

Psychology-Informed Reinforcement Learning for Situated Virtual Coaching in Smoking Cessation

Albers, N.

DOI

[10.4233/uuid:5e1d1500-1c45-4869-b995-6ed7db8f8b3b](https://doi.org/10.4233/uuid:5e1d1500-1c45-4869-b995-6ed7db8f8b3b)

Publication date

2025

Document Version

Final published version

Citation (APA)

Albers, N. (2025). *Psychology-Informed Reinforcement Learning for Situated Virtual Coaching in Smoking Cessation*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:5e1d1500-1c45-4869-b995-6ed7db8f8b3b>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Psychology-Informed Reinforcement Learning

for Situated Virtual Coaching
in Smoking Cessation

NELE ALBERS



**PSYCHOLOGY-INFORMED REINFORCEMENT
LEARNING FOR SITUATED VIRTUAL COACHING
IN SMOKING CESSATION**

Nele ALBERS

PSYCHOLOGY-INFORMED REINFORCEMENT LEARNING FOR SITUATED VIRTUAL COACHING IN SMOKING CESSATION

Dissertation

for the purpose of attaining the degree of doctor
at Delft University of Technology,
by the authority of Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
chair of the Board for Doctorates,
to be defended publicly on
Thursday, 27 February 2025, at 12:30 o'clock

by

Nele ALBERS

Master of Science in Computer Science,
Delft University of Technology, the Netherlands,
born in Lüdinghausen, Germany.

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus	chairperson
Dr. ir. W.-P. Brinkman	Delft University of Technology, <i>promotor</i>
Prof. dr. M.A. Neerincx	Delft University of Technology, <i>promotor</i>

Independent members:

Prof. dr. T. Bosse	Radboud University
Prof. dr. C.M. Jonker	Delft University of Technology
Prof. dr. L.G.M.T. Keijsers	Erasmus University Rotterdam
Prof. dr. S.A. Murphy	Harvard University, United States of America
Prof. dr. M.T.J. Spaan	Delft University of Technology

The thesis has been funded by the multidisciplinary research project Perfect Fit, which is supported by several funders organized by the Netherlands Organization for Scientific Research (NWO), program Commit2Data - Big Data & Health (project number 628.011.211).

SIKS Dissertation Series No. 2025-05.

The research reported in this thesis has been carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.



Keywords: Reinforcement Learning, Smoking, Behavior Change, eHealth, Virtual Coach, Conversational Agent, Persuasion, Ethics, Psychology-Informed Algorithm

Printed by: Proefschriftspecialist | <https://www.proefschriftspecialist.nl/>

Cover by: Jay Patel

Style: TU Delft House Style, with modifications by Moritz Beller

Copyright © 2025 by Nele Albers

ISBN 978-94-6518-010-6

An electronic version of this dissertation is available at <http://repository.tudelft.nl/>.

CONTENTS

Summary	vii
Samenvatting	xiii
1 Introduction	1
2 Users' needs for a digital smoking cessation application	13
3 <i>How</i> – Reinforcement learning algorithm for persuading to quit smoking	43
4 <i>How</i> – Using states and user characteristics to predict behavior	73
5 <i>What</i> – Reinforcement learning for proposing smoking cessation activities	91
6 <i>Who</i> – The impact of human support in a chatbot-based smoking cessation intervention	111
7 Conclusion	137
A Appendix – Users' needs for a digital smoking cessation application	153
B Appendix – Reinforcement learning algorithm for persuading to quit smoking	173
C Appendix – The impact of human support in a chatbot-based smoking cessation intervention	191
Bibliography	221
SIKS Dissertations	261
Acknowledgments	271
Curriculum Vitae	275
List of Publications	277
List of Relevant Datasets and Analysis Code	281
List of Relevant Preregistrations	283
List of Relevant Virtual Coach Implementations	285

SUMMARY

Despite recent decreases in tobacco smoking rates in high-income countries such as the Netherlands, smoking still is the second largest risk factor for early death and poor health worldwide, killing more than 8 million people every year. Together with other unhealthy behaviors such as physical inactivity, smoking places a substantial strain on the financial, staffing, and societal sustainability of healthcare. eHealth applications, especially those that include support from a virtual coach, have the potential to help people change behaviors both effectively and at reduced financial and staffing costs. While these eHealth applications often suffer from dropout and a lack of engagement, adapting the support they offer to the variety of environmental, cognitive, and social factors influencing human behavior is a promising way to increase satisfaction and engagement. It is especially promising to adapt the support to the state a person is in, which is described by components such as risk perception, self-efficacy, and skills. A person's state is an interesting basis for adaptive support because of initial evidence that it influences the effectiveness of different types of support, such as differently framed messages, for changing behaviors such as unhealthy eating and smoking. However, types of support can, in turn, also affect a person's future state. For instance, the impact of messages on the future state component of self-efficacy can vary depending on whether they are framed positively or negatively. From the many computational methods used to adapt eHealth applications, one method that allows us to account for people's current and future states is Reinforcement Learning (RL). Recent studies show the potential and even medical efficacy of using RL with a consideration of people's current states in adapting behavior change support in contexts such as stress and pain management. Given initial positive results also for accounting for both current and future states (i.e., *full RL*), we investigated in this thesis further how full RL can be used to make smoking cessation support more effective.

To create effective smoking cessation support, we first wanted to establish what kind of support smokers need. Then we could design support that meets these needs. To this end, we conducted a first study in which 671 daily smokers interacted with a virtual coach in up to five sessions spread over about two weeks (Figure 1). Based on their self-reported experiences with interactions during the study and views on envisioned interaction scenarios thereafter, we obtained 14 main needs. These needs pertain to people's behaviors related to the eHealth application, personal characteristics, the involvement of other parties such as general practitioners or virtual coaches, and the environmental context in which behaviors occur. The most commonly mentioned need was that behaviors are perceived as useful because they provide motivation, encouragement, help, advice, or learning opportunities.

To address smokers' need for motivation and encouragement, our first application of RL was to adapt *how* people are motivated to do activities for quitting smoking. We hypothesized that a personalized RL algorithm would be effective. This personalized RL algorithm considered people's current states, their future states, and their similarity based

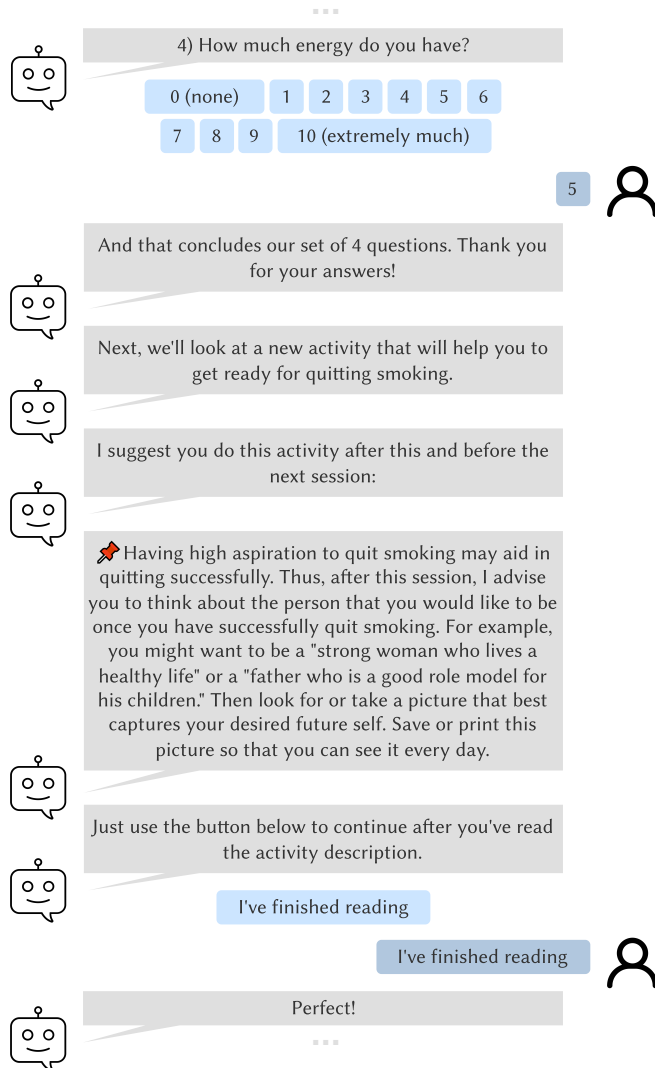


Figure 1: Example excerpt from a dialog with a virtual coach. After asking questions to determine the user's state, the virtual coach proposes a new preparatory activity.

on personal characteristics such as personality and stage of change. The similarity was used to give a larger weight to interaction samples from more similar people. To study the effect of using more algorithm components (i.e., current states, future states, similarity) on the perceived motivational impact of the virtual coach interactions and effort spent on activities, we split participants of our first study into four groups. Each group was motivated based on a different algorithm complexity level, with higher levels using more algorithm components. Based on 2,366 interaction samples, we obtained some support that perceived motivational impact and effort are positively affected by including more algorithm components, with the latter becoming apparent only after some time. Looking at the algorithm complexity levels separately, it seems that the level considering people's current and future states but not their similarity is most effective. Support for this is even stronger when focusing on individuals who were most involved in their activities. To further assess the generalizability of models for selecting motivational strategies, we compared the most effective strategies identified for smoking cessation activities with those for increasing physical activity. This was possible because half of the activities in the study focused on increasing physical activity as a strategy to facilitate quitting smoking. Our findings suggest that even though there is some agreement between the most effective strategies for the two activity types, it is relatively low.

While we did not find that considering personal characteristics *in addition* to states is helpful, it could still be that considering personal characteristics *alone* is more effective than considering states. This would be beneficial because personal characteristics data only needs to be collected once at the start of the intervention rather than before each interaction, as is the case with states. However, performing cross-validation based on the 2,366 collected interaction samples, we found that we can better predict the effort people spend on preparatory activities after being motivated with different strategies based on states than personal characteristics measurable before an intervention. Factors conceptually closer to an advocated behavior can thus better predict it: states based on components such as motivation are closer than personal characteristics such as personality.

While motivation and encouragement are among the needs we identified in our first study, the most prevalent need was that smokers perceive advocated behaviors as useful. And these usefulness perceptions do not necessarily match the ones experts have. For example, smokers do not always agree with experts that physical activity can facilitate quitting smoking. Thus, if we simply propose activities that experts find useful, smokers may not do them and hence never build the competencies required for quitting. On the other hand, if we just propose activities that smokers find useful, we might not cover all the competencies. Our question for our second study was thus how we can build an RL model to choose activities in a way that trades off these two perspectives. This means that having looked at *how* people are motivated in the previous study, we now focused on building an RL model to adapt *what* people are asked to do. To create this model, we first established competency-building activities as those that experts typically propose to prepare for quitting smoking. Afterward, we used two repertory grid studies, one with experts and one with smokers, to determine which competencies experts and smokers think are built by these activities. We found nine smoker-identified competencies, such as *self-efficacy*, *awareness of negative outcomes*, and *knowledge of how to maintain/achieve well-being*. From experts, we obtained six competencies, including *clear future identity*,

motivation, and *insights into personal weaknesses*. Next, we created nine belief-changing activities to convince smokers of the usefulness of the nine smoker-identified competencies. The reason for including these activities was that we might have to change smokers' usefulness beliefs for them to do activities related to a competency. And finally, we designed and trained the RL model using data from a new crowdsourcing study, this time with 542 daily smokers interacting with a virtual coach in up to five sessions. Using simulations to evaluate the model, we found that proposing activities based on the model can allow smokers to build 91% of expert-identified competencies within five interactions. All model components contribute to this. However, smokers' current state based on their usefulness beliefs, energy, and degrees of having built the expert-identified competencies contributed the most. The contribution of the transition to next states, on the other hand, is small. One reason for this might be the positive but small effects of belief-changing activities. So if not the *future*, at least the *current* views of smokers and experts are helpful to consider in the RL model.

Following our investigation into how RL can be used to adapt the *how-* and *what-* dimensions of support, we next studied adapting *who* provides the support. This was motivated by having observed in our first study that smokers were concerned about the accountability and companionableness of other parties involved in an intervention, such as virtual coaches and general practitioners. As also expressed by participants of our first study, both accountability and companionableness can be influenced by whether another party is human or AI. Previous work indicates that including support from a human coach can make people feel more accountable and satisfied. Therefore, we examined how effective adding human support to a virtual coach-based smoking cessation intervention is. Adding extensive human support would, however, undermine the goal of eHealth applications, which aim to minimize reliance on scarce and expensive healthcare professionals. Thus, we focused on relatively low-cost human support, in our case feedback messages written by Master's students in Psychology. Using data from a third crowdsourcing study with 679 daily smokers and vapers, we analyzed psychological, economic, and ethical factors that play a role when allocating this support. Our simulations show that providing human feedback more often generally leads to a higher effort spent on activities over time. However, human feedback messages sent to people in about half of the states can be removed without a large drop in effort. When providing human feedback is expensive, the largest long-term increase in effort can be obtained by giving feedback to people who perceive preparing to quit as not that important and at the same time have high self-efficacy for preparing to quit. Notably, while it may seem intuitive to allocate limited human feedback to those who would benefit most, it is, in fact, only one of several ethical principles for allocating scarce medical resources. Alternatives include allocating feedback to those who want it most or to individuals who have the lowest likelihood of successfully quitting without such support. The choice of allocation principle reflects a moral stance on who should have the best chance at successful smoking cessation and improved health outcomes. Our analysis of the "standard" benefit-maximizing model reveals that it prioritizes people who want feedback and are already doing well. However, other allocation principles can be included in the model to favor other smoker subgroups. Yet, these principles, we found, are often interdependent. For example, since we observed that people who are not doing well tend not to want human feedback and benefit less from it, giving more feedback to them means

doing worse in respecting people's autonomy and maximizing total benefit. Since different smoker subgroups benefit depending on the chosen allocation principles, these findings show that moral decisions are unavoidable.

Together, our findings suggest that considering current and future states increases the effort smokers spend on smoking cessation activities and helps them build quitting-related competencies over time. Since our models were based on behavior change theories, this highlights the potential of using psychology-informed RL to develop effective long-term support for smoking cessation.

SAMENVATTING

Ondanks recente dalingen in het tabaksgebruik in hoge-inkomenslanden zoals Nederland, is roken wereldwijd nog steeds de op één na grootste risicofactor voor vroegtijdige sterfte en slechte gezondheid. Jaarlijks overlijden hier meer dan 8 miljoen mensen aan. Samen met ander ongezond gedrag, zoals lichamelijke inactiviteit, verhoogt roken de kosten van de gezondheidszorg, de druk op het personeel en de uitdagingen van de maatschappij voor het leveren van inclusieve en duurzame zorg. eHealth-toepassingen, met name die met een virtuele coach, hebben het potentieel om mensen effectief te ondersteunen bij gedragsverandering, terwijl ze tevens de financiële en personele kosten verlagen. Hoewel huidige eHealth-toepassingen vaak nog weinig gebruikt worden, zal een afstemming op de verschillende omgevings-, cognitieve en sociale factoren van gedragsbeïnvloeding dit gebruik naar verwachting substantieel vergroten. Het is vooral veelbelovend om de ondersteuning aan te passen aan de toestand waarin iemand zich bevindt, die wordt beschreven door componenten zoals risicoperceptie, zelfredzaamheid en vaardigheden. De toestand van een persoon is een interessante basis voor adaptieve ondersteuning omdat er aanwijzingen zijn dat deze de effectiviteit beïnvloedt van verschillende ondersteuningsvormen, zoals verschillend geformuleerde berichten, voor het veranderen van gedrag zoals ongezond eten en roken. Maar ondersteuningsvormen kunnen op hun beurt ook de toekomstige toestand van een persoon beïnvloeden. De impact van berichten op de toekomstige toestandcomponent van zelfredzaamheid kan bijvoorbeeld variëren afhankelijk van een positieve of negatieve formulering. Van de vele computationele methoden die worden gebruikt om eHealth-toepassingen aan te passen, is er één methode waarmee we rekening kunnen houden met de huidige en toekomstige toestand van mensen: Reinforcement Learning (RL). Recente studies tonen het potentieel en zelfs de medische effectiviteit aan van het gebruik van RL met inachtneming van de huidige toestand van mensen bij het aanpassen van ondersteuning voor gedragsverandering in contexten zoals stress- en pijnmanagement. Gezien de eerste positieve resultaten van het gebruik van zowel huidige als toekomstige toestanden (d.w.z. *volledige* RL), hebben we in dit proefschrift verder onderzocht hoe volledige RL gebruikt kan worden om ondersteuning bij stoppen met roken effectiever te maken.

Om effectieve ondersteuning bij het stoppen met roken te creëren, wilden we eerst vaststellen wat voor soort ondersteuning rokers nodig hebben. Daarna konden we ondersteuning ontwerpen die aan deze behoeften voldeed. Hiertoe voerden we een eerste onderzoek uit waarin 671 dagelijkse rokers interactie hadden met een virtuele coach in maximaal vijf sessies verspreid over ongeveer twee weken (Figuur 2). Op basis van hun zelfgerapporteerde ervaringen met deze coach en hun meningen over de beoogde interactiescenario's daarna, verkregen we 14 basisbehoeften. Deze behoeften betreffen het gedrag van mensen met betrekking tot de eHealth-toepassing, persoonlijke kenmerken, de betrokkenheid van andere partijen zoals huisartsen of virtuele coaches, en de omgevingscontext

waarin het gedrag plaatsvindt. De meest genoemde behoefte was dat gedrag als nuttig wordt ervaren omdat het motivatie, aanmoediging, hulp, advies of leermogelijkheden biedt.

Om tegemoet te komen aan de behoefte van rokers aan motivatie en aanmoediging, paste de eerste RL-toepassing zich aan *hoe* mensen gemotiveerd zijn om activiteiten te ondernemen om te stoppen met roken. We stelden dat een gepersonaliseerd RL-algoritme effectief zou zijn. Dit gepersonaliseerde RL-algoritme hield rekening met de huidige toestand van mensen, hun toekomstige toestand en hun gelijkenis op basis van persoonlijke kenmerken zoals persoonlijkheid en fase van verandering. De gelijkenis werd gebruikt om een groter gewicht toe te kennen aan interactie samples van mensen die meer op elkaar leken. We verdeelden de deelnemers aan ons eerste onderzoek in vier groepen om het effect te bestuderen van het gebruik van meer algoritme-componenten (d.w.z. huidige toestanden, toekomstige toestanden, gelijkenis) op de waargenomen motiverende impact van de virtuele coachinteracties en de moeite die werd besteed aan activiteiten. Elke groep werd gemotiveerd op basis van een ander algoritme-complexiteitsniveau, waarbij hogere niveaus meer algoritme-componenten gebruikten. Op basis van 2.366 interactie samples vonden we enig bewijs dat de waargenomen motiverende impact en moeite positief worden beïnvloed door het gebruik van meer algoritme-componenten. Als we de algoritme-complexiteitsniveaus afzonderlijk bekijken, lijkt het erop dat het niveau dat rekening houdt met de huidige en toekomstige toestand van mensen, maar niet met hun gelijkenis, het meest effectief is. Dit wordt nog sterker ondersteund wanneer we ons richten op individuen die het meest betrokken waren bij hun activiteiten. Om de generaliseerbaarheid van modellen voor het selecteren van motivatiestrategieën verder te beoordelen, vergeleken we de meest effectieve strategieën voor stoppen met roken activiteiten met die voor het verhogen van lichamelijke activiteit. Dit was mogelijk omdat de helft van de activiteiten in het onderzoek zich richtte op het verhogen van lichamelijke activiteit als strategie om het stoppen met roken te vergemakkelijken. Onze bevindingen suggereren dat, hoewel er enige overeenstemming is tussen de meest effectieve strategieën voor de twee typen activiteiten, deze relatief laag is.

Hoewel we niet vonden dat het in beschouwing nemen van persoonlijke kenmerken *als aanvulling* op toestanden nuttig was, zou het nog steeds zo kunnen zijn dat het in beschouwing nemen van persoonlijke kenmerken *alleen* effectiever is dan het in beschouwing nemen van toestanden. Dit zou gunstig zijn omdat gegevens over persoonlijke kenmerken maar één keer hoeven te worden verzameld aan het begin van de interventie in plaats van voor elke interactie, zoals het geval is bij toestanden. Door cross-validation uit te voeren op basis van de 2.366 verzamelde interactie samples, vonden we echter dat we de moeite die mensen besteden aan voorbereidende activiteiten nadat ze gemotiveerd zijn met verschillende strategieën beter kunnen voorspellen op basis van toestanden dan op basis van persoonlijke kenmerken die gemeten kunnen worden vóór een interventie. Factoren die conceptueel dichter bij het gewenste gedrag staan, kunnen dit gedrag dus beter voorspellen.

Hoewel motivatie en aanmoediging tot de behoeften behoorden die we in ons eerste onderzoek identificeerden, was de meest voorkomende behoefte dat rokers aanbevolen gedrag als nuttig ervaren. En deze percepties van nut komen niet noodzakelijk overeen met die van experts. Rokers zijn het bijvoorbeeld niet altijd eens met experts dat lichamelijke activiteit het stoppen met roken kan vergemakkelijken. Dus, als we simpelweg activiteiten



Figuur 2: Voorbeeldfragment uit een dialoog met een virtuele coach. Nadat er vragen zijn gesteld om de toestand van de gebruiker te bepalen, stelt de virtuele coach een nieuwe voorbereidende activiteit voor.

voorstellen die experts nuttig vinden, zullen rokers ze misschien niet doen en dus nooit de competenties opbouwen die nodig zijn om te stoppen. Aan de andere kant, als we alleen activiteiten voorstellen die rokers nuttig vinden, dekken we misschien niet alle competenties. Onze vraag voor onze tweede studie was dus hoe we een RL-model kunnen bouwen om activiteiten te kiezen op een manier die deze twee perspectieven combineert. Dit betekent dat we, nadat we in de vorige studie hadden gekeken naar *hoe* mensen gemotiveerd zijn, ons nu richtten op het bouwen van een RL-model om aan te passen *wat* mensen gevraagd wordt te doen. Om dit model te creëren, hebben we eerst competentie-opbouwende activiteiten vastgesteld, zoals de activiteiten die experts gewoonlijk voorstellen om mensen voor te bereiden op het stoppen met roken. Daarna gebruikten we twee repertory grid onderzoeken, één met experts en één met rokers, om te bepalen welke competenties experts en rokers denken dat deze activiteiten opbouwen. We vonden negen door rokers geïdentificeerde competenties, zoals *zelfredzaamheid*, *bewustzijn van negatieve uitkomsten* en *kennis over hoe welzijn te behouden/bereiken*. Van experts vonden we zes competenties, waaronder *een duidelijke toekomstige identiteit*, *motivatie* en *inzicht in persoonlijke zwakheden*. Vervolgens creëerden we negen overtuigingsveranderende activiteiten om rokers te overtuigen van het nut van de negen door rokers geïdentificeerde competenties. De reden voor het opnemen van deze activiteiten was dat we misschien de overtuigingen over het nut van rokers moesten veranderen om hen activiteiten te laten doen die gerelateerd waren aan een competentie. Tot slot hebben we het RL-model ontworpen en getraind met behulp van gegevens uit een nieuw crowdsourcing-onderzoek, dit keer met 542 dagelijkse rokers die interactie hadden met een virtuele coach in maximaal vijf sessies. Met behulp van simulaties om het model te evalueren, ontdekten we dat het voorstellen van activiteiten op basis van het model rokers in staat kan stellen om 91% van de door experts geïdentificeerde competenties op te bouwen binnen vijf interacties. Alle modelcomponenten dragen hieraan bij. De huidige staat van rokers op basis van hun overtuigingen over het nut, hun energie en de mate waarin ze de door experts geïdentificeerde competenties hadden opgebouwd, droegen echter het meest bij. De bijdrage van de overgang naar volgende toestanden is daarentegen klein. Eén reden hiervoor zou kunnen liggen in de positieve maar kleine effecten van activiteiten die overtuigingen veranderen. Dus als het niet de *toekomst* is, dan zijn in ieder geval de *huidige* opvattingen van rokers en experts nuttig om mee te nemen in het RL-model.

Na ons onderzoek naar hoe RL kan worden gebruikt om de *hoe-* en *wat-*dimensie van ondersteuning aan te passen, onderzochten we *wie* de ondersteuning biedt. De reden hiervoor was dat we in ons eerste onderzoek hadden gezien dat rokers zich zorgen maakten over de aanspreekbaarheid en vriendelijkheid van andere partijen die betrokken waren bij een interventie, zoals virtuele coaches en huisartsen. Zoals deelnemers aan ons eerste onderzoek ook aangaven, kunnen zowel verantwoordelijkheid als kameraadschap verschillend beïnvloed worden bij een menselijke versus AI coach. Eerder werk geeft aan dat mensen zich meer verantwoordelijk en tevreden voelen als ze ondersteuning krijgen van een menselijke coach. Daarom hebben we onderzocht hoe effectief het is om menselijke ondersteuning toe te voegen aan een stoppen met roken interventie die gebaseerd is op een virtuele coach. Het toevoegen van uitgebreide menselijke ondersteuning zou echter het doel van eHealth toepassingen ondermijnen de afhankelijkheid van schaarse en dure zorgprofessionals te minimaliseren. Daarom hebben we ons gericht op relatief goedkope

menselijke ondersteuning, in ons geval feedbackberichten geschreven door masterstudenten Psychologie. Met behulp van gegevens uit een derde crowdsourcingstudie met 679 dagelijkse rokers en vapers, analyseerden we psychologische, economische en ethische factoren die een rol spelen bij het toekennen van deze steun. Onze simulaties laten zien dat het vaker geven van menselijke feedback over het algemeen leidt tot een grotere inspanning voor activiteiten na verloop van tijd. Echter, menselijke feedbackberichten aan mensen in ongeveer de helft van de toestanden kunnen worden verwijderd zonder een grote daling in inspanning. Wanneer het geven van menselijke feedback duur is, kan de grootste toename in inspanning op de lange termijn worden verkregen door feedback te geven aan mensen die het voorbereiden om te stoppen als niet zo belangrijk ervaren en tegelijkertijd een hoge zelfredzaamheid hebben voor het voorbereiden om te stoppen. Hoewel het intuïtief lijkt om beperkte menselijke feedback toe te wijzen aan degenen die er het meeste baat bij hebben, is dit in feite slechts één van de ethische principes voor het toewijzen van schaarse medische middelen. Alternatieven zijn onder andere het toewijzen van feedback aan degenen die dat het meest willen of aan individuen die de minste kans hebben om succesvol te stoppen zonder dergelijke ondersteuning. De keuze van het toewijzingsprincipe weerspiegelt een moreel standpunt over wie de beste kans zou moeten hebben op succesvol stoppen met roken en verbeterde gezondheidsresultaten. Onze analyse van het “standaard” voordeel maximaliserend model laat zien dat het prioriteit geeft aan mensen die feedback willen en het al goed doen. Er kunnen echter andere toewijzingsprincipes in het model worden opgenomen om andere subgroepen rokers te bevoordelen. We ontdekten echter dat deze principes vaak onderling afhankelijk zijn. Bijvoorbeeld, omdat we hebben gezien dat mensen die het niet goed doen meestal geen menselijke feedback willen en er minder van profiteren, betekent het geven van meer feedback aan hen dat het slechter gaat met het respecteren van de autonomie van mensen en het maximaliseren van het totale voordeel. Aangezien het profijt voor verschillende subgroepen rokers afhangt van de gekozen toewijzingsprincipes, tonen deze bevindingen aan dat morele beslissingen onvermijdelijk zijn.

Al met al suggereren onze bevindingen dat rokers meer inspanning leveren bij activiteiten voor het stoppen van roken die hen op den duur helpen om stop-gerelateerde competenties op te bouwen, wanneer er rekening gehouden wordt met de huidige en toekomstige toestand. Aangezien onze modellen gebaseerd waren op gedragsveranderings-theorieën, benadrukt dit het potentieel van het gebruik van psychologie-geïnformeerde RL om effectieve lange termijn ondersteuning voor stoppen met roken te ontwikkelen.

1

INTRODUCTION

1

MOTIVATION

Projections show that by 2060, Dutch healthcare spending, which is already the second highest per capita in the European Union [84], would need to rise from the current 10% to 18% of the Gross Domestic Product (GDP) to take care of the aging population [1]. Moreover, the number of people working in healthcare would need to increase from one in seven to one in three - a demand unlikely to be met even if the employment rate were to grow by such a margin [105]. A large part of these costs can be attributed to unhealthy behavior. Already in 2000, 40% of premature deaths in the United States were caused by unhealthy behavior [253]. And looking at the overall disease burden, a 2018 report for the Netherlands states that 18.5% of this burden stems from unhealthy behavior, with 9.4% linked specifically to smoking [324]. This issue is not confined to the US or the Netherlands; it is a global concern. Smoking, for instance, is the largest behavioral risk factor for early death and poor health worldwide, and the second largest risk factor overall after high systolic blood pressure [171]. Additionally, unhealthy behavior strains not only the financial and staffing sustainability of healthcare but also its societal sustainability. There is an increasing reluctance to cover costs for others' lifestyle-related conditions caused by smoking and obesity [105]. Therefore, we need approaches to effectively change unhealthy behaviors at low cost, including a low need for human healthcare staff.

One promising approach are eHealth applications, which provide elements of healthcare over the Internet or connected technologies such as apps and text messaging [290]. This includes applications for communication between healthcare professionals and clients, lifestyle monitoring, and medication dispensing [403]. Being able to provide support anywhere anytime, eHealth applications have been shown to be effective in contexts such as smoking cessation [126, 215] as well as physical activity, sedentary behavior, and unhealthy eating and sleep [343]. To increase engagement, discuss relevant information, and form a connection with people [158, 255], these applications often include support from a conversational agent taking the role of a virtual coach. Examples include the Quit Coach [33], the StopCoach Suzanne [237], and Dejal@bot [278]. These virtual coaches can, for example, offer information on the health effects of smoking, tips for quitting, and motivational messages [33], and provide practical and motivational support [237]. Since human behavior depends on a variety of environmental, social, and cognitive factors [231, 240], adapting eHealth applications [180] and in particular interactions with these virtual coaches [199] to (groups of) individuals is a promising approach to improving satisfaction, engagement, and health outcomes.

One way to adapt the support provided by eHealth applications is based on enduring user characteristics. For example, Zalake et al. [398] saw that the relative effectiveness of persuasive strategies by Cialdini [93] used to promote coping skills for mental health depended on personality, Oyibo et al. [289] found the relative effectiveness of the same strategies to depend on cultural background, and Steward et al. [351] observed that the effectiveness of positively versus negatively framed smoking cessation messages depended on people's need for cognition. Yet, the effects of such adaptations on behavior tend to be small (e.g., [106, 183]). Since many of the environmental, social, and cognitive factors influencing human behavior [231, 240] are not static, a more effective alternative may be to consider the current state a person is in. This state refers to an individual's condition or status at a specific moment in time, characterized by relative stability in its components

[37]. Examples of state components used for adaptation include risk perception [371], awareness, motivation, and commitment [198], and pain intensity, past physical activity, sleep, skills, and intervention phase [300]. However, not only can a person's state influence the effectiveness of different forms of support, but the support can, in turn, also affect a person's future state. For example, self-efficacy can influence how effective differently framed messages are [54], and these messages can again differ in their effects on self-efficacy [351]. We thus have a cyclic chain of effects: *self-efficacy* → *effectiveness of differently framed messages* → *self-efficacy* ... Such recursive effects are also expressed in behavior change models such as the COM-B model in which capability (C), opportunity (O), and motivation (M) influence and are influenced by behavior (B) [387]. This suggests that if we strive to improve the long-term effectiveness of behavior change support [158, 180], we should account for both current and future user states.

Various computational methods have been used to adapt eHealth interventions to users, including classification-based methods and algorithms comparing user behavior to norms and guidelines [172]. One method that allows us to adapt to current and future user states is the machine learning paradigm of Reinforcement Learning (RL) [357]. In RL, an intelligent agent learns what to do to maximize a numerical reward signal. That is, by interacting with a dynamic environment, the agent over time learns which actions to take in a certain situation to maximize a reward. The dynamic environment in our case is the person we want to support in changing their behavior, actions are different forms of support, and the reward is a measure of support effectiveness, either actual behavioral change or a more proximal measure such as motivation or engagement with the intervention. RL not only differs from other machine learning paradigms in being able to capture the effects that actions taken in certain states have on future states, but also in not requiring a labeled training set. This makes RL a promising tool for adaptive behavior change support. In addition, if the agent continues to try different actions during a live behavior change intervention to see which actions yield the highest reward, a process called exploration, the agent can adapt to changes in action effects.

A recent review by Weimann and Gißke [384] highlighted the potential and even medical efficacy of RL for adaptive health behavior change support. Paredes et al. [292] showed that an RL approach that accounts for current user states when selecting stress management interventions enabled a larger stress reduction after four weeks than proposing random interventions. And Piette et al. [300] saw that choosing therapist feedback types based on an RL approach accounting for current user states significantly reduced pain outcomes while using less therapist time compared to delivering the most intensive type of therapist feedback. However, despite the increasing popularity of RL, only a relatively small number of the studies reviewed by Weimann and Gißke [384] account for both current and future states (i.e., *full RL*¹). Examples include timing running notifications [381], suggesting step goals [140], and selecting messages for diabetes prevention [195]. Since initial results comparing full RL to random baselines (e.g., [140, 381]), simpler RL models (e.g., [140]) and not precisely defined static baselines (e.g., [195]) are promising, we investigate in this thesis how full RL can be used to adapt smoking cessation support. This leads to the following

¹This is also frequently referred to as *non-myopic RL*, distinguishing it from *myopic RL* which only tries to maximize immediate rewards [357]. However, in this thesis, we adopt the term *full RL*, consistent with the review by Weimann and Gißke [384].

overarching research question:

How can reinforcement learning be used to make the support in virtual coach-based smoking cessation interventions more effective?

In answering this research question, we will first investigate user needs for such support and then look at three dimensions of adaptive support: 1) *how* people are persuaded (i.e., with different persuasive strategies), 2) *what* people are asked to do (i.e., different activities for quitting smoking), and 3) *who* they are supported by (i.e., only a *virtual* coach or also a *human* coach). Figure 1.1 provides a conceptual overview of the support and corresponding research questions and hypothesis.

RESEARCH QUESTIONS AND HYPOTHESES

Before thinking about adaptivity, we must establish what kind of support it is that we want to make adaptive (green elements in Figure 1.1). While it is difficult to systematically assess the effect of different features of virtual coach-based smoking cessation interventions on user engagement due to their variety [158], previous studies have identified several combinations of factors that play a role. Alphonse et al. [33], for example, saw in an interview study that the virtual coach's interaction style and format, participants' support needs, and the anthropomorphism of and accountability to the virtual coach influenced users' experiences with the Quit Coach. Moreover, Meijer et al. [237] noted that age and support from one's social environment may affect the usability and usefulness of the StopCoach. These examples illustrate the importance of taking a holistic approach that considers factors related to the technology, the user, and the user's environment [373]. In light of the wide variety of possibly relevant factors, a more thorough understanding of user needs is welcome. This is especially the case in light of the often poor adherence, lack of engagement, and even abandonment of eHealth applications [147, 191], which suggest that current eHealth applications insufficiently meet user needs. These considerations led to our first research sub-question:

RQ1: What are users' needs for a virtual coach-based smoking cessation intervention?

Having examined human factors for designing the support, this thesis subsequently looks at how RL can be used to make the support adaptive. The first dimension of support we examine is *how* people are persuaded (blue elements in Figure 1.1). Specifically, when persuading users to do activities such as making a plan, reflecting on the previous week, or going for a run, eHealth applications commonly make use of different persuasive strategies. For example, an application may say that experts recommend going for a run to reduce one's vulnerability to cardiovascular disease, thus implementing the principle of *authority* by Cialdini [93]. Previous work has stated the importance of considering user states [54, 60, 137] as well as shown that persuasive strategies, in turn, can influence user states [80, 286, 287, 351], suggesting that there is a benefit to considering current and future states when choosing persuasive strategies. In addition, more enduring user characteristics such as the need for cognition [351] can affect how effective different persuasive strategies

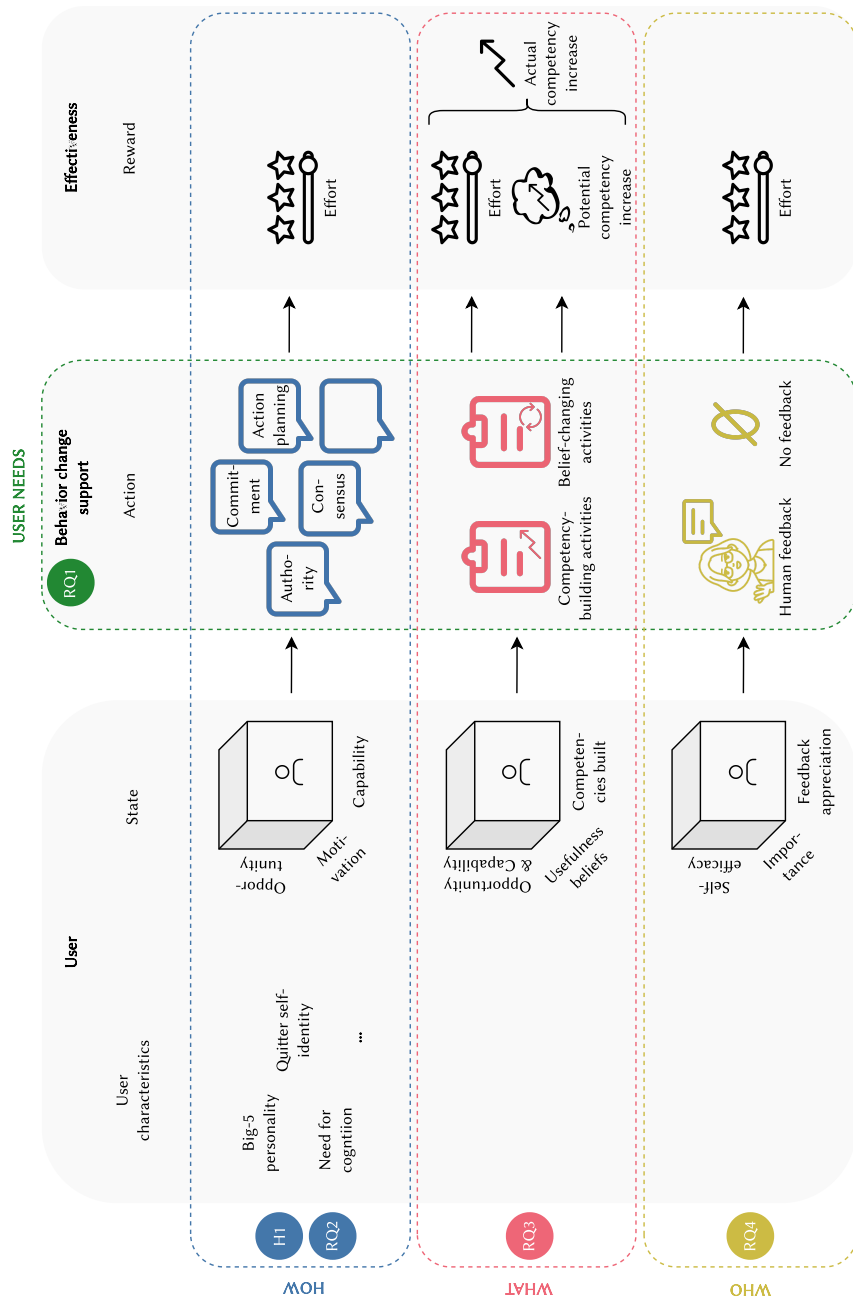


Figure 1.1: Conceptual overview of the behavior change support provided in a virtual coach-based smoking cessation intervention and corresponding research questions and hypothesis this thesis sets out to answer.

chosen in a certain state are. Hence, we expect that individuals with more similar user characteristics are more likely to respond similarly to persuasive strategies. Taking the consideration of states, future states, and user characteristic-based similarity together, this thesis hypothesizes the following:

H1: Subsequently incorporating 1) states, 2) the consideration of future states, and 3) the weighting of samples based on the similarity of people into an algorithm that selects the best persuasive strategy is more effective than not incorporating the respective element.

If the effect of incorporating user characteristic-based similarity *alongside* current and future states is limited, accounting for user characteristics *alone* could still be more effective than considering user states. In contrast to states, fixed user characteristics need to be inferred only once at the start of the intervention rather than before each persuasive attempt. This is potentially advantageous since asking people many questions may be more cognitively demanding and thus seen as more effortful and liked less [302], which is negatively associated with technology use [377]. Implicit data collection methods such as sensors could be used instead, however, these do not yet succeed at collecting high-quality data unobtrusively [395]. Given that user characteristics are thus easier to infer than states, our second sub-question compares how well we can predict people's behavior after persuading them based on different persuasive strategies using either user characteristics or states to choose strategies. If we can predict this behavior better, we can make a more informed choice of persuasive strategies that lead to favorable behavior. Our second sub-question thus is

RQ2: How does predicting behavior based on user characteristics compare to doing so based on states?

Besides adapting *how* people are persuaded, applications can also adapt *what* people are asked to do (red elements in Figure 1.1). The motivation for doing this is that while health experts can create a set of activities that help people build the competencies (e.g., knowledge, skills, mindsets, thought patterns) [118] needed for successful behavior change, users may not engage with them if they do not find them useful [340, 377]. So instead of simply proposing activities that experts find useful, we might also want to consider which competencies users find useful. Moreover, if we cannot build all competencies because users find all related activities unhelpful, we can try to change users' usefulness beliefs. Again there is a recursive relationship in which activity selection may depend on an individual's usefulness beliefs and degree of having built different competencies, both of which are themselves affected by the activity chosen. We can thus formulate this as an RL problem where the goal is to maximize an individual's competency increase over time. The state thereby captures the views of users (i.e., the usefulness beliefs) and experts (i.e., the degree to which users have built expert-identified competencies). Our third sub-question is how we can create an RL model for this problem:

RQ3: How can we create an RL model for building human competencies that combines the views of experts and users?

The third dimension of support we can adapt is *by whom* people are supported (yellow elements in Figure 1.1). More precisely, we can adapt whether we combine an eHealth application with support from a *human coach*, which may increase people's engagement [115], accountability [210, 250] and satisfaction [327]. Adding large amounts of professional human support would undermine the goal of eHealth applications to reduce reliance on scarce and expensive healthcare staff. However, relatively low amounts of human support (e.g., at most three daily text messages in a text-messaging intervention for people with schizophrenia or 36 minutes of therapist support in a ten-week web-based treatment of tinnitus) can already lead to good results [337], with professional qualifications of the human coaches not necessarily playing a large role [47]. It is relatively low-cost human support that we are interested in adding to a virtual coach-based application. Since the effectiveness of human support may depend on a person's state (e.g., intrinsic motivation [252]), which in turn may be influenced by human support [252], the decision of when to allocate human support also lends itself to an RL formulation. While initial results on using RL with a consideration of current states to allocate human support in eHealth applications [300, 301] are promising, also incorporating future states to capture the effect human support can have on user states deserves further exploration. This can especially help in understanding the long-term effects of human support. Using RL with a consideration of current and future states to allocate human support, our fourth research sub-question thus is

RQ4: How effective is adding human support to a virtual coach-based behavior change intervention in the long term?

RESEARCH APPROACH AND THESIS STRUCTURE

Overall, this thesis takes an empirical approach using data from real smokers to assess the effects of algorithms. However, doing this in a complete behavior change intervention is complicated due to the complexity of the intervention [45]. The many different, and possibly over time changing, components make it difficult to determine what is driving an observed effect [45]. Can the effect be attributed to gamification, goal-setting, connection to other users, interface design, or really the algorithm for adapting persuasive strategies? And what if the other components also change over time - goals are adapted based on user feedback, and the other users in the intervention that people can connect to vary? Therefore, to better assess the effects of our algorithms, all our analyses are based on minimal interventions excluding any components not needed for our algorithms. For example, we did not include goal-setting even though this can positively affect behavior change [121]. To quickly and meaningfully evaluate the effects of our algorithms [45], we thereby used proximal outcome measures (e.g., motivation or engagement with the intervention) rather than relying on more distal outcomes (e.g., smoking abstinence). Moreover, we focused on *preparing* people for quitting smoking. This is because both a

sub-optimal algorithm and a minimal intervention could lead to insufficient support. In the case of an intervention for actual smoking cessation, this might be less ethical as it could mean a failed smoking cessation attempt and thus adverse health outcomes. The behavior change preparation we focus on instead is a less risky but still relevant setting: a preparation phase is often included in behavior change applications (e.g., [243, 266, 374]) to increase the chance of successful behavior change thereafter. While our analyses are thus based on minimal interventions in the context of preparing for quitting smoking, approaches shown to be successful in our studies can later be translated to and tested in full smoking cessation interventions following the stages in the development of technological health interventions defined by Brinkman [68].

We approached our investigation of user needs (*RQ1*) and the *how*-dimension of support (*H1* and *RQ2*) with a joint study on the online crowdsourcing platform Prolific. In this study, 671 daily smokers who were contemplating or preparing to quit smoking interacted with the text-based virtual coach Sam in up to five sessions, which were three to five days apart. In each session, Sam assigned participants a new preparatory activity for quitting smoking (e.g., envisioning one's desired future self, visualizing smoking as a fighting match, or learning progressive muscle relaxation) together with one of five persuasive strategies (Figure 1.2). In the next session, participants reported the effort spent on and experience with their activity. After the five sessions, participants further filled in a post-questionnaire on their involvement in the activities, the barriers and motivators they had for doing their activities, and their views on interaction scenarios for a virtual coach (e.g., whether they would want to receive motivational messages). Examples of participants' views on interaction scenarios are given in Figure 1.3.



Figure 1.2: Examples of messages implementing the persuasive strategies *authority* and *consensus*. To increase the processing of these messages, they were supplemented with reflective questions on their personal relevance.

To examine user needs for the first research question (*RQ1*), we performed a mixed-method analysis, combining a thematic analysis of participants' experiences with their activities, barriers, motivators, and views of interaction scenarios with triangulation with literature and quantitative data (e.g., ratings of the interaction scenarios, user characteristics such as quitter self-identity). Such triangulation of multiple methods and data sources helps to get a comprehensive understanding when performing qualitative analyses [82]. This analysis led to 14 themes describing user needs, based on which we formulated literature-based recommendations for designing a virtual coach-based smoking cessation

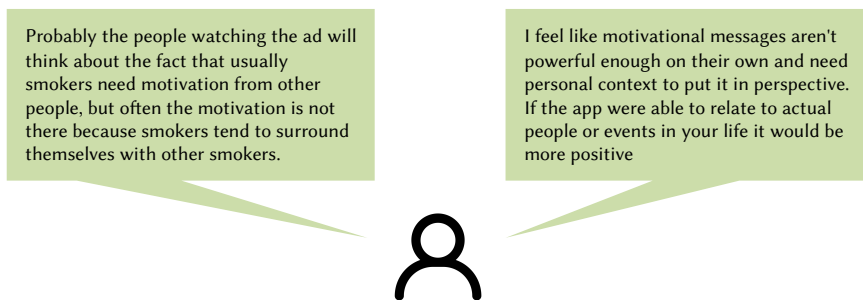


Figure 1.3: Examples of participants' views on receiving motivational messages.

intervention. Since the analysis is based on users' experiences and views expressed after interacting with such a virtual coach-based system for a few sessions where the novelty effect [101, 333] has worn off and all unsatisfied users have not yet dropped out [123], we set out to gain a thorough and more realistic assessment of user needs and how to address them. Chapter 2 presents this assessment.

To test our hypothesis on the effects of different algorithm elements (*H1*), we designed the study as a double-blind mixed-design study with two within-subject factors and one between-subject factor. The within-subject factors were the session in which a persuasive attempt was made (4 levels: sessions 1-4) and algorithm activeness (2 levels: off/on for sessions 1-2/3-4). The between-subject factor was the algorithm complexity used to choose a persuasive strategy after session 2. This factor had four levels with successively more elaborate optimization strategies. Ordered by complexity, the algorithm levels look for the highest value of either: 1) the average reward, 2) the average reward in a person's state, 3) the expected cumulative sum of rewards in a person's state over time, or 4) the similarity-weighted expected cumulative sum of rewards in a person's state over time. This means that starting from sending a persuasive strategy with the highest average reward, we progressively added the consideration of states, future states, and the weighting of samples based on the similarity of people. To answer *H1*, we fit two Bayesian multi-level models to assess the impact of the algorithm complexity level on the effort and the perceived motivational impact. Our approach and results for *H1* and related analyses are presented in Chapter 3.

Our analysis of the effects of user characteristics versus states for our second research question (*RQ2*) is based on performing simulations on the dataset also used to assess our first hypothesis (*H1*). Specifically, we performed leave-one-out cross-validation on the collected data to compute how well we can predict the effort-based reward using either user characteristics or states. Having collected data on more than 30 user characteristics (e.g., personality, need for cognition, quitter self-identity) and 8 possible state features, we performed the comparison based on three user characteristics and three state features, selected so that they capture much of the variation in the reward. Using three user characteristics and three state features ensures a fair comparison and that a realistic number of questions would need to be asked to users in a live application. Our baselines were predicting the reward based on 1) the overall mean reward, 2) the mean reward per

persuasive strategy (i.e., action), 3) and the mean reward per action and state. Our findings for *RQ2* and further related analyses can be found in Chapter 4.

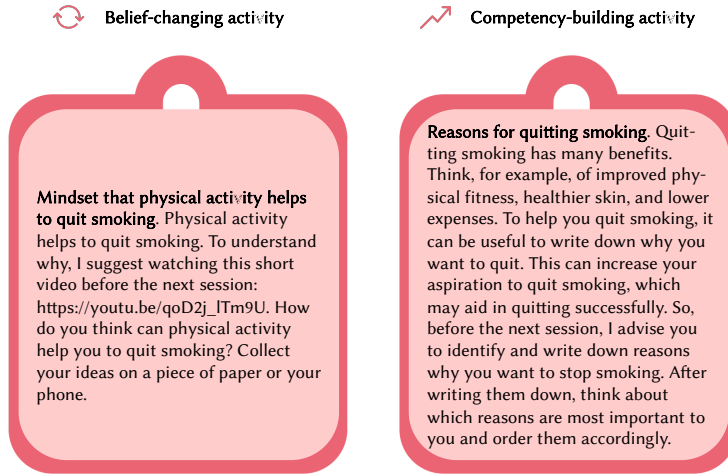


Figure 1.4: Examples of belief-changing and competency-building activities for quitting smoking.

After having investigated adapting *how* people are persuaded, we next studied adapting *what* people are asked to do. More precisely, for our third research question (*RQ3*) we wanted to build an RL model for choosing competency-building activities for quitting smoking that accounts for the views of health experts and smokers. Therefore, we first had to determine the views of experts and smokers. To this end, we conducted two repertory grid studies [134], one with experts and one with smokers. Based on personal construct theory [192], the goal of the repertory grid technique is to explore personal construct systems, or, in other words, see the world as other people see it [134]. In our case, we specifically wanted to know how experts and smokers view competency-building activities for quitting smoking (i.e., which competencies they think are built by these activities). Suppose we know which competencies the activities build according to *experts*. In that case, our model can keep track of the extent to which smokers have already built the different competencies and choose activities that help build missing competencies. On the other hand, by knowing which competencies *smokers* think are built by the activities, our model can consider which competencies smokers find useful when choosing activities. Having established the six competencies forming the view of experts and the nine competencies forming the view of smokers, the next step was to train an RL model that accounts for these views when choosing activities. Thus, we performed another longitudinal study with more than 500 daily smokers interacting with a text-based virtual coach in up to five sessions. In each session, participants received a new randomly chosen activity, either one that was meant to build their competencies or one that was meant to change their usefulness beliefs (Figure 1.4). This setting can be seen as a micro-randomized design [196]

in which each participant receives a random intervention option at each pertinent decision point. To assess how much each model component contributes to allowing smokers to build competencies for quitting smoking, we performed human data-based simulations with different ablated² model versions. Such simulations are a common way of evaluating RL models [384] as they allow testing many different parameter settings [217] before choosing one or more promising settings to study in more costly real-world experiments. Our entire pipeline for building the RL model as well as our evaluation approach and results can be found in Chapter 5.

Having studied the *how*- and *what*-dimensions of support, the third and last dimension of support we studied is *who* people are supported by. To answer *RQ4* about the long-term effectiveness of human support for quitting smoking, we performed a third longitudinal study with a micro-randomized design on Prolific, this time with 679 daily smokers and vapers. Participants again completed up to five sessions with a text-based virtual coach in each of which they were assigned a new preparatory activity for quitting smoking. This time, however, participants also had a 20% chance of receiving a feedback message (Figure 1.5) from human coaches between each pair of sessions. Using the data from this study, we fit an RL model that chooses when to allocate human support. To assess the long-term effects of this model, we again performed human data-based simulations. In these simulations, we compared the RL model to static baselines that allocate human support with varying frequencies. In addition, we assessed the effect of different economic costs for providing human support on human support allocation and the resulting effort-based reward over time. More information on this and related analyses of the effects of human support are provided in Chapter 6.

Hello, it's great that you have made the decision to stop smoking! You're very right, if we want to achieve a goal we need to be consistent. However, it is important to give yourself some compassion, too. We cannot expect ourselves to be perfect at something from the very beginning. Be gentle with yourself through this process and celebrate the small successes. Exercise really is a great way to divert yourself from wanting to smoke and ignoring that little voice in your head that one more cigarette won't be so bad. I hope that the next activity will help you with getting some motivation to get more physically active! I can see that you are motivated to stop smoking and are confident in your ability to do so, your mindset is already there! Keep up the good work and remember to be kind to yourself.

Best wishes,
Karina & Goda on behalf of the Perfect Fit Smoking Cessation Team



Figure 1.5: Example of human feedback sent to participants of our study on the long-term effectiveness of human support.

²In machine learning, ablation is the removal of a model component. It is commonly used to study the contribution of model components to the performance of the overall model.

1

The final chapter of this thesis, Chapter 7, summarizes the main findings and contributions of this thesis. To better understand these conclusions, it also reflects on possible limitations and the responsibility of the research. Directions for future work as well as final remarks are also provided.


2

USERS' NEEDS FOR A DIGITAL SMOKING CESSATION APPLICATION AND HOW TO ADDRESS THEM: A MIXED-METHODS STUDY

Background. *Despite their increasing prevalence and potential, eHealth applications for behavior change suffer from a lack of adherence and from dropout. Advances in virtual coach technology provide new opportunities to improve this. However, these applications still do not always offer what people need. We, therefore, need a better understanding of people's needs and how to address these, based on both actual experiences of users and their reflections on envisioned scenarios.*

Methods. *We conducted a longitudinal study in which 671 smokers interacted with a virtual coach in five sessions. The virtual coach assigned them a new preparatory activity for quitting smoking or increasing physical activity in each session. Participants provided feedback on the activity in the next session. After the five sessions, participants were asked to describe barriers and motivators for doing their activities. In addition, they provided their views on videos of scenarios such as receiving motivational messages. To understand users' needs, we took a mixed-methods approach. This approach triangulated findings from qualitative data, quantitative data, and the literature.*

Results. *We identified 14 main themes that describe people's views of their current and future behaviors concerning an eHealth application. These themes relate to the behaviors themselves, the users, other parties involved in a behavior, and the environment. The most prevalent theme*

 **Nele Albers**, Mark A Neerincx, Kristell M Penforinis, and Willem-Paul Brinkman. Users' needs for a digital smoking cessation application and how to address them: A mixed-methods study. *PeerJ*, 10:e13824, 2022. doi: 10.7717/peerj.13824.

was the perceived usefulness of behaviors, especially whether they were informative, helpful, motivating, or encouraging. The timing and intensity of behaviors also mattered. With regards to the users, their perceived importance of and motivation to change, autonomy, and personal characteristics were major themes. Another important role was played by other parties that may be involved in a behavior, such as general practitioners or virtual coaches. Here, the themes of companionableness, accountability, and nature of the other party (i.e., human vs. AI) were relevant. The last set of main themes was related to the environment in which a behavior is performed. Prevalent themes were the availability of sufficient time, the presence of prompts and triggers, support from one's social environment, and the diversity of other environmental factors. We provide recommendations for addressing each theme.

Conclusions. *The integrated method of experience-based and envisioning-based needs acquisition with a triangulate analysis provided a comprehensive needs classification (empirically and theoretically grounded). We expect that our themes and recommendations for addressing them will be helpful for designing applications for health behavior change that meet people's needs. Designers should especially focus on the perceived usefulness of application components. To aid future work, we publish our dataset with user characteristics and 5074 free-text responses from 671 people.*

INTRODUCTION

When creating an eHealth application for behavior change, one is confronted with many choices. The first one relates to behavior change techniques, for which Michie et al. [244] alone formulated 93 options, including coping planning, self-talk, and social support. Second, one has to decide how to implement these behavior change techniques. For example, should users create coping plans regularly, or only when they feel the need? Third, it gets more complicated when another party, such as a virtual coach, general practitioner, or somebody from the social environment, is involved. When should these parties be included, and how? And lastly, all of these choices should be made so that users use and continue to use the application. For people to use an application, it has to meet their needs. So what are users' needs for using a behavior change application, and what does this imply for somebody creating such an application?

Recent years have seen a surge of eHealth applications with 78,000 new ones in major app stores in 2017 alone [320]. These applications can be easy to use, available at all times, scalable, cost-effective, and can facilitate tailoring of the intervention [218]. These characteristics make such applications beneficial for people wishing to change their health behavior, which can be difficult without help. For instance, more than two-thirds of adult smokers in the United States want to quit smoking [44], but most unassisted quit attempts fail [99]. However, despite their potential, users commonly do not adhere to eHealth applications or abandon them entirely [56, 147, 191]. Thus, there appears to be a mismatch between what the applications offer and what users need.

To improve behavior change applications, users' needs must be better understood. Thereby, it is crucial to take a holistic approach that considers not only the technology itself but also the user and their environment [373]. Previous work in the context of quitting smoking has, for example, found that the intuitiveness of the user interface [204], users' experience with computers [142], the appreciation expressed by a conversational agent

[204], and support from one's social environment [354] play a role. This illustrates the diversity of factors that may need to be considered in eHealth applications.

Studies for getting input on the system design from users differ in two ways. First, they employ systems of differing maturity, ranging from mere design ideas to complete applications. Each of the two extremes has an advantage: the former allows one to test multiple design options more easily; the latter helps users to more accurately identify benefits and barriers to using the application [79]. Second, users interact with a system for varying amounts of time before being asked for their input. Both very short and very long uses are at risk of resulting in an overly positive evaluation of a system: the former because people's initial curiosity and excitement about a novel system tend to fade as they become more aware of the system's limitations [101, 333], and the latter because people for whom the system does not work tend to drop out according to the law of attrition [123]. Thus, data should be collected in the middle range, where the novelty effect has worn off and average users have not yet dropped out. This is to allow one to more accurately assess users' needs.

This study aims to get a more accurate assessment of users' needs for eHealth applications for behavior change. To this end, we collected data on both the use of an application and views on multiple design ideas from this middle time range. More precisely, we conducted a longitudinal study in which 671 smokers interacted with a text-based virtual coach. Virtual coaches or conversational agents have been receiving a lot of attention in the health context due to their potential ability to increase engagement, provide and discuss relevant information, and form a connection with users [255]. Participants of our study interacted with such a virtual coach in up to five sessions spread over at least nine days. In each session, participants were assigned a new preparatory activity for quitting smoking or increasing physical activity, with the latter possibly aiding the former [152, 368] and vice versa [291]. To gain a comprehensive understanding of participants' needs for using the application, we conducted a mixed-methods analysis. This analysis was based on participants' characteristics such as their physical activity identity, their feedback on their activities as well as barriers and motivators and thus information on actual behavior, their views on videos of interaction scenarios described after completing the five sessions and thus information on experience-based behavioral intentions for multiple design options, and findings from the literature. We found a comprehensive set of 14 themes that describe users' needs. We used these themes to formulate recommendations to support designers of future health behavior change applications. To further aid future research on understanding user needs, we publish our data together with this article.

MATERIALS & METHODS

We conducted a longitudinal study from 20 May 2021 until 30 June 2021. The Human Research Ethics Committee of Delft University of Technology approved the study (Letter of Approval number: 1523), and we preregistered the study in the Open Science Framework (OSF) [14].

STUDY DESIGN

The study followed a mixed design with five sources of information. We collected the first four from participants: **###** their characteristics (e.g., physical activity identity), **↩** their feedback on their previously assigned preparatory activity for quitting smoking or increasing physical activity, **🚧** barriers and motivators they had for doing their activities, and **🗣️** their views on interaction scenarios for a virtual coach (e.g., whether participants would like to receive motivational messages). The user characteristics were quantitative, the barriers and motivators were qualitative, and the activity feedback and views on interaction scenarios were quantitative and qualitative. Each participant saw a random selection of two interaction scenarios with the goal of presenting each scenario the same number of times across the sample population. Figure 2.1 illustrates how we gathered data from participants. Finally, **📚** previous studies provided information to triangulate the findings from the other four sources of information. Triangulation of multiple data sources or methods has been described as a way to examine the validity of qualitative research and to obtain a comprehensive understanding of a phenomenon [82]. A successful example of triangulating qualitative findings with previous studies as part of the analysis is the work of Nahar et al. [262] in the context of software engineering, which we took as an inspiration for this study.

MATERIALS

We used the online crowdsourcing platform Prolific to recruit, invite, and communicate with participants, Qualtrics to host the questionnaires and instructions for the sessions, Google Compute Engine to host the virtual coach and the sessions using Rasa X, and YouTube to host the videos shown for the interaction scenarios.

The virtual coach used for the sessions was implemented in Rasa [61] and had the name Sam. Sam introduced itself as wanting to help people to prepare to quit smoking and become more physically active, with the latter possibly facilitating the former. The code of Sam can be accessed online [9]. Sam proposed a new preparatory activity related to quitting smoking or increasing physical activity in each session. The virtual coach randomly drew these activities from a pool of 24 activities, 12 each for quitting smoking and increasing physical activity. The activities were based on components of the smoking cessation app StopAdvisor [243] and future-self exercises [236, 294], and reviewed by a psychologist and smoking cessation expert. Examples of activities are formulating a rule for not smoking or tracking one's physical activity. Table A.1 shows the complete list of activities. An example of a conversation with Sam is shown in Figure A.8. Based on their acceptance of Sam measured in the post-questionnaire with six items on scales from -5 to 5 and with 0 being neutral [18], participants had an overall positive attitude toward Sam ($M = 2.50$, $SD = 1.68$, $95\% \text{ HDI} = [2.32, 2.68]$).

In the post-questionnaire, each participant saw 2 out of 13 interaction scenarios in video form. Each video presented an imaginary persona alongside her situation and described an interaction for this persona. The video ended with a question about whether the viewer would engage in the interaction if they were the persona. The topics for the scenarios (Table A.2) were drawn from the literature and discussions within the consortium of the multidisciplinary Perfect Fit project [238] that this study is a part of. This project aims to develop an app that helps smokers quit smoking and become more physically active.

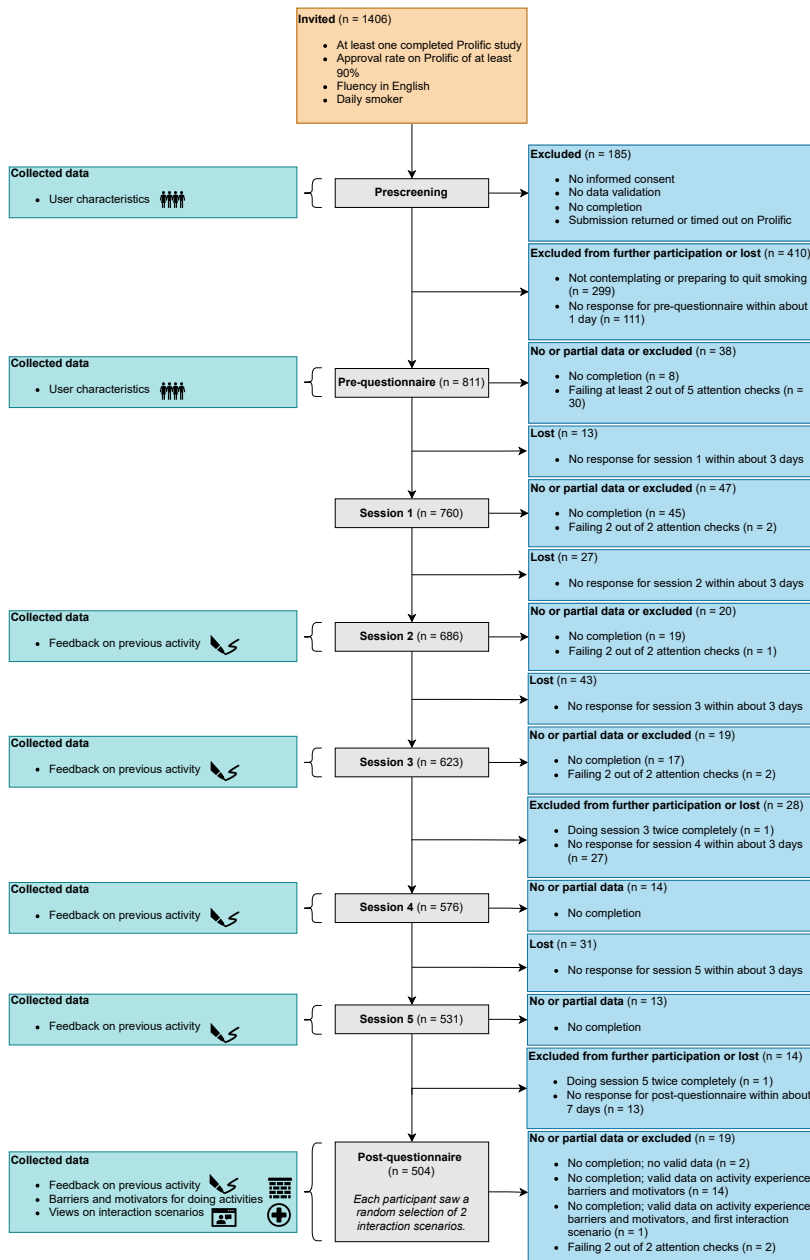


Figure 2.1: Study design. Design of the study, including the study components, collected data, and participant flow. Icons illustrate the four types of data we collected from participants: characteristics, feedback on preparatory activities, barriers and motivators for doing the activities, and views on interaction scenarios for a virtual coach. The numbers next to the study components indicate how many participants started the respective component. For the post-questionnaire, we show which data we collected from participants who did not complete it.

There were two versions for each video, one with a male and one with a female persona. Male and female participants saw a video with a persona whose gender matched their own; participants with a different gender saw one with a persona whose gender was chosen randomly. The information in the videos was presented using text and voice-over. Table A.3 provides links to the videos on YouTube.

2

MEASURES

We used the following measures in our analysis:

✎ **Activity effort and experience.** Using an adaptation of the scale from Hutchinson and Tenenbaum [170], participants were asked the amount of effort they spent on their activity from the previous session. Moreover, we asked participants about their experience with their activity through a free-text question. After describing their experience, participants could provide modifications in a second free-text response. Table A.4 provides details on these three measures.

☒ **Barriers and motivators for doing the activities.** We asked participants about their barriers and motivators for doing their assigned activities using two free-text questions (Table A.4).

☒ **Views on interaction scenarios.** Each interaction scenario ended with a question about whether participants would engage in the shown interaction if they were the persona from the video. Participants were asked to provide a rating on a scale from -5 to 5 and a free-text response after the prompt "Why do you think so?". Table A.5 shows the question and scale endpoints for each interaction scenario.

☒ **User characteristics.** We measured several user characteristics to explore their effect on the other measures. This included quitter and non-smoker self-identity measured with three items each based on Meijer et al. [235] and physical activity identity based on an adaptation of the exercise identity questionnaire by Anderson and Cychosz [38] to physical activity. All identity-related items were measured on 5-point Likert scales. In addition, we measured the Transtheoretical Model (TTM)-stage for becoming physically active based on an adaptation of the question by Norman et al. [271] to physical activity, and people's Big-Five personality based on the 10-item questionnaire by Gosling et al. [146]. The 10-item questionnaire by Gosling et al. [146] was chosen due to its brevity and use in previous work on individual differences in behavior (e.g., [182]). Despite its brevity, its convergent correlations compared to longer questionnaires such as the 44-item Big-Five Inventory (see [178]) have been found to be substantial [146]. We also gathered information from participants' Prolific profiles. This included their age range (e.g., 21 – 25), smoking frequency, weekly exercise amount, household size, and their highest completed education level. We used the education level as a measure of socioeconomic status, as is commonly done in smoking research [235].

PARTICIPANTS

Eligible participants were those who were fluent in English, smoked tobacco products at least once per day, were contemplating or preparing to quit smoking [112], were not part of another intervention to quit smoking, and provided informed consent. Further, we aimed to increase the quality of the responses by requiring participants to have at least one completed study and an approval rate of at least 90% for their completed studies on

Prolific. 1406 participants started the prescreening questionnaire, and 485 of the 922 eligible participants successfully responded to both interaction scenarios in the post-questionnaire. Participants had about one day to respond to their invitation to the pre-questionnaire, three days for the sessions, and seven days for the post-questionnaire. The participant flow is presented in Figure 2.1.

Participants who successfully completed a study component were paid based on the minimum payment rules on Prolific (5 pound sterling per hour). Since some participants faced difficulties accessing the videos of the interaction scenarios, participants who completed everything but part of the scenario questions in the post-questionnaire were also paid ($N = 15$). Participants were told that whether they did and how they reported on their assigned preparatory activities would not affect their payment. This was to account for self-interest and loss aversion biases. Self-interest bias can come into play when there are incentives that motivate participants to respond in a certain way; loss aversion bias can arise when participants suspect that they may not get paid fairly and thus choose not to participate or to drop out [117].

Participants on Prolific were nationals of or lived in member countries of the Organisation for Economic Co-operation and Development (OECD) with the exception of Turkey, Lithuania, Colombia and Costa Rica and the addition of South Africa [310]. Of the 671 participants with at least one valid free-text response, 349 were female, 310 were male, and 12 indicated a different gender or provided no information. The youngest participant was 18 and the oldest 74. With regards to smoking behavior, participants could be characterized as smoking once (5.37%), 2 – 5 times (24.59%), 6 – 10 times (31.74%), 11 – 19 times (28.32%), or more than 20 times (9.54%) per day. Moreover, 78.69% of the participants indicated having previously quit smoking for at least 24 hours. An overview of these and further participant characteristics is provided in Table A.6.

While sample sizes are less relevant for Bayesian analyses like ours than for frequentist ones [90], we conducted a power analysis to get an idea of the statistical power of the quantitative part of our analysis in which we compute Spearman correlation coefficients. Following the Monte Carlo approach described by Kruschke [202], we used 1000 simulations of two standardized variables with a medium correlation of 0.3 according to Cohen [97]. For each simulation, we computed the 95% Highest Density Interval (HDI) for the correlation, with an HDI being "the narrowest interval containing the specified probability mass" [234]. The power was then calculated as the fraction of simulations in which the lower bound of the HDI was greater than zero. The result was a power of 0.68 for a sample size of 71, a power of 0.95 for a sample size of 148, and a power of >0.99 for a sample size of 300. These sample sizes are the smallest, median and largest number of samples we obtained for a group of interaction scenarios used in our quantitative analysis.

PROCEDURE

Participants meeting the qualification criteria, passing the prescreening, and successfully completing the pre-questionnaire were invited to the first of five sessions with the virtual coach Sam. Those participants who successfully completed all five sessions were invited to the post-questionnaire. The post-questionnaire first asked participants about their effort spent on and experience with their last activity, then asked them about their motivators and barriers for doing their activities, and finally showed them two interaction scenarios.

Before each scenario, participants were told that they would see a video and asked to turn on their audio. Underneath the video, we provided a link to the video on YouTube in case participants could not see the video in Qualtrics. Once the duration of the video had passed, participants could proceed to the next page to provide a rating and a free-text response for the scenario. Invitations to the next session or post-questionnaire were sent about two days after completing the previous session. Showing the interaction scenarios after participants had interacted with the virtual coach in five sessions spread over at least nine days ensured that participants had personal experience of interacting with a virtual coach. Using an operational system has been described as crucial to be able to see possible benefits of health information technology [79].

DATA PREPARATION AND ANALYSIS STRATEGIES

DATA PREPARATION

We preprocessed the gathered data by 1) using only data from sessions and the post-questionnaire if participants passed at least one attention check in the respective component, 2) using the first recorded submission for a study component if participants did the component more than once, 3) removing ratings and free-text responses for the interaction scenarios for people who wrote in their free-text responses that they could not see the video ($N = 2$), and 4) anonymizing free-text responses by removing potentially identifying or sensitive information. In addition, we computed the reliability of the items corresponding to the quitter, non-smoker, and physical activity identity measures. Since the reliability was sufficiently high for quitter (Cronbach's $\alpha = 0.76$, $N = 671$), non-smoker (Cronbach's $\alpha = 0.69$), and physical activity identity (Cronbach's $\alpha = 0.89$), we used the means of the items as index measures. We also reversed the scale for the TTM-stage for becoming physically active such that a higher value denotes a higher stage of change.

ANALYSIS

We took a mixed-methods approach and proceeded in four steps to analyze the data. These steps were the thematic analysis steps described by Braun and Clarke [64] with the addition of triangulation based on literature and quantitative results. We used two types of triangulation: method and investigator triangulation [82]. Method triangulation was performed using data on both people's actual behavior from their activity experiences and efforts as well as their views on possible behaviors from their free-text responses and ratings for the interaction scenarios. We also used data on user characteristics (e.g., physical activity identity) and findings from the literature. Regarding investigator triangulation, two researchers with different backgrounds were involved in all parts of the analysis. The result are the analysis steps that we now describe in detail.

Preparation of coding scheme. To create our coding scheme, the first author (NA) with a background in artificial intelligence and eHealth first familiarized herself with the data by reading all free-text responses and noting initial inductive codes. These codes were further refined deductively by looking through literature on technology acceptance and use, human motivation and behavior, and perceptions of virtual agents and robots. This included the two versions of the Unified Theory of Acceptance and Use of Technology (UTAUT) [376, 377] (including their extensions with autonomy [194, 206], self-efficacy [160], and characteristics of the technology, situation, task, individual and other humans

[70]), self-determination theory [108], the Capability-Opportunity-Motivation-Behavior (COM-B) model of behavior [242], barriers to behavior [31], the findings by de Graaf et al. [104] on users' experiences with a social robot, and the Ability-Benevolence-Integrity model of trustworthiness [232]. A draft coding scheme was discussed with SV who has a background in interaction design and had also read responses and formulated initial codes. The final coding scheme consisted of three levels, with 4 codes at the highest level, 15 codes at the second level, and 86 codes at the third level. Codes thereby captured both semantic and latent meanings of the responses [65]. The coding scheme is shown in Figure A.7.

Coding of free-text responses. All free-text responses were coded by NA based on the developed coding scheme. Multiple codes were used if relevant. We assessed the reliability using double coding. The second coder SV was further trained by independently coding six sets of ten responses and discussing the coding with NA after each set. Then, SV coded 100 responses. These 100 responses were chosen randomly but such that there were at least six responses per question (i.e., the 13 interaction scenarios, barriers, motivators, and activity experiences). The number of double-coded responses was selected to allow for an error margin of at most 10% to be obtained when calculating percent agreement [150]. We obtained moderate agreement (Cohen's $\kappa = 0.41$) [209] at the third coding level. Due to its more robust nature [151], we also computed the Brennan-Prediger κ [67] for a value of 0.97. Since participants primarily corrected spelling and grammar errors in their modifications of their activity experience answers, these modifications were excluded from further analysis.

Triangulation with literature and quantitative results. As NA and the third author (KP), with a background in psychology, gained insights from the coded free-text responses, literature and quantitative results were used to triangulate the qualitative results. Relevant literature came from diverse research fields such as eHealth, behavior change theories, human-robot interaction, and various application domains. Moreover, we incorporated two types of quantitative data. First, we computed means and Bayesian credibility intervals for the ratings per interaction scenario. The credibility intervals we report are the 95% HDIs. Second, we computed Spearman correlation coefficients between user characteristics (e.g., physical activity identity) on the one hand and users' activity efforts and ratings for groups of interaction scenarios on the other hand. Note that we combined scenarios about similar interactions into groups to facilitate their discussion, as shown in Table A.2. We conducted Bayesian tests for the correlations using the Bayesian First Aid R-package [43] and report the median values and 95% HDIs. We classified the size of the resulting correlations using the guidelines by Cohen [97]. Furthermore, we calculated the posterior probability that a positive correlation is greater than zero and evaluated the probability using the guidelines by Chechile [90].

Search, review and definition of themes, and production of the report. NA and KP examined the results to identify overarching themes. A final set of themes was obtained using multiple rounds of discussion. To produce the report, which is the last thematic analysis step described by Braun and Clarke [64], NA selected participant responses that illustrate the themes. Participants are referred to by numbers (e.g., P123).

RESULTS

We depict the frequencies of the most frequent codes from our coding scheme in Figure 2.2 and those of all codes in Figure A.9. Figure 2.3 further presents the ratings for the interaction scenarios. In addition, we show the correlations between participants' activity efforts and ratings for the interaction scenario groups on the one hand and user characteristics on the other hand in Figure 2.4. We will refer to these figures throughout our discussion of the themes. In this discussion, we move from the smallest unit of analysis, a behavior, to the user who performs a behavior, to another party that may be involved in a behavior, to the largest unit of analysis, which is the environment (Figure 2.5). This approach follows the idea of distinguishing micro, meso, and macro elements of behavior [176] as similarly done in previous work (e.g., [335]).

2

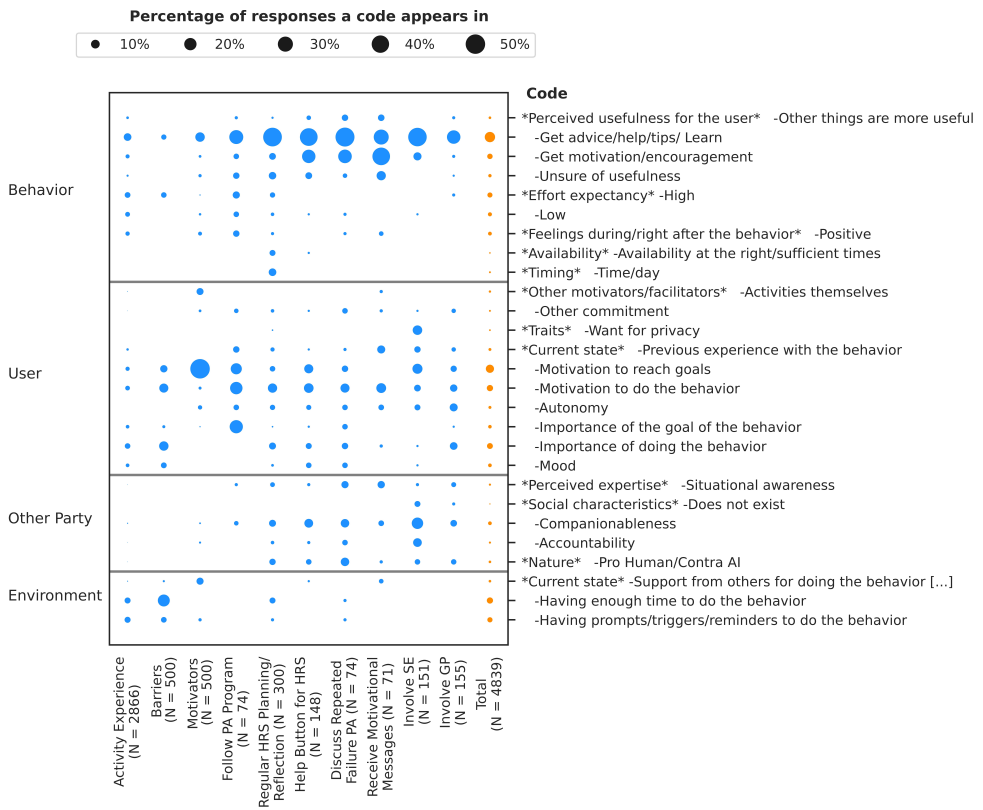


Figure 2.2: Percentage of times that codes from the coding scheme appear in each response type as well as across all response types together. We show only the percentages of those codes that appear in at least 4% of the responses for at least one response type. The response types are the activity experiences, barriers, motivators, and the groups of interaction scenarios.




Figure 2.3: Means and 95% HDIs for the intentions to engage in the interactions from the interaction scenarios. Abbreviations: PA, Physical activity; HRS, High risk situation; SE, Social environment; SO, Significant other; GP, General practitioner.

BEHAVIOR

PERCEIVED USEFULNESS

The most frequent topics both overall and for the interaction scenarios and activity experiences revolved around the perceived usefulness of the behavior (Figure 2.2).

Getting help, advice or tips and learning. Thinking that they would get help, advice, tips, or learn something by engaging in a behavior was the most frequent topic overall (13.97%) and for all interaction scenario groups except for the scenario about receiving motivational messages (Figure 2.2). Participants' concerns included whether the behavior would help to reach their goals (e.g., P283), teach them how to deal with cravings (e.g., P92), or serve as a prompt to reflect in general (e.g., P274) or on their current behavior (e.g., P507). Several participants also stated that they thought a behavior was (not) helpful without providing specific reasons for this evaluation (e.g., P224, P151). For example, some participants who were against involving their General Practitioner (GP) noted that they did not see any way in which their GP could help them (e.g., P345, P639):

No i wouldn't [consult my GP], i don't think my GP could do anything to help.

(☒ Consult GP in case of smoking relapse, P345)

Obtaining information or knowledge has previously been identified as a theme in participants' thoughts on using a self-regulation-based eHealth intervention to increase

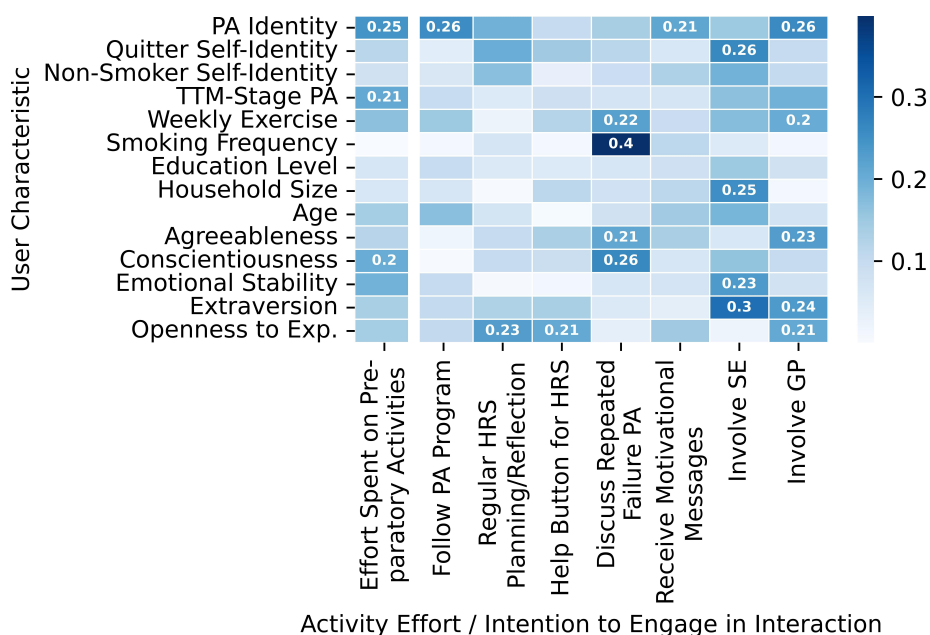


Figure 2.4: Overview of Spearman correlation coefficients between participant characteristics on the one hand, and the effort participants spent on their activities and their intentions to engage in the interactions from the scenario groups on the other hand. Value labels are provided for all coefficients with an absolute value of at least 0.2. The color scheme is based on the absolute values of the coefficients. Abbreviations: PA, Physical activity; TTM, Transtheoretical model; Exp., Experience; HRS, High risk situation; SE, Social environment; GP, General practitioner.

physical activity and intake of fruit and vegetables [303]. It also plays a role in the context of eHealth applications for other domains, including self-management of chronic conditions such as chronic pain [348] and type 2 diabetes [219]. The scoping review of Wilson et al. [391] also showed that the opportunity to learn new information is a motivator for the use of eHealth tools by older adults. It has even been argued that gaining knowledge is such a crucial motivation for using online activities and applications that it makes users active consumers and producers of health knowledge [226].

Getting motivation or encouragement. One element participants were looking for in the behaviors was help in the form of motivation or encouragement, which was with 3.82% the second most frequent topic overall for "behavior" and with 40.85% the most frequent one for the interaction scenario about receiving motivational messages (Figure 2.2). Concerns about receiving motivational messages included whether the messages would be tailored to the user and situation at hand (e.g., P212, P497), help to increase or maintain self-confidence (e.g., P158), or serve as a reminder of (reasons for) quitting smoking (e.g., P6, P25):

I don't generally respond much to motivational messages, but in this instance, anything that can serve as a reminder for why I want and need to do this so much, is a good thing. Being asked to reflect on our reasons for quitting is

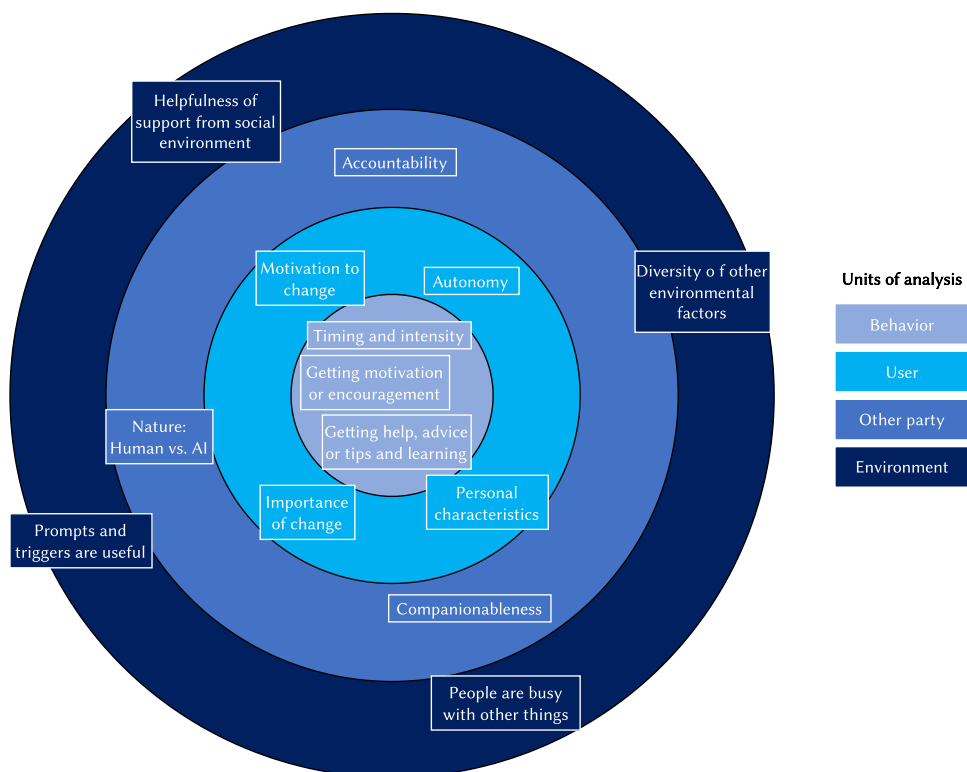


Figure 2.5: Overview of the main discussed themes for each of the four hierarchical units of analysis. The four units are behavior, user (who performs a behavior), another party (that may be involved in a behavior), and environment.

definitely a good thing and something I personally would benefit from. (☑ Receive motivational messages, P25)

Similar findings came to light regarding using a help button for High Risk Situations (HRSs). After pressing such a help button, the virtual coach would provide support for dealing with cravings or difficulties to do planned physical activity. Here 23.65% of responses referred to getting motivation or encouragement. Participants mentioned wanting to get motivation in general (e.g., P283, P617), be reminded of (reasons for) quitting smoking (e.g., P52, P291), or get the strength to resist cravings (e.g., P79). Some participants also explicitly mentioned the importance of being motivated by someone or something else (e.g., P57, P417, P636):

... sometimes I need a bit more motivation than what's going on in my head and need a little kick or nudge in the right direction (☑ Help button for PA HRSs, P636)

These findings coincide with work by Poppe et al. [303], which found that the opportunity to be motivated by being reminded of one's goals was a reason for participants to prefer

a mobile application instead of a website for increasing physical activity and the intake of fruit and vegetables. Moreover, Kulhánek et al. [204] found motivation strengthening to be a frequently mentioned benefit of a conversational agent that assists with quitting smoking.

2

Improving feelings or mood, novelty, and comparison with others. Besides getting help and motivation, several other topics related to the perceived usefulness of the behavior. This included, for instance, whether engaging in the behavior improved one's feelings or mood (e.g., P50, P244), whether the behavior was novel (e.g., P236, P464), or whether the behavior made a comparison with others possible (e.g., P310). Giving users the option to compare their performance to others, for example, can be an effective motivation strategy in persuasive games [285], just as novelty can be motivating [101, 333]. According to this novelty effect, people are initially curious about a new system or technology and have high expectations regarding its usefulness. Yet, these perceptions decrease over time as the system's limitations become apparent. The novelty effect was implied by one participant who gave the following response when asked to describe their motivators for doing their preparatory activities:

curiosity at first, but that waned (☉ Motivators, P338)

Recommendations for addressing the perceived usefulness of the behavior. The central role perceived usefulness plays for the acceptance of a system has been modeled by the UTAUT [376], which posits that the extent to which a person thinks that using a system will lead to personal performance improvement influences the intention to use that system. This effect on behavioral intention has been illustrated in studies of technologies in diverse contexts such as an app for insomnia treatment [130] and a socially assistive robot [136]. One approach to increase the perceived usefulness is to tailor advice or content in general to users' preferences. Previous work has, for example, attempted to automatically optimize suggested activities for older adults [100], physical activities for adults [313], or breast cancer screening recommendations [42]. Yet, future work on such automated approaches should take the novelty effect into account. An alternative is to use a participatory approach in which potential users can contribute to the design [103]. Yet, our participants often drew conclusions about the usefulness of an intervention component without having tried it. This suggests that other elements such as telling people that the content is meant specifically for them [114] or explaining in more detail why and how something helps them [166] may be necessary.

Lastly, since getting motivation or encouragement was frequently sought out in our study, special attention should be given to how participants can be motivated or encouraged. Receiving motivational messages was overall seen positively in our study (Figure 2.3), but we saw earlier that several people were concerned about the format of these messages. In addition, participants were also looking for motivation in other interactions, such as a discussion with their virtual coach about their repeated failure to reach their physical activity goals (Figure 2.2).

TIMING AND INTENSITY

The timing of the behavior is interesting in that it appears to play a role primarily for one type of interaction, namely, regular planning and reflection for HRSs. There appeared to

be no consensus regarding an ideal time for this. For instance, some participants liked the proposed day (e.g., P470, P487) for physical activity- or time of the day (e.g., P121, P493, P627) for smoking-related planning and reflection. But others found the suggested timing for planning and reflection to be inappropriate:

Maybe [I would make a plan with my virtual coach] on a different day, e.g. friday for the week ahead starting on Monday. I do like the planning of the week and think ahead of the hurdles to avoid them. (☒ Plan for PA HRSs on Sundays, P262)

I think it's great idea [to make a plan with my virtual coach,] just not sure if I would have enough time in the morning to do so (I start work at 6.30 a.m. so that would have to be quick plan) (☒ Plan for smoking HRSs in the mornings, P113)

Besides timing, participants were also concerned about the frequency (e.g., P158, P178) and duration (e.g., P362, P552) of (potential) application components.

Finally, while some participants favored the regularity of the planning and reflection interactions (e.g., P462), others expressed their preference for on-demand rather than regular support:

I think consulting the virtual coach might be helpful - but the option to let the app know whenever I have a craving would be more convenient. (☒ Reflect on smoking HRSs in the evenings, P100).

Recommendations for addressing the timing and intensity of the behavior. Previous work on needs, barriers, and facilitators for older adults to use eHealth applications found instant availability of help through such applications to be a facilitator [378]. However, given the diversity of user needs and preferences, there appears to be a need for tailored timing and intensity of intervention components. The importance of personalizing these elements of health behavior change applications has previously been pointed out by Dijkstra [114]. Approaches to achieving such personalization include both letting participants choose themselves and automatically determining opportune moments [167]. Letting participants choose acknowledges that users may know best what time suits them and that supporting autonomy and competence may increase motivation and performance [332]. In contrast, determining opportune moments for the user aims to account for the fact that people's self-knowledge is generally not very accurate [119, 375]. A middle ground may be to support users in systematically finding out what suits them [186].

USER

IMPORTANCE OF AND MOTIVATION TO CHANGE

Importance of change. When it comes to physical activity, several participants pointed out that becoming (more) physically active was not important to them. Reasons included already being physically active (e.g., P143, P533) and not seeing how physical activity helps to quit smoking (e.g., P657). Some participants, for instance, thought that working simultaneously on becoming physically active and quitting smoking was too difficult:

It's a bit condescending. Giving up smoking is hard enough without having to do a fitness regime also. (☒ Follow PA program while quitting smoking, P464).

Notably, this was even though participants were informed about the potential positive effects of physical activity on quitting smoking at the start of the study as well as every time they were assigned an activity for physical activity increase. Participants who were in favor of physical activity-related behaviors, on the other hand, frequently pointed out the benefits of physical activity both in general (e.g., P245, P335) and for quitting smoking specifically:

I really think that physical activity could help me to quit smoking forever (☉ Motivators, P543)

The TTM provides a framework for assessing participants' views on the importance of changing a behavior. Stage one, the precontemplation stage, is one in which people do not aim to change their behavior within the next six months and are often un- or under-informed about the consequences of their behavior [309]. As people progress through the stages, they become increasingly aware of the consequences of their behavior(s) and ready to take action. The relationship between TTM-stage and taking action was observable in our study. More precisely, we observed a small correlation of 0.21 between people's TTM-stage for becoming physically active and the effort they spent on their activities (Table 2.1). Hence, participants in higher stages overall spent somewhat more effort on their activities. Notably, we find this association between the TTM-stage for becoming physically active and the effort spent on activities even though only half of the activities were targeted at physical activity.

Table 2.1: Results of Bayesian analyses of Spearman correlations between user characteristics on the one hand and activity efforts and intentions to engage in the interactions from the scenario groups on the other hand. For all seven correlations, at least 99.8% of the posterior distribution for the mean correlation was greater than zero. This leads to at least a very strong bet that the mean correlations are greater than zero [90].

User Characteristic	Effort/Scenario Group Rating	Median [HDI]
Conscientiousness	Effort	0.20 [0.17, 0.24]
Extraversion	Rating for "Involve SE"	0.31 [0.15, 0.45]
	Rating for "Involve GP"	0.24 [0.08, 0.38]
Household size	Rating for "Involve SE"	0.26 [0.10, 0.40]
PA identity	Effort	0.25 [0.21, 0.28]
Smoking frequency	Rating for "Discuss repeated failure PA"	0.41 [0.20, 0.59]
TTM-stage PA	Effort	0.21 [0.17, 0.24]

Abbreviations: HDI, Highest density interval; SE, Social environment; GP, General practitioner; PA, Physical activity; TTM, Transtheoretical model.

Motivation to change. Once participants have become aware of the importance of changing their behavior, they also need to be motivated to do so. In our study, people's motivation to reach their goals was with 50.00% the most commonly reported motivator

for completing their activities, making it the overall most frequently mentioned topic for the "user" unit (Figure 2.2). Goals which motivated participants included quitting smoking (e.g., P403, P463) and becoming more physically active (e.g., P126, P142), but also more general goals such as improving one's health (e.g., P224, P390) and very individual ones such as being able to hold one's breath for longer (P624). Sometimes, participants explicitly linked these other goals to quitting smoking or becoming more physically active:

I was motivated by exploring all the positive health, family and psychological benefits that comes with quitting. And it fuels my drive to do the assigned activities. (☉ Motivators, P521)

my current health is not good so to try to improve this i need to give up smoking and do more exercise. This is my main motivation (☉ Motivators, P493)

At the same time, not wanting to quit smoking was mentioned by some participants as a barrier to doing their activities (e.g., P455, P591). Moreover, several participants were not confident in their ability to reach their goals (e.g., P111). These findings resonate with ones by Milcarz et al. [247], according to which difficulties in quitting smoking and a lack of willingness to quit were the most commonly mentioned barriers.

Besides conscious or reflective motivation, automatic motivation also plays a role. Automatic motivation consists of automatic processes involving, for example, emotional reactions and reflex responses such as feeling excited at the prospect of going running in the evening [245]. One way to capture people's automatic motivation is to look at their self-identity [242, 386], which has been shown to predict people's (intention and motivation to perform a) behavior [306, 393]. In our results, we find some influence of self-identity on behavior based on a small correlation of 0.25 between physical activity identity and the effort people spent on their activities (Table 2.1). Thus, it appears that participants with stronger physical activity identities spent more effort on their activities. In addition, physical activity identity was also positively related to people's willingness to follow a physical activity program besides quitting smoking, receive motivational messages, and involve their GP in the quit smoking process (Figure 2.4). This is interesting, as some of these behaviors do not involve physical activity. One explanation could be that people with stronger physical activity identities felt more involved in the intervention as a whole due to the combination of quitting smoking and becoming physically active. One participant, for instance, pointed out the importance of being willing to become more physically active to be a suitable participant for the intervention:

I'm fairly positive about this question because considering some of the activities were physical it would be important that the person was willing to be more physically active. (☒ Follow PA program while quitting smoking, P204)

This suggests that participants may need to be separately motivated to reach all goals an intervention puts forward. Notably, the association between quitter or non-smoker self-identities and the effort participants spent on their activities is much weaker than in the case of physical activity identity (Figure 2.4). This likely is the case because participants had to be in either the contemplation or the preparation stage of the TTM for quitting smoking to be eligible for the study. As such, their identity as quitter or non-smoker

may have been quite weak because not (yet) fitting with their (readiness to change their) smoking behaviors.

Recommendations for addressing the perceived importance of change and motivation to change. Our results show that as a first step, people need to be able to link a behavior to a desired outcome. For example, in this study, the benefits of physical activity for quitting smoking may have needed to be clearer. Our findings resonate with ones by Poppe et al. [303], who showed that participants did not always agree with provided information on the positive effects of physical activity and fruit and vegetable intake. Participants who are not aware of the consequences of their behavior are often described as being in the precontemplation stage of the TTM. Important processes of change in this stage include receiving information about the behavior, evaluating how the problem affects one's environment, and experiencing emotions about one's situation and problems [309]. Once people have learned about the consequences of their behavior, a next process of change may be self re-evaluation or realizing that the behavior change matches one's identity better [309]. Several participants of our study mentioned, for example, that an activity asking them to think about their feared or desired future self concerning quitting smoking or becoming more physically active was a major motivator (e.g., P243, P263):

The task where I was to think about what I would be in the future gave me a huge boost for the rest of the tasks (⊕ Motivators, P592)

Thinking about oneself in this way may weaken identity related to a problem behavior and strengthen identity related to one's non-problem behavior (i.e., non-smoking if the problem behavior is smoking). Consequently, as identity represents a strong form of automatic motivation, such a change in identity can lead to successfully changing (health) behaviors such as smoking [293, 390]. In addition, people need to gain confidence to change their behavior. Ways to increase confidence include allowing users to make small wins [35], letting them observe a relatable other succeed [46], verbally persuading them [353], or improving their mood [189].

AUTONOMY

A topic that appeared several times in participants' responses was autonomy (Figure 2.2). For example, one participant perceived asking for help in a physical activity HRS as being a sign of excessive dependence on others, especially since they regarded the situation as not serious enough to require help:

If I needed help like this just because I wasn't exercising one evening I would feel very concerned that I could no longer stand on my own two feet in life. This is going too far. We are adults and this kind of hand holding is not healthy or helpful. (☒ Help button for PA HRSs, P133)

While a need for autonomy was put forward in several different interactions, the percentage of times this topic appeared was with 8.39% notably high in the context of involving one's GP (Figure 2.2). Participants mentioned wanting to decide for themselves whether to involve their GP instead of following the advice of the virtual coach (e.g., P429, P528). They also mentioned wanting to keep trying to change their behavior themselves (e.g., P472,

P660), or only contacting their GP in severe cases (e.g., P85, P539). One participant also emphasized that seeing a GP is not helpful unless one is also motivated to quit smoking:

... I believe that quitting smoking comes entirely from self motivation. Even if I would go to my GP, if I'm not personally committed it would be just a waste of time and money (☒ Consult GP at start of quit attempt, P624)

According to self-determination theory, people need to feel in control of their behaviors and goals to initiate behavior [332]. As such, the proposal to involve a GP likely violated this need. Besides our study, other eHealth interventions have identified autonomy and independence as needs for long-term weight maintenance [40], healthy living for cardiovascular disease prevention and rehabilitation [66], and self-management of chronic pain [348]. The latter study also explicitly reported the difficulty of finding a balance between asking for help and being independent. Interestingly, however, our participants' overall stance on involving the GP was also much more negative than for the other proposed interactions, with the mean of the credible interval even being less than zero for contacting the GP at the start of the quitting process (Figure 2.3). This contradicts findings of previous work, which show that smokers tend to accept unsought conversations about smoking with their GP [212]. Therefore, it was likely not the involvement of a GP alone that concerned participants in our study, but rather the way and personal situation in which involving a GP was proposed.

Recommendations for addressing users' autonomy. Our results suggest that when a virtual coach recommends help for participants, their need for autonomy can be violated. Thereby, it appears to be crucial to not only consider what is recommended but also how and when. Interesting work in this regard has been conducted by Tielman et al. [363]. Their model for referring patients to human care is based on a combination of willingness to see a human and severity of the situation. If the situation is not severe and a patient is not in favor of seeing a human, it may be better not to actively try to persuade a person to see a human. The reasoning behind this is based on social judgment theory [339], which posits that any recommendation made to a person who is against a suggestion will likely fail and make the user more opposed to the idea. Besides considering people's willingness to contact a human and the severity of the situation, it may also be useful to use a different formulation for the recommendations. For example, more emphasis could be placed on persuading people using, for instance, testimonials [114]. Alternatively, one could explain how the recommendation is in line with a user's values [114], or formulate utterances more carefully, as suggestions and less as commands [135].

PERSONAL CHARACTERISTICS

Besides people's self-identity and TTM-stages, we also looked at the effects that their smoking frequency, weekly exercise amount, age, education level, or personality may have (Figure 2.4). While previous work suggests a relationship between socioeconomic status and smoking [127, 162] and physical activity behaviors [334], as well as between age and smoking behaviors [205, 214], our results do not suggest a strong effect of education level and age on people's activity efforts and ratings of the interaction scenarios (Figure 2.4). What we did observe is a moderate correlation of 0.41 between smoking frequency and willingness to discuss repeated failures to reach physical activity goals (Table 2.1). One

explanation for this observation could be that as heavier smokers are more likely to seek help in quitting smoking [87], they may be more open to receiving support. In addition, we find several small to moderate correlations between people's Big-Five personality dimensions and their activity efforts and interaction scenario ratings (Figure 2.4). For example, there is a small correlation of 0.20 between conscientiousness and activity effort (Table 2.1). This matches the observation that the code "conscientiousness" appeared several times in relation to the preparatory activities (e.g., P320, P380):

I approach the activities with clear focus and intense dedication and discipline
(☞ Activity experience, P521)

Moreover, extraversion correlated with people's willingness to involve both their social environment and GP as shown in Table 2.1. This is in line with the observation that, especially with regards to involving their social environment, several participants expressed their desire for privacy (Figure 2.2). While some participants wanted privacy in general (e.g., P495, P573), others were concerned explicitly with not wanting to be seen failing:

I would prefer to keep it to myself until im confident that i can kick this habit for good (☞ Tell SE about quit attempt, P111)

... I have done that [i.e., tell my social environment about my quit attempt] before and felt pressure to stop. Then you feel like a failure if you dont succeed
(☞ Tell SE about quit attempt, P79)

Overall, however, the effects of personality were relatively small. This is in line with previous work, which found primarily small effects of personality on physical activity [322, 323] and smoking behaviors [258, 358].

Recommendations for addressing users' personal characteristics. Our results are indicative of personal characteristics playing a role in people's preferences for using an eHealth application. The consequence is that it may be harder to convince some participants of the merits of certain behaviors, such as involving their social environment. However, our observed effects were primarily small, and the topics "want for privacy" and "conscientiousness" appeared relatively infrequently in the free-text responses (Figure A.9). Therefore, rather than directly tailoring application components to people's characteristics, the most straightforward approach may be to simply acknowledge that people differ in their preferences and leave room for these differences. Interestingly, we find that, even when not explicitly asked to, people tend to follow their preferences. For example, several people modified a preparatory activity in which they were asked to visualize smoking or becoming more physically active as a fighting match. Instead of a fighting match, they imagined a bike race (P104), a soccer match (P539), or a verbal fight (P87). While not all people provided reasons for doing so, P539, for instance, mentioned imagining a soccer match because of being a soccer fan.

OTHER PARTY

Now we zoom out one more unit to the other party involved in an interaction. This includes the virtual coach, social environment, and the GP, as well as people featured in educational videos that were part of some preparatory activities.

Companionableness. Appearing in 1.82% of the 4839 responses, companionableness was the most frequently mentioned topic relating to other parties involved in a behavior (Figure 2.2). To participants, this included having another person or friend to talk to (e.g., P260, P645), not feeling alone (e.g., P442, P513), finding the other party condescending or patronizing (e.g., P274, P423), feeling close or connected to the other party (e.g., P89, P95, P476), being able to relate to the other party (e.g., P68), being able to share accomplishments (e.g., P528, P635), and feeling supported and not pressured (e.g., P416, P547). Support from others was frequently mentioned as a motivator for doing the activities (Figure 2.2), and companionableness of the other party appeared to be important for support to be helpful:

Having an extra support is really helpful and if it comes through an important person it has more impact in my choices (☑ Discuss with an SO how they can support the quit attempt, P593).

Depends on what happens after I press the button. Will the ai try to act as my friend and scold me? (☑ Help button for smoking HRSs, P263)

Further support to this is given by the observation provided in Table 2.1 that participants with a larger household, and thus likely a larger and closer social environment, were more willing to involve their social environment.

Previous work has described the importance of companionableness, such as being able to count on a social robot [104], trust one's primary healthcare providers as well as the received health-related information in the context of an eHealth application for cardiovascular disease and dementia [7], and feel supported in eHealth applications for both long-term weight maintenance [40] and knee osteoarthritis [267]. The pilot study for the smoking cessation intervention txt2stop also found that people disliked messages that were seen as patronizing [135]. Similar recommendations were formulated by Michie et al. [243] for the internet-based smoking cessation intervention StopAdvisor. Moreover, Henkemans et al. [159] showed in the context of a robot playing a self-management education game with children with type 1 diabetes that the children answered more questions correctly and perceived the interactions as more pleasurable when the robot was designed to account for the children's needs for, among others, relatedness.

Nature: Human vs. AI. Importantly, (lack of) companionableness was ascribed to both the virtual coach and humans. However, the nature of the other party was referred to by some people. For example, some people were entirely against using a virtual coach:

If it was an actual person, I could probably consult, but being a virtual coach, I would not be as interested. (☑ Reflect on PA HRSs on Sundays, P573)

I would never consult a virtual coach (☑ Reflect on smoking HRSs in the evenings, P330)

However, more commonly, people expected certain characteristics or abilities to (not) be present in a virtual coach compared to a human. One such characteristic was situational awareness, or the ability to understand and tailor to the individual user and their situation:

... I'd love to see tips, but the reason why is much deeper and requires a human. (☑ Reflect on PA HRSs on Sundays, P452)

Other aspects mentioned by participants with regards to the nature of the other party included a lack of empathy (e.g., P630) and finding a virtual coach less motivating than a human:

I don't think I'd consult any virtual coach on this issue, as I find real people to be much more motivating. (☒ Discuss repeated failure of reaching PA goals, P416)

People who did like the virtual nature of the coach, on the other hand, mentioned fearing or not liking to contact real people (P222, P225) as well as feeling less embarrassed to admit failures to a virtual coach than to a human (P269).

Several of these factors have also been found to play a role in previous work. Issom et al. [173], for instance, saw that the empathy expressed by a conversational agent was particularly appreciated, and de Graaf et al. [104] observed that the expression of human-like emotions by a social robot was perceived as important. Expressing empathy has also been shown to help a virtual agent to form and maintain a relationship with a user [57], which can support behavior change [400]. In addition, Issom et al. [173] observed that the anonymous nature of conversing with a conversational agent was valued. Regarding situational awareness, both de Graaf et al. [104] and Issom et al. [173] obtained similar findings in that participants preferred the social robot to understand more than just pre-programmed commands and that users of conversational agents requested more flexible answer choices. Notably, there can also be differences between AI embodiments. For instance, Sinoo et al. [345] saw that children's feelings of friendship were stronger toward a physical robot than an avatar.

Accountability. The fourth most frequent topic with regards to the "other party" unit, after companionableness, nature, and situational awareness, was with 0.60% accountability (Figure 2.2). People felt accountable to the virtual coach (e.g., P31, P466) and their social environment (e.g., P475, P638), although accountability was perceived as stronger when coming from humans:

With my experience with Sam I realise that a virtual assistant can really help and I also think about the fact that if I fail at some point it's "more ok" to let down a fake person than someone real that I care about a lot. (☒ Reflect on smoking HRSs in the evenings, P593)

As a result, accountability was sometimes seen as too strong when coming from the social environment and too weak when coming from the virtual coach:

Adding peer-pressure to an already stressful situation would not be useful (☒ Tell SE about quit attempt, P264)

It's easy to dismiss a virtual coach, maybe it works for people who are very committed to quit and would be a reminder (☒ Plan for smoking HRSs in the mornings, P273)

The importance of feeling accountable to somebody comes forward in work by Nelligan et al. [267], who found accountability to be part of the primary themes describing attitudes

and experiences in the context of an eHealth application. The relevance of accountability was also observed by Lie et al. [219], who saw that individuals with type 2 diabetes felt more accountable to regular health consultations than virtual ones.

Recommendations for the other party. Our results show that companionableness is a key ingredient in interactions with another party. While addressing the perceived companionableness of the social environment or GP may lie outside the reach of an eHealth application, improving the one of a virtual coach does not. As a start, it is important to be aware that it is not only possible for people to form a relationship with a machine or computer [101], but that people also tend to treat their computers as social beings [263]. Relatively simple strategies can help to improve such a relationship with a system. This includes giving a virtual coach a name to increase its social presence [400], avoiding repetitiveness and predictability to improve engagement, enjoyment of the interactions, and motivation to perform an advocated behavior [57, 101], and trying to avoid responses that may be seen as too enthusiastic [135]. Other aspects such as learning from individual conversations, building on and referring to previous conversations, and conveying in-depth information on various topics as humans commonly do, however, remain open challenges [101]. Nevertheless, paying close attention to the relationship between a user and a virtual coach is likely to pay off, as a good relationship can support behavior change [400]. Notably, however, improving the relationship between a user and a virtual coach should not come at the expense of transparency: the user needs to be aware that they are interacting with a virtual coach and not a human [228].

ENVIRONMENT

DIFFICULTY OF INTEGRATING (HEALTH) BEHAVIORS INTO PEOPLE'S BUSY LIVES

People are busy with other things. The most frequently mentioned topic with regards to the environment was with 5.04% whether participants had enough time to perform a behavior, and especially to complete their preparatory activities (Figure 2.2). Notably, 18.80% of barriers to completing the activities involved the availability of time, making it the most commonly mentioned barrier (Figure 2.2). It turned out that participants tended to be busy with their daily lives, including work, child care, and daily chores, and that these tasks left no time (e.g., P111, P432), caused people to be too tired (e.g., P66, P140, P262), or made it difficult for people to focus on their preparatory activities (e.g., P495, P600, P642):

I run out of time with home life taking a focus so didnt get time to complete
(☞ Activity experience, P432)

It was hard to plan the exercise in my daily planning, as the days are full and i'm very tired at the end of the day (☞ Activity experience, P262)

I thought about it [i.e., the person I would like to be once I have successfully quit smoking] for a couple minutes with the intention to write it down, but got distracted with other things. It did remind me that I did want to quit though.
(☞ Activity experience, P642)

When participants had no time to do their activities, they reported having spent their time on other priorities such as home life in the example of P432. A similar phenomenon

was found in Lie et al. [219] where people dropped out from an eHealth intervention for self-management of type 2 diabetes because the daily life took the front stage. Similarly, a scoping review by Wilson et al. [391] saw that the inability to incorporate a behavior into one's routine was a barrier to using eHealth tools. In addition, the statement by P262 (i.e., being too tired to do a behavior) suggests that the participant expected the behavior to require a considerable amount of effort, which was also one of our codes for the "behavior" unit (Figure 2.2). The expectation of effort or effort expectancy is one of the predictors of the intention to perform a behavior according to the UTAUT [376], and several studies have confirmed this relationship in the context of eHealth tools (e.g., [62, 130, 312]).

Prompts and triggers are helpful. The second most frequent topic for "environment" was with 3.49% of all responses whether participants had prompts or triggers for doing a behavior, and especially to complete a preparatory activity (Figure 2.2). Participants commonly reported that they forgot to do (part of) their preparatory activity (e.g., P269, P331). Reasons included being busy (e.g., P43, P527), and that some activities (e.g., tracking one's smoking behavior) required one to remember to do something at specific times (e.g., P182, P227):

I approached the activity with a positive thought but found myself forgetting to record the time i had a cigarette. This was mainly due to smoking when i had opportunity for a quick break so was always rushing. (☞ Activity experience, P182)

Participants also commonly started an activity but then stopped to do something else and forgot to get back to it:

I watched 1 minute but I started different activity and forgotten about it. (☞ Activity experience, P281)

The importance of prompts and triggers is further illustrated by the fact that several participants who did complete their activities reported making use of them. Participants mentioned completing their activity right after the session had ended (e.g., P186, P480), once they received the reminder message we sent after the session on Prolific (e.g., P180, P393, P417), or based on a reminder they had set themselves (e.g., P7, P346):

I set a remind on my smartphone to recall me the activity, so yesterday, in my bed before sleeping, I thought about who I want to be once I have quit smoking (☞ Activity experience, P346)

A scoping review by Wilson et al. [391] found a lack of reminders or alerts to be a barrier for older individuals to use eHealth tools and their presence to be a motivator. Similarly, Horsch et al. [166] saw that people favored the use of reminders to help with forgetting in the context of insomnia treatment. Participants of the study by Horsch et al. [166] also emphasized, however, that users themselves should set reminders. This links to the earlier discussed topic of autonomy in that users want to be in control of application components.

Recommendation for addressing the difficulty of integrating (health) behaviors into people's busy lives. According to the COM-B model of behavior [242], one predictor

of behavior is whether people have the opportunity to perform the behavior. This includes sufficient time and prompts or triggers to remind them. Our results suggest that both of these factors tend to be lacking. A straightforward way to help people who lack time is to create action plans. Action plans specify where, when, and how one plans to do something to create a link between a cue and a behavioral response [153]. Action plans have been effective in changing behaviors such as physical activity, smoking, and alcohol consumption [153, 346]. To further account for sudden barriers to doing a behavior such as being tired, one can specify coping plans to create a link between a possible risk situation and a feasible way of coping with it [347].

Another strategy could be to prompt participants to create reminders themselves or determine suitable times for sending automatic reminders [167]. This was actually recommended by one participant (☞ Activity experience, P247). It turns out that reminders are already one of the most common persuasive components of eHealth applications [213]. However, it is important to keep in mind that a high effort expectancy and other more relevant priorities likely also play a role for somebody who is too tired or has no time to do a behavior. This shows that characteristics of the environment can be intertwined with ones of the behavior and the user.

Lastly, the topics "having enough time" and "having prompts or triggers" primarily appeared in participants' statements about their actual behavior (i.e., their activity experiences and barriers) rather than their views on possible behaviors in the form of the interaction scenarios (Figure 2.2). Thus, these factors are less evident to people when they are just asked about their views on possible behaviors. One likely explanation for this is the optimism bias, according to which people tend to be overly optimistic about themselves and their future [385]. For example, the study of Horsch et al. [166] showed that people tended to be rather optimistic about their future adherence to an eHealth application for insomnia treatment. Reasons may include relying too much on future willpower and ignoring things that could go wrong [166]. This underlines the importance of having participants interact with a system to get a thorough assessment of their needs [79]. On the other hand, other topics, namely ones related to the other party, primarily appeared in the views on the interaction scenarios rather than statements about actual behavior (Figure 2.2). Thus, combining data on actual and potential behaviors offers a clear benefit.

HELPFULNESS OF SUPPORT FROM SOCIAL ENVIRONMENT

We have already touched upon the role of the social environment in the context of characteristics of the other party that influence people's views on interaction scenarios. However, our results also suggest the general importance of support from one's social environment. In fact, 6.60% of participants reported support from their social environment to be a motivator for doing their preparatory activities (Figure 2.2). The social environment supported these participants in their wish to reach their goals such as quitting smoking (e.g., P212, P475) and helped them to complete the activities:

Wrote down on a list last night and discussed with partner. Helpful and motivating.
 (☞ Activity experience, P133)

An important form of support is not just verbal but also behavioral. Some participants felt less motivated because their social environment did not live according to their own

behavioral goal of not smoking (e.g., P127, P207, P315), or because they did not feel part of a group that worked toward the same behavioral goal (e.g., P262):

Probably the people watching the ad will think about the fact that usually smokers need motivation from other people, but often the motivation is not there because smokers tend to surround themselves with other smokers. (☒ Receive motivational messages, P207)

... As well as the motivation of other participants. When feeling you are a part of a group that want to achieve the same thing i feel that this is more motivated. (☒ Barriers, P262)

Previous work by Willemsen et al. [389] found social pressure from the social environment of Dutch employees to be a predictor of intention to quit smoking, and Breeman et al. [66] concluded that involving the social environment was a desired core attribute for an eHealth application for healthy living. In addition, Meijer et al. [235] note that the support and social norms present in social environments can shape identities, which in turn can affect behavior as discussed previously. For example, according to the social identity model of recovery [55], a person's recovery identity in the context of addiction can become stronger if it is shared with other people who favor recovery.

Recommendations to address the helpfulness of support from one's social environment. We find that support from one's social environment can be motivating and can help to perform activities that are part of a behavior change intervention. A straightforward way to include social support in an eHealth intervention is to prompt participants to either tell their social environment about their behavior change process or discuss with a significant other about how they can support it. Both of these elements were generally seen positively by our participants (Figure 2.3). However, it is important to keep in mind the earlier discussed personal characteristics such as a want for privacy and characteristics of the social environment that may influence whether people want to involve their social environment. Taking a similar approach to Tielman et al. [363] and taking people's willingness to involve a human and situation severity into account may be beneficial. Another approach may be to connect people who work toward the same behavioral goal. Promising results can be obtained with relatively simple solutions such as a WhatsApp group [342]. Such a group has the advantage that it is easy to implement and accessible due to the omnipresence of WhatsApp [342].

DIVERSITY OF OTHER ENVIRONMENTAL FACTORS

Several other environmental factors such as not having access to the Internet on certain days (P233), not using one's phone when at home or on specific days (e.g., P376, P624), not being able to access the videos contained in the activities due to one's location (P668), poor weather (P468), or restrictions related to the COVID pandemic (e.g., P4, P631) were mentioned. Given the diversity of environmental factors that can play a role, it is likely difficult, if not infeasible, to specifically address all of them. Since such other environmental factors appeared relatively infrequently (Figure A.9), the most important insight may thus be to design an eHealth application in such a way that it leaves room for individual barriers and gives users resources to try to cope with these barriers themselves.

DISCUSSION

Through a thematic analysis based on qualitative data, quantitative data, and literature, we have discussed 14 main themes that are associated with people's actual behavior and views on potential interactions in the context of a virtual coach for quitting smoking and becoming more physically active. These 14 themes can be structured by assigning them to four hierarchical units of analysis. These units are "behavior," "user" (who performs a behavior), "other party" (involved in a behavior), and "environment." While these units provide a convenient frame for analysis, it is important to note that the observed themes often span multiple units or depend on themes in other units. For instance, the environment, user, and behavior are involved in observing that people are often too tired from their busy daily lives to perform a behavior. This is in line with previous work that has highlighted the interdependence of factors from the environment, user, and the technology [373].

Aligning time, perceived usefulness, and users' goals. The most common topics for the "environment" unit were having enough time and prompts or triggers for doing something. One could address this by making suggestions at convenient times (e.g., [167]) or helping people create action plans [153]. Yet, whether people have time for something and remember to do it likely depends on how useful they find it. For instance, some participants set their own reminders for doing preparatory activities. Those people likely perceived the activities as so useful that they wanted to ensure they did them. Perceived usefulness was the most common theme in participants' responses, but it is also connected to another topic. More precisely, somebody who does not see the link between a behavior and their goals is likely to find the behavior less useful. Recall that the motivation to reach their goals was the most common motivator for people to do their preparatory activities. Those who mentioned this motivator likely saw how the preparatory activities aligned with their goals. Ways to strengthen the link between perceived usefulness and users' goals include referring to people's goals and beliefs when giving advice [2]. Another approach, which motivated several participants, is to think about one's feared or desired future self with regards to a behavior [236, 294]. Given that these topics of having enough time, perceived usefulness, and users' goals appear to be connected, addressing them together may be beneficial.

We have formulated recommendations for how to address each observed theme as part of our analysis. Besides these recommendations, we find that the following challenges would benefit from more attention in the future:

Is there a set of standard factors that are generalizable across domains? Our study was conducted in the context of smoking cessation and physical activity increase and with participants that were enrolled on the online crowdsourcing platform Prolific and hence had at least some experience with using digital services. Thus, the question regarding the generalizability of our findings to other types of behavior change, or even eHealth more generally, and users with less experience with digital services arises. In addition, as our participants were contemplating or preparing to quit smoking and were paid for completing the conversations with the virtual coach, it is not clear how our findings generalize to a setting with people who are not yet contemplating to change and do not receive such a payment. We think that there are two important steps. The first one involves carefully describing the study context to more easily compare it across studies. While progress has been made for systematically reporting behavior change techniques [244], others

elements of the context such as the environment and virtual coach are often described less extensively and without clear guidelines. One useful direction in this regard is the work by Fitrianie et al. [129], which aims to create a questionnaire that reports characteristics of virtual agents in a standardized fashion. A second step is to examine which factors affect user needs and how these factors change as the study context varies. For instance, while several of our themes coincided with earlier studies on eHealth applications for other domains such as self-management of chronic conditions [219, 348], more research is needed to determine how findings can be generalized. Other important characteristics of the study context include the embodiment of and way of interacting with an AI (e.g., [148, 345]) as well as whether the intervention is blended (e.g., [237]).

How to get input from (many) users? Some themes, such as lacking prompts and triggers or having no time to perform a behavior, primarily appeared in people's descriptions of their actual behavior as opposed to their views on possible interaction scenarios presented in videos. This was the case even though people provided their views on the interaction scenarios after having experienced actual interactions in the form of doing suggested preparatory activities. The disparate theme distribution is likely due to the optimism bias, according to which people tend to be overly optimistic about their future [385]. This underlines the need for carefully choosing a method and using multiple ones if a comprehensive understanding of users' needs is desired. As done in our study, crowdsourcing can facilitate reaching a large and diverse number of people once one has chosen a method. Yet, it is very time-consuming to conduct a thematic analysis of many free-text responses manually. Thus, crowdsourcing needs to be supplemented with automatically extracting codes from text to allow large-scale thematic analyses to be adopted more widely. Promising results for identifying predefined codes exist in specific application areas such as cognitive therapy [72] or news [274]. However, as novel codes may appear in the data, such approaches need to be combined with ones for generating codes. First approaches exist (e.g., [211, 220]), but challenges remain regarding explainability and trust, among others [91]. While those challenges persist, automated approaches may be beneficial as an adjunct to qualitative analysis by informing the creation of codes, checking the accuracy of coding, or pointing out ambiguity [91, 211].

How to tailor application components? In several of our themes, there was an apparent need for tailoring application components to individuals or groups of users. For example, while some people favored involving their social environment, others expressed a need for privacy. However, it is not clear how such tailoring can best be accomplished. Options include letting users choose themselves to support autonomy and competence [167], automatically tailoring to users to account for people's lack of self-knowledge [167], and helping users to self-experiment [186]. Moreover, each of these general approaches can be implemented in many ways, and they can also be combined. Ranjbartabar et al. [317], for instance, employed users' preferences as a starting point for subsequent automatic tailoring. Further research is needed to determine which approaches to tailoring are effective in increasing adherence and under which conditions. One promising way to test many tailoring approaches and configurations may be micro-randomized trials, which allow participants to be randomized hundreds of times during a single study [196].

CONCLUSIONS

In light of the dropout and lack of adherence common to eHealth application for behavior change, we need a better understanding of user needs and how to address them. We thus conducted a thematic analysis of people's experiences with actual and views on potential behaviors. The context was a text-based virtual coach for quitting smoking and becoming more physically active.

We found that users' needs are often interconnected and include characteristics of the behavior, the user, other parties such as the social environment, and the environment. We identified 14 main themes that describe users' needs: of these, the perceived usefulness of behaviors is most prominent and relates to environmental characteristics such as having sufficient time and the user's state such as their motivation to reach their goals. We publish our dataset with user characteristics and 5074 free-text responses from 671 people to aid future work on understanding the interplay between users' needs and characteristics. This dataset can also be used to improve preparatory activities for quitting smoking and becoming more physically active, as it contains 2866 descriptions of experiences with 24 such activities.


Based on this analysis, we formulated recommendations for how users' needs can be addressed in eHealth applications for behavior change. Besides the specific recommendations we provide for each need, we suggest that associated needs should be addressed together. Adherence could, for example, be strengthened by referring to users' goals and their beliefs when giving advice on quitting smoking and increasing physical activity.

3

3

ADDRESSING PEOPLE'S CURRENT AND FUTURE STATES IN A REINFORCEMENT LEARNING ALGORITHM FOR PERSUADING TO QUIT SMOKING AND TO BE PHYSICALLY ACTIVE

Behavior change applications often assign their users activities such as tracking the number of smoked cigarettes or planning a running route. To help a user complete these activities, an application can persuade them in many ways. For example, it may help the user create a plan or mention the experience of peers. Intuitively, the application should thereby pick the message that is most likely to be motivating. In the simplest case, this could be the message that has been most effective in the past. However, one could consider several other elements in an algorithm to choose a message. Possible elements include the user's current state (e.g., self-efficacy), the user's future state after reading a message, and the user's similarity to the users on which data has been gathered. To test the added value of subsequently incorporating these elements into an algorithm that selects persuasive messages, we conducted an experiment in which more than 500 people in four conditions interacted with a text-based virtual coach. The experiment consisted of five sessions, in each of which participants were suggested a preparatory activity for quitting smoking or increasing physical activity together with a persuasive message. Our findings suggest that adding more elements to the algorithm is effective, especially in later sessions and for people who thought the activities were useful. Moreover, while we found

 **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Addressing people's current and future states in a reinforcement learning algorithm for persuading to quit smoking and to be physically active. *PLoS ONE*, 17(12):e0277295, 2022. doi: 10.1371/journal.pone.0277295.

some support for transferring knowledge between the two activity types, there was rather low agreement between the optimal policies computed separately for the two activity types. This suggests limited policy generalizability between activities for quitting smoking and those for increasing physical activity. We see our results as supporting the idea of constructing more complex persuasion algorithms. Our dataset on 2,366 persuasive messages sent to 671 people is published together with this article for researchers to build on our algorithm.

INTRODUCTION

3

Imagine a woman called Janine who wants to motivate her friend Martha to become more physically active. Janine could motivate Martha to go for a run because that has worked for her other friends. However, this likely only works if Martha has running shoes. If she does not have any, just asking Martha to go for a walk might be more successful. So the success of the motivation may depend on the state Martha is currently in. In addition, if Janine cares about the overall success of all her motivational attempts, she should probably begin by telling Martha how to buy running shoes. This may cause Martha not to work out this week, but future attempts to motivate her to work out are much more likely to be successful once Martha has running shoes. So Janine should also consider the future states of Martha. And, lastly, people differ in whether they prefer to walk or run, no matter if they have running shoes. So Janine should also consider what type of person Martha is. Since Janine is not always available to motivate Martha, we want to create a virtual coach. Can this virtual coach do what Janine does?

Changing behavior such as becoming more physically active is crucial to improving health and reducing premature death. For example, 40% of deaths in the United States are brought about by unhealthy behavior [253, 336]. In addition, changing one behavior can make changing another one easier. For instance, becoming more physically active may facilitate quitting smoking [152, 368] and vice versa [291]. However, while many people want to change their behavior, doing so without help can be difficult. For example, more than two-thirds of adult smokers in the United States want to quit smoking [44], but most unassisted quit attempts fail [99]. One promising way to support people in changing their behavior are eHealth applications [290], which provide elements of healthcare over the Internet or connected technologies such as apps and text messaging. However, while such applications can be easy to use, available at all times, scalable, cost-effective, and can facilitate tailoring [218], adherence to them remains low [56, 191]. Adherence refers to whether and how thoroughly people do the activities suggested by the application.

We, therefore, aim to develop persuasion algorithms that successfully encourage people to adhere to their behavior change intervention. A one-size-fits-all approach to persuasion is unlikely to be effective [53, 128], as behavior change theories [240, 307] suggest many factors that affect personal behavior. However, these factors can be used as a starting point for designing algorithm-driven persuasion. Algorithm-driven persuasion is persuasion that is determined by programming code, with the advantage that it can use behavioral user data, target individuals or groups, and be adaptive [328]. Previous work on persuasion algorithms has shown that one can use data gathered on other people [183, 184], similar people [106, 165] or a single individual [179, 183, 184, 249, 331, 383] to choose a persuasion type (e.g., advice from peers vs. experts). However, it is essential also to consider the context of a persuasive attempt [34, 277, 397]. One way to define a context is by describing

the current state a persuadee is in. For example, Bertolotti et al. [54] show that the success of different messages to reduce red meat consumption depends on the persuadee's self-efficacy. In addition, persuasion types depend not only on the persuadee's state for their success, but they in turn also influence the persuadee's state for future persuasive attempts. For instance, messages for quitting smoking differ in their impact on self-efficacy [351]. Thus, if we want to maximize the effectiveness of persuasive attempts over time, we need to consider both current and future states.

One framework that allows us to formulate an adaptive and data-driven algorithm that considers both current and future states is Reinforcement Learning (RL). There are first results for applying RL with consideration of people's states to adapting the framing of messages for inducing healthy nutritional habits [81] or the affective behavior of a social robot teacher for children [145]. In our approach, we investigate whether states are also helpful in persuading people to do preparatory activities for quitting smoking, such as writing down and ranking reasons for quitting smoking. In addition, we go a step further by also taking the similarity of people into account. The reason is that previous work has shown that characteristics such as the stage of behavior change [106] and personality [32, 106, 154, 182, 288, 359, 360, 398] affect the effectiveness of different persuasion types. The result is a personalized RL algorithm for choosing persuasive messages.

To systematically assess the value of subsequently adding the consideration of states, future states, and the similarity of people, we conducted a longitudinal experiment. Since the effects of these algorithm elements are difficult to assess in a complete behavior change intervention in which many other components such as goal-setting and progress feedback can play a role (e.g., see Brinkman et al. [69] in the context of usability testing), we created a minimal intervention in which people were only coached to *prepare* for changing their behavior. In this intervention, a conversational agent served as a virtual coach that suggested and persuaded people to do preparatory activities for quitting smoking. Since becoming more physically active may facilitate quitting smoking [152, 368], half of the activities addressed preparing for increasing physical activity.

HYPOTHESES

The objective of this study was to test a personalized RL approach to persuading people to do preparatory activities for quitting smoking and increasing physical activity. The complete algorithm considers a person's state, future states, and the similarity of people when choosing a persuasion type. The goal thereby is that people do their activities more thoroughly, which is supposed to facilitate quitting smoking (Figure 3.1). Therefore, our first hypothesis is that subsequently incorporating the elements of the personalized RL algorithm is more effective with respect to how thoroughly people do their activities. Furthermore, our algorithm does not distinguish between preparatory activities for quitting smoking and ones for increasing physical activity because both types of activities serve the same behavioral goal of quitting smoking. This leads to our second hypothesis, which is that the best persuasion strategy is similar if we use data collected on both types of activities compared to using data collected on solely one type of activity. We now motivate each hypothesis in turn.




Figure 3.1: **Overarching goal of our work.** The goal of our persuasion algorithm is that people do their preparatory activities more thoroughly, which is supposed to facilitate quitting smoking.

3

H1: ALGORITHM EFFECTIVENESS

In the introductory example, Martha wanted to become physically active. In the simplest case, the virtual coach could send her the persuasion type that has led people to do their activities most thoroughly in the past. As a measure for this thoroughness, the virtual coach could use the self-reported effort people put into assigned activities. Of course, if no such data is available yet, the virtual coach would need to choose either randomly or based on other sources of information such as experts due to the cold-start problem. Assuming such data is available and the effort for three persuasion types is as shown in Figure 3.2A, for instance, the virtual coach would choose persuasion type $P2$. However, intuitively, Martha’s reaction to persuasive attempts might differ based on the state she is currently in. For instance, if she has no running shoes, just motivating her to go for a walk ($P1$) might be better than motivating her to go running (Figure 3.2B). Previous work has posited the importance of considering the context of a persuasive attempt when striving to create effective persuasion [34, 277, 397], for example, by defining the current state of the persuadee. This should be done so that knowing the persuadee’s state allows one to predict the effectiveness of different persuasive messages. One such characteristic of a persuadee’s state is the presence of barriers, such as Martha’s lack of running shoes. Alfaifi et al. [31], for example, distinguish health, environmental, psychological, personal, and social barriers. Another potential state feature is self-efficacy, as it influences which health messages are more effective [54]. Moreover, how a person processes messages changes based on their mood [60, 131]. We, therefore, posit that choosing a message based on a persuadee’s state is more effective than choosing the overall most effective message.

The effectiveness of persuasive attempts might depend on the persuadee’s state, but a persuasive attempt in turn might also affect the state and thus the effectiveness of future persuasive attempts. Sending instructions on buying running shoes when Martha does not have any, for example, may cause Martha to buy some and thus remove the corresponding barrier. Future persuasive attempts that aim to increase Martha’s motivation to run may then be more successful. Thus, even though informing Martha about buying running shoes with $P3$ may lead to less effort at the current time step than motivating her to go for a walk with $P1$, the former may allow the virtual coach to more successfully motivate Martha in the future (Figure 3.2C). To estimate the overall effectiveness of $P3$, we can compute the discounted sum or Q-value of the efforts after $P3$ at the current time step and the most effective persuasion type at the next time step. Discounting thereby means that we give a lower weight to efforts in the more distant future due to the importance of initial small wins [35]. In the example in Figure 3.2C, the discount factor is set to 0.85, and the discounted sum or Q-value is with 11.8 the largest if we choose $P3$ at the current step. Regarding the previously mentioned state features from the literature, Steward et al. [351]

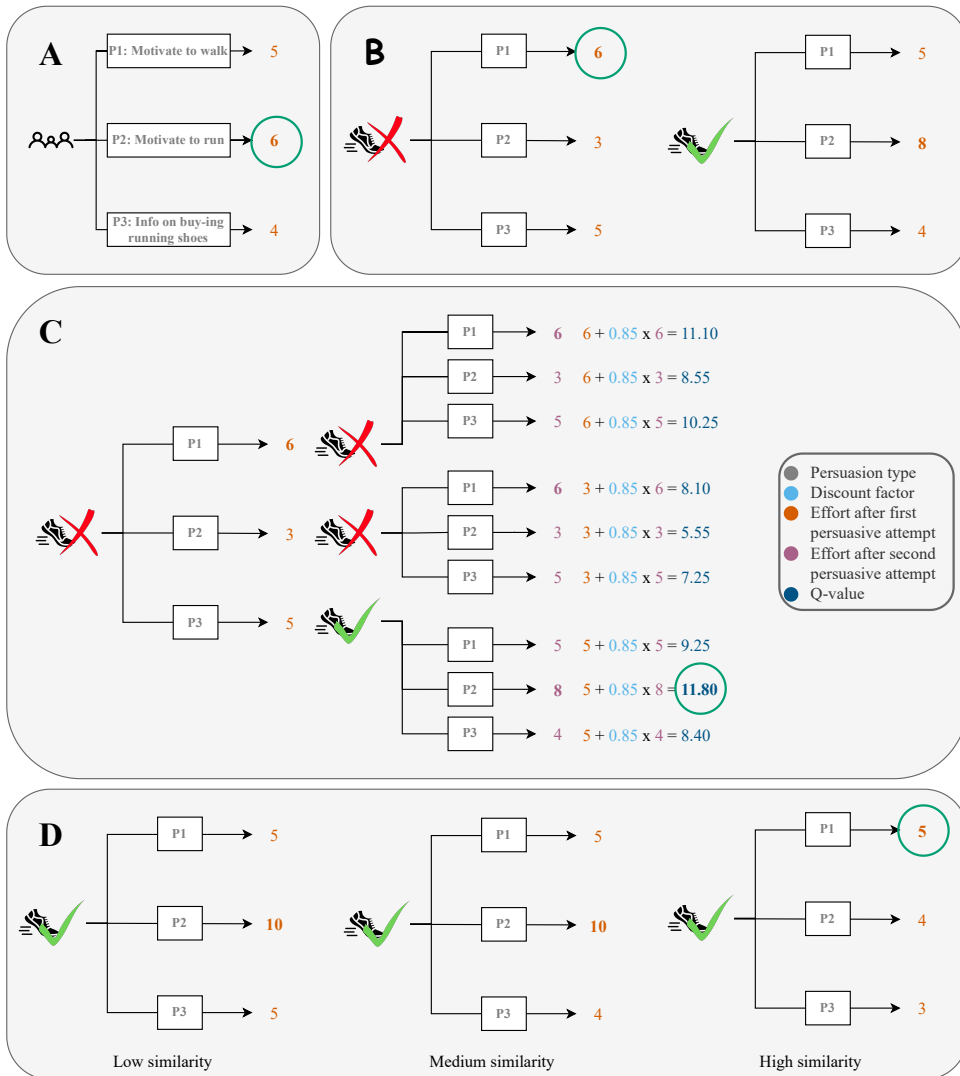


Figure 3.2: **Illustration of the algorithm components.** Illustration of our proposed algorithm components. To the baseline of sending the most effective persuasion type (A) we add the consideration of states (B), next states (C), and the similarity of people (D). Circles indicate the most effective persuasion type for the example person Martha described in the text.

found that differently framed messages vary in their impact on self-efficacy. Given that self-efficacy determines how effective different health messages are [54], a message choice at this time point thus determines the effectiveness of messages in the future. Besides self-efficacy, the type of message might also affect a person’s intention to act, anticipated regret, and attitude toward behavior [80]. We thus hypothesize that selecting a message based on both the present and the future states of a persuadee is altogether more effective

than considering only the persuadee's present state and choosing the overall most effective message.

A person's state can change frequently, so we need to infer it each time we make a persuasive attempt. However, there are also relevant characteristics of a person which change, if at all, very slowly. For example, the impact of message types on self-efficacy depends on a person's need for cognition [351]. Other variables that may affect the success of different messages include the stage of behavior change [106], personality [32, 106, 154, 182, 288, 359, 360, 398], age and gender [182, 257], cultural background [289], how people approach pleasure and pain [85, 221], self-construal or the perceived relationship between the self and others [355], and in the context of quitting smoking the experience with previous quit attempts [351]. Thus, we suppose that people who are more similar concerning such characteristics are more likely to respond similarly to persuasive attempts. When deciding how to persuade somebody, we thus want to weigh the data observed from other people based on how similar they are to the person at hand. For example, we may find that for people like Martha, it is better to motivate them to go for a walk ($P1$) than to go for a run ($P2$) once they have running shoes (Figure 3.2D). We hence posit that considering a persuadee's similarity to other people besides their current and future states when choosing a persuasive message is more effective than not taking the similarity to other people into account. Overall, we thus hypothesize the following:

H1: Subsequently incorporating 1) states, 2) the consideration of future states, and 3) the weighting of samples based on the similarity of people into an algorithm that selects the best persuasive message type is more effective than not incorporating the respective element.

H2: SIMILARITY OF OPTIMAL PERSUASION STRATEGIES

Previous work on persuasion algorithms claims the need for considering the domain. For example, Alslaity and Tran [34] found that the impact of persuasion types varies between domains such as e-commerce and movie recommendations. Intuitively, it is possible to continuously split domains into sub-domains such as e-commerce for clothes and e-commerce for books. Nevertheless, this is not done by persuasion approaches such as the ones by Alslaity and Tran [34] and Kaptein et al. [183]. The underlying assumption is that there is a certain level of domain granularity at which one can meaningfully generalize from one persuasive attempt to another. We, therefore, assume that we can persuade people similarly for preparatory activities for quitting smoking and those for increasing physical activity, as they serve the same behavioral goal of quitting smoking. Thus, we hypothesize the best persuasion strategy (i.e., policy) to be similar if we use data collected on both types of activities compared to using data collected on only one type of activity, or more formally:

H2: The optimal policy is similar when learned based on a combined data set of activities for smoking cessation and increasing physical activity, and when learned based on a data set of activities for smoking cessation and on a data set of activities for increasing physical activity separately.

METHODS

To test our hypotheses stated above, we conducted a longitudinal experiment from 20 May 2021 until 30 June 2021. The Human Research Ethics Committee of Delft University of Technology granted ethical approval for the research (Letter of Approval number: 1523). Before the collection of data, the experiment was preregistered in the Open Science Framework (OSF) [14].

EXPERIMENTAL DESIGN

The experiment consisted of a prescreening to determine the eligibility of participants, a pre-questionnaire, five sessions in which a virtual coach attempted to persuade participants to do a new preparatory activity for quitting smoking or increasing physical activity, and a post-questionnaire. Participants were persuaded with a random persuasion type in the first two sessions and a persuasion type chosen by a persuasion algorithm after that.

Figure 3.3 shows the experimental design of the study. It was set up as a double-blind mixed-design study with two within-subject factors and one between-subject factor. The within-subject factors were the session in which a persuasive attempt was made (4 levels: sessions 1–4) and algorithm activeness (2 levels: off/on for sessions 1–2/3–4). The between-subject factor was the algorithm complexity used to choose a persuasion type after session 2. This factor had four levels with successively more elaborate optimization strategies. Ordered by complexity, the algorithm levels look for the highest value of either: 1) the average reward, 2) the average reward in a person's state, 3) the Q-value in a person's state, or 4) the similarity-weighted Q-value in a person's state. This means that starting from sending a persuasion type with the highest average reward, we progressively added the consideration of states, future states, and the weighting of samples based on the similarity of people. Table 3.1 provides an overview of the four complexity levels, whose components will be explained in the next section.

After session 2, we randomly assigned participants to one of the four algorithm complexity levels. Thereby, we aimed to balance the four groups regarding the potential covariates gender, Big-Five personality, stage of change for becoming physically active, and the effort participants put into their first activity. We used block randomization for the categorical gender variable, and adaptive covariate randomization for the other variables. Adaptive covariate randomization considers both previous assignments and covariates of people to balance condition assignments within covariate profiles [222]. Our approach to adaptive covariate randomization was a modification of the algorithm put forward by Xiao et al. [394].

ALGORITHM

We created a virtual coach embedded in a conversational agent that attempted to persuade people to do small activities. For each persuasive attempt, the virtual coach selected a persuasion type based on its learned policy. After two to five days, the user provided the virtual coach with feedback by reporting the effort put into their activity. The virtual coach used this feedback to update its policy. Formally, we can define our approach as a Markov Decision Process (MDP) $\langle S, A, R, T, \gamma \rangle$. The action space A thereby consisted of different persuasion types, the reward function $R : S \times A \times S \rightarrow [-1, 1]$ was determined by the self-reported effort, $T : S \times A \times S \rightarrow [0, 1]$ described the transition function, and

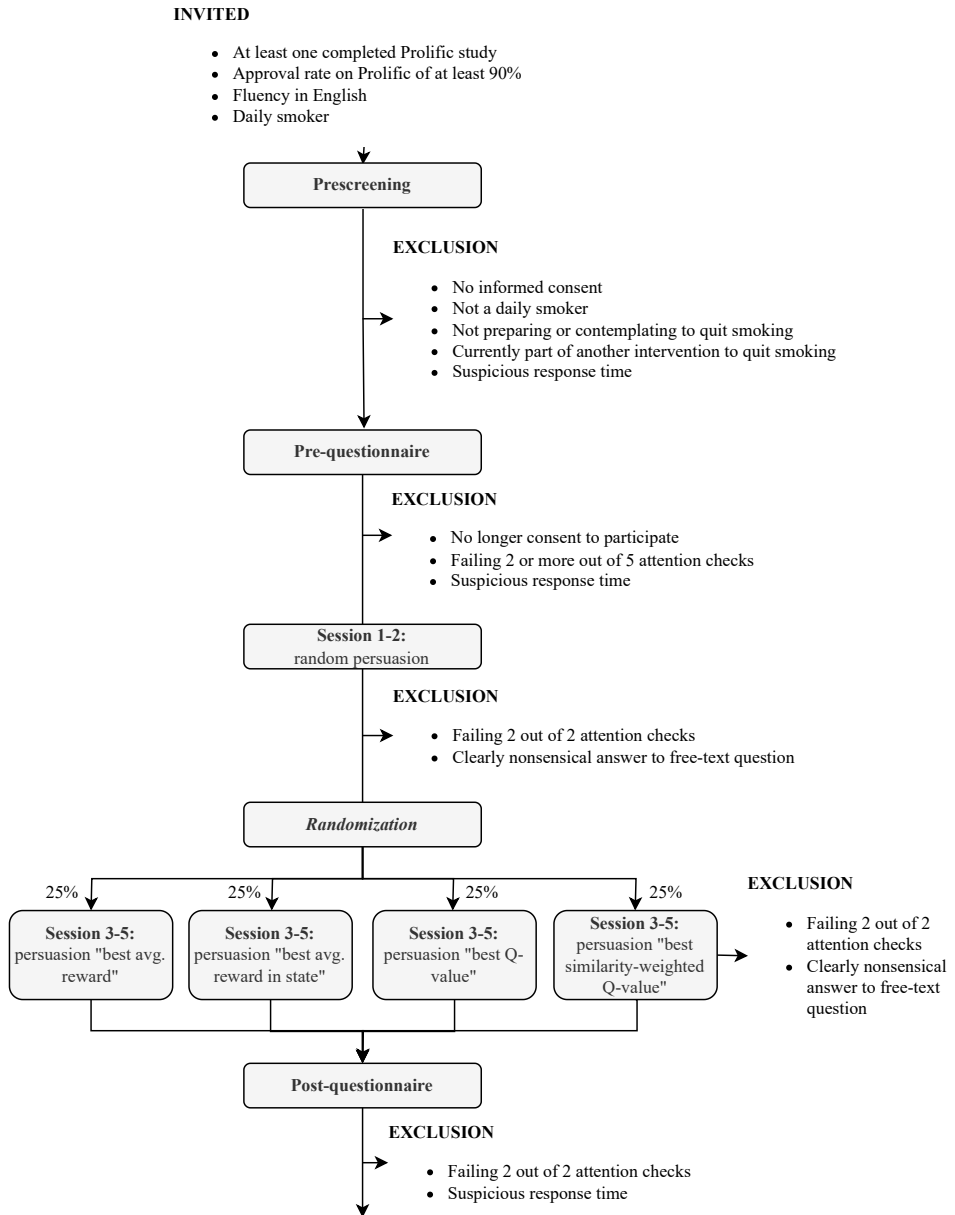


Figure 3.3: **Experimental design.** Design of the experiment, including the study components and in- and exclusion criteria for participants.

the discount factor γ was set to 0.85 to favor rewards obtained in the near future over rewards obtained in the more distant future. The intuition behind this value for γ was

Table 3.1: Chosen persuasion type for each algorithm complexity level.

Chosen persuasion type
1: BEST AVG. REWARD The persuasion type a with the overall highest average reward: $\max_{a \in A} \{R(a)\}$.
2: BEST AVG. REWARD IN STATE The persuasion type a with the highest average reward in a person's state s : $\max_{a \in A} \{R(s, a)\}$.
3: BEST Q-VALUE The persuasion type a with the highest Q-value in a person's state s : $\max_{a \in A} \{Q^*(s, a)\}$.
4: BEST SIMILARITY-WEIGHTED Q-VALUE The persuasion type a with the highest similarity-weighted Q-value in the state s of person i : $\max_{a \in A} \{Q_i^*(s, a)\}$.

Abbreviations: A , Action space; $R(a)$, Average reward for taking action a ; $R(s, a)$, Average reward for taking action a in state s ; $Q^*(s, a)$, Expected cumulative discounted reward for taking action a in state s and an optimal action in all subsequent states.

that while we wanted to persuade a user over multiple time steps successfully, a failed persuasive attempt in the near future could cause a user to become less receptive to future ones or even to drop out entirely: early success might encourage people to continue [35]. The finite state space S described the state a user was in and was captured by answers to questions about a user's capability, opportunity, and motivation to perform an activity [245]. The goal of an agent in an MDP is to learn an optimal policy $\pi^* : S \rightarrow \Pi(A)$ that maximizes the expected cumulative discounted reward $\mathbb{E}[\sum_t \gamma^t r_t]$ for acting in the given environment. The expected cumulative discounted reward for taking action a in state s and an optimal action in all subsequent states is given by the Q-value function $Q^* : S \times A \rightarrow \mathbb{R}$. To incorporate the similarity of people, the virtual coach maintained a policy π_i for each user i . When updating π_i , an observed sample from user j was weighted based on how similar i and j were. We provide an overview of the algorithm component definitions in Table 3.2. In the following, we describe each algorithm component in detail.

Table 3.2: Overview of the algorithm components and their definitions.

Algorithm component
STATE SPACE $S = \{000, 001, 010, 011, 100, 101, 110, 111\}$, using three binary features based on the COM-B model (see the "State space"-section for more information).

Table 3.2: (continued)

Algorithm component
<p>ACTION SPACE $A = \{\text{Commitment (CM)}, \text{Consensus (CN)}, \text{Authority (AU)}, \text{Action planning (AP)}, \text{No persuasion (NP)}\}$</p>
<p>REWARD</p> $r = \begin{cases} -1 + \frac{e}{\bar{e}} & \text{if } e < \bar{e} \\ 1 - \frac{10-e}{10-\bar{e}} & \text{if } e > \bar{e} \\ 0 & \text{otherwise,} \end{cases}$ <p>where $e \in [0, 10]$ is an effort response and \bar{e} the mean effort.</p>
<p>REWARD FUNCTION $R : S \times A \times S \rightarrow [-1, 1]$ such that $R(s, a, s')$ is the instant reward for taking action a in state s and arriving in state s'.</p>
<p>TRANSITION FUNCTION $T : S \times A \times S \rightarrow [0, 1]$ such that $T(s, a, s') = Pr(s' s, a)$ is the probability of arriving in state s' after taking action a in state s.</p>
<p>DISCOUNT FACTOR 0.85</p>
<p>SIMILARITY COMPUTATION</p> <ol style="list-style-type: none"> 1. Scale TTM-stage and five personality dimensions to the interval $[0, 1]$. 2. Compute for a person i her similarity to all other people j based on the Euclidean distance between their six-dimensional trait vectors, whereby the largest distance is mapped to a similarity of 0 and the smallest distance to a similarity of 1. 3. Compute the weight w_{ij} of a sample from person j for person i as follows: $w_{ij} = \max\left(\frac{s_{ij}}{\sum_k s_{ik}}, 0.0001\right),$ <p>where s_{ij} is the similarity of i and j.</p>

Abbreviations: COM-B model, Capability-Opportunity-Motivation-Behavior model; TTM, Transtheoretical model.

STATE SPACE

Each session, participants answered ten questions on a 5-point Likert scale. Seven of these questions were based on the Capability-Opportunity-Motivation-Behavior (COM-B) self-evaluation questionnaire [245]. According to the COM-B model, capability, opportunity, and motivation together determine one's behavior, which in turn influences one's capability, opportunity, and motivation [387]. This made using capability, opportunity, and motivation as state variables for our RL approach appealing. We supplemented the seven questions from the self-evaluation questionnaire with people's self-efficacy due to the impact of self-efficacy on the effectiveness of different types of persuasive messages [54]. In addition, we asked people about their smoker and physical activity identities as according to Prime theory, self-identity can be a reliable predictor of behavior [386]. These additional questions fall under motivation in the COM-B model of behavior [168].

To lower the required amount of training data from the two training sessions, we subsequently reduced the size of the state space in two ways. First, we turned each state feature into a binary feature based on whether a value was greater than or equal to the feature mean (1) or less than the feature mean (0). Second, we selected three features in a way that was inspired by the G-algorithm [88]. Features were selected based on average rewards for level 2 and Q-values for levels 3 and 4 of algorithm complexity. The result of this state-space reduction was a state space of size $2^3 = 8$ (see Table 3.5 for the chosen features).

ACTION SPACE

Five persuasion types defined the action space. These were authority, commitment and consensus from Cialdini [93], action planning [153], and no persuasion. For each persuasion type, we formulated a set of message templates that were filled in for each activity to obtain persuasive messages. In the case of Cialdini's persuasion types, we created these templates by removing the domain-specific information from the validated healthy eating messages by Thomas et al. [360]. Due to the importance of self-identity in behavior [386], we also added two identity-based message templates for the commitment persuasion type. For action planning, we created templates based on the formulation by Sniehotta et al. [346]. However, rather than asking participants to enter their action plans in a table, the virtual coach prompted them to create an if-then plan of the form "If <situation>, then I will <do activity>" based on Chapman et al. [89]. In addition, the virtual coach provided an example of such an if-then plan as recommended by Hagger and Luszczynska [153]. Table B.1 lists examples of the resulting templates and persuasive messages.

The virtual coach asked the participants to type their action plans into the chat, which indicated whether the participants had read the message. In the case of the three persuasion types from Cialdini, however, simply showing the persuasive messages may have meant that participants did not centrally process or even read the messages [338]. However, such central processing was desirable. As the elaboration likelihood model indicates, high-effort central processing of messages leads to attitudes that are more likely to be persistent over time, resistant to counterattack, and influential in guiding thought and behavior [298]. Therefore, we attempted to increase in-depth central processing of the persuasive messages in three ways. First, the virtual coach printed the persuasive messages in boldface [259] to reduced distraction [298]. Second, the virtual coach asked participants to answer

reflective questions to increase self-referencing [74, 298] (Table 3.3). Third, we repeated the persuasion type [298] by adding reminder questions to the reminder messages participants received after each session (Table 3.4). We based these reminder questions on the ones used by Schwerdtfeger et al. [338] to remind people of their action plans and sent them for Cialdini’s persuasion types as well as action planning.

Table 3.3: Reflective questions for authority, commitment, and consensus.

Persuasion type	Reflective question: "Please tell me what you think: ..."
Authority	Which other experts, whose opinion you value, would agree with this?"
Commitment	In what way does doing this activity match your decision to successfully quit smoking?"
Commitment (Identity)	In what way does doing this activity match your decision to become somebody who has successfully quit smoking?"
Consensus	How would people like you, in a situation like yours, agree with this?"

Table 3.4: Examples of reminder question templates. Examples of templates for the reminder questions that are added to the reminder messages people receive.

Persuasion type	Example of reminder question template
Action planning	Keep in mind your rule for <i>⟨doing activity⟩</i> before the next session!
Authority	Do you remember which experts, whose opinion you value, would argue that <i>⟨doing activity⟩</i> may help to <i>⟨positive impact of activity⟩</i> ?
Commitment	Recall how <i>⟨doing activity⟩</i> is in line with your decision to successfully quit smoking!
Consensus	Don’t forget how people like you, in a situation like yours, would testify that <i>⟨doing activity⟩</i> may help to <i>⟨positive impact of activity⟩</i> !

REWARD

In sessions 2–5, participants were asked about the overall effort they put into their previously assigned activity on a scale from 0 to 10. Based on the mean effort \bar{e} computed after session 2¹, the reward $r \in [-1, 1]$ for an effort response e was computed as follows:

¹The mean effort was computed at the sample population level.

$$r = \begin{cases} -1 + \frac{e}{\bar{e}} & \text{if } e < \bar{e} \\ 1 - \frac{10-e}{10-\bar{e}} & \text{if } e > \bar{e} \\ 0 & \text{otherwise.} \end{cases}$$

The idea behind this reward signal was that an effort response that was equal to the mean effort was awarded a reward of 0, and that rewards for efforts greater and lower than the mean were each equally spaced.

REWARD AND TRANSITION FUNCTIONS

The reward and transition functions were estimated based on the samples collected from the first batch of people ($N = 516$) who successfully completed session 2². No updates to these training samples were made afterward as more data was collected.

SIMILARITY COMPUTATION

Rather than choosing the same persuasion type for each person in a state, the virtual coach maintained a separate policy π_i for each user i . When computing π_i , an observed sample from user j was weighted based on how similar i and j were. The virtual coach computed the similarity based on people's Big-Five Personality [146] and Transtheoretical Model (TTM)-stage [308] for becoming physically active. We chose these variables due to extensive previous work showing their impact on the success of different forms of persuasion [32, 106, 154, 182, 288, 359, 360, 398]. We did not consider the TTM-stage for quitting smoking, as participants had to be in one of two specific stages (i.e., contemplation or preparation) to be eligible for the study. For the similarity computation, the virtual coach first scaled the TTM-stage and the five personality dimensions to the interval $[0, 1]$ so that the features had the same scale. Next, the virtual coach computed for a person i her similarity to all other people j based on the Euclidean distance between their six-dimensional trait vectors. Thereby, the virtual coach mapped Euclidean distances to similarities so that the similarity for the smallest Euclidean distance was 1, and the similarity for the highest Euclidean distance was 0. Lastly, the virtual coach computed the weight w_{ij} of a sample from person j for person i as follows:

$$w_{ij} = \max\left(\frac{s_{ij}}{\sum_k s_{ik}}, 0.0001\right), \quad (3.1)$$

where s_{ij} is the similarity of i and j , k denotes a person on which samples were gathered, and the addition of 0.0001 was to ensure that no sample was given a weight of 0.

These similarity-based sample weights affected how the reward and transition functions were estimated for a person. For example, given a training set with one sample of the form $\langle s_k, a_k, r_k, s'_k \rangle$ from each of K people, the reward $R_i(s, a, s')$ for person i was computed as so:

$$R_i(s, a, s') = \frac{\sum_{k \in K, s_k=s, a_k=a, s'_k=s'} w_{ik} r_k}{\sum_{k \in K, s_k=s, a_k=a, s'_k=s'} w_{ik}}. \quad (3.2)$$

²We had aimed for at least 10 samples for each of the $8 \times 5 = 40$ state-action combinations.

ALGORITHM TRAINING

The persuasion algorithm on all four complexity levels was trained based on the data gathered in sessions 1 and 2 for the first batch of people ($N = 516$) who successfully completed session 2 (see Table 3.5 for the resulting policies). No samples were later added to this training set of 516 samples so that the policies for all people were trained on the same number of samples and hence comparable. For algorithm complexity levels 3 and 4, Q-values were computed via value iteration based on the estimated reward and transition functions.

Table 3.5: The learned policy used in sessions 3–5 for each algorithm complexity level. The state feature selection and training of all policies were based on the data gathered in sessions 1 and 2 for the first batch of people ($N = 516$) who successfully completed session 2. This training set of 516 samples was not updated thereafter as more data was gathered.

1: BEST AVG. REWARD

Commitment

2: BEST AVG. REWARD IN STATE

State feature	State							
	1	2	3	4	5	6	7	8
Feeling like wanting to do an activity (F5)	0	0	0	0	1	1	1	1
Feeling like being part of a group that is doing these kinds of activities (F4)	0	0	1	1	0	0	1	1
Thinking they can do an activity (F8)	0	1	0	1	0	1	0	1
<i>Action</i>	<i>AP</i>	<i>AP</i>	<i>CN</i>	<i>NP</i>	<i>CN</i>	<i>NP</i>	<i>NP</i>	<i>CM</i>

3: BEST Q-VALUE

State feature	State							
	1	2	3	4	5	6	7	8
Thinking that it would be a good thing to do an activity (F7)	0	0	0	0	1	1	1	1
Thinking they can do an activity (F8)	0	0	1	1	0	0	1	1
Knowing why it is important to do an activity (F1)	0	1	0	1	0	1	0	1
<i>Action</i>	<i>CN</i>	<i>CN</i>	<i>CM</i>	<i>NP</i>	<i>AU</i>	<i>CM</i>	<i>CN</i>	<i>CM</i>

4: BEST SIMILARITY-WEIGHTED Q-VALUE

Table 3.5: (continued)

State feature	State							
	1	2	3	4	5	6	7	8
Thinking that it would be a good thing to do an activity (F7)	0	0	0	0	1	1	1	1
Thinking they can do an activity (F8)	0	0	1	1	0	0	1	1
Knowing why it is important to do an activity (F1)	0	1	0	1	0	1	0	1
<i>Action</i>	<i>CN</i>	<i>CN</i>	<i>CM</i>	<i>NP</i>	<i>NP</i>	<i>AU</i>	<i>CN</i>	<i>CM</i>

The policy above is an example for one person (there was a separate policy for each person).

Abbreviations: Avg., Average; F, Feature; AP, Action planning; CN, Consensus; NP, No persuasion; CM, Commitment; AU, Authority.

The number of samples used to train the algorithm complexity levels was based on the guidelines by Cohen [97] for multiple regression analysis and a medium effect, an alpha of 0.05, and three independent variables (i.e., the three state features that describe the state space). This resulted in a sample size of 76. Since we have five actions, we multiplied this sample size by five for a value of 380. We used with 516 more than 380 samples as we had more people in the first batch of people who successfully completed session 2. Moreover, we analyzed the impact of sample sizes ranging from 25 to 2,300 on the Q-value estimation and optimality of chosen actions. Specifically, we estimated the reward function, transition function, and the resulting "true" Q-values and optimal policy based on all 2,366 samples gathered in our study. We then randomly drew different numbers of samples from these 2,366 samples and computed the mean L_1 -error for predicting the true Q-values based on 100 repetitions per sample size. We obtained a mean L_1 -error of 0.68 for our sample size of 516, which is a reduction by more than two thirds of the mean L_1 -error for a sample size of 25. In addition, the mean L_1 -error for the true Q-values of the estimated optimal actions compared to the true optimal actions per state is only 0.08. This shows that the optimal action chosen based on 516 samples is only slightly worse than the true optimal action. Figure B.9 provides further information on this.

MATERIALS

We used four online services in this study: Prolific for recruiting, inviting, and communicating with participants, Qualtrics for hosting the questionnaires and instructions for the conversational sessions, and Google Compute Engine to host the virtual coach and the sessions via Rasa X. In addition, some activities assigned in the sessions involved watching a video on YouTube.

VIRTUAL COACH

The virtual coach used for the sessions was implemented in Rasa [61]. It had the name Sam, which may help to increase its social presence [400]. Sam presented itself as being there to help people to prepare to quit smoking and become more physically active as the latter may aid the former. In its responses, Sam used techniques from motivational interviewing [59] such as giving compliments for putting much effort into assigned activities and expressing empathy otherwise. Empathy can thereby also help to form and maintain a relationship with a user [57], which can support behavior change [400]. Based on discussions with smoking cessation experts, Sam maintained a generally positive and encouraging attitude while trying to avoid responses that may be perceived as too enthusiastic [135]. To make the conversations accessible for people with low literacy levels, large chunks of text were broken up into multiple messages, in between which participants had to indicate when to proceed. In addition, participants communicated mainly by clicking on buttons with possible answer options. Only when free-text input was crucial, such as when writing about the experience with an assigned activity, were buttons not used. Lastly, to avoid repetitiveness, Sam randomly chose from several different formulations for an utterance. This is important, as repetitiveness can negatively influence the engagement with a system and motivation to perform an advocated behavior [57]. The implementation of the virtual coach can be found online [9].

PREPARATORY ACTIVITIES

In each session, the virtual coach asked participants to complete a new preparatory activity for next time that related to quitting smoking or increasing physical activity, such as writing down and ranking reasons for quitting smoking:

Having high aspiration to quit smoking may aid quitting successfully. So, before the next session, I advise you to identify and write down reasons why you want to stop smoking. After writing them down, think about which reasons are most important to you and order them accordingly.

The virtual coach selected the activities from a pool of 24 activities of similar duration, 12 each for quitting smoking and increasing physical activity. The activities for quitting smoking were based on components of the StopAdvisor smoking cessation intervention [243] and future-self exercises [236, 294]. The ones for increasing physical activity were generated by adapting the smoking cessation activities. Each activity formulation included reasoning for why the activity could help to prepare to quit smoking. A psychologist and smoking cessation expert read through the activity formulations to ensure they were suitable and clear. The virtual coach proposed one activity for quitting smoking and one for increasing physical activity in the first two and the subsequent two sessions to each participant. The virtual coach chose the type of activity in the fifth session randomly. It selected an activity for an activity type uniformly at random while avoiding repetitions of the same and very similar activities (e.g., creating a personal rule for not smoking and creating a personal rule for becoming more physically active). So participants were never asked to do an activity more than once, as the goal was not to create habits. The formulations of the activities are provided in Table B.8.

MEASURES

PRIMARY MEASURES

To assess the effectiveness of subsequently adding the algorithm components for our first hypothesis, we used the following primary measures:

Effort. The virtual coach measured the effort by asking participants about the effort they put into their previously assigned activity on a scale from 0 ("Nothing") to 10 ("Extremely strong"). The scale was adapted from Hutchinson and Tenenbaum [170]. Note that this effort measure also served as a basis for choosing persuasion types in the four algorithm complexity levels.

Perceived motivational impact. The virtual coach measured the perceived motivational impact of the sessions by asking participants "Please rate the impact of our last 2 conversations on your motivation to do your previous assigned activities" at the beginning of their third and fifth sessions. The virtual coach prompted participants to enter any number between -10 and 10, with -10 being "Very negative," 0 being "Neutral" and 10 being "Very positive."

SECONDARY MEASURES

Algorithm input measures. We measured several variables as input for the persuasion algorithms. This included ten possible state features, seven of which were adapted from the COM-B self-evaluation questionnaire [245] (e.g., "I know why it is important to do the activity") and answered on a 5-point Likert scale. The other three features were based on measuring self-efficacy [346] on a 5-point Likert scale, smoker identity [78] with the additional answer option "Smoker," and physical activity identity. For the latter, we adapted the item with the highest factor loading from the Exercise Identity Scale [38] to physical activity and asked participants to rate it on a 5-point Likert scale. The highest algorithm complexity level further required computing how similar people were. We accomplished this by using people's Big-Five personality based on the 10-item questionnaire by Gosling et al. [146] and TTM-stage for becoming physically active based on an adaptation of the question by Norman et al. [271] to physical activity.

Activity involvement. For exploration purposes, we further measured participants' involvement in their assigned activities in the post-questionnaire. We, therefore, asked participants to rate three items to assess whether they found their assigned activities interesting, personally relevant, and involving. The three items were based on Maheswaran and Meyers-Levy [229] and answered on a scale from -5 ("Disagree strongly") to 5 ("Agree strongly").

Potential covariates. We collected data on potential covariates for the first hypothesis in the pre-questionnaire. Besides the variables discussed above, this included quitter self-identity measured with three items based on Meijer et al. [235], the need for cognition based on the three items from Cacioppo et al. [77] used by Steward et al. [351], and physical activity identity based on an adaptation of the Exercise Identity Scale by Anderson and Cychosz [38] to physical activity. All items were rated on 5-point Likert scales.

Data for future research. We measured several other variables for future research. These variables are not discussed in this chapter but are described in our OSF preregistration form [14].

PARTICIPANTS

Prior to the experiment, we computed a conservative estimate of the sample size required for evaluating the effectiveness of the four algorithm complexity levels using the Monte Carlo simulation described in Chapter 4.9.2 in Chechile [90]. We ran the simulation based on the ability to find an effect size (Cohen's g) of 0.1, which is halfway between a small ($g=.05$) and a medium ($g=.15$) effect size according to Cohen [97], with reliable Bayes factor values of 19 or more for four conditions and a binary response variable. The result was a sample size of 132 per algorithm complexity level, resulting in a total of 528 participants. This estimate was conservative, as we had interval dependent variables instead of binary ones. In reaching the sample size, we were constrained by a budget limit of 5,000 euros.

Moreover, we conducted a Bayesian power analysis based on a Monte Carlo approach. We used 500 simulations of two conditions with a medium difference of 0.3 [97] between their standard normally distributed means. For each simulation, we computed the Bayes factor for the hypothesis that the mean of the second condition is higher than the one of the first condition. The power was then calculated as the fraction of simulations in which the Bayes factor was at least 19. The result was a power of 0.78 for 129 samples per condition, which is the lowest number of samples we obtained for an algorithm complexity level for the last session.

To be eligible, participants had to be fluent in English, smoke tobacco products daily, contemplate or prepare to quit smoking [112], not be part of another intervention to quit smoking, and provide informed consent. In addition, we used the quality measures on Prolific to choose people who had completed at least one previous study and an approval rate of at least 90% for their previously completed studies. 1406 participants started the prescreening questionnaire, and 521 of the 922 eligible participants successfully reported on all their assigned activities in sessions 2 to 5. Participants were not invited to a subsequent study component when doing an entire component twice or failing two or more attention checks. In addition, participants had to respond to a study component invitation within about one day for the pre-questionnaire, three days for the sessions, and seven days for the post-questionnaire. The participant flow is depicted in Figure B.2. Participants who completed a study component were paid based on the minimum payment rules on Prolific, which require a payment rate of five pounds sterling per hour. Participants were informed that their payment was independent of how they reported on their suggested preparatory activities to account for self-interest and loss aversion biases [117]. Self-interest bias can arise when incentives exist that motivate participants to respond in a certain way; loss aversion bias can occur when participants choose to not participate or to drop out when suspecting that they may not be paid fairly. Participants who failed two or more attention checks in a study component were not reimbursed.

Table B.3 lists participant characteristics such as age, gender, TTM-stage for quitting smoking, and the existence of previous quit attempts for each algorithm complexity level. We compared Bayesian models with and without the algorithm complexity level as a predictor for each characteristic to test for systematic differences between the four levels. We found that based on the Widely Applicable Information Criterion (WAIC), the model with the algorithm complexity level as a predictor always performed worse than the model without this predictor, therefore providing no indication of systematic differences between the algorithm complexity levels for these characteristics. Participants were nationals of

diverse countries such as the United Kingdom, Portugal, the Russian Federation, the United States, Chile, South Africa, Nigeria, Turkey, India, Malaysia, and Australia.

PROCEDURE

We opened the intake for the prescreening questionnaire on Prolific on each of seven different days between 20 May and 8 June 2021. Each time, we invited as many people as we could afford to reach but not exceed our budget limit in case of no further dropout. Participants meeting the qualification criteria could access the prescreening questionnaire on Prolific, and those who passed the prescreening were invited to the pre-questionnaire about one day later. In the pre-questionnaire, we collected participant data such as the Big-Five personality and TTM-stage for becoming physically active. One day after completing the pre-questionnaire, we invited participants to the first of five sessions in which the virtual coach Sam assigned them a preparatory activity together with a persuasion type. Participants received instructions on interacting with Sam in Qualtrics before being directed to a website for the conversation. The structure of the conversations is depicted in Figure B.4 and two excerpts of actual conversations are shown in Figure B.5. Each session lasted about five to eight minutes, and invitations to a subsequent session were sent about two days after having completed the previous one. The study ended with a post-questionnaire, to which participants were invited about two days after completing the last session.

DATA PREPARATION AND ANALYSIS STRATEGIES

First, we corrected entry errors from state and attention check questions in the sessions that participants messaged us about on Prolific ($N = 4$). The corrections for state questions ($N = 2$) pertained to the question on smoker identity, with two participants correcting their entries from "non-smoker" and "ex-smoker" to "smoker" for session 1. As participants were persuaded randomly in session 1, these entry errors did not affect the conversations. Entry errors for attention check questions ($N = 2$) had no effect on the conversations irrespective of the session. Next, we preprocessed the gathered data by 1) using only data from sessions and the post-questionnaire if people passed at least one attention check during the respective component, and 2) using the first recorded submission for a study component if people did the component more than once. In the following, we describe our data and analysis strategies for each hypothesis in detail. All data and analysis code can be found online [19].

H1: ALGORITHM EFFECTIVENESS

We conducted a multi-level (i.e., hierarchical) Bayesian analysis of the data.

Further data preparation. We removed the data of people who did not complete session 2 and were therefore not assigned to a condition ($N = 5$). To make an exploratory analysis of subgroups based on activity involvement possible, we computed the reliability of the corresponding three items. As the reliability was sufficiently high (Cronbach's $\alpha = 0.89$), we used the mean of the items as an index measure.

Statistical models. We created three statistical models. For both dependent variables, the effort people put into their activities and the perceived motivational impact of the sessions, we fit models that contained a general mean, a random intercept for each participant, a fixed effect for algorithm activeness, a fixed effect for the algorithm complexity level, and

a fixed interaction effect between the algorithm complexity level and algorithm activeness. For the effort, we also fit a second model that additionally included a main fixed effect for the session, as well as fixed interaction effects between the session and the other two factors. We fit all three models with diffuse priors based on the ones used by McElreath [234]. In addition, we performed a prior sensitivity analysis to assess the impact of using different priors. We found only a limited effect on the posterior probability for a hypothesis being true, as it changed by at most 0.02. A t -distribution was fit for the dependent variable in each model.

3

Covariates. We explored potential covariates such as the type of the assigned activities (i.e., quitting smoking or increasing physical activity), physical activity identity, and quitter self-identity. Adding these variables did not change the conclusions drawn about our hypothesis, and therefore we did not include the variables in the final models.

Inference criteria. For each of the three statistical models, we computed the posterior probability that our hypothesis was true based on samples drawn from the estimated model. This means that we evaluated for each model the posterior probability that the relevant parameter was greater than 0. For the first two models, this parameter was the fixed two-way interaction effect between the algorithm complexity level and algorithm activeness. For the third model it was the fixed three-way interaction effect between the algorithm complexity level, algorithm activeness, and the session. We interpreted posterior probabilities using the guidelines by Chechile [90] and their extension to values below 50% based on Andraszewicz et al. [39]. We also report the 95% Highest Density Intervals (HDIs) for the parameters, with an HDI being "the narrowest interval containing the specified probability mass" [234]. In addition, we used the Region of Practical Equivalence (ROPE) method [203] as a secondary method to evaluate the results. This method allows one to accept or reject a hypothesis or to withhold a decision. As the results of this method were all inconclusive, we do not report them.

Implementation. All analyses were carried out in R with the rethinking package [234]. We provide code to reproduce the analyses in a Docker container as recommended by van de Schoot et al. [372].

Exploratory subgroup analysis based on activity involvement. We divided participants into subgroups based on whether their activity involvement was greater than or equal to the median ($N = 269$) or less ($N = 231$). The analyses for the three models were then repeated separately for each subgroup. Note that since we measured the activity involvement in the post-questionnaire, this analysis only included participants for whom we had data on at least one passed attention check from the post-questionnaire ($N = 500$).

H2: SIMILARITY OF OPTIMAL PERSUASION STRATEGIES

Data preparation. We compared the optimal policies computed based on all collected data to using only the data on activities for either quitting smoking or increasing physical activity. Therefore, we distributed the gathered data over three datasets based on the activity type. This resulted in 1,175 samples for quitting smoking, 1,191 for increasing physical activity, and 2,366 for both activity types together.

Analysis plan. We computed the optimal policies for each non-baseline algorithm complexity level (i.e., levels 2-4) for each dataset. To use equal amounts of data for both activity types, we randomly drew 1,000 samples from each activity type. This means that

we used 1,000 and 2,000 samples, respectively, when computing a policy based on a single activity type and both activity types together. To account for the impact of this random selection, the sampling and subsequent optimal policy computation were conducted 100 times. Afterward, we concatenated the optimal policies for the 100 repetitions into a single list for each data type. In the case of multiple best actions for a state, one of the best actions was chosen uniformly at random.

Inference criteria. For each non-baseline algorithm complexity level, we calculated Cohen’s κ between the list of optimal policies based on activities for both quitting smoking and increasing physical activity, and the list of optimal policies computed based on only samples in which participants were advised to do a preparatory activity for either quitting smoking or increasing physical activity. The outcomes were interpreted based on the guidelines by Landis and Koch [209] shown in Table B.6. We also determined Cohen’s κ between the optimal policies computed based on different samples drawn from the same data type for exploratory purposes. This allowed us to draw conclusions about the consistency of policies computed on a certain data type.

Implementation. We provide code to reproduce the analyses in Python.

RESULTS

Table 3.6 and Table 3.7 provide overviews of the mean effort and perceived motivational impact per algorithm complexity level and measurement moment. As the reward was computed based on the effort, the former gives an indication of the reward obtained by the virtual coach for the four algorithm complexity levels. Furthermore, Table B.7 shows the mean effort per activity and activity type. To give some intuition for the size of the observed differences in Table 3.6 and Table 3.7, we can divide the largest change between the first and last measurement moment for a complexity level by the standard deviation. The resulting effect size is 0.14 based on an observed change of at most 0.40 for the effort, and 0.20 based on an observed change of at most 0.66 for the perceived motivational impact. Both effect sizes are at most small according to Cohen [97].

Table 3.6: Mean effort per algorithm complexity level and measurement moment.

Complexity level	Measurement moment			
	1	2	3	4
1: Best avg. reward	5.55 (2.52)	5.49 (2.74)	5.49 (2.94)	5.59 (2.88)
2: Best avg. reward in state	5.54 (2.54)	5.20 (2.84)	5.20 (3.09)	5.46 (3.01)
3: Best Q-value	5.48 (2.56)	5.40 (2.94)	5.32 (2.99)	5.88 (2.86)
4: Best similarity-weighted Q-value	5.45 (2.47)	5.13 (2.81)	5.29 (3.00)	5.11 (3.12)

Standard deviations are provided in parentheses.

Effort was measured on a scale from 0 to 10.

Abbreviations: avg., average.

Table 3.7: Mean perceived motivational impact per algorithm complexity level and measurement moment.

Complexity level	Measurement moment	
	1	2
1: Best avg. reward	5.01 (3.42)	5.20 (3.48)
2: Best avg. reward in state	4.83 (3.61)	5.23 (3.97)
3: Best Q-value	4.75 (3.22)	5.40 (3.22)
4: Best similarity-weighted Q-value	4.64 (3.41)	5.05 (3.43)

Standard deviations are provided in parentheses.

Perceived motivational impact was measured on a scale from -10 to 10.

Abbreviations: avg., average.

3

H1: ALGORITHM EFFECTIVENESS

Between the two baseline sessions and the two sessions in which the algorithms were active, the largest increase in effort was observed in complexity levels 1 and 3 (Figure 3.4A). Quantifying these observations based on our Bayesian analysis, Table 3.8 reveals that it is not worth betting against higher complexity levels leading to a larger increase in effort, with the mean of the credible values showing a decrease of 0.05 in effort between complexity levels 1 and 4 when the algorithms are active. The HDI thereby ranges from -0.43 to 0.33, with only 39% and thus less than half of the credibility mass favoring higher complexity levels leading to a larger increase in effort. However, a detailed examination of Figure 3.4A suggests that there are differences between the two active sessions, which are sessions 3 and 4. Specifically, complexity level 3 exhibits a change from an effort similar or lower compared to level 1 in session 3 to the highest effort in session 4. Complexity level 4, on the other hand, shows a decrease in effort between the two active sessions.




Figure 3.4: Mean effort per session and algorithm complexity level.

These observations are found back when the fit model is extended with the session as a predictor and specifically a three-way interaction effect between algorithm complexity, algorithm activeness, and session. This fit model assigns a posterior probability of 0.70 to the hypothesis that the increase in effort between the two active sessions is larger for higher complexity levels (Table 3.8). As a result, more than half of the credibility mass are

Table 3.8: Results of Bayesian analyses of effort and perceived motivational impact.

DV	Parameter	Mean [HDI] (SD)	Post	Evaluation
Effort	alg. level × alg. activeness	-0.05 [-0.43, 0.33] (0.19)	0.39	Not worth betting against
	alg. level × alg. activeness × session	0.20 [-0.57, 0.97] (0.40)	0.70	Not worth betting on
	alg. level × alg. activeness	0.09 [-0.34, 0.54] (0.22)	0.67	Not worth betting on

Abbreviations: DV, Dependent variable; HDI, Highest density interval; SD, Standard deviation; Post, Posterior probability that a parameter's value is greater than 0; PMI, Perceived motivational impact.

Table 3.9: Results of Bayesian analyses of effort and perceived motivational impact for people with low and high activity involvement.

DV	Parameter	Mean [HDI] (SD)	Post	Evaluation
LOW ACTIVITY INVOLVEMENT				
Effort	alg. level × alg. activeness	0.18 [-0.55, 0.94] (0.39)	0.67	Not worth betting on
	alg. level × alg. activeness × session	-0.45 [-1.83, 0.94] (0.71)	0.26	Not worth betting against
	alg. level × alg. activeness	-0.45 [-1.36, 0.44] (0.46)	0.16	Only a casual bet against
HIGH ACTIVITY INVOLVEMENT				
Effort	alg. level × alg. activeness	-0.20 [-0.67, 0.26] (0.24)	0.20	Only a casual bet against
	alg. level × alg. activeness × session	0.80 [-0.10, 1.69] (0.46)	0.96	Good bet - too good to disregard
	alg. level × alg. activeness	0.26 [-0.20, 0.71] (0.23)	0.87	Only a casual bet

Abbreviations: DV, Dependent variable; HDI, Highest density interval; SD, Standard deviation; Post, Posterior probability that a parameter's value is greater than 0; PMI, Perceived motivational impact.

in favor of this hypothesis. Similarly, we observe in Figure 3.5A that complexity level 3 is accompanied by the largest increase in perceived motivational impact. Despite the increase being lower for complexity level 4 than for level 3, there is hence additional support for the first hypothesis. A Bayesian analysis confirms this (Table 3.8). More precisely, the posterior probability that the increase in perceived motivation impact is larger for higher complexity levels is 0.67. In other words, more than half of the credibility mass support this.




Figure 3.5: Mean perceived motivational impact when the algorithms are off/on per algorithm complexity level.

We also conducted separate analyses for people with high and low involvement in their activities for exploratory purposes. The mean effort is higher for people with high involvement (Figure 3.4C) than for those with low involvement (Figure 3.4B). For the high involvement subgroup, complexity level 1 shows the largest increase in effort between the two baseline and the two active sessions (Figure 3.4C). Quantitatively, the mean credible value is -0.20 with a posterior probability of only 0.20 that this increase is larger for higher complexity levels (Table 3.9). Again, however, complexity level 3 is associated with the largest increase in effort between the two active sessions (Figure 3.4C). This matches the quantitative results, according to which the posterior probability in favor of higher complexity levels leading to a larger increase in effort between the two active sessions is 0.96, a "good bet - too good to disregard" (Table 3.9). For people with low involvement, on the other hand, complexity level 1 is associated with the largest drop in effort between the two baseline and the two active sessions (Figure 3.4B). While this quantitatively leads to a posterior probability of 0.67 for higher complexity levels leading to a larger increase in effort (Table 3.9), none of the algorithm levels is very effective for this subgroup. Regarding the increase in effort between the two active sessions, complexity level 2 performs best but is still hardly effective in session 4 (Figure 3.4B). Therefore, based on our Bayesian analyses, it is not worth betting against higher complexity levels leading to a larger increase in effort between the two active sessions for this subgroup due to a posterior probability of 0.26.

Next, we conducted an exploratory analysis of the perceived motivational impact in the two subgroups. For people with high involvement (Figure 3.5C), the perceived motivational impact is much higher than for people with low involvement (Figure 3.5B). Thereby, complexity levels 3 and 4 show a larger increase in perceived motivational impact than levels 1 and 2 for the high involvement subgroup (Figure 3.5C). Given that this is the case only for complexity level 3 when both subgroups together are considered (Figure 3.5A), there seems to be more support for higher complexity levels leading to a larger increase

in perceived motivational impact for people with high involvement. Quantitatively, the posterior probability in favor of this is 0.87, which can be qualified as a casual bet (Table 3.9). However, in contrast to the high involvement subgroup, we do not find much support for this for the low involvement subgroup. Neither of the four complexity levels is associated with an apparent increase in perceived motivational impact and instead levels 1 and 4 suggest even a slight decrease (Figure 3.5B). The posterior probability of 0.16 confirms this. In other words, it qualifies as a casual bet against higher complexity levels leading to a larger increase in perceived motivational impact for this subgroup (Table 3.9).

H2: SIMILARITY OF OPTIMAL PERSUASION STRATEGIES

Table 3.10 shows fair to moderate agreement between the policies computed based on both activity types together on the one hand and activities solely for quitting smoking or increasing physical activity separately on the other hand for the non-baseline algorithm complexity levels. The agreement thereby tends to be much stronger for algorithm complexity level 3 and for physical activity also for algorithm complexity level 4. For reference, an upper limit of agreement was calculated by examining the agreement between policies computed from samples drawn from the same data set. This upper limit is moderate agreement for all three complexity levels.

Table 3.10: Cohen's κ for algorithm complexity levels 2 to 4 after computing the optimal policies based on different types of data. The types of data result from splitting the data based on the activity type, or using the data on both activity types together.

Data type		Complexity level		
		2	3	4
MAIN ANALYSES				
Smoking cessation	Both	0.33	0.36	0.26
Physical activity increase	Both	0.29	0.49	0.48
REFERENCE ANALYSES				
Both	Both	0.60	0.54	0.58
Smoking cessation	Smoking cessation	0.54	0.56	0.42
Physical activity increase	Physical activity increase	0.51	0.57	0.59

DISCUSSION AND CONCLUSION

The presented longitudinal study examined the effectiveness of subsequently adding the consideration of states, future states, and the similarity of people to a personalized RL algorithm for persuading people to do preparatory activities for quitting smoking and increasing physical activity. The findings provide some support that people's reported motivation is positively affected by using higher algorithm complexity levels. The effort people spent on the activities also provides some support. Here, however, the overall advantage of using more algorithm elements becomes apparent only after some time, and initially, there seems to be no positive impact. This is reflected by the three-way interaction

effect in Table 3.8 and also visible in the increase in effort between session 3 and 4 for complexity level 3 in Figure 3.4A. Looking at the algorithm complexity levels separately, the level that considers current and future states by choosing a persuasion type with the highest Q-value seems most successful in moving people to future states in which they can be persuaded better. Support for this is even stronger for people who found the suggested activities most useful. An explanation may be that the persuasive messages have a stronger and more persistent impact on people with high activity involvement. According to the elaboration likelihood model, high involvement in an issue makes it more likely that messages are processed in detail [229]. Such in-depth processing in turn is more likely to have a persistent impact [298].

Extending the algorithm by weighting observed data samples based on the similarity of people did not perform well in this study. Specifically, the results suggest that the fourth algorithm complexity level that additionally considers the similarity of people based on their TTM-stage for becoming physically active and Big-Five personality is associated with a lower effort spent on the activities than the third level. This shows that increased personalization can be harmful, even if it is informed by literature, as in our case. One reason could be the necessity of using more domain-specific similarity variables such as quitter self-identity [235]. Future work can use our published data to determine whether such similarity variables are relevant in our domain. Moreover, while we computed similarity based on the Euclidean distances between vectors of user characteristics, other distances such as the cosine distance could be used (see Ontañón [283] for an overview).

Another interesting observation is that the impact of using higher algorithm complexity levels for people with low activity involvement appears to be not zero but in fact negative for the increase in effort between the two active sessions and the perceived motivational impact. This suggests that choosing persuasion types based on higher algorithm complexity levels is worse for this subgroup than doing so based on lower ones. The reason might be a novelty effect [101, 333]. A novelty effect arises because people are initially curious about a new system or technology and have high expectations. However, this curiosity and perceived usefulness fade over time as people become aware of the system's limitations. Applying this novelty effect to our study, participants likely had high expectations about the system's capability to help them prepare to quit smoking at the beginning. Afterward, the perceived usefulness of the approach may have decreased for some people as their expectations were not met. However, since we used the data gathered in the first two sessions as training data for the persuasion algorithms, the algorithms were trained mainly based on people who thought the system was useful. This likely lowered the performance of our algorithms overall, but especially the performance of higher algorithm complexity levels for people with low activity involvement. This is the case because these higher algorithm complexity levels were fit more tightly to the data gathered from people with high activity involvement in the first two sessions. An important implication for future work is that it may be relevant to consider when a data sample was gathered during the behavior change process. Especially when persuasive attempts are made over a long time, might it be beneficial to give a lower weight to samples collected at the beginning of the interaction with the system. Furthermore, since people's preferences can also change over time [382], the weights for later samples need to be chosen carefully as well.

As to our second hypothesis, we see some agreement between the persuasion types

chosen based on all collected data and the ones chosen based on data for only quitting smoking or increasing physical activity. This lends some support to transferring knowledge between these two activities types. Yet, we find that the agreement tends to be higher between optimal policies computed based on different samples of the same data type. This suggests that the algorithm could be improved by considering the activity type in the optimal policy computation. One reason for the lower agreement between optimal policies computed based on different data types might be that the involvement in the two activity types differed. We observed, for example, that people put overall less effort into activities for physical activity increase than ones for smoking cessation (Table B.7). Since the effort in our study was lower for people with low involvement in the activities, the processing of messages for the two activity types may have been different. Potentially, the link between becoming more physically active and quitting smoking could be made more evident. Besides higher agreement between optimal policies computed on the same data type, we also find that higher agreement between optimal policies computed based on different data types is achieved for algorithm complexity level 3 and for physical activity also for level 4. Thus, incorporating future states in algorithm complexity levels 3 and 4 and using people's TTM-stage for becoming physically active to weigh the observed samples in algorithm complexity level 4 appears to have helped capture the difference between the two activity types. It would be interesting to see in future work if other state or similarity variables could further improve upon this.

Besides the ideas mentioned above, there are many further directions for extending our work. First, it is interesting to think about the choice of reward signal. We want that people do their preparatory activities more thoroughly, so that they are better prepared for quitting smoking, so that they better achieve and maintain abstinence from smoking. While the more distal outcome measures in this chain capture the actual behavior we want to see, using them as reward signal to compare the four algorithm complexity levels leads to several challenges. This includes the time until we receive the signal (e.g., it may take several months or years before we know whether somebody has maintained their abstinence) and the signal's noisiness (e.g., a cancer diagnosis a year from now can also affect abstinence maintenance). In the case of being prepared for quitting smoking, an additional challenge is how this "preparedness" can be measured given that the activities differ in what they are meant to achieve, be it increasing self-efficacy or removing smoking cues. Certainly, several questions could be asked, but the number of questions should be kept low in light of the already low adherence rates for eHealth applications. Since the links in the chain from thoroughly doing preparatory activities to successful maintenance of smoking abstinence have already been supported by other literature (e.g., [243]), we thus chose the effort people spent on their activities as a more proximal reward signal. Notably, however, an even more proximal reward signal could be added. This is because our results suggest that with motivation, one of the predictors of behavior is increased to a greater extent if future states are taken into account. A combination of effort and perceived motivational impact could thus be used. Hiraoka et al. [161], for instance, use a reward signal that combines user satisfaction, the success of persuasion, and the naturalness of system utterances. It may also be worthwhile to add a more objective measures of behavior than self-reported ones such as the effort in our study. This may, however, not be feasible for some (parts of) activities, such as placing a rule for not smoking in a place one can see

every day.

Another direction for future work is to use our learned policy as a starting point and subsequently adapt to single individuals by focusing on their personal samples (rather than samples from similar people as in our fourth complexity level). The reason is that research has shown that the way people respond to a persuasive attempt is a good predictor of how they will respond to the same attempt in the future [183, 249, 331, 383]. Moreover, it may be important to consider the impact repeated actions can have. Repeatedly sending the same persuasion type may make it less effective [360], but could also help to strengthen the link between cue and response for action planning [338] or to scrutinize arguments objectively [76]. One interesting study in this regard is the one by Mintz et al. [249], which considers the effects both of repeating a message and of subsequently not sending the message for some time. Lastly, another avenue for future work is to ensure that the algorithms are ethical. For instance, it may not be ethical to choose a persuasion type that is predicted to be effective while at the same time lowering a person's self-efficacy. One way to incorporate such values or norms may be to learn a separate constrained policy from ethical examples [269] or expert preferences [92]. Other relevant issues are user trust, user privacy, and low bias [400].

On a higher level, our results show us that the impact of message-based persuasion algorithms on predictors of behavior and behavior itself is small. For example, the mean perceived motivational impact when the algorithms are active does not differ by more than 0.35 (Cohen's $d=0.10$) between the four algorithm complexity levels (Figure 3.5A). This qualifies as less than a small effect size according to Cohen [97]. Similarly, the mean effort for session 3 does not differ by more than 0.29 (Cohen's $d=0.10$) between the four algorithm complexity levels (Figure 3.4A). While other persuasive messages could have a larger effect, our findings are in line with other work. Kaptein et al. [183], for instance, found that the difference between a random and a tailored persuasive message with regards to the number of daily consumed snacks is 0.08 for a single persuasive attempt. Similarly, de Vries [106] saw that self-reported physical activity increases over time for both a tailored and a random message condition, with the physical activity being slightly but not significantly higher for the tailored condition. While the results of de Vries [106] and the snacking study of Kaptein et al. [183] were based on relatively small sample sizes of overall 47 and 73 participants, respectively, our results now show that the impact of persuasion algorithms on behavior is small even when conducting a large-scale experiment with at least 129 participants per condition and a resulting power of at least 0.78. Arguably, the impact of persuasion algorithms has been found to increase over time in both our experiment and the snacking study of Kaptein et al. [183]. More research is needed to test whether and how this increase in effectiveness occurs when persuasive attempts are made over more extended periods such as weeks or months. Even though the dangers of message amplification have been pointed out in various contexts such as social media [328], it is not yet well understood how and to which extent it influences actual behavior.

Given the so far limited impact of persuasive messages, an alternative may be to strategically persuade people through an entire dialog (e.g., [161], [329]). In addition, one could aim to further increase the processing of the persuasion by using multi-modal forms of persuasion (e.g., [379, 382, 388]). This choice of modalities can also be learned [132]. Alternatively, one could optimize the suggested behavior rather than the persuasive message.

Our results show, for example, that the mean effort for the most effective complexity level in session 3 is by 2.84 (Cohen's $d=0.95$) higher for people with high than for people with low involvement in their activities (Figure 3.4). In contrast to the difference between complexity levels, this qualifies as a large effect size [97]. Relevant approaches in this regard are ones to optimize suggested step goals [5, 402], activities for elderly people [100], physical [313] or learning [95] activities, or breast cancer screening recommendations [42].


In conclusion, we have presented a personalized RL-based persuasion algorithm and systematically tested the effectiveness of the algorithm components. Our results support the importance of taking states and future states into account to persuade people again. We expect that future work can build on these results to improve persuasion algorithms further. We make the dataset on 2,366 persuasive messages sent to 671 people publicly available to facilitate this. Given the sparsity of public datasets in this field and the expensive nature of collecting data on human behavior, we think this helps those wishing to develop new algorithms or to test existing ones. For the field of behavior change, our dataset provides the effectiveness of different activities based on the effort people spent on them. This shows, for example, that the link between increasing physical activity and quitting smoking needs to be made more evident for participants. In addition, our results lend support to the COM-B model of behavior change, as state variables derived from this model showed to help predict behavior. Thus, our study can be seen as a successful example of combining computer science and behavior change theories to test behavior change theories in a large-scale experiment [246]. Given the fruitful insights for both fields, we encourage further work at their intersection.

4

PERSUADING TO PREPARE FOR QUITTING SMOKING WITH A VIRTUAL COACH: USING STATES AND USER CHARACTERISTICS TO PREDICT BEHAVIOR

4

Despite their prevalence in eHealth applications for behavior change, persuasive messages tend to have small effects on behavior. Conditions or states (e.g., confidence, knowledge, motivation) and characteristics (e.g., gender, age, personality) of persuadees are two promising components for more effective algorithms for choosing persuasive messages. However, it is not yet sufficiently clear how well considering these components allows one to predict behavior after persuasive attempts, especially in the long run. Since collecting data for many algorithm components is costly and places a burden on users, a better understanding of the impact of individual components in practice is welcome. This can help to make an informed decision on which components to use. We thus conducted a longitudinal study in which a virtual coach persuaded 671 daily smokers to do preparatory activities for quitting smoking and becoming more physically active, such as envisioning one's desired future self. Based on the collected data, we designed a Reinforcement Learning (RL)-approach that considers current and future states to maximize the effort people spend on their activities. Using this RL-approach, we found, based on leave-one-out cross-validation, that considering states helps to predict both behavior and future states. User characteristics and especially involvement in the activities, on the other hand, only help to predict behavior if used in combination with states rather

 **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Persuading to prepare for quitting smoking with a virtual coach: Using states and user characteristics to predict behavior. In Noa Agmon, Bo An, Alessandro Ricci, and William Yeoh, editors, *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 - 2 June 2023*, pages 717–726. ACM, 2023. doi: 10.5555/3545946.3598704.

than alone. We see these results as supporting the use of states and involvement in persuasion algorithms. Our dataset is available online.

INTRODUCTION

Recent years have seen a surge of eHealth applications for behavior change (e.g., [125, 227, 237]), which provide behavior change support over the Internet or connected technologies such as apps and text messaging. Such applications often ask their users to do activities such as setting a goal, planning a running route, or watching an educational video. Persuasive messages are commonly used to motivate users to do these activities. For example, users may be reminded that doing an activity is in line with their decision to change their behavior. However, the effect of single persuasive attempts on behavior tends to be small (e.g., [20, 106, 183]).

Several studies have tried to increase the effectiveness of a persuasive attempt. One way is to consider the current state people are in (e.g., confidence, knowledge, motivation). Such a state describes a person's condition or status at a certain time that is relatively stable with regards to its elements [37]. Di Massimo et al. [111] and Klein et al. [198], for instance, account for people's self-efficacy when selecting messages for behavior change. Doing so is in line with behavior change theories, which posit that behavior is influenced by people's current state (e.g., [6, 242]). Yet, behavior in turn can also influence people's states. For example, verbally persuading people [353] or improving their mood [189] may increase their self-efficacy. Intuitively, we want to persuade people in such a way that they move to a state in which they are more likely to be successfully persuaded again. One framework that allows one to consider both current and future states is Reinforcement Learning (RL). RL with consideration of states has been applied to adapting the framing of messages for inducing healthy nutritional habits [81] or the affective behavior of a social robot teacher [145]. However, it is not yet sufficiently clear how persuasive attempts affect behavior and future states, especially after a sequence of these attempts.

An alternative to considering people's current state when choosing a persuasive strategy is to consider their characteristics such as gender, personality, and involvement in an issue. While previous work has found such characteristics to play a role (e.g., [106, 182, 229]), little work has comprehensively compared the use of user characteristics to the one of states. In addition, it may be helpful to combine these two approaches: behavior after applying a persuasive strategy in a state may differ based on user characteristics.

Our goal thus is to shed light on the effects of considering algorithm components such as states, user characteristics, or both when choosing a persuasive strategy. While previous work has tested algorithms with such components (e.g., [145, 165]), we do not yet understand the effects of individual algorithm components in practice. Therefore, rather than developing a new algorithm and comparing it to existing ones, we want to first get a better understanding of the practical impact of algorithm components. This can enable informed decisions on which components to include, which is desirable due to the larger amount of human data that needs to be collected when more components are used. Collecting more human data is costly and places a burden on users of eHealth applications that is unlikely to benefit the already low adherence rates to these applications. If data is explicitly collected by means of questions, people are likely to stop using the application if many questions are asked. For example, Pommeranz et al. [302] saw that more cognitively

demanding preference elicitation methods were seen as more effortful and liked less, which is negatively associated with technology use [377]. Moreover, while implicit data collection methods such as sensors have the potential to collect high-quality data less obtrusively, they often do not yet succeed at this. Yang et al. [395], for instance, found in the context of smoking cessation that improvements in sensing technology are needed to obtain higher data quality, lower the burden to users, and increase adherence.

Thus, to get a better understanding of the effects of algorithm components, we conducted a study in which smokers interacted with the text-based virtual coach Sam in up to five sessions. In each session, Sam assigned people a new preparatory activity for quitting smoking together with a persuasive strategy. The goal of these activities was to prepare people for change, which is typically done at the start of a behavior change intervention to increase the likelihood of successful change. Half of the activities targeted becoming more physically active as this may facilitate quitting smoking [152, 368]. In the next session, Sam asked about the effort people spent on their activity to measure their behavior. To determine people's states, Sam asked questions about people's capability, opportunity, and motivation to do an activity. Each pair of states from consecutive sessions forms a transition sample that we used to predict states after persuasive attempts. Moreover, we measured 32 characteristics covering demographics, smoking and physical activity, personality, and involvement in the activities. Based on the resulting 2366 transition samples from 671 people, we compared the effectiveness of considering states, user characteristics, or both for predicting behavior after persuasive attempts. In addition, we used simulations to assess the long-term effects of optimally persuading people based on an RL-approach that considers current and future states to maximize the effort people spend on their activities.

This chapter's contribution is evidence supporting the use of states derived from behavior change theories as well as people's overall involvement as components in persuasion algorithms. Following the stages in the development of technological health interventions defined by Brinkman [68], this justifies research on including these components in a full intervention as a next step.

BACKGROUND

PERSUASIVE STRATEGIES

Several sets of persuasive strategies have been defined. For example, Oinas-Kukkonen and Harjumaa [276] distinguish persuasive strategies such as social comparison and competition, Cialdini [93] defines six persuasive strategies such as authority, Fogg [131] differentiates between persuasive strategies related to "technology as a tool" (e.g., self-monitoring) and those related to "technology as a social actor" (e.g., language cues), and Consolvo et al. [98] describe nine persuasive strategies such as credibility. Such persuasive strategies are meant to directly influence people's motivation [242]. In addition, there are strategies that are meant to influence motivation indirectly by, for example, restructuring a person's environment. Examples include action and coping planning [347]. Notably, many of these persuasive strategies can be implemented in several ways. For instance, there are different ways of framing messages (e.g., [83, 351]) and communication modalities (e.g., [379, 382]). In this work, we focus on persuasive strategies that can be implemented in a text-based virtual coach that supports a single person in their behavior change process, without

requiring external elements such as sensor data or peers. We thereby interpret the term persuasion broadly to also include strategies that influence motivation indirectly.

STATES

Persuasive strategies are not equally effective in all circumstances: the context of a persuasive attempt matters [34, 277]. One way to describe the context is the state a persuadee is in. For example, the effectiveness of different health messages depends on a persuadee's self-efficacy [54], and the processing of messages depends on a persuadee's mood [60, 131]. Several of these state features have been formalized as influencing behavior in behavior change theories. One such theory is the behavior change wheel [242], at whose center lies the Capability-Opportunity-Motivation-Behavior (COM-B) model of behavior. This COM-B model is an overarching causal model of behavior, according to which a person's capability, motivation, and opportunity determine their behavior. Capability includes having the necessary knowledge and skills, motivation considers the brain processes influencing behavior, and opportunity captures factors outside of an individual such as support from one's social environment. The COM-B model is overarching in the sense that components of other behavior change theories can be mapped to it. For example, Fogg's behavior model specifies that ability, motivation, and a trigger need to come together for behavior to happen [131]. Ability can be mapped to "Capability" in the COM-B model, motivation to "Motivation," and the trigger to "Opportunity." The COM-B model thus provides an indication of which information about a persuadee's state needs to be considered to predict behavior after persuasive attempts. One question we pose hence is:

Q1: How well can states derived from the COM-B model predict behavior after persuasive attempts?

FUTURE STATES

In the COM-B model, a person's capability, opportunity, and motivation influence their behavior, and the behavior in turn influences their capability, opportunity, and motivation. Thus, behavior influences people's future states. This effect of behavior on a person's state has also been studied in the context of persuasion. For instance, Steward et al. [351] found that the framing of messages influences their effect on self-efficacy, and Carfora et al. [80] saw that the message type affects a person's intention to act, anticipated regret, and attitude towards behavior. Thus, persuasive strategies differ in their effect on a persuadee's state. Ideally, we would choose a persuasive strategy that positively influences a persuadee's state by, for example, increasing motivation. To do so, we need to be able to predict not just the behavior, but also the state after a persuasive attempt. We thus investigate the following question:

Q2: How well can states derived from the COM-B model predict states after persuasive attempts?

Ideally, a persuasive attempt moves a person to a future state in which they are more likely to be successfully persuaded again. Since capability, opportunity, and motivation determine behavior, the goal is that each person ultimately moves to, and then stays in, a

state with high values for these predictors of behavior. We, therefore, want to examine what happens to people's states after a sequence of persuasive attempts in the ideal case. The ideal case is that we always use the optimal persuasive strategy:

Q3: What is the effect of (multiple) optimal persuasive attempts on persuadees' states?

Being able to predict states may help to choose effective sequences of persuasive strategies, but how is behavior affected by using sequences of optimal persuasive strategies? And importantly, how much does it matter what a virtual coach says? Hence, we pose the following question:

Q4: How do optimal and sub-optimal persuasive attempts compare in their effect on behavior?

USER CHARACTERISTICS

Considering people's states is one way to capture their differing responses to persuasive strategies - considering user characteristics is another. With user characteristics we mean information about a user that changes, if at all, very slowly and irrespective of persuasive attempts. Kaptein and Eckles [182], for instance, showed that age, gender, and personality may influence which of the persuasive strategies by Cialdini [93] is most effective. Several other works have confirmed the influence of user characteristics such as the stage of behavior change [106], personality [32, 106, 154, 288, 398], age and gender [257], cultural background [289], and how people approach pleasure and pain [85]. Another potentially important user characteristic is involvement. According to the Elaboration Likelihood Model (ELM) [298], messages are more likely to be processed in detail when people are highly involved in an issue [229]. Such in-depth processing in turn is more likely to have a persistent effect [298]. Predicting the effectiveness of persuasive attempts based on user characteristics has the advantage that we need to collect data less often: in contrast to states, we do not need to gather this data before each persuasive attempt. We thus pose the following question:

Q5: How does predicting behavior based on user characteristics compare to doing so based on states?

Rather than *replacing* states with user characteristics, one may also *use both* states and characteristics. For instance, Steward et al. [351] showed that a person's need for cognition influences the effect of message types on self-efficacy. Thus, user characteristics may have an effect on the states after persuasive attempts. Intuitively, one would expect people who are more similar with regard to these user characteristics to respond more similarly to persuasive attempts. We, therefore, investigate the following question:

Q6: How does incorporating users' similarity based on characteristics, besides the consideration of states, improve the prediction of behavior?

METHODOLOGY

To answer our research questions, we developed the virtual coach Sam that persuaded people to do preparatory activities for quitting smoking based on an RL-algorithm. This algorithm for choosing persuasive strategies aimed to maximize the effort people spend on their activities over time. Data for the algorithm was collected in a longitudinal study. The data and analysis code underlying this chapter as well as the Appendix can be found online [24].

VIRTUAL COACH

We implemented the text-based virtual coach Sam that helped people prepare for quitting smoking and becoming more physically active in conversational sessions. In each session, Sam randomly proposed to users a new preparatory activity for quitting smoking or becoming more physically active such as tracking one's smoking behavior. These activities were based on components of the StopAdvisor smoking cessation intervention [243] and future-self exercises [236, 295]. After proposing the activity, Sam asked questions to determine a user's current state. This state was used as input for choosing how to persuade the user to do the activity. In the next session, Sam asked about users' experience with their activity and the effort they spent on it. Throughout the dialog, Sam used techniques from motivational interviewing [59] such as giving compliments for spending a lot of effort on activities and otherwise expressing empathy. Empathy can also facilitate forming and maintaining a relationship with a user [57], which can support behavior change [400]. Moreover, based on discussions with smoking cessation experts, Sam maintained a generally positive and encouraging attitude while trying to avoid responses that may be perceived as too enthusiastic [135]. The implementation of the virtual coach, based on Rasa and Rasa Webchat, can be found online [9]. The structure and an example of the conversational sessions as well as examples of the activities are available in the Appendix.

PERSUASION ALGORITHM

For each persuasive attempt, Sam chose a persuasive strategy based on its learned policy. In the next session, the user provided Sam with feedback by reporting the effort they spent on their activity. Formally, we can define our approach as a Markov Decision Process (MDP) $\langle S, A, R, T, \gamma \rangle$. The action space A consisted of different persuasive strategies, the reward function $R : S \times A \times S \rightarrow [-1, 1]$ was determined by the self-reported effort, $T : S \times A \times S \rightarrow [0, 1]$ described the transition function, and the discount factor γ was set to 0.85 to favor rewards obtained in the near future over rewards obtained in the more distant future. The intuition behind this value for γ was that while we wanted to persuade a user over multiple time steps successfully, a failed persuasive attempt in the near future could cause a user to become less receptive to future ones or even to drop out entirely: early success might encourage people to continue [35]. The finite state space S described the state a user was in and was captured by answers to questions about a user's capability, opportunity, and motivation to perform an activity [245]. The goal of an agent in an MDP is to learn an optimal policy $\pi^* : S \rightarrow \Pi(A)$ that maximizes the expected cumulative discounted reward $\mathbb{E}[\sum_t \gamma^t r_t]$ for acting in the given environment. The value function $V^\pi : S \rightarrow \mathbb{R}$ describes the expected cumulative discounted reward for executing π in state s and all subsequent states. V^* denotes the value function if $\pi = \pi^*$. Figure 3 in the Appendix

illustrates the algorithm idea.

State space. In each session, users provided answers to questions about their capability, opportunity, and motivation to do preparatory activities (e.g., “I feel that I need to do the activity”) on 5-point Likert scales. These questions were based on the COM-B self-evaluation questionnaire [245] with an additional question about self-efficacy based on Sniehotta et al. [346] to assess motivation (see Table 2 in the Appendix). To use the time and effort of users efficiently, we only asked those questions that we envisioned to differ between people for our domain. We transformed the questions to binary features based on whether a value was greater than or equal to the feature mean (1) or less than the feature mean (0). To further reduce the size of the state space, we used our collected data to select three out of eight features in a way that was inspired by the G-algorithm [88]. This involved iteratively selecting the feature for which the Q-values were most different when the feature is 0 compared to when the feature is 1. Besides the reduction in state space size, this feature selection also has the benefit that fewer questions would need to be answered by users in practice. The three chosen features were 1) whether users felt like they wanted to do an activity, 2) whether they had things that prompted or reminded them to do an activity, and 3) whether they felt like they needed to do an activity. The resulting state space had a size of $|S| = 2^3 = 8$. We denote states with binary strings such as 001 (here the first and second features are 0 and the third feature is 1).

Action space. Five persuasive strategies formed the action space: authority, commitment, and consensus from Cialdini [93], action planning [153], and no persuasion. The first three persuasive strategies consisted of a persuasive message (e.g., “Experts recommend *<doing activity>* to *<positive impact of activity>*.”) and a subsequent reflective question (e.g., “Which other experts, whose opinion you value, would agree with this?”). The latter was meant to increase the in-depth central processing of the persuasive message. According to the ELM, such high-effort central processing of messages leads to attitudes that are more likely to be persistent over time, resistant to counterattack, and influential in guiding thought and behavior [298]. Persuasive messages were based on the validated messages from Thomas et al. [360]. For action planning, users were asked to create an if-then plan for doing their activity based on the formulation by Sniehotta et al. [346]. Yet, rather than asking users to enter their action plans in a table, the virtual coach prompted them to create an if-then plan of the form “If *<situation>*, then I will *<do activity>*” based on Chapman et al. [89]. For the first four persuasive strategies, a message that reminded people of their new activity after the session also contained a question based on the persuasive strategy. These reminder questions were based on the ones by Schwerdtfeger et al. [338]. Repeating a persuasive attempt can also increase in-depth central processing [298]. Examples of persuasive messages and reflective questions are given in the Appendix.

REWARD

In sessions 2–5, participants were asked about the overall effort they spent on their last activity on a scale from 0 to 10, adapted from Hutchinson and Tenenbaum [170]. Based on the mean effort \bar{e} , the reward $r \in [-1, 1]$ for an effort e was computed as follows:

$$r = \begin{cases} -1 + \frac{e}{\bar{e}} & \text{if } e < \bar{e} \\ 1 - \frac{10-e}{10-\bar{e}} & \text{if } e > \bar{e} \\ 0 & \text{otherwise.} \end{cases}$$

The idea behind this reward signal was that an effort that was equal to the mean was awarded a reward of 0, and that rewards for efforts greater and lower than the mean were each equally spaced.

DATA COLLECTION

Study. We conducted a longitudinal study in which people interacted with Sam in up to five conversational sessions between 20 May 2021 and 30 June 2021. The Human Research Ethics Committee of Delft University of Technology granted ethical approval for the research (Letter of Approval number: 1523). Before the collection of data, the study was preregistered in the Open Science Framework (OSF) [14]. Participants were recruited from the online crowdsourcing platform Prolific. Eligible were people who were contemplating or preparing to quit smoking [112], smoked tobacco products daily, were fluent in English, were not part of another intervention to quit smoking, had an approval rate of at least 90% and at least one previous submission on Prolific, and provided informed consent. Participants were persuaded randomly in the first two sessions. Afterward, participants were split into four groups, each of which was persuaded based on a different policy. We provide details on these policies in Table 5 in the Appendix. 760 people started the first session, and 518 people successfully completed session 5 (see Figure 3 in the Appendix). Participant characteristics such as age and education level are shown in Table 6 in the Appendix.

Data. We gathered 2366 $\langle s, a, r, s' \rangle$ -samples from 671 people, where s is the state, a the action, r the reward, and s' the next state. Besides these transition samples, we also collected data on user characteristics. This includes 31 pre-characteristics (i.e., characteristics measured before any persuasive attempt) covering demographics, smoking, physical activity, personality, and need for cognition. Moreover, we measured users' overall involvement in their assigned activities after the five sessions. Due to dropout, we obtained involvement data for only 500 participants. The Appendix provides more information on the user characteristics we measured.

RESULTS

We now investigate each of our six research questions. For each research question, we first describe our setup, followed by our findings and the resulting answer to the research question.

Q1: How well can states derived from the COM-B model predict behavior after persuasive attempts?

Setup. Knowing the state a persuadee is in may help to predict their behavior after persuading them with different persuasive strategies (i.e., actions). The behavior in our

case is the effort people spend on their preparatory activities, which is captured by our reward function. We compared two approaches for predicting the reward: 1) the mean reward per action, and 2) the mean reward per action and state. We used leave-one-out cross-validation for the 671 participants with at least one transition sample to compare the two approaches based on the mean L_1 -error and its Bayesian 95% credible interval (CI) [279] per state. In contrast to classical confidence intervals, Bayesian CIs provide information on the most likely values (i.e., a likely range) [164]. We regard non-overlapping 95% CIs as a credible indication that values are different, both for this research question and the subsequent ones.




Figure 4.1: Left axis: Mean L_1 -error with 95% CIs for predicting rewards based on 1) the mean reward per action and 2) the mean reward per action and state. Right axis: Mean reward overall and per state.

Findings. Considering the state tends to result in lower L_1 -errors for predicting the reward than not considering the state (Figure 4.1). This makes sense, as the mean reward strongly differs between states. For example, while state 000 has a mean reward of -0.52, state 111 has one of 0.25 (see the red line in Figure 4.1). In such states with mean rewards much lower or higher than the overall mean reward, the advantage of considering states for the reward prediction is pronounced with the 95% CIs for the two approaches not overlapping. This provides a credible indication that considering states performs better. For states with mean rewards more similar to the overall mean reward, on the other hand, the 95% CIs for the two approaches tend to overlap. So there is no credible indication that one of the two approaches is better for those states.

Answer to Q1. Considering the state a persuadee is in helps to predict the effort they spend on an activity, as long as the state is one in which people tend to spend much less or more effort on activities than on average. Using features derived from the COM-B model, we obtained such states.

Q2: How well can states derived from the COM-B model predict states after persuasive attempts?

Setup. Ideally, we want to persuade a person in such a way that they move to a state in which they are likely to again be persuaded to spend a lot of effort on an activity. Therefore, we need to be able to predict the state after a persuasive attempt. Using leave-one-out cross-validation, we compared three approaches for predicting the next states for the samples from the left-out person: 1) assigning an equal probability to all states, 2) predicting that people stay in their current state, and 3) using the transition function estimated from the training data. We compared the three approaches based on the mean likelihood of the next state and its 95% CI per state. A higher likelihood suggests that next states can be predicted better.

4




Figure 4.2: Comparison of three approaches to predicting next states with regards to the mean likelihood of next states with 95% CIs for each state.

Findings. Figure 4.2 shows that considering the current state, by either predicting that people stay in their current state or assigning a probability to next states based on the estimated transition function, leads to a higher mean likelihood of next states than assigning an equal probability to all next states. This shows that state transitions do not occur uniformly at random. Notably, predicting that people stay in their current state leads to the highest mean likelihood of next states in three of the eight states. These states are states 000, 010, and 111. In each of these states, the mean for predicting that people stay in their current state is highest and the corresponding 95% CI does not overlap with the ones for the other two approaches. This shows the high probability of staying in those three states, which are states with either very low or very high mean rewards (Figure 4.1).

Answer to Q2. Our results show that considering the current state a persuadee is in helps to predict their next state after a persuasive attempt. For persuadees who are in states in which people tend to spend very little or very much effort on their activities, this next state tends to be the same as the current one. This means that if we just persuade people as we

did in the study used to collect data, we will have limited success in moving people from low-effort to higher-effort states. Though once people are in higher-effort states, they are likely to stay.

Q3: What is the effect of (multiple) optimal persuasive attempts on persuadees' states?

Setup. We would like that people ultimately move to the states in which they are most likely to be persuaded to spend a lot of effort on activities. Starting from an equal distribution of people across the states, we calculated the percentage of people in each state after following the optimal policy π^* for a certain number of time steps. π^* was computed via value iteration based on all gathered samples. Table 7 in the Appendix shows π^* .

Findings. Figure 4.3 depicts the transition function under π^* . It is evident that people tend to move to better states or stay in the best state (blue lines). With better states we mean states with higher V^* . In fact, for each state, there is a probability of at least $\frac{1}{|S|}$ that a person moves to a better state. And once people have reached the best state, which is state 111, there is a high probability of 0.8 that they stay there. However, there are some red lines in Figure 4.3 as well. These lines show that people sometimes move to worse states or stay in the worst state after being persuaded based on π^* . This happens especially for states with lower V^* such as states 000 and 010. For both of these states, there is also a relatively high probability of staying in them. For example, there is a probability of 0.41 that people stay in state 000 once there. Yet, people can also move from states with relatively high V^* to states with low V^* . For state 011, for instance, there is a probability of 0.22 that people move to the lower-value state 010.

Besides the short-term effects of following π^* , we are also interested in the long-term effects when using multiple persuasive attempts. The results of simulating transitions for applying π^* for up to 20 time steps are shown in Figure 4.4. It is evident that compared to the initial state distribution with an equal number of people in each state, more people are in state 111 and fewer people in all other states after 20 time steps. Given that state 111 is the state with the highest value, people thus tend to move to the best state. In fact, 62.61% of people and thus more than half are in state 111 after 20 time steps. However, there are always some people in the states with lower values. For example, 6.63% of people are in state 000, the state with the lowest value, after 20 time steps.

Answer to Q3. While persuading people optimally multiple times allows most of them to move to and stay in the state in which they are expected to spend the *most* effort on their activities, a few people remain in the state in which they are expected to spend the *least* effort on their activities.

Q4: How do optimal and sub-optimal persuasive attempts compare in their effect on behavior?

Setup. Once we are able to predict states, we would like to choose effective sequences of persuasive strategies. Yet, it is not clear how much the choice of persuasive strategy matters




Figure 4.3: Transition probabilities under π^* . Only transitions with a probability of at least $\frac{1}{|S|}$ are shown. We distinguish transitions to a state with a higher or highest V^* (blue), lower or lowest V^* (red), and the same V^* (black). A thicker line denotes a higher probability.

when it comes to the effort people spend on their activities over time. We calculated the mean reward per transition over time when following 1) the optimal policy π^* , 2) the worst policy π^- , and 3) the average policy π^- . π^- is a theoretical policy for comparison purposes in which each action is taken $\frac{1}{|A|}$ times for each person at each time step, where $|A|$ is the number of actions. We considered two initial state distributions, namely, the distributions across states in the first session of our study based on a) all people and b) only those people whose first reward was in the lowest 25%-percentile of all first rewards. Distributions are from our study's first session to represent a general population of people who have never been persuaded to do preparatory activities. We further specifically look at people who initially spend very little effort on their activities when persuaded randomly as at the start of our study, because it is more beneficial to coach people who are not yet performing well.

Findings. The mean reward for π^* is highest for all time steps and increases over time for an initial state distribution that is based on all people (Figure 4.5). After 100 time steps, the mean reward per transition is 0.17 and therewith above the 50%-percentile of rewards for the first session of 0.13. This means that the mean reward is increased compared to the actual mean reward we observed in session 1. In contrast, the mean reward drops for the other two policies and is only 0.02 for π^- and -0.13 for π^- after 100 time steps. The former falls in the 40–50%-percentile of rewards for the first session and the latter in the 30–40%-percentile. Hence, the difference in mean reward between the three policies

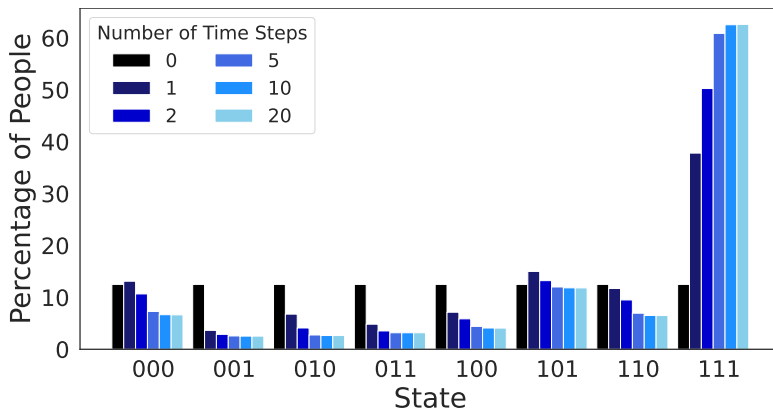


Figure 4.4: Percentage of people in each state after following π^* for varying numbers of time steps.

4

increases over time. We also observe this if we consider the initial state distribution of only those people with low rewards for the first session. For example, the difference between π^* and π^{\sim} increases from 0.08 to 0.15 and thus almost doubles.

Answer to Q4. These findings show that it matters, both for people overall and for people who are not performing well initially, how we persuade them to do preparatory activities for quitting smoking. Choosing how to persuade people based on an optimal RL-policy thereby performs better than doing so based on a worst or an average RL-policy.




Figure 4.5: Mean reward per transition over time for three policies. The initial populations are the state distribution of all people (solid line) or of only the people with a reward in the lowest 25%-percentile (dashed line) for the first session.

Q5: How does predicting behavior based on user characteristics compare to doing so based on states?

Setup. An alternative to using states to predict behavior is using user characteristics. This alternative has the advantage that data on such characteristics do not need to be collected before each persuasive attempt. To compare the use of user characteristics to the one of states, we selected three user characteristics in a similar fashion as the three state features. More precisely, we first turned the user characteristics into binary variables based on whether their value was greater than or equal to the mean (1) or less than the mean (0). Then we iteratively selected the variable with the largest difference in reward when the variable is 0 compared to when it is 1. This is because when the reward is very similar for both values of a variable, it does not improve the reward prediction very much to consider the value of the variable. We considered two different sets of candidate variables. First, we considered only the pre-characteristics and thus data that we can collect from people without having to provide any information about the activities (i.e., we excluded people's involvement in their activities). Second, we considered all characteristics (i.e., we also included involvement). The selected characteristics in the first case were the Transtheoretical Model (TTM)-stage for becoming physically active, conscientiousness, and smoking status; the ones in the second case were involvement, physical activity identity, and smoking status. For each case, we created a user characteristic state space of size $2^3 = 8$ analogously to the case of state features. Based on these state spaces, we computed the mean L_1 -error for predicting the reward using leave-one-out cross-validation. Our baselines were predicting the reward based on 1) the overall mean reward, 2) the mean reward per action, 3) and the mean reward per action and state.

Findings. Figure 4.6 shows that predicting rewards based on user characteristics in addition to actions outperforms predicting the overall mean reward. Of the two ways of predicting rewards based on user characteristics, the one that includes people's involvement in their assigned activities leads to a lower L_1 -error. More precisely, the mean L_1 -error is 0.43 for user characteristics with involvement, and 0.45 when excluding involvement. The two 95% CIs thereby do not overlap, providing a credible indication that the mean L_1 -error is lower for the former than for the latter. However, none of the two ways of predicting rewards based on user characteristics performs better than using states, with the latter leading to a mean L_1 -error of 0.41. While the 95% CI for predicting rewards based on states overlaps with the one for predicting rewards based on user characteristics including involvement, it does not overlap with the one for using only user pre-characteristics.

Answer to Q5. These results provide a credible indication that using states allows us to better predict the effort people spend on their activities than using only user characteristics that we can collect data on without having to tell people about the activities. If we include people's involvement in the activities as a user characteristic, however, there is no longer a credible indication that using states outperforms using user characteristics.

Q6: How does incorporating users' similarity based on characteristics, besides the consideration of states, improve the prediction of behavior?




Figure 4.6: Mean L_1 -error for predicting the reward with 95% CIs when considering different components for the reward prediction. *None* denotes that we predicted the reward based on the mean overall reward. Abbreviations: UPC, User pre-characteristic; UC, User characteristic; Inv., Involvement.

Setup. While user characteristics *alone* may not help to predict behavior compared to states, they may do so *in combination with states*: people with different characteristics may respond differently to a persuasive attempt in a certain state. We thus examine the effect of incorporating people’s similarity, based on user characteristics, on our ability to predict the effort people spend on their activities. We do so by weighting observed samples differently for each persuadee, whereby a larger weight is given to samples from people more similar to the persuadee. Using different user characteristics and weighting parameters, we tried a total of 68 configurations for weighting samples based on similarity (see Appendix). We here report the results for the configuration with the lowest mean L_1 -error based on leave-one-out cross-validation. This best configuration used people’s involvement in their activities to measure similarity.

Findings. Even though the mean L_1 -error is lower for incorporating users’ similarity than for the original approach without similarity, the 95% CIs overlap (see the two rightmost bars in Figure 4.6).

Answer to Q6. Incorporating users’ similarity besides the consideration of states appears to offer some improvement, but there is no credible indication that it allows us to better predict the effort users spend on preparatory activities after persuasive attempts.

DISCUSSION AND CONCLUSION

The presented study examined the use of states and user characteristics to predict the effort people spend on preparatory activities for quitting smoking after being persuaded by a virtual coach. States were based on the COM-B model and captured people’s capability, opportunity, and motivation to do an activity. Our findings suggest that states derived from the COM-B model help to predict behavior: the effort people spend on their activities clearly differs between states (Q1). In addition, considering states also helps to predict next

states (Q2). This may aid in choosing persuasive strategies that move people to future states in which they are more likely to be successfully persuaded again to spend a lot of effort on their activities. With regards to long-term effects, we find based on simulations that people tend to move to better states or stay in the best state when they are persuaded optimally (Q3). With good states we mean states in which people are expected to spend a high amount of effort over time when persuaded optimally. However, some people are always in states in which little effort tends to be spent on activities. Our simulation further shows that it matters how we persuade people (Q4). More precisely, people tend to spend more effort on activities if they are persuaded optimally based on an RL-algorithm compared to being persuaded based on the worst or an average persuasive strategy. The difference in mean effort per persuasive attempt between the three strategies increases as more persuasive attempts are made before ultimately plateauing.

4

Using user characteristics to predict behavior did not perform as well in this study. Compared to using states, we observed worse results when using user pre-characteristics alone (Q5). This is the case even though we performed experiments with 31 pre-characteristics that capture a wide range of information about demographics, smoking and physical activity, personality, and need for cognition. Additionally considering users' overall involvement in their activities led to slightly better predictions than considering pre-characteristics alone, but the predictions were still not better than for states. In line with findings by Kaptein [181] in the context of persuasive marketing messages, this suggests that predictions of behavior improve if the predictors are conceptually closer to the behavior. While pre-characteristics such as quitter self-identity may say something about the effort a person is willing to spend to prepare to quit smoking, the person's involvement in such activities is conceptually closer. And states derived from the COM-B model are even closer: they specify theoretically grounded predictors of behavior before each activity. Notably, we find that considering user characteristics *in addition* to states does offer some benefit (Q6). But even here, characteristics that are conceptually closer to the behavior we want to predict are most useful, with involvement performing best. However, it may not always be clear how to measure such conceptually closer characteristics. Involvement in our study was, for example, only measured after the persuasive attempts and could thus not inform the selection of persuasive attempts. Asking people to rate prototypes of activities in advance may be a way to address this. As involvement can change, it could also be measured in each session.

Limitations and directions for future work. The main limitation of our work is the data it is based on. While we did gather data from human subjects, we did not assess the effects of our approaches on the actual behavior or states of these humans. Instead, we performed leave-one-out cross-validation and simulations. The primary reason is that this allowed us to test a large number of approaches while staying within a reasonable budget. The best-performing approaches can then be tested in the wild in the future. When doing so, however, several additional factors may need to be addressed. This is because all of our approaches assume the transitions between states and the effort people spend on their activities to be stationary. Stationary here means that the transition probabilities and the mean effort people spend for combinations of states and actions do not change. But intuitively, such changes may occur. For instance, repeatedly sending the same persuasive

strategy may make it less effective [360], but could also help to strengthen the link between cue and response for action planning [338] or to scrutinize arguments objectively [76]. One approach to address the effects of such repetitions is the work by Mintz et al. [249] on non-stationary bandits. Moreover, once people move beyond preparatory activities and start to actually change their behavior, habits may form after several weeks [138]. Such habits may reduce the cognitive effort and awareness required to do a behavior [138]. One could address this by including information on habits in the state description (e.g., [399]).

A more general limitation of our work is the way we defined our problem. First, our state description is based on the COM-B self-evaluation questionnaire and only a subset of the questions therein. While this is a good starting point as our results show, other features, potentially derived from other theories, could be useful. For example, physical capability may play a role when people are to be persuaded to do more complex tasks such as going for a run. Importantly, however, not all people may be willing to answer many questions in each session. So it may be beneficial to either limit the number of questions for all people, or to give people the option to answer additional questions for more precise tailoring (e.g., [165]). Second, our results are based on five widely used persuasive strategies that we deemed to be applicable in our context. Given the large number of other strategies, it is possible that user characteristics play a more important role in explaining the effectiveness of other strategies. Notably, however, there is also ample literature suggesting the importance of user characteristics for the persuasive strategies we used (e.g., [288, 360, 398]). Third, we measured people's response to persuasive attempts based on the self-reported effort they spent on their activities. It would be interesting to see whether our findings also hold when a more objective measure of behavior is used. Lastly, another interesting direction to improve our model is to use Bayesian RL. This allows one to incorporate prior information about the dynamics in a flexible manner as well as to consider the uncertainty in the learned parameters when making decisions [141, 330]. For example, one can model relations between state features using a dynamic Bayesian network [330]. This may be useful, as behavior models such as COM-B specify relations between predictors of behavior.

Conclusion. We want to make informed decisions on which components to use in persuasion algorithms for eHealth applications for behavior change that are effective as well as more cost-effective and user-friendly by reducing the amount of required human data. Therefore, a better understanding of the components' individual effects on predicting behavior after persuasive attempts is welcome. We have thus compared the use of states and user characteristics, and a combination thereof, in predicting behavior after persuasive attempts in the context of preparing for quitting smoking with a virtual coach. Our results lend support to the idea of considering states and the user characteristic "involvement" in persuasion algorithms for behavior change. Research on smoking cessation can directly build on these insights and examine the use of these components in a full application. Moreover, both components seem to be domain-independent measures that could also be used in eHealth applications for other behaviors.


5

REINFORCEMENT LEARNING FOR PROPOSING SMOKING CESSATION ACTIVITIES THAT BUILD COMPETENCIES: COMBINING TWO WORLDVIEWS IN A VIRTUAL COACH

5

Background. *Reaching personal goals typically requires building competencies (e.g., insights into personal strengths), but expert health professionals and non-expert clients often think differently about which competencies are needed. Just having a virtual coach advise activities for “expert-devised” competencies may not motivate clients to carry them out, while advising only “non-expert devised” activities may not result in all required competencies being built.*

Methods. *We integrated the client and health expert worldviews in our modeling method for informing the activity selection by a virtual coach: We created a pipeline to build a reinforcement learning model for proposing activities in the context of quitting smoking. This model considers smokers’ current and future levels for expert-devised competencies as well as their beliefs about the usefulness of different competencies when choosing activities. To train the model, we conducted a study in which 542 smokers interacted with a virtual coach in five sessions spread over at least nine days. Using the data from this study, we performed simulations to systematically assess the impact of the different model components on the competencies built by smokers. Moreover, we performed paired Bayesian *t*-tests to determine the effect of persuasive activities on smokers’ usefulness beliefs.*

 **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Reinforcement learning for proposing smoking cessation activities that build competencies: Combining two worldviews in a virtual coach. *Under review.*

Results. *Our simulations show that smokers' current levels for the expert competencies and their usefulness beliefs are important to consider when building expert competencies. In fact, we saw improvements of up to 22% when considering current competencies, and an additional 13% when also accounting for usefulness beliefs. Furthermore, although we found credible evidence that persuasive activities changed smokers' usefulness beliefs, the effects might be too small to contribute in an optimal strategy for building competencies.*

Conclusion. *The worldviews of both health experts and smokers are important to consider when proposing activities for quitting smoking. We have presented an RL model that combines these worldviews and we hope that our work can be an example of incorporating different worldviews in a reinforcement learning model for building competencies. Our dataset is publicly available.*

INTRODUCTION

5 Considering that 18.5% of the disease burden in the Netherlands stems from unhealthy behavior [324], coupled with the projection that by 2060, one in three Dutch workers will need to work in healthcare to cater to the aging population [105], eHealth applications have a large potential in supporting people in changing behaviors such as physical inactivity [22], alcohol consumption [149], and unhealthy eating [303]. Since smoking alone causes 9.4% of the Dutch disease burden [324], applications supporting smoking cessation are especially welcome. To increase engagement, discuss relevant information, and form a connection with people [158, 255], such eHealth applications commonly integrate conversational agents that take the role of virtual coaches guiding people through the behavior change intervention. This guidance often includes proposing activities such as envisioning one's desired future self after quitting smoking, tracking one's smoking behavior, or creating a motivational slogan. Here, we investigate how the virtual coach should decide which activities to propose.

The virtual coach ultimately wants to propose activities that allow people to reach their behavioral goals. This often requires building competencies, such as being able to perform a breathing exercise, knowing what constitutes a healthy diet, or having self-confidence. These competencies are characteristics (e.g., knowledge, skills, mindsets, thought patterns) that when used, alone or together, result in successful behavior [118]. Following a means-end problem-solving approach, when people select subgoals (i.e., "means") to reach their goals [268, 270], the competencies that they perceive to be important are obvious "means" candidates. A person's conscious (sub)goals affect the actions they take [224] and, consequently, knowing them helps to predict the subjective usefulness and effort investment in the related actions (e.g., if a person thinks that practical knowledge will help them more to reach their goal of quitting smoking than physical fitness, they are likely to spend more effort on knowledge building than on physical activity). So knowing a person's subgoals can help a virtual coach propose actions (i.e., activities) that the person is likely to spend effort on, and that will thus build the person's competencies.

However, people do not always know which competencies are required for reaching a goal. According to the Dunning-Kruger effect [201], for example, people with little experience or knowledge regarding a task tend to overestimate their competence (e.g., because similar competencies to reach a goal are also needed to assess one's performance

[201]). Thus, when selecting subgoals, people may select different ones than experts would. This means that if the virtual coach would simply propose activities that people regard as useful, the people might never build the competencies that experts consider relevant.

Our aim was thus to develop a model for informing a virtual coach's selection of activities that build people's competencies from the perspective of experts while accounting for the fact that if people do not find an activity useful, they are unlikely to do it thoroughly. People may need to first be convinced of its usefulness. Our model hence needs to consider what people find useful and the degree to which they have built the expert competencies, both currently and in the future. One framework that allows us to formulate a model accounting for such current and future states is Reinforcement Learning (RL). RL for adaptive behavior change support [384] with consideration of current and future states has previously been applied to send running notifications [381], suggest step goals [113, 140], recommend diabetes coaching interventions [110], or choose persuasive strategies for preparing for quitting smoking [20]. Here, we investigate how RL can be used to combine two worldviews to build human competencies. Our overarching research question thus is

How can we build an RL model for building human competencies for quitting smoking that combines the views of experts and smokers?

Our pipeline for creating such an RL model consisted of the five steps shown in Figure 5.1: 1) Establishing competency-building activities as the actions that health experts recommend to reach the goal of quitting smoking, 2) obtaining the views of health experts and smokers that describe which competencies they think are built by these activities (i.e., expert-identified vs. smoker-identified competencies), 3) creating persuasive activities that can persuade smokers of the usefulness of smoker-identified competencies, 4) designing an RL model for proposing activities that optimizes the degree to which smokers build the expert-identified competencies while considering that the effort smokers spend on activities depends on which smoker-identified competencies they perceive as useful, and 5) training the model with data from a crowdsourcing study. Afterward, we evaluated the model by examining its effectiveness in building expert competencies and in changing smokers' usefulness beliefs using human data-based simulations, which is a common way to evaluate RL models [384].

This chapter contributes insights into the effects of subjective usefulness beliefs and of possibilities to change them with short persuasive activities. These highlight the importance of accounting for people's current worldviews rather than trying to change them when striving to build people's competencies. Furthermore, we provide a model for proposing competency-building activities for quitting smoking which combines the views of health experts and smokers. This model alone is not a complete behavior change intervention. Instead, it can be used to personalize elements of both face-to-face and digital smoking cessation interventions, specifically the recommendation of activities. To facilitate this, we have made the dataset used to train our model and our activities publicly available in the online repository accompanying this chapter [28]. Lastly, we hope that other researchers wishing to incorporate different worldviews in a reinforcement learning model for building competencies can use our pipeline as inspiration.




Figure 5.1: Pipeline for creating an RL model for proposing smoking cessation activities that build competencies for quitting smoking by accounting for the views of health experts and smokers.

BACKGROUND

PERSUASIVE STRATEGIES IN EHEALTH APPLICATIONS FOR BEHAVIOR CHANGE

5

Providing behavior change support over the Internet or connected technologies such as apps and text messaging, eHealth applications for behavior change commonly ask their users to do activities such as designing motivational slogans, learning about nicotine replacement therapy, or reflecting on the past week. Persuasive strategies are often used to motivate people to do these activities. Several sets of persuasive strategies have been outlined. Oinas-Kukkonen and Harjumaa [276], for instance, identify persuasive strategies such as social learning and cooperation. Cialdini [93] introduces six persuasive strategies such as consensus, while Fogg [131] distinguishes between persuasive strategies associated with “technology as a tool” (e.g., self-monitoring) and those linked to “technology as a social actor” (e.g., language cues). Consolvo et al. [98] provide nine persuasive strategies, including aesthetics. It is worth noting that many of these persuasive strategies can be applied in various ways, such as framing messages differently (e.g., [83, 351]) and using different communication methods (e.g., [379, 382]).

ALGORITHMS FOR ADAPTIVE PERSUASIVE ATTEMPTS

When applying these persuasive strategies, using a one-size-fits-all approach is unlikely to have a large effect on behavior [53, 128], as behavior change theories [240] suggest that many personal factors influence behavior. Using these factors as inspiration, previous work has developed algorithms for adapting *how* people are persuaded, *when*, and *by whom*. Work on the former includes adapting persuasive strategies to dynamic factors (e.g., people’s states derived from the COM-B model [20, 27], self-efficacy [111]) as well as more stable personal characteristics (e.g., personality, gender, and stage of change [107], age, gender, and personality [182]). Algorithmic techniques thereby range from RL (e.g., [20, 140, 299]) to recommender systems (e.g., [165]) and logistic regression (e.g., [362]). Dynamic factors have also been considered to optimize the timing and sender of persuasive attempts, for example in RL models for sending notifications for physical activity [381] and oral self-care [367] or deciding on the degree of human involvement in an intervention for chronic pain [300]. Yet, the effects of these approaches on behavior are typically small (e.g., [20, 106, 183]).

PROPOSING USEFUL ACTIVITIES

One reason for these small effects is that people do not necessarily find useful what is proposed to them. For example, Albers et al. [20] observed a large effect for personal relevance, involvement, and personal interest on the effort spent on activities for quitting smoking, in contrast to a small effect of adapting *how* people were persuaded to do the activities. Moreover, recommendations for eHealth applications for people with low socioeconomic status by Faber et al. [124] include that such applications should be perceived as useful. This is in line with the algorithmic acceptance model [340], which posits that besides convenience, usefulness predicts people's attitude toward an algorithm system and thus its actual use. Moreover, the related notion of performance expectancy is also one of the main predictors of the intention to use technology in the Unified Theory of Acceptance and Use of Technology (UTAUT) [376].

Several previous works have thus optimized *what* is proposed to people. Costa et al. [100], for instance, select activities for elderly people by generating arguments in support of the activities and deciding which one would be preferred by a person based on data from previous interactions. And Klein et al. [198] address a person's bottlenecks for behavior change (e.g., attitude) based on urgency and the degree to which they can be changed. Yet, these approaches consider only what is useful objectively or from the perspective of *experts*, not what is useful from the perspective of *users*. Users' usefulness beliefs, however, do not necessarily match the ones of experts. Even though physical activity can make it easier to quit smoking [152, 305], for example, smokers do not necessarily regard physical activity as useful for quitting smoking [21]. Thus, while we ultimately want users to do activities that are perceived as useful by experts (i.e., build the competencies experts consider relevant), we need to account for users' perceptions of usefulness. Given that RL allows us to consider people's degrees of having built the expert competencies as well as their usefulness beliefs both currently and in the future, our first analysis question, therefore, is the following:

AQ1: How effective is an RL model that combines the views of experts and smokers in building expert-identified competencies?

CHANGING BELIEFS

Rather than just *considering* people's usefulness beliefs, one can also try to *change* them. This is especially important when the virtual coach can otherwise not build all expert competencies (e.g., because people find all related activities not useful). Yet, changing beliefs can be difficult because people attribute importance to their beliefs and are hence prepared to act on and hold to these beliefs even when presented with conflicting evidence, especially when the beliefs are strong [260]. From the perspective of conceptual change, learners bring conceptions constructed from their own experiences with them that are potentially incorrect from the standpoint of established knowledge and thus hinder learning [260]. When such misconceptions exist, learning requires changes in learners' personal mental models or representations. This is because information that does not fit the learners' mental models is ignored or misunderstood [208]. Any of the persuasive strategies defined earlier can in principle be used to try to change beliefs. One theoretical framework that appears especially suitable to the health context is Protection Motivation Theory (PMT) [326]. PMT posits that a threat's severity and vulnerability on the one hand and response efficacy and

self-efficacy on the other hand influence whether people take a recommended health action. Applied to people's beliefs about the usefulness of competencies for quitting smoking, it is thus the severity of and vulnerability to the consequences of not building a competency as well as the effectiveness of and self-efficacy for building the competency that influence whether people want to build the competency. Using PMT to create persuasive activities, our second analysis question is as follows:

AQ2: How effective are persuasive activities in changing usefulness beliefs?

In the following, we describe our five pipeline steps for building our RL model for proposing smoking cessation activities as shown in Figure 5.1.

METHODS

STEP 1: ESTABLISHING COMPETENCY-BUILDING ACTIVITIES

The first step was to understand which activities are currently used by health experts to prepare smokers for quitting smoking. These activities build the competencies for quitting smoking that health experts consider relevant, even if the competencies have not been standardized. Based on discussions with health experts, the activities by Albers et al. [21], the behavior change techniques by Michie et al. [241], and smoking cessation material by organizations such as the National Cancer Institute and the Dutch Trimbos Institute, we obtained 44 preparatory activities (e.g., envisioning quitting smoking as a fighting match, thinking of past successes, or writing a positive diary). Some activities addressed becoming more physically active since this can make it easier to quit smoking [152, 305]. A health psychologist and smoking cessation expert checked the activities to ensure they were suitable and clear. Table S1 in the Appendix [28] lists all preparatory activities.

STEP 2: OBTAINING THE VIEWS OF HEALTH EXPERTS AND SMOKERS

Having established the 44 preparatory activities, the next step was to determine how health experts and smokers view them. The “views” in our case are the two sets of competencies for quitting smoking that experts and smokers think are built by the activities. If we know which competencies the activities build according to *experts*, our model can keep track of the extent to which smokers have already built the different competencies (yellow rectangles in Figure 5.2) and choose activities that help build missing competencies. On the other hand, by knowing which competencies *smokers* think are built by the activities, our model can consider which competencies smokers find useful when choosing activities (red rectangles in Figure 5.2).

To this end, we conducted two repertory grid studies, one with experts and one with smokers. Based on personal construct theory [192], the goal of the repertory grid technique is to explore personal construct systems, or, in other words, see the world as other people see it [134]. The people whose world one would like to see, in our case experts and smokers, were given three preparatory activities and asked to divide them into two groups based on considering how two activities are alike in some way but different from the third activity. After providing a label for each resulting group, participants rated each of the 44 activities on a seven-point scale from “not at all related to *<label>*” to “strongly related to *<label>*”

for each label.

For each repertory grid study, these ratings served as input for an exploratory factor analysis with an oblique rotation, minimum residuals as extraction method as recommended by Izquierdo et al. [174], and the common cutoff value of 0.4 [225]. For each possible number of factors identified with the scree method and parallel analysis, we examined the resulting factors according to their theoretical and practical plausibility. The final factors describing the views of experts and smokers were chosen based on this examination by two researchers. All our factors satisfy the recommendation that independent of the sample size, factors are reliable as long as the average of the four largest loadings is greater than 0.60 [350].

The studies were preregistered in the Open Science Framework [15] and approved by the Human Research Ethics Committee of Delft University of Technology (Letter of Approval number: 2338, date: 27 June 2022). All participants gave digital informed consent. Below we provide more information on the two studies and their results.

VIEW OF HEALTH EXPERTS

First, we investigated which competencies the preparatory activities build according to experts.

Approach. To account for some variation in the perceptions of smoking cessation experts, we conducted a repertory grid study with four smoking cessation experts who described their backgrounds as “psychology” ($N = 2$), “health and medical psychology” ($N = 1$), and “general practitioner” ($N = 1$). On a scale from 0 (“No expertise at all”) to 10 (“Extremely strong expertise”), the experts reported having strong expertise in coaching for behavior change ($M = 8.25$, $SD = 0.96$), coaching for quitting smoking ($M = 8.25$, $SD = 0.50$), and coaching for becoming more physically active ($M = 7.50$, $SD = 1.00$). Each expert was asked to do the task four times, each time with a new set of preparatory activities, using the question “When it comes to competencies for quitting smoking that smokers build by doing the activities, how are two activities alike in some way but different from the third activity?” To ensure that the experts understood the question, they had to pass an attention check on the question’s meaning after being provided both a definition and an example.

Results. The exploratory factor analysis on the 32 items (i.e., labels and corresponding activity ratings) led to three factors: 1) practical skills – clear future identity, 2) motivation – knowledge, and 3) insights into personal strengths – insights into personal weaknesses. The six factor endpoints gave us six individual competencies for quitting smoking (yellow rectangles in Figure 5.2). Table 5.1 provides examples of the labels and explanations by experts mapped to the first factor. For examples for all three factors refer to Table S2 in the Appendix.

VIEW OF SMOKERS

Next, we explored smokers’ views on competencies for quitting smoking built by preparatory activities.

Approach. Aiming for 4 participants per combination of values for age range (3 levels), gender (2 levels), weekly exercise amount (3 levels), and smoking frequency (2 levels), we conducted an online crowdsourcing study with $4 \times (3 \times 2 \times 3 \times 2) = 144$ daily smokers who were contemplating or preparing to quit smoking [112]. Each participant received two sets of three preparatory activities. 76 participants were instructed to divide the activities in a

Table 5.1: Factor loadings, labels, and explanations of the three items with the most positive and negative factor loadings for two of the factors found through the repertory grid studies (i.e., one factor from the repertory grid study with health experts and one factor from the study with smokers). The labels and explanations are direct, uncorrected quotes from participants.

FACTOR FOR EXPERT-IDENTIFIED COMPETENCIES: (+) PRACTICAL SKILLS – CLEAR	
FUTURE IDENTITY (–)	
0.85	practical skills: preparing practically for activities
0.80	strategies: through social learning the individual might find out strategies to successfully change behavior themselves
0.78	problem solving: This helps smokers to think ahead and come up with solutions for barriers
-0.63	Identity: These activities help to strengthen feared and ideal future selves
-0.63	identity: These activities help to envision the ideal and feared future selves
-0.64	Future-self: Future-selves can act as powerful motivators
FACTOR FOR SMOKER-IDENTIFIED COMPETENCIES: (+) SELF-EFFICACY – PRACTICAL	
KNOWLEDGE (–)	
0.92	Self motivation: Activities that focus on motivation and planning
0.85	Mindset: Activities that will help you with the right mindset needed to quit smoking
0.79	Mindsets: getting into the right mindset
-0.75	Problem solving: Quitting smoking and being more active can some times be hard and some barriers may show up, so being able to find a way around those is important
-0.84	Knowledge: Consuming educational content, gaining knowledge
-0.86	Practical: Learning real techniques to quit smoking

set into two groups based on the question “When it comes to competencies for quitting smoking that smokers build by doing the activities, how are two activities alike in some way but different from the third activity?” The other participants were asked to divide the activities based on what people have to do for an activity (e.g., visualize, record) for future research. For each resulting group, participants provided a label as well as an explanation of the label. To increase the validity of the data, participants had to pass a multiple-choice attention check question on the meaning of competencies or doing something for an activity after seeing both an explanation and an example. Based on our observations from two small pilot studies, we suspected that not all participants followed the instructions. Therefore, the first author coded all obtained labels as 1) competency ($N = 153$), 2) a way of doing a preparatory activity ($N = 254$), or 3) unclear ($N = 169$) by looking at both the labels and their explanations. To examine the reliability of the coding, we made use of a second coder. The first author trained this second coder by explaining the coding of 12 example labels and giving feedback on six rounds of coding ten labels. Based on the subsequent independent coding of 100 labels by the second coder, we obtained a Cohen’s κ of 0.55 and a Brennan-Prediger κ of 0.56, indicating moderate agreement [209]. Finally, in the exploratory factor analysis, we included only those activity ratings whose labels the first

coder had coded as competency.

Results. We obtained five factors from whose endpoints we created nine competencies (red rectangles in Figure 5.2)¹. For example, we created the competencies “self-efficacy” and “practical knowledge” from the first factor (Table 5.1).

STEP 3: CREATING PERSUASIVE ACTIVITIES

If the virtual coach knows which smoker-identified competencies a person finds useful, it can propose competency-building activities that they find more useful. However, the virtual coach can also try to change the person’s usefulness beliefs about these competencies, especially when building all expert-identified competencies is otherwise not possible because the person finds all related activities not useful. We thus designed nine persuasive activities to each persuade smokers of the usefulness of one of the nine smoker-identified competencies. As we worked with two different worldviews (i.e., of smokers and experts) we accepted to some extent that smokers might do the “right” thing for the wrong reasons. Still, we first verified that each smoker-identified competency could be mapped to one or more expert-identified competencies (e.g., “self-efficacy” could be mapped to “motivation” and “insights into personal strengths”). This ensured that the content of the persuasive activities was also grounded in the views of the experts. Each persuasive activity was then built to persuade people of the usefulness of one smoker-identified competency by addressing elements from PMT (e.g., see Table S4 in the Appendix). A health psychologist and smoking cessation expert read through all activities to ensure that they were suitable and clear. The activities can be found in Table S5 in the Appendix.

STEP 4: DESIGNING THE MODEL

Next, we designed a model that a virtual coach can use to choose activities. We can define our approach as a Markov Decision Process (MDP) $\langle S, A, R, T, \gamma \rangle$. The action space A consisted of 53 activities (i.e., 44 preparatory and 9 persuasive activities), the reward function $R : S \times A \rightarrow [0, 6]^2$ was determined by the self-reported effort spent on activities and by the activities’ contributions to the expert-identified competencies, $T : S \times A \times S \rightarrow [0, 1]$ was the transition function, and the discount factor γ was set to 0.85 to favor rewards obtained earlier over rewards obtained later. The finite state space S described the state a person was in and was captured by their beliefs about the usefulness of smoker-identified competencies, their capability and opportunity, and their levels for expert-identified competencies. The goal of an agent in an MDP is to learn an optimal policy $\pi^* : S \rightarrow \Pi(A)$ that maximizes the expected cumulative discounted reward $\mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r_t]$ for acting in the environment. The optimal Q-value function $Q^* : S \times A \rightarrow \mathbb{R}$ describes the expected cumulative discounted reward for executing action a in state s and π^* in all subsequent states. Figure 5.2 visualizes the model.

STATE SPACE

The state space had three components: 1) people’s beliefs about the usefulness of the nine smoker-identified competencies, 2) their capability and opportunity, and 3) their levels for

¹As the endpoints of one factor were negations of each other, we created a single competency for that factor.

²In practice, no activity fully contributes to all six expert competencies, which means that a reward of 6 does not occur.

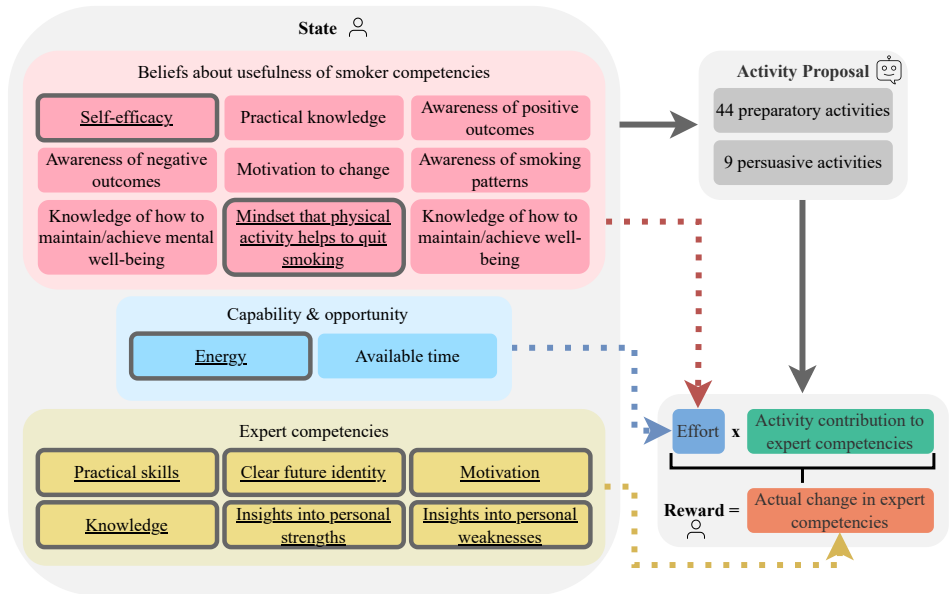


Figure 5.2: Visualization of the RL model. The state features with thick borders are used in the final trained model.

the expert-identified competencies. We included people’s capability and opportunity as they predict behavior according to the COM-B model [242] alongside motivation, which was captured by the usefulness beliefs.

To infer a person’s state, the virtual coach would ask questions during its interaction with them. For the usefulness beliefs, people would answer nine questions after the prompt “Please rate how you think the following 9 factors affect quitting smoking. Answer on a scale from -10 to 10, where -10 indicates that quitting smoking is made a lot harder and 10 indicates that quitting smoking is made a lot easier. 0 indicates ‘neutral’” in each session with the virtual coach. To ensure that the questions are understandable for smokers, we used the terminology smokers used in the repertory grid study together with some specific examples they gave (Table S6 in the Appendix). For example, for the competency “practical knowledge,” we used the formulation “practical preparation (e.g., learning how to relieve stress, knowing effects of nicotine, getting organized).”

To measure people’s capability and opportunity to do preparatory activities, the virtual coach would further ask people about their energy and available time on 11-point scales.

Lastly, people’s levels for the six expert-identified competencies (i.e., degrees of having built these competencies) were initialized to 0 and subsequently updated to a value in the set $\{0, 0.33, 0.67, 1\}$ to obtain a reasonably sized state space. The updating process is described in more detail for the transition function.

ACTION SPACE

There were 53 actions: the 44 preparatory activities for quitting smoking and the 9 activities meant to persuade people of the usefulness of the smoker-identified competencies for quitting smoking.

REWARD

The intuition behind the reward signal is that people should ultimately build the competencies identified by experts, that these competencies are only built if people do their activities thoroughly, and that there is an upper limit to building each competency (i.e., at some point, a competency has been fully built). The idea thus is that the reward captures the *actual* increase in these competencies. Therefore, the reward was, accounting for an upper limit of 1 for each expert-identified competency, the product of two measures: 1) an activity's contribution to the expert-identified competencies, and 2) the effort people spent on the activity.

For the first measure, we computed the contribution of each preparatory activity to the expert-identified competencies based on the factor loadings from the repertory grid study with experts, scaled to the interval $[0, 1]$ (Table S7 in the Appendix). The contributions of the persuasive activities were set to 0 as these activities do not build any competencies but only aimed at changing usefulness beliefs. For the second measure, people were in each session asked about the overall effort they spent on their last activity on a scale from 0 to 10, adapted from Hutchinson and Tenenbaum [170] as also used by Albers et al. [20]. The effort responses were also scaled to the interval $[0, 1]$, with the mean effort mapped to 0.5 so that values for efforts greater and lower than the mean were each equally spaced. To reduce the amount of required data, we grouped the preparatory activities into five clusters to predict the effort. To this end, we performed k-means clustering using the smokers' ratings of the preparatory activities' contribution to smoker-identified competencies from the repertory grid study. This means that preparatory activities seen as contributing similarly to the smoker-identified competencies were grouped together (Table S1 in the Appendix).

Given a maximum value of 1 for each expert-identified competency, the *actual* increase in the expert-identified competencies for person i after spending effort $e_{a,i}$ on activity a was then calculated as $\sum_{j=0}^5 \min\{pc_{a,j,i}, 1 - comp_{j,i}\}$, where $pc_{a,j,i}$ is the possible contribution $cont_{a,j} \times e_{a,i}$ mapped to the possible levels for the expert competencies $\{0, 0.33, 0.67, 1\}$, $cont_{a,j}$ is the contribution of activity a to expert-identified competency j , and $comp_{j,i}$ is the current level of competency j for person i (Figure 5.3). Note that since the preparatory activities are only clustered for predicting one component of the reward (i.e., the effort), the reward can be different for each of the 44 preparatory activities.

TRANSITION FUNCTION

The transitions between values for the user-inquired state features (i.e., the usefulness beliefs and people's capability and opportunity) were learned from data, whereby transitions for one state feature were considered independent of the values of other features. To reduce the amount of data required to reliably predict the transitions, the preparatory activities were grouped into the same clusters as for the effort prediction when predicting the next usefulness beliefs and people's capability and opportunity. People's levels for the six expert-identified competencies, on the other hand, were updated up to a maximum value of 1 based on 1) an activity's contribution to the expert-identified competencies according to the data from the repertory grid study with experts and 2) the effort people spent on the activity, mapped to the interval $[0, 1]$. The raw updated level of person i for an expert-identified competency j after spending effort $e_{a,i}$ on activity a was then computed as $comp_{j,i,t+1} = comp_{j,i,t} + e_{a,i} \times cont_{a,j}$, where $cont_{a,j}$ is the contribution of

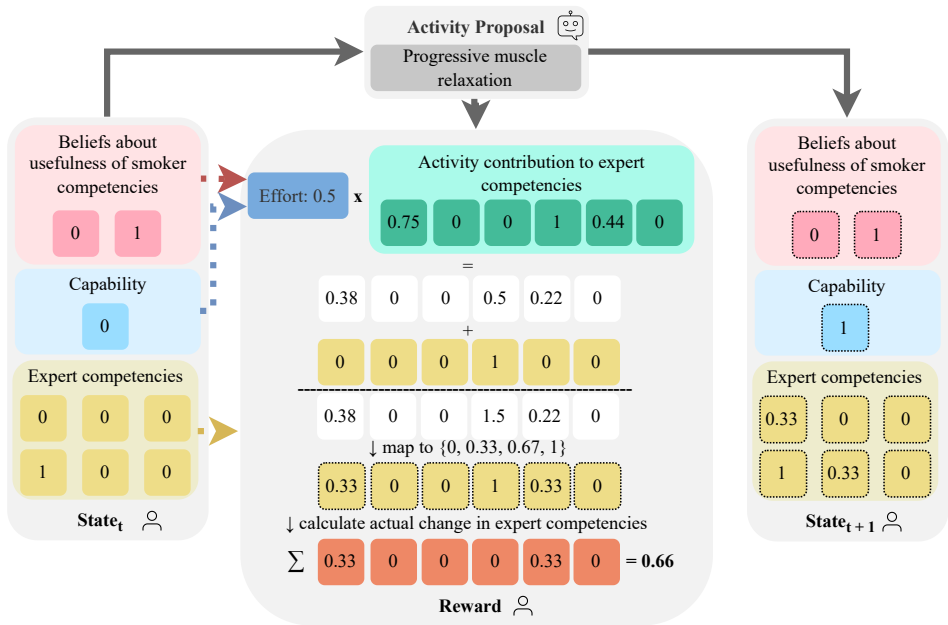


Figure 5.3: Example of how the reward of 0.66 and the next expert competency levels are computed after a person spent an effort of 0.5 on the activity “progressive muscle relaxation.”

activity a to expert-identified competency j (Figure 5.3). The resulting raw value was then mapped to the closest value in $\{0, 0.33, 0.67, 1\}$ to get the next value for the state feature. The contributions of all preparatory activities to the expert-identified competencies are provided in the Appendix. The nine persuasive activities do not contribute to the expert-identified competencies. Note that the change in levels for the expert-identified competencies also informed the reward computation.

STEP 5: TRAINING THE MODEL

To train our model, we conducted a study in which daily smokers interacted with the virtual coach Mel in five sessions spread over at least nine days. While Mel proposed randomly chosen activities in this study to facilitate training the model with the collected data, we envision Mel ultimately using our trained model to choose activities.

Virtual coach. We implemented the text-based virtual coach Mel. Mel introduced itself as wanting to prepare people for quitting smoking and becoming more physically active, with the latter possibly facilitating the former. In each session, Mel determined people’s current state by asking about their beliefs regarding the usefulness of the competencies for quitting smoking identified by smokers as well as their available time and energy. Afterward, Mel proposed a new preparatory or persuasive activity. In the next session, which participants were invited to about two days later, Mel asked about the effort people spent on their activity as well as their experience with it. In its conversation structure and style, Mel was closely based on the virtual coach Sam [9], whose scripted dialogs were

developed for another smoking cessation study and were overall perceived positively by its users [21, 23]. Like Sam, Mel gave compliments for spending a lot of effort on activities, expressed empathy otherwise, and kept a generally positive and encouraging attitude. The Rasa-based implementation of the virtual coach [10] as well as a demo video [11] are available online. The conversation structure is depicted in the Appendix.

Study. We conducted a study in which people interacted with Mel in up to five conversational sessions between 21 July and 27 August 2023. The Human Research Ethics Committee of Delft University of Technology granted ethical approval for the research (Letter of Approval number: 2939) on 31 March 2023. Before data collection, the study was preregistered [16]. Participants were recruited from the crowdsourcing platform Prolific Academic. Eligible were people who were contemplating or preparing to quit smoking [112], smoked tobacco products daily, were fluent in English, were not part of another intervention to quit smoking, had not participated in our repertory grid studies, and gave digital informed consent. To choose a new activity, Mel first randomly chose from the five preparatory activity clusters and nine persuasive activities. If a cluster was chosen, Mel then randomly selected one of the activities mapped to it. For completing each study part, participants were paid based on the minimum payment rules on Prolific (i.e., six GBP per hour). They were also informed that their payment was independent of them completing their activities. 682 people started the first session and 349 people completed session 5 (Figure S2 in the Appendix). Participant characteristics such as age, gender, and smoking frequency at the start of the study are shown in Table S8 in the Appendix. Two days ($T1$, $N = 324$) and eight weeks ($T2$, $N = 245$) after the last session, participants' smoking frequency was lower ($T1 - T0$: $M = -0.67$, 95%-HDI = $[-0.96, -0.38]$; $T2 - T0$: $M = -0.96$, 95%-HDI = $[-1.30, -0.63]$; 8-point scale) and quitter self-identity³ [235] higher ($T1 - T0$: $M = 0.21$, 95%-HDI = $[0.15, 0.27]$; $T2 - T0$: $M = 0.10$, 95%-HDI = $[0.01, 0.19]$; 5-point scale) than at the start of the study ($T0$) (see Table S9 in the Appendix).

Collected data. We gathered 1710 $\langle s, a, r, s' \rangle$ -samples from 542 people, where s is the state, a the action, r the reward, and s' the next state. Participants spent an average effort of 5.58 ($SD = 2.86$) on their activities, with the mean effort per preparatory activity cluster ranging from 5.30 ($SD = 2.85$) to 6.14 ($SD = 2.72$) and the one per persuasive activity from 5.19 ($SD = 2.94$) to 5.95 ($SD = 2.66$) (Table S10 in the Appendix). In sessions 2–5, participants were asked about their likelihood of having returned to the session in case of an unpaid smoking cessation program on a scale from -5 (“definitely would have quit the program”) to 5 (“definitely would have returned to this session”). The mean of these responses was 1.44 ($SD = 2.74$) in session 2 and 1.80 ($SD = 2.94$) in session 5, with responses from the full range of the scale in each session.

State space reduction. Even when using only binary state features, using all nine usefulness beliefs and both capability and opportunity features in our model would lead to $2^{11} = 2048$ possible values for those state features that influence dynamics components that we need to estimate from data (i.e., the effort and the transitions between these features). To reduce the size of the state space and hence the amount of required data, we transformed the usefulness beliefs and the capability and opportunity features into binary features based on whether a value was greater than or equal to the sample mean (1) or less than the mean

³Since the reliability was sufficiently high for $T1$ (Cronbach's $\alpha = 0.67$, $N = 324$), $T2$ (Cronbach's $\alpha = 0.78$, $N = 324$), and $T3$ (Cronbach's $\alpha = 0.83$, $N = 245$), we used the mean of the three items as an index measure.

(0). Moreover, we used our data to select three features in a way that was inspired by the G-algorithm [88]. This involved iteratively selecting the feature for which the effort-based Q-values were most different when the feature is 0 compared to when the feature is 1. Besides the reduction in state space size, this selection also has the benefit that users would need to answer fewer questions in practice. The selected features were: 1) belief about the usefulness of “self-efficacy,” 2) belief about the usefulness of the competency “mindset that physical activity helps to quit smoking,” and 3) energy. The final model had $2^3 = 8$ different values for those state features that influence dynamics components we need to estimate from data as well as $4^6 = 4096$ different values for the expert competency features that influence the dynamics deterministically. The entire state space hence had size $|S| = 8 \times 4096 = 32768$. Figure S3 in the Appendix shows the mean effort per combination of values for the three selected user-inquired features.

Model training. Using the reward and transition functions estimated from our data, we computed an 0.001-optimal policy and corresponding Q^* with Gauss-Seidel value iteration from the Python MDP Toolbox. If an optimal activity had already been proposed to a person in the past, an activity with the next highest Q^* was proposed.

5

RESULTS

We now investigate our two analysis questions. For each analysis question, we first describe our approach, followed by our findings and the resulting answer. The data and analysis code underlying this chapter are available online [28].

AQ1: BUILDING EXPERT COMPETENCIES

Setup. To examine how each of our model components contributes to building people’s competencies, as seen by experts, we compared the effects of optimal policies of ablated versions of our model that included or excluded specific components. For this, we analyzed results from human data-based simulations, examining each time how 1000 simulated people would progress in their competency development over multiple interactions with a virtual coach that bases its activity advice on a specific policy. To obtain a realistic population, these simulated people were initially distributed across the user-inquired state features following the distribution we observed in the first session of our data collection study. We created ablated versions of our model by removing increasingly more components from the model: first the learned transitions to the next user-inquired feature values ($-uf'$), then the transitions to the next expert competency levels ($-ec'$), then the current user-inquired feature values ($-uf$), and so forth. The first five optimal policies we compared, we computed based on these five model versions: 1) the full model (π^*), 2) assuming that all next user-inquired feature values are equally likely ($\pi^{-uf'}$), 3) not considering any future states ($\pi^{-ec',uf'}$), 4) considering only a person’s current value for the expert competencies ($\pi^{-ec',uf',uf}$), and 5) considering only current user-inquired feature values by randomly picking one of the activities in the preparatory activity cluster with the highest expected effort ($\pi^{-ec',uf',ec}$). A sixth policy was choosing preparatory activities uniformly at random (π^r). Two policies are thus derived from a full RL model (π^* and $\pi^{-uf'}$), three policies from models that are contextual bandits ($\pi^{-ec',uf'}$, $\pi^{-ec',uf',uf}$, and $\pi^{-ec',uf',ec}$), and one policy from a simple baseline model. The five learned optimal policies all differ in some




Figure 5.4: Fraction of expert competencies built after different numbers of proposed activities when using policies based on different models to choose activities. The lines for the first two policies overlap almost completely.

states. For example, the optimal activity indices in the eight possible starting states are 31, 16, 31, 31, 31, 4, 16, 5 for π^* , 31, 16, 31, 4, 31, 4, 16, 21 for $\pi^{-u_{f'}}$, and 8, 9, 9, 32, 31, 4, 32, 21 for $\pi^{-ec', u_{f'}}$. It was striking that none of the five learned policies included a persuasive activity.

Results. Using more model components generally allows the expert competencies to be built more quickly (Figure 5.4). However, removing the learned transitions to the next user-inquired feature values (*dashed yellow line*) and the transitions to the next expert competencies (*dotted green line*) each leads to at most a small deterioration (1%). Proposing a random preparatory activity or a preparatory activity that people are expected to spend the most effort on based on their current user-inquired feature values performs worst. After five proposed activities, using π^* has allowed people to build 91% of the competencies, $\pi^{-u_{f'}}$ 90%, $\pi^{-ec', u_{f'}}$ 89%, $\pi^{-ec', u_{f'}, u_{f'}}$ 76%, $\pi^{-ec', u_{f'}, ec}$ 56%, and π^r 54%. The largest drops in performance result from removing either the current user-inquired feature values (13%, effect size Cohen's $h = 0.34$) or the current expert competency levels (32%, Cohen's $h = 0.76$). These are small to medium effects according to the classification guidelines by Cohen [96].

Answer to AQ1. With just five proposed activities, an RL model that combines the views of experts and smokers allows smokers to build 91% of expert-identified competencies. 54% could be attributed to assigning any preparatory activity, 22% to considering current levels of expert competencies, 13% to also considering smokers' current usefulness beliefs and energy, and 1% each to further considering future levels of expert competencies and smokers' future usefulness beliefs and energy.

AQ2: CHANGING USEFULNESS BELIEFS

Setup. Our analysis for AQ1 showed that considering smokers' current usefulness beliefs helps to choose activities that build expert competencies more quickly. This suggests that people's usefulness beliefs impact the effort they spend on preparatory activities (see also Figure S3 in the Appendix). It would thus be beneficial if we could change people's usefulness beliefs so that people spend more effort on activities. In AQ1, we already saw

that optimal policies propose only preparatory and not persuasive activities. So persuasive activities are likely not as effective in changing people’s usefulness beliefs as we envisioned them to be. To investigate whether there is any effect of the persuasive activities on the usefulness beliefs, we performed paired Bayesian t -tests using the Bayesian First Aid package [43], comparing the usefulness belief corresponding to a persuasive activity before and after people were assigned the activity.

Table 5.2: Impact of the nine persuasive activities on the corresponding usefulness beliefs based on paired Bayesian t -tests.

Competency	Mean (SD)	95% HDI	Prob > 0
Self-efficacy	0.35 (0.17)	[0.01, 0.69]	0.98
Practical knowledge	0.17 (0.18)	[-0.18, 0.52]	0.84
Awareness of positive outcomes	0.31 (0.15)	[0.00, 0.61]	0.98
Awareness of negative outcomes	1.39 (0.52)	[0.45, 2.42]	> 0.9995
Motivation to change	0.58 (0.28)	[0.05, 1.13]	0.99
Knowledge of how to maintain/ achieve mental well-being	0.40 (0.18)	[0.04, 0.76]	0.99
Mindset that physical activity helps to quit smoking	0.43 (0.20)	[0.04, 0.81]	0.99
Awareness of smoking patterns	1.16 (0.27)	[0.63, 1.70]	> 0.9995
Knowledge of how to maintain/ achieve well-being	0.23 (0.21)	[-0.18, 0.63]	0.87

Abbreviations: SD, Standard deviation; HDI, Highest density interval.

Results. Table 5.2 shows that all persuasive activities positively impact the corresponding usefulness beliefs with a posterior probability of at least 0.84. For two activities, this probability is even > 0.9995, which can be evaluated as “nearing certainty” that the effect is positive [90]. Effect sizes (Cohen’s d) thereby range from 0.01 for “practical knowledge” to 0.45 for “awareness of negative outcomes” and are hence less than small to at most small according to the classification guidelines by Cohen [97].

Answer to AQ2. While persuasive activities overall do have a positive effect on the usefulness beliefs, the effects seem to be too small for optimal policies to suggest them instead of actual preparatory activities. This might at least be the case when user-inquired features are binary as in the case of our simulation.

DISCUSSION

We have presented a five-step pipeline for creating an RL model for building human competencies for quitting smoking that combines the views of health experts and smokers. To train the model, we conducted a crowdsourcing study with 542 daily smokers doing randomly chosen preparatory and persuasive activities for quitting smoking in up to five sessions. It is interesting to point out that even though participants only did *random* activities, their quitter self-identity was somewhat higher and their smoking frequency lower than before the study both two days and eight weeks after the last session. There is

some evidence from waitlist conditions in random waitlist-controlled trials that quitter self-identity [295] and smoking frequency [73, 185, 392] remain relatively constant if there is no intervention. While this suggests that already doing randomly chosen preparatory activities might increase quitter self-identity and reduce smoking frequency in smokers, caution is required as external factors such as time could have contributed to the observed effect.

Based on the data from the study, we performed simulations to assess the benefit of each RL model component in building expert-identified competencies in smokers. Within just five interactions with the virtual coach, proposing activities based on the full model can allow smokers to build 91% of expert competencies (AQ1). All model components contribute to this. People's current state based on both their levels for the expert competencies and their usefulness beliefs and energy is most important. In fact, we saw improvements of up to 22% when considering current competencies, and an additional 13% when also accounting for current usefulness beliefs and energy. The contributions of the learned transitions to the next user-inquired feature values and the transitions to the next expert competencies, on the other hand, are small. This confirms the value of considering, if not the *future*, at least the *current* views of smokers and experts and thus the benefit of a psychology-informed model which, analogously to physics-informed algorithms incorporating physical laws to facilitate learning [187], includes psychological information.

Given that considering the transitions to next values for the user-inquired features hardly contributes to building expert competencies, it seems that the effect of *preparatory activities* on usefulness beliefs is small. Our analysis of the effects of *persuasive activities* indicates the same for these activities (AQ2). Specifically, while there is a high probability that all persuasive activities positively impact the corresponding usefulness beliefs, the effects are too small for an optimal policy to suggest persuasive instead of preparatory activities. This is in line with the often small effects of individual persuasive attempts on behavior (e.g., [20, 183, 349]). Multiple persuasive attempts might hence be needed to clearly change a usefulness belief. As long as users' current usefulness beliefs still make them do activities that eventually build the expert competencies, however, it might be more effective for the virtual coach to focus on proposing activities that people already regard as useful than trying to change usefulness beliefs. At least in our simulations, people still succeed in building the expert competencies. Notably, however, our data is from a study in which participants were paid for completing the sessions in which they were assigned activities. Even though participants were informed that their payment was not contingent on completing the activities, they might have felt at least some obligation to do the activities. In a real-world application without such payments, participants who do not find the assigned activities useful might simply drop out and thus never build the expert competencies. This is supported by the observation that in each session some of our participants said that they would definitely have quit the program if it was unpaid. Future work could incorporate these dropout responses into the reward signal. Furthermore, our simulation was based on the average effort spent on activities in certain states. However, there might be individuals who sometimes or generally spend very little effort on activities. Given our binned expert competency levels, such individuals might only very slowly or never build any expert competencies in our model (e.g., see Figure S5 in the Appendix).

Besides the reliance on crowdsourced data, our work has several other data-related

limitations. First, due to the high cost of collecting human data like ours, we obtained a relatively limited dataset of 1710 samples. We thus turned our user-inquired features into binary features and used only a subset of usefulness beliefs as state features. It could be that using more values for more usefulness beliefs could capture the small positive effects that persuasive activities have on usefulness beliefs. However, a larger dataset is necessary to reliably capture such effects. Notably, using more usefulness beliefs in the model would require asking users more questions in each session, which might be more effortful [302] and thus affect technology use negatively [376]. Second, we grouped preparatory activities perceived similarly by smokers into clusters to more reliably predict the effort and the transitions to values of user-inquired features. A larger dataset could also allow one to remove this clustering and instead capture the effects of individual preparatory activities. Further capturing individual (e.g., [183, 249, 331]) or trait-based (e.g., [20, 27]) differences between smokers might also be worthwhile. Lastly, since we took an offline RL approach, our insights are dependent on our dataset [251]. While human data-based simulations are a common way to assess RL models [384], future work should compare policies trained based on different model components in a random controlled trial with activities assigned to real people to see how well our insights from human data-based simulations generalize.

5

Regarding our model formulation, one limitation further is that our model did not capture delayed effects of activities beyond the next state that could arise because it takes people more time to thoroughly reflect on the activities and change their usefulness beliefs accordingly. Defining surrogate rewards could be a way to address this (e.g., [367]). Moreover, we used domain knowledge to incorporate structure into our RL model and create a relational decomposition that specifies relations between model components [251]. This reduces the amount of data needed to train the model but limits what can be learned. For example, while we specified that the effort does not depend on the expert competency levels, it could be that building one competency depends on other competencies (e.g., as in educational systems [316] or games [58]). Future work should examine how well our modeling assumptions hold. Furthermore, while constructivism posits that each individual has their own personal construct system with which they see the world [157], we defined a joint construct system to capture the view of all smokers on preparatory activities for quitting smoking. Intuitively, however, the construct systems of individual smokers might differ as they are shaped by personal and smoking-specific experiences (e.g., previous quit attempts). Examining these differences in the future would be interesting. Lastly, our model did not account for some activities having a logical order (e.g., first tracking one's behavior before considering what to change).

CONCLUSIONS

To help a virtual coach propose effective activities, we have presented an RL model for building human competencies for quitting smoking that combines the worldviews of health experts and smokers. Simulations based on data from a multi-part study with 542 daily smokers support the use of both worldviews in the model, with small to medium effects for smokers' current usefulness beliefs and energy as well as their current levels for expert competencies. Moreover, while it is possible to positively affect smokers' usefulness beliefs using short persuasive activities, the effect of these persuasive activities is too small for them to be considered instead of activities that directly aim to build competencies. These


findings suggest that it might be more effective to look for the most competency-building activities among the activities people find useful than to try to persuade people of the usefulness of other activities.

6

THE IMPACT OF HUMAN FEEDBACK IN A CHATBOT-BASED SMOKING CESSATION INTERVENTION: AN EMPIRICAL STUDY INTO PSYCHOLOGICAL, ECONOMIC, AND ETHICAL FACTORS

6

Integrating human support with chatbot-based behavior change interventions raises three challenges: 1) attuning the support to an individual's state (e.g., motivation) for enhanced engagement, 2) limiting the use of the concerning human resources for enhanced efficiency, and 3) optimizing outcomes on ethical aspects (e.g., fairness). Therefore, we conducted a study in which 679 smokers and vapers had a 20% chance of receiving human feedback between five chatbot sessions. We find that having received feedback increases retention and effort spent on preparatory activities. However, analyzing a Reinforcement Learning (RL) model fit on the data shows there are also states where not providing feedback is better. Even this "standard" benefit-maximizing RL model is value-laden. It not only prioritizes people who would benefit most but also those who are already doing well and want feedback. We show how four other ethical principles can be incorporated to favor other smoker subgroups, yet, interdependencies exist.

 **Nele Albers**, Francisco S Melo, Mark A Neerincx, Olya Kudina and Willem-Paul Brinkman. The impact of human feedback in a chatbot-based smoking cessation intervention: An empirical study into psychological, economic, and ethical factors. *Under review.*

INTRODUCTION

Suppose Alice, Bob, Charlie, and 163 others are trying to quit smoking with an eHealth application. This is a sensible approach as meta-analyses on eHealth applications show that they are effective for quitting smoking [126, 215] and changing other lifestyle behaviors [343] by providing support anywhere anytime [230]. Now in such an eHealth application, the *virtual* coach Kai can support Alice, Bob, Charlie, and the other people wishing to quit. Next to Kai, the *human* coach Hannah can give additional feedback to increase people's adherence to the virtual coach intervention. Since Hannah has agreed with her manager that she only has time to give six-minute feedback to around 60 clients per day, she every day needs to choose clients to help. To increase everybody's chance of successfully quitting, it might be ideal to give feedback to all three clients alternately. However, Hannah expects, for example, that Alice would most benefit from the feedback, so she is tempted to give more feedback to Alice. On the other hand, Alice is already doing very well, whereas Bob is struggling with the intervention. So even though Bob would benefit less than Alice, Hannah feels some obligation to help him so that he also has a chance to succeed at quitting smoking. But to make it more complicated, Bob told Hannah that he does not appreciate human feedback, so she feels like she should respect his wish. This leaves Charlie: he expressed high appreciation for human feedback but Hannah thinks that the human feedback will mainly distract him from the virtual coach intervention. So who should Hannah give feedback to?

6

The question about good allocation of health professionals' time is not a trivial one with projections indicating that by 2060, healthcare expenses will need to rise to 18% of the Dutch GDP [1] and one in three people will need to work in healthcare to support the aging population [105]. There is thus an evident need to make healthcare more scalable and cost-effective. One promising way are eHealth applications [290] which provide elements of healthcare over the Internet or connected technologies such as apps and text messaging and thus reduce the need for scarce and costly human healthcare staff. Since 9.4% of the disease burden in countries such as the Netherlands results from smoking [324], eHealth applications for quitting smoking are especially welcome. While these applications commonly integrate conversational agents that take the role of *virtual* coaches [158], combining virtual with *human* support can be effective. Such human support can complement the strengths of virtual coaches, not only in terms of responsibility, risk, and oversight [63], but also by providing more tailored support [21, 250], addressing things other than health in people's lives [250], and being more empathetic [21, 250]. Human support can also lead to higher credibility [188], which may make application features such as personalization more effective [223]. Moreover, people may be more engaged [115] and feel more accountable [21, 210, 250] and satisfied [327] when a human coach is involved, which may help address the dropout eHealth applications often suffer from [123, 147]. However, too much human support can reduce motivation, one's sense of self-worth, autonomy, and opportunities for learning [337]. In light of these considerations, we want to examine the effects of human support in a chatbot-based intervention for quitting smoking (*RQ1*). We are specifically interested in its effects on engagement due to the central role engagement plays in intervention effectiveness [396]. Thereby, before delving into long-term effects, we will first concentrate on the short-term effects.

Whether human support is effective may depend on how motivated, confident, or

tired a person is (i.e., their state). This state refers to a person's condition or status at a specific moment in time, characterized by relative stability in its components [37]. For example, people with high intrinsic motivation who are already adhering to and engaging with the intervention might perceive human support as controlling or questioning their ability or competence [252]. At best providing human support to people with high intrinsic motivation only makes it ineffective, but it can also undermine people's intrinsic motivation and thus lead to lower adherence in the future [252]. So whether human support is given in a person's *current* state can also affect a person's *future* state and thus the effectiveness of future human support. One approach that allows us to consider both current and future states is Reinforcement Learning (RL) [357]. While RL has previously been used to allocate human support in eHealth applications, the algorithms tend to not consider people's future states (e.g., [300, 301]) and current states (e.g., [132, 133]). With such a consideration of current and future states, RL for adaptive behavior change support [384] has previously been applied to various domains, such as timing running notifications [381], suggesting step goals [113, 140], selecting messages for diabetes prevention [195], and choosing persuasive messages for preparing for quitting smoking [20]. Here we investigate whether RL with a consideration of current and future states can also be used to allocate human support for long-term effectiveness. Our second research question thus concerns how effective human support for quitting smoking is in the long term (RQ2), again with focus on engagement.

One crucial difference when allocating human support compared to adapting other elements of eHealth applications is that human support is limited: Since one of the main motivators for creating eHealth applications is that they reduce the need for scarce and costly healthcare staff, adding large amounts of human support to eHealth applications defeats this purpose. Current eHealth applications commonly provide human support on demand (e.g., [51, 52, 115, 321]). While this does not explicitly limit the amount of human support, many people do not use optional human support [337] due to reasons such as preference for self-management [21] or lack of perceived usefulness [21, 52] or time [52]. Therefore, since the requested amount remains relatively low in practice, limits such as maximum amounts of support per person (e.g., up to three text messages per day [52]) may not be necessary. A downside of this approach is that people who do not ask for support may still benefit from it [337]. To address this, some applications supplement the on-demand support with a certain minimum level of human support per person (e.g., [51, 52, 115]). This, in turn, has the disadvantage that support might be allocated to people who do not benefit or even are put off by it. Current RL algorithms for allocating human support [132, 133, 300, 301] do allocate human support to those who would most benefit from it by optimizing for measures such as meeting calorie goals [132] or reducing opioid analgesic misuse risk [301]. However, they do not necessarily respect people's autonomy by not assigning support to people who do not want it, which brings us back to the idea of providing human support on demand. Moreover, unlike applications providing a certain amount of support to each person, current RL algorithms do not ensure equal treatment by assigning the same amount of support to everybody.

Evidently, allocating limited human smoking cessation support requires moral considerations regarding who gets to benefit from human support, and thus who may increase their chance of successful smoking cessation and positive health outcomes. More generally, the question of allocating limited human support can be seen as a problem of allocating scarce

medical resources, for which Persad et al. [297] formulated four categories of ethical principles: 1) treating people equally, 2) favoring the worst-off, 3) maximizing the total benefit of the client population, and 4) promoting and rewarding social usefulness (Table 6.1). Each of these categories can be implemented in different ways. For example, treating people equally could mean allocating support randomly or on a first-come, first-served basis [297]. In addition, resource allocation differs in whether it respects people's autonomy. According to self-determination theory, the satisfaction of autonomy together with competence and relatedness enhances motivation and well-being [332]. Moreover, autonomy is, besides justice, non-maleficence, and beneficence, one of the four main principles of biomedical ethics [49]. Goodman and Houk [144] argue that a patient should have the ultimate say in whether to proceed with a treatment they are offered if their autonomy is to be respected. Applied to our context, this could mean that a person who does not want human support should not be given support.

Persad et al. [297] claim that no single principle is sufficient to include all morally relevant considerations. However, it is not obvious how the principles should be combined. Even in the intensive care unit (ICU) triage context, which has been well-studied during the COVID-19 pandemic, guidelines differ between countries [3]. For example, while maximizing benefits is a central triage ground in many countries, countries disagree on whether some priority should be given to younger people regardless of medical arguments [3]. Before deciding how to allocate human support for quitting smoking, it would help to first better understand the consequences of using different ethical allocation principles, and combinations thereof, for various subgroups. Importantly, the focus here is on first *understanding* the effects of including a wide range of different principles rather than *proposing* that certain principles should be used. Therefore, our third research question is what ethically (ir)responsible consequences may occur from using the learning algorithm and can be mitigated with algorithmic refinements (RQ3).

The context in which we investigate our three research questions is *preparing* for quitting smoking or vaping with a virtual coach. Specifically, we envision a virtual coach that prepares people for quitting smoking or vaping by assigning them preparatory activities such as visualizing one's desired future self or thinking of strategies for dealing with cravings. The goal of these activities is to prepare people for change, which is often done at the start of a behavior change intervention (e.g., [266, 278, 374]) to increase the likelihood of successful change thereafter. We focus on this first part of a behavior change intervention since feedback effects are more difficult to assess in a complete behavior change intervention with many other (adaptive) elements such as goal-setting or social support [45]. In the few days between sessions with the virtual coach, smokers may sometimes receive a feedback message from a human coach to motivate and keep them engaged (Figure 6.1). To assess the effect of the feedback, the virtual coach asks smokers about the effort spent on their activities in the next session. The choice of who receives feedback thereby is guided by psychological (i.e., factors describing an individual's state that influence the effects of human feedback), economic (i.e., cost of human feedback), and ethical (i.e., which ethical principles are used) concerns.

Since the effectiveness of eHealth applications for behavior change hinges on user engagement [396], we here examine whether receiving human feedback increases people's engagement with the preparatory activities proposed by the virtual coach, specifically, the

Table 6.1: Allocation principles by Persad et al. [297], with the addition of autonomy [49, 332], and corresponding examples in the context of allocating human support for preparing for quitting smoking.

Allocation principle	Example for preparing for quitting smoking
TREATING PEOPLE EQUALLY	
Lottery	- Random allocation
First-come, first-served	- Longest time since last human support
FAVORING THE WORST-OFF: PRIORITARIANISM	
Sickest first	- Least likely to successfully prepare for quitting smoking without human support - Most likely to experience negative consequences of smoking in the future without human support
Youngest first	- Youngest first
MAXIMIZING TOTAL BENEFITS: UTILITARIANISM	
Prognosis	- Largest increase in chance of successfully preparing for quitting smoking because of the support - Largest reduction in negative consequences of smoking in the future because of the support
PROMOTING AND REWARDING SOCIAL USEFULNESS	
Instrumental value	- Largest value to society in the future (e.g., healthcare staff, workers producing influenza vaccine, people who agree to improve their health and thus use fewer resources in the future)
Reciprocity	- Past usefulness or sacrifice (e.g., past organ donors, people who participated in vaccine research, people who made healthy lifestyle choices that reduced their need for resources in the past)
RESPECTING AUTONOMY	
Autonomy	- Highest appreciation of human support

effort spent on the activities and the likelihood of returning to the next session. Therefore, the effects of the intervention on actual smoking cessation lie outside the scope of our study. For our analyses, we use 2,326 interaction samples from a longitudinal study in which 679 daily smokers and vapers interacted with a text-based virtual coach in up to five sessions and sometimes received human feedback between sessions. Moreover, we perform human data-based simulations with an RL model to examine the long-term effects of human feedback given in different states and under varying cost settings. Even this “standard” RL model is value-laden. It prioritizes people who are already doing well and want feedback. We conclude by showing that building an RL model for allocating limited human feedback necessitates making ethical decisions and illustrating how different ethical

Hello, it's great that you have made the decision to stop smoking! You're very right, if we want to achieve a goal we need to be consistent. However, it is important to give yourself some compassion, too. We cannot expect ourselves to be perfect at something from the very beginning. Be gentle with yourself through this process and celebrate the small successes. Exercise really is a great way to divert yourself from wanting to smoke and ignoring that little voice in your head that one more cigarette won't be so bad. I hope that the next activity will help you with getting some motivation to get more physically active! I can see that you are motivated to stop smoking and are confident in your ability to do so, your mindset is already there! Keep up the good work and remember to be kind to yourself.

Best wishes,
Karina & Goda on behalf of the Perfect Fit Smoking Cessation Team



Figure 6.1: Example of human feedback sent to participants in our intervention.

6

principles can be incorporated to favor different smoker subgroups. For this, we use the principles by Persad et al. [297] with the addition of autonomy to capture a wide range of concrete principles used in practice. Despite the breadth of these principles, we do not exclude the possibility that other principles could be relevant in this context. To combine principles, we can use weights. For example, a total priority score could be computed as $0.8 \times \textit{prognosis} + 0.2 \times \textit{youngest first}$, thereby assigning a weight of 0.8 to the benefit someone would receive from human support, a weight of 0.2 to their age-based priority, and a weight of 0 to their priorities based on any other principle. Rather than proposing optimal weights, our goal is to show the influence of different weights and, in particular, weights based on smokers' preferences. For the latter, we use the preferences of the 449 participants of our post-questionnaire. We hope that our work helps eHealth application designers make the ethical decisions needed for allocating human support. To make our findings accessible to a broader audience, a lay summary can be found in the Appendix.

RESULTS

We collected 2,326 interaction samples from 679 people. On a scale from 0 to 10, participants reported spending an average effort of 5.74 ($SD = 2.75$) on their activities, ranging from 4.80 ($SD = 2.72$, $N = 71$) for the activity "Role model for others by quitting smoking" to 6.62 ($SD = 2.42$, $N = 21$) for the activity "How friends and/or family will receive one's desired future self after quitting smoking"¹. In sessions 2–5, participants were asked about their likelihood of having returned to the session in case of an unpaid smoking cessation program on a scale from –5 ("definitely would have quit the program") to 5 ("definitely would have returned

¹Refer to Table C.1 for the mean effort for each of the 37 preparatory activities.

to this session”). The mean of these responses was 1.57 ($SD = 2.73$, $N = 679$) in session 2 and 2.11 ($SD = 2.68$, $N = 504$) in session 5, with responses from the full range of the scale in each session. Participants seem to have read most of the messages as they clicked on the reading confirmation links for 81.72% of the 380 interaction samples with human feedback. Moreover, of the 270 people in the post-questionnaire who received at least one human feedback message, 82.59% said they noticed the human feedback messages, and 81.85% that they read the human feedback messages either sometimes (11.48%) or always (70.37%). Receiving human feedback appears to have not influenced the actual return to the next session. For example, the percentage of people answering at least one state question in session 2 is 87.26% for people who received feedback after session 1 and 86.09% for those who did not.

RQ1: SHORT-TERM EFFECTS OF HUMAN FEEDBACK ON ENGAGEMENT

Direct effect of human feedback. We can frame posterior probabilities as “bets” we can place with varying confidence levels [90]. Here, we can place a casual bet that human feedback increases the effort people spend on their activities ($b = 0.08$, 95% HDI = [-0.13, 0.29], $P(b > 0) = 0.76$, Cohen’s $d = 0.05$), whereas it is not worth betting on a positive effect for the return likelihood ($b = 0.03$, 95% HDI = [-0.15, 0.22], $P(b > 0) = 0.64$, Cohen’s $d = 0.02$). The effect sizes can be evaluated as less than small according to the guidelines by Cohen [97].

Delayed effect of human feedback. People spend more effort on their activities ($b = 0.39$, 95% HDI = [0.17, 0.62], $P(b > 0) > 0.9995$, Cohen’s $d = 0.25$) and are more likely to return to the next session ($b = 0.29$, 95% HDI = [0.08, 0.50], $P(b > 0) > 0.995$, Cohen’s $d = 0.18$) when they have received human feedback in the past (e.g., two sessions ago). The posterior probabilities can be classified as nearing certainty that the effect of having received feedback is positive for effort and a very strong bet that it is positive for the return likelihood. The effect size is small for effort and less than small for the return likelihood [97].

Delayed effect of multiple human feedback messages. We can place a casual bet that having received multiple human feedback messages increases the effort people spend on their activities ($b = 0.14$, 95% HDI = [-0.21, 0.50], $P(b > 0) = 0.80$, Cohen’s $d = 0.09$), whereas it is not worth betting on a positive effect for the return likelihood ($b = 0.06$, 95% HDI = [-0.24, 0.37], $P(b > 0) = 0.66$, Cohen’s $d = 0.04$). Both effect sizes are less than small [97].

RQ2: LONG-TERM EFFECTS OF OPTIMALLY ALLOCATED HUMAN FEEDBACK ON ENGAGEMENT

Long-term effects of unlimited human feedback. Figure 6.2 depicts the mean effort-based reward per activity assignment for four policies that provide different amounts of human feedback. Comparing these policies in Figure 6.2 shows that providing more human feedback *generally* leads to a higher mean reward per activity assignment. The mean rewards per activity assignment for the policies of never, half the time, and always providing feedback are 0.53, 0.54, and 0.55 after a single time step, and 0.57, 0.61, and 0.64

after 100 time steps. The latter three correspond to efforts of 6.32, 6.68, and 6.97 and thus to increases of the mean effort of 5.74 by 10.18%, 16.41%, and 21.52%. Always providing feedback thus ultimately leads to an effort that is by 0.65 scale points higher than never providing feedback. Looking at the optimal policy $\pi^{*,0}$ in Figure 6.2, however, reveals that providing more human feedback is not *always* better: the optimal policy does not always assign human feedback but ultimately leads to a higher mean reward per activity assignment than always providing human feedback. Specifically, there are two states where it is optimal not to give human feedback. In both of these states, the perceived importance of preparing to quit is high, and either the self-efficacy for preparing to quit or the human feedback appreciation low. The mean reward per activity assignment for the optimal policy after 100 time steps is 0.66, which corresponds to an effort of 7.08 and an increase of the mean effort by 23.36%.




Figure 6.2: Mean reward per previous activity assignment over time for four policies. People are initially distributed across the 12 base states as in the first session of our study.

Long-term effects of limited human feedback. Table 6.2 shows the states in which people would receive human feedback when different human feedback costs are used. States with low or medium perceived importance and high self-efficacy are those where human feedback has the most positive effect on engagement in the long run, as those states are still allocated feedback for the highest costs. Increasing the cost generally leads to only a small drop in reward (Figure 6.3a) even if a lot less human feedback is given (Figure 6.3b). The notable exception is increasing the cost from 0.07 to 0.09, which results

in a clear drop in the mean effort spent on activities because then people who have high perceived importance, self-efficacy, and feedback appreciation (i.e., those in state 211) no longer receive feedback (Table 6.2).

Table 6.2: States with human feedback (✓) for optimal policies $\pi^{*,c}$, computed based on different costs c . We refer to the 12 states with three-digit strings representing the values of the three state features: 1) perceived importance, 2) self-efficacy, and 3) human feedback appreciation.

Policy	Low importance				Medium importance				High importance			
	000	001	010	011	100	101	110	111	200	201	210	211
$\pi^{*,0}$	✓	✓	✓	✓	✓	✓	✓	✓	✓			✓
$\pi^{*,0.02}$	✓		✓	✓	✓	✓	✓	✓	✓			✓
$\pi^{*,0.03}$			✓	✓	✓	✓	✓	✓	✓			✓
$\pi^{*,0.05}$			✓	✓	✓	✓	✓	✓	✓			✓
$\pi^{*,0.055}$			✓	✓	✓	✓	✓	✓				✓
$\pi^{*,0.07}$			✓	✓	✓	✓	✓	✓				✓
$\pi^{*,0.09}$			✓	✓			✓	✓				✓
$\pi^{*,0.1}$			✓				✓	✓				
$\pi^{*,0.102}$							✓	✓				
$\pi^{*,0.12}$								✓				
$\pi^{*,0.18}$												

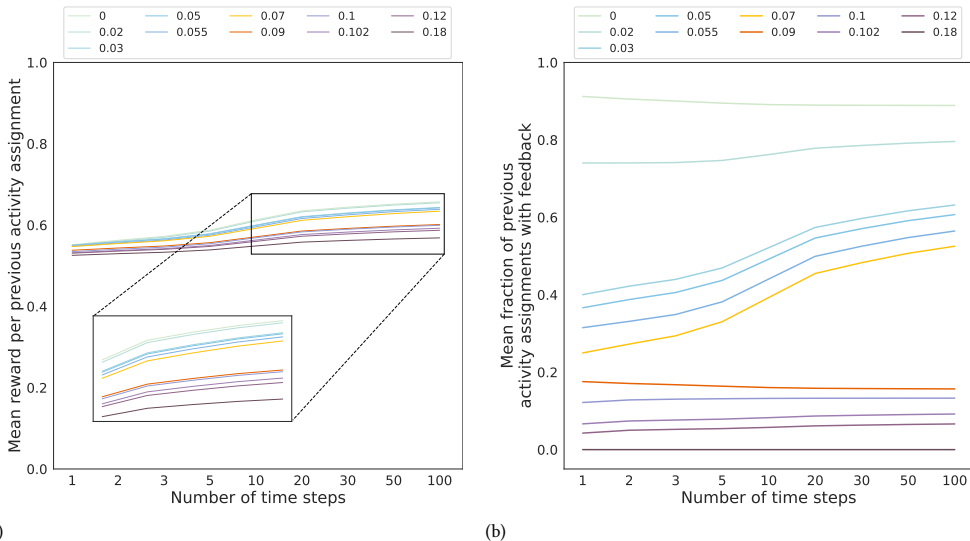


Figure 6.3: Mean (a) reward per previous activity assignment and (b) fraction of previous activity assignments with human feedback for different human feedback costs. Up to costs of 0.07, large amounts of human feedback can be saved (b) without large drops in reward (a).

RQ3: EFFECT OF DIFFERENT ETHICAL ALLOCATION PRINCIPLES ON HUMAN FEEDBACK RECEIVED BY SMOKER SUBGROUPS

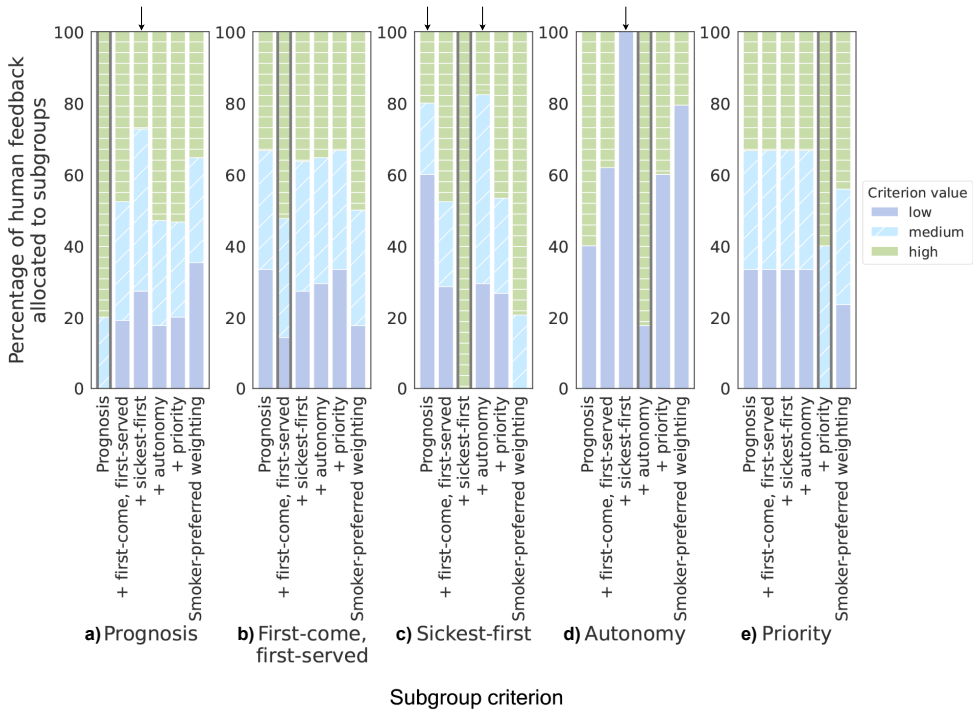


Figure 6.4: Percentage of human feedback allocated to smoker subgroups with high, medium, and low values for the criteria a) *prognosis*, b) *first-come, first-served*, c) *sickest first*, d) *autonomy*, and e) *priority* when using six different policies to allocate feedback. These policies are based on either a single reward (*prognosis*), the weighted sum of two rewards (e.g., *prognosis + autonomy*), or the weighted sum of all five rewards using the smoker-preferred weighting shown in Table 6.4 (*smoker-preferred weighting*). There are only low and high values for the criterion *autonomy*. For each criterion, the bar of the policy that specifically addresses only that criterion besides *prognosis* is highlighted with a thick border. Arrows indicate conflicts between ethical principles.

Figure 6.4 depicts the percentage of human feedback allocated to smoker subgroups, distinguished based on the criteria *prognosis*, *first-come, first-served*, *sickest first*, *autonomy*, and *priority*, when using different policies to allocate feedback. For example, when we allocate feedback using only a *prognosis*-based reward, around 80% of the feedback is given to people with a high value for the criterion *prognosis* (i.e., those most benefitting from the feedback, first bar in Figure 6.4a). On the other hand, if we use a weighted combination of *prognosis* and *autonomy* to allocate feedback, only about 50% of the feedback is given to people who would benefit most from it (fourth bar in Figure 6.4a). In general, Figure 6.4 shows that the reward functions included in the RL model influence the way smoker subgroups are allocated feedback. While adding to the *prognosis*-based reward a single

auxiliary reward based on another ethical principle allows more feedback to be allocated according to that principle (see bars with thick borders in Figure 6.4), ethical principles can be conflicting. Specifically, allocating feedback according to *sickest first* leads to much worse performance for *autonomy* and *prognosis* and vice versa (see the arrows in Figure 6.4). This suggests that people who are not doing well are less likely to want and benefit from human feedback than people who are doing well. Considering smokers' allocation preferences (Table 6.4), we observe that an optimal policy based on smoker-preferred principle weighting allocates less feedback to those who would most benefit from it (*prognosis*, Figure 6.4a) and value it the most (*autonomy*, Figure 6.4d), compared to an optimal *prognosis*-based policy. On the other hand, considering smokers' preferences means allocating *more* feedback to people who have spent a lot of time since the last feedback (*first-come, first-served*, Figure 6.4b), are expected to spend the lowest effort on their activities without feedback (*sickest first*, Figure 6.4c), and have an individual characteristic-based priority (*priority*, Figure 6.4e). This shows that smokers' preferences differ from what is optimal when we just strive to optimize population-level health outcomes. Free-text descriptions of smokers' allocation preferences from our post-questionnaire underline this. For example, one participant wrote, "Help those first who need the most help but also be equal like give feedback at least once to each person." Table C.7 shows for each principle a quote from a participant who afterward indicated a relatively high preference for the principle.

DISCUSSION

Short-term effects on engagement. The presented study examined the psychological, economic, and ethical factors that arise when combining human support in the form of human feedback messages with a chatbot-based smoking cessation intervention. While our inferential analyses provide only limited support for a positive *direct* effect of human feedback on the effort people spend on their activities and their return likelihood, they do provide strong support for a positive *delayed* effect, albeit one that is at most small (*RQ1*). That is, people who have received human feedback at some point in the past (e.g., two sessions ago) spend more effort and are more likely to return to the next chatbot session. Such small positive effects of human support have also been found in other contexts such as mental health [47, 361]. Our observed delayed effect instead of a direct effect makes sense in the tradition of primarily seeing human support as enhancing accountability or adherence [252, 337]: human feedback received in the past also affects accountability at least in the short-term future. It would be interesting to see whether more direct effects are observed for different kinds of support, such as one that encourages curiosity [143] or aims to deepen the skills or knowledge taught in an activity [337].

Beyond the effect of having received human feedback, we only find limited support for a positive effect of having received *multiple* feedback messages. This is in line with previous work in the context of an Internet-based intervention for panic disorder [197], a web- and mobile app-based mental health promotion intervention [319], and a computerized psychological intervention for comorbid mental health and substance use problems [190]. Thus, more human support beyond some minimal level of human involvement might not have additional benefits. Notably, even people who never received human feedback in our study spent a somewhat higher effort ($M = 5.83$, $SD = 2.43$, $95\% \text{ HDI} = [5.49, 6.17]$) than

participants of our two previous chatbot-based smoking cessation studies without human involvement [14, 16] ($M = 5.60$, $SD = 2.56$, 95% HDI = [5.43, 5.78]) on activities common to all three studies. Therefore, future research on the effect of merely knowing that a human *could* give feedback is warranted.

Long-term effects on engagement. Regarding long-term effects, simulations with an RL model show that while providing more human feedback generally leads to a higher effort-based reward over time, there are also two states where it is better in the long run not to provide feedback ($RQ2$). These are states where the perceived importance is high, and either self-efficacy or the appreciation of human feedback is low. This underlines that whether providing human support positively affects engagement in the long run depends on a person's state described by psychological factors. While our examined psychological factors primarily relate to engagement, future research could also examine factors related to other reasons why people may fail to benefit from a behavior change application and hence benefit from human support (e.g., lack of knowledge on how to use the intervention) [337]. Interesting is also that even in states where human feedback appreciation is low, providing human feedback is often beneficial in the long run (Table 6.2). While not receiving one's preferred type of support does not necessarily negatively influence adherence and effectiveness [319], it could also be that people who are strongly opposed to certain types of feedback drop out at the very start of the intervention [319]. Such self-selection might also have occurred in our study. Since we observed a median human feedback appreciation rating of 6 on a scale from -10 to 10 , this seems likely.

6

While it is optimal to give human feedback in almost all states, about 50% of human feedback can be removed without a large drop in effort spent on activities (Figure 6.3). When the human feedback cost is high, the largest long-term increase in effort can be obtained by giving feedback to people with low or medium perceived importance of preparing for quitting and high self-efficacy, which suggests an importance-enhancing effect of human feedback (Figure C.2). While the cost of our human support was already relatively low because our human coaches were Master's students in Psychology who spent only five to ten minutes per message, previous work suggests other types of support that could be used to potentially further reduce cost. For example, given that the qualification of human coaches providing support does not seem to play a large role in internet-based mental health interventions [47], it has been proposed that technicians [365] or expert patients [327] could also provide feedback. Other alternatives include tips and personal stories from other users [115] as well as group sessions [207], even though the greater effort for users and loss of anonymity associated with the latter might lead to higher dropout [319]. Furthermore, besides changing the type of support, one could also reduce costs by examining why around 20% of people did not read feedback messages. It could be that some of them needed less feedback. Lastly, one could investigate how human coaches can write feedback more quickly. For example, it was more difficult for our human coaches to write feedback when people did not provide much information in their introduction texts and activity experiences.

More generally, however, sharing user data with human coaches raises privacy and transparency concerns [210]. For example, while our participants were told in the ethics board-approved informed consent form that human coaches could potentially read their

anonymized introduction texts, activity experiences, and background information (e.g., baseline smoking/vaping frequency), we need to keep in mind that participants might not actually (fully) read or understand informed consent forms [296]. Our participants were reminded during the virtual coach sessions only that their activity experiences could be read by human coaches, so not all participants might have been aware of *all* the information that the human coaches could access. One way to address this is to let users see which data their human coach has access to and explicitly share data with them [115]. This, however, could mean that users share very little data, which could come at the expense of effective feedback as the experiences of our human coaches and also the therapists in the study by Doherty et al. [115] suggest. Explaining to users the benefit of sharing more information could help. Future research should also investigate the effect of feedback quality.

Effect of different ethical allocation principles. When it comes to allocating limited human feedback, the RL model that optimizes the effort spent on activities is essentially value-laden [200], because it takes a stand on the ethical question of who should receive feedback when resources are limited. Specifically, in a setting where only around 35% of people receive feedback, the model gives less feedback to people who have spent the longest time since the last feedback (*first-come, first-served*), are least likely to engage with the intervention without feedback (*sickest first*), or should be prioritized based on individual characteristics such as age or future or past value to society (*youngest first, instrumental value, reciprocity*) than the 449 participants of our post-questionnaire would prefer (Table 6.4) (*RQ3*). Our results show that we can define additional reward functions to favor those people. Notably, however, we find that favoring one principle can come at the expense of another. For example, since our participants with low engagement often did not want feedback, giving more feedback to them means doing worse on *autonomy*. These conflicts between ethical principles raise the question of how allocation decisions can be made in practice. In a medical context, it has often been advocated that multidisciplinary teams should decide how to allocate limited medical resources [3, 177]. In case these teams are unsure of their preferences over different allocation principles, an explicitly multi-objective approach can be taken where teams are presented with and guided in choosing from the set of allocations that are optimal under different weights given to allocation principles [94, 156]. To increase the acceptance of the selected allocation, the underlying rationale should be explained [4, 200]. Individual allocation outcomes could further be explained with directions toward how other outcomes can be obtained [344]. Lastly, in case the model is continuously updated with new data collected during the intervention, approaches for normative monitoring of the model (e.g., [281, 282]) might be useful to determine and potentially adjust the degrees to which allocation principles are followed.

Limitations. In addition to the limitations related to the type and quality of feedback as well as the possible self-selection of participants who favor receiving human feedback, our work is further limited in several ways. First, our RL-based analysis of the long-term effects of human feedback is based on human data-based simulations. Although this is a common way to assess RL models [384], future work should compare the long-term effects of different ways of allocating human feedback in a randomized controlled trial to see how well our findings generalize. Such a trial could also integrate our preparatory

activities into a full behavior change intervention to confirm whether higher engagement with preparatory activities is associated with more successful smoking cessation. Just like the engagement with these activities, their effectiveness may also depend on people's states. We are currently investigating this in a separate study [16]. Analyzing the data on people's experiences with their activities that we publish with this chapter [30] can further provide insights on when and for whom preparatory activities are effective. Additional field studies could also assess the predicted ethical implications of our analyses. Second, while our analysis of short-term effects for *RQ1* shows a delayed effect of providing feedback, it is not clear whether the state transitions in our RL model are sufficient to capture these delayed effects. Future works should investigate this. Third, since participants were paid for completing the conversational sessions, they might have felt at least some obligation to complete the activities even though they were informed that their payment was not contingent on their reported completion. As such, there might already have been some accountability to the intervention, which might have limited the additional effect of human feedback. Interestingly, feedback is allocated in fewer and different states when using the return likelihood instead of the effort as the basis for the reward (Table C.3), which underlines the importance of testing the effect of human feedback in an intervention without payments. Fourth, as our participants were relatively young and well-educated (Table C.9), our findings might not generalize to an older and less educated sample. Such a sample might, for example, benefit from more support, particularly that which also addresses low eHealth literacy and limited technology skills [8, 318]. Future research could examine how our findings generalize to people who are underrepresented in our sample. Furthermore, there is potentially limited generalizability of our findings to other clinical settings. Even though similar analyses could be performed, the specific results (e.g., states in which human feedback is optimal, ways in which ethical principles are conflicting) would probably differ. Since human feedback has also been shown to be effective in other contexts such as mental health [47, 361], it is promising to investigate this further. Lastly, our findings depend on the specific ways in which we operationalized the ethical principles in our context. This operationalization alone can have ethical consequences (e.g., see Obermeyer et al. [272] in the context of using historical health costs as a proxy for health needs).

6

Conclusion. Based on data from our longitudinal study in which 679 daily smokers and vapers interacted with a text-based virtual coach in up to five sessions and sometimes received human feedback between sessions, we demonstrate that people who have received human feedback spend more effort on the activities proposed by the virtual coach and are more likely to return to the next session. This suggests that it would be beneficial to have a human coach check in at least once with people who are preparing to quit smoking with a virtual coach. Moreover, concerning long-term effects, simulations with an RL model show that while providing more human feedback generally leads to a higher effort, there are also states where it is better not to provide feedback. When only very few resources for providing feedback are available, the highest effort spent on activities over time can be obtained by giving it to people with high self-efficacy and low or medium perceived importance of preparing for quitting smoking/vaping. Third, even the "standard" benefit-maximizing RL model is value-laden, prioritizing people who are already doing well and

want feedback. This is noteworthy in times when the increasing pressure on the healthcare system leads to calls to focus more on the cost-effectiveness of healthcare [105]. We further show how the RL model can be extended to incorporate other ethical principles such as favoring the worst-off or treating people equally, which we find to influence which smoker subgroups receive feedback. Since there is thus no value-free allocation of human support, moral decisions on who gets human support cannot be avoided. Given the complexity and dependencies between ethical principles, determining the consequences of different moral decisions is crucial. We hope that our work facilitates this and thus helps in making moral allocation decisions.

METHODS

We conducted an online crowdsourcing study in which participants interacted with the virtual coach Kai in up to five sessions between 1 February and 19 March 2024. The Human Research Ethics Committee of Delft University of Technology granted ethical approval for the study (Letter of Approval number: 3683). We preregistered the study in the Open Science Framework (OSF) [17] and no changes were made compared to the preregistration.

STUDY DESIGN

We performed a longitudinal study with a micro-randomized design [196], which entails assigning an intervention option at random to each participant at each pertinent decision point. The two intervention options were providing and not providing human feedback, which were chosen with probabilities of 20% and 80%, respectively. The four decision points were the days between each pair of five sessions with the virtual coach. To assess the effect of the intervention options, participants reported their effort spent on the activity assigned by the virtual coach as well as their return likelihood in case of an unpaid intervention in sessions 2–5 (Figure 6.5). Based on the collected data, we performed inferential statistics to determine the effect of human feedback on the effort and return likelihood (*RQ1*). Moreover, we trained an RL model that optimizes the effort people spend on their activities over time. Using this model, we ran human data-based simulations to assess the long-term effects of human feedback under varying settings for the cost of providing feedback (*RQ2*). Such human data-based simulations are a common way to assess RL models [384]. Lastly, we compared the optimal policies of RL models that not only optimize the effort spent on activities (i.e., *prognosis*) but also account for other ethical principles (Table 6.4) concerning the human feedback allocated to different smoker subgroups (*RQ3*). The weights assigned to the different ethical principles are thereby also based on smokers' preferred principles for allocating human feedback from our post-questionnaire (Table 6.4).

MATERIALS

Virtual coach. We implemented the text-based virtual coach Kai by closely following the implementation of the virtual coach Sam [9], which was developed for another smoking cessation study and overall perceived positively by smokers [21, 23]. There were two versions of Kai, one for smokers and one for vapers. Below we describe the smoking version².

²The only difference of the vaping version is that smoking-related terms in the dialogs were replaced with their vaping counterparts (e.g., "smoker" was replaced by "vaper").

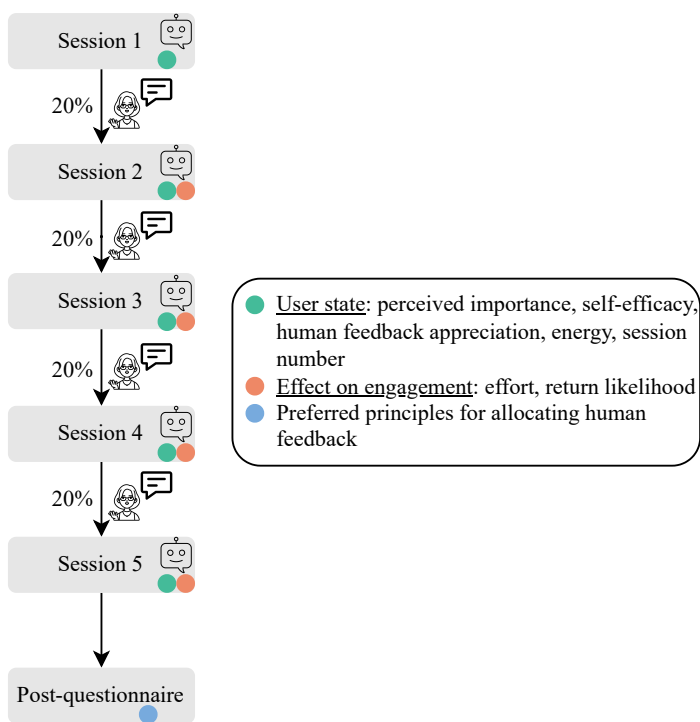


Figure 6.5: Study design, including the collected data. Between each pair of sessions, participants had a 20% chance of receiving a human feedback message.

After introducing itself as wanting to prepare people for quitting smoking and becoming more physically active, with the latter possibly aiding the former [152, 305], Kai explained that one of two human coaches could sometimes send a feedback message between sessions. These human coaches were described as having a background in Psychology, including knowledge of how to help people change their behavior. In each of up to five sessions, Kai collected information on an individual's current state by asking about their importance and self-efficacy for preparing for quitting, human feedback appreciation, and energy. Afterward, Kai proposed a new preparatory activity. In the next session, Kai asked about the effort people spent on their activity and their experience with it as well as their likelihood of returning to the session if it was unpaid. People were told that one of the human coaches could read their experience description to write a feedback message and that more specific descriptions would help write more specific feedback. Kai informed people when they were chosen to receive human feedback after the session. At the end of the session, participants received a reminder message with their activity on Prolific (Figure C.4). Like Sam, Kai gave compliments for spending a lot of effort on activities, expressed empathy otherwise, and maintained an encouraging attitude. The Rasa-based implementation of Kai [12] and a demo video [13] are available online. The conversation structure is shown in Figure C.5.

Preparatory activities. In each session, Kai proposed a new preparatory activity. This activity was randomly chosen from a set of 37 short activities (e.g., past successes for quitting smoking, role model for others by quitting smoking, visualizing becoming more physically active as a battle) created based on discussions with health experts, the activities of the smoking cessation applications by Michie et al. [243] and Albers et al. [21], the behavior change techniques by Michie et al. [241], and smoking cessation material by organizations such as the National Cancer Institute and the Dutch Trimbos Institute. Since becoming more physically active can make it easier to quit smoking [152, 305], 17 activities addressed becoming more physically active. One example of an activity is given in Table 6.3 and all activities can be found in Table C.6.

Table 6.3: Title and formulation of 1 of the 37 preparatory activities for quitting smoking used in the study.

Reasons for quitting smoking. Quitting smoking has many benefits. Think, for example, of improved physical fitness, healthier skin, and lower expenses. To help you quit smoking, it can be useful to write down why you want to quit. This can increase your aspiration to quit smoking, which may aid in quitting successfully. So, before the next session, I advise you to identify and write down reasons why you want to stop smoking. After writing them down, think about which reasons are most important to you and order them accordingly.

Human feedback. The human feedback messages were written by one of two human coaches, who were Master's students in Psychology. Following the model by op den Akker et al. [284], the human coaches were instructed to write messages that contained the following components: feedback, argument, and suggestion or reinforcement. They also received the general guidelines to refer to things in people's lives to build rapport, show understanding if people have low confidence, and reinforce people when they are motivated. When writing the feedback, the human coaches had access to anonymized data on people's baseline smoking and physical activity behavior (i.e., smoking/vaping frequency, weekly exercise amount, existence of previous quit attempts of at least 24 hours, and the number of such quit attempts in the last year), introduction texts from the first session with the virtual coach, previous preparatory activity (i.e., activity formulation, effort spent on the activity and experience with it, return likelihood), current state (i.e., self-efficacy, perceived importance of preparing for quitting, human feedback appreciation), and new activity formulation. All feedback messages ended with a disclaimer that they were not medical advice. A screenshot of how we sent human feedback messages to participants is provided in Figure 6.6. All 523 written messages are available online [29].

MEASURES

PRIMARY MEASURES

Effort. The virtual coach asked participants about the effort they put into their previously assigned activity on a scale from 0 ("Nothing") to 10 ("Extremely strong"), adapted from Hutchinson and Tenenbaum [170] as also done by Albers et al. [20].

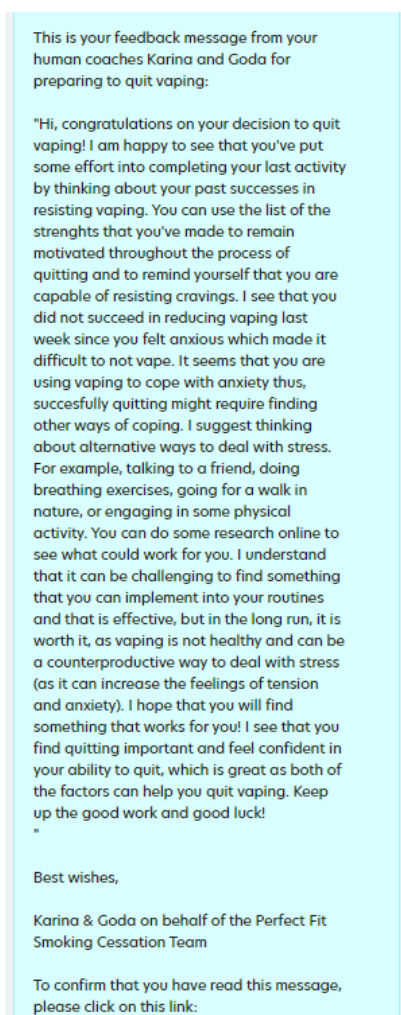


Figure 6.6: Screenshot of a human feedback message sent to a participant on the crowdsourcing platform Prolific Academic. The message ended with a disclaimer that it was not medical advice.

Return likelihood. Participants were asked “Currently you are taking part in a paid experiment. Imagine this was an unpaid [smoking/vaping] cessation program. How likely would you then have quit the program or returned to this session?”, rated on a scale from -5 (“definitely would have quit the program”) to 5 (“definitely would have returned to this session”). 0 was labeled as “neutral.”

SECONDARY MEASURES

State features. We measured five variables (i.e., features) that describe a person’s state in each session: 1) the perceived importance based on the question “How important is it to you to prepare for quitting [smoking/vaping] now?”, adapted from Rajani et al.

[314] and rated on a scale from 0 ("not at all important") to 10 ("desperately important"), 2) self-efficacy based on the question "How confident are you that you can prepare for quitting [smoking/vaping] now?", adapted from the Exercise Self-Efficacy Scale by McAuley [233] and rated on a scale from 0 ("not at all confident") to 10 ("highly confident"), 3) human feedback appreciation based on the question "How would you view receiving a feedback message from a human coach after this session?", rated on a scale from -10 ("very negatively") to 10 ("very positively"), with 0 labeled as "neutral," 4) energy based on the question "How much energy do you have?", rated on a scale from 0 ("none") to 10 ("extremely much"), and 5) the session number.

Preferred principles for allocating human feedback. Participants were asked to distribute 100 points across 11 allocation principles after the question "Based on which principles/rules should the virtual coach decide when a human coach should give feedback to people who are preparing to quit [smoking/vaping]? Assign 100 credits to the principles below, where more credits mean that you are more in favor of a principle." 9 principles were derived from the ones presented by Persad et al. [297], adapted to the smoking cessation context (Table C.7). We supplemented these principles with one further formulation of treating people equally (i.e., *least amount of human feedback so far*) and with the principle of respecting people's autonomy by prioritizing people who most appreciate receiving human feedback.

PARTICIPANTS

Participants were recruited from the crowdsourcing platform Prolific Academic. Eligible were people who smoked tobacco products or vaped daily, were fluent in English, and had not participated in the conversational sessions of our two previous studies with similar preparatory activities [14, 16]. Participants further had to give digital informed consent, confirm smoking/vaping daily, and indicate being contemplating or preparing to quit smoking/vaping [112] and not being part of another intervention to quit smoking/vaping to pass the prescreening questionnaire. The study was framed as preparation for quitting smoking/vaping for people recruited as daily smokers/vapers. Out of 852 people who started the first conversational session, 500 completed all five sessions, and 449 provided their preferences for allocating human feedback based on different principles in the post-questionnaire. To increase the chance that participants would read the human feedback messages, they were told they might be asked to confirm having read a received message to be invited to the next session. Participants who failed more than one attention check in the prescreening questionnaire were not invited to the first session. Moreover, participants had to respond to the invitations to the sessions and post-questionnaire within two days. The participant flow is shown in the Appendix. Participants who completed a study component were paid based on the minimum payment rules on Prolific, which require a payment rate of six pounds sterling per hour. Participants were informed that their payment was independent of how they reported on their preparatory activities to account for self-interest and loss aversion biases [117]. Participants who failed more than one attention check in the prescreening or post-questionnaire were not compensated for that respective questionnaire. Participants were from countries of the Organization for Economic Co-operation and Development (OECD), excluding Turkey, Lithuania, Colombia, and Costa

Rica, but including South Africa [311]. Of the 679 participants with at least one interaction sample, 330 (48.60%) identified as female, 335 (49.34%) as male, and 14 (2.06%) provided another gender identity. The age ranged from 19 to 71 ($M = 36.30$, $SD = 11.21$). Further participant characteristics (e.g., education level, smoking/vaping frequency) can be found in Table C.9.

PROCEDURE

Participants meeting the qualification criteria could access the prescreening questionnaire on Prolific, and those who passed the prescreening were invited to the first session with Kai about one day later. Invitations to a subsequent session were sent about three days after having completed the previous one. Between sessions, participants each time had a 20% chance of receiving a human feedback message. About three days after completing the last session, participants were invited to a post-questionnaire in which they were asked about their preferred principles for allocating human feedback, first by means of an open question and then by distributing points across given principles.

DATA PREPARATION AND ANALYSIS STRATEGIES

We collected all interaction samples of pairs of sessions in which people answered at least the effort, return likelihood, and the first state feature question (i.e., perceived importance) in the next session. Missing values in interaction samples ($N = 5$) for the remaining state features were imputed with the corresponding feature's sample population median. Below we explain our analysis strategies for each of the research questions. Our data and analysis code are publicly available [30].

RQ1: SHORT-TERM EFFECTS OF HUMAN FEEDBACK ON ENGAGEMENT

First, we wanted to assess whether human feedback positively affects engagement in the short term. For this, we performed Bayesian inferential analyses.

Direct effect of human feedback. To determine the direct effect of human feedback on the effort people spend on their activities and their return likelihood, we compared samples where people received human feedback to samples where they did not. For each of the two dependent variables (i.e., effort and return likelihood), we fit a model containing a general mean, a random intercept for each participant, and a binary fixed effect for human feedback received after the previous session. We fit both models with diffuse priors based on the ones used by McElreath [234] and assessed them by interpreting the posterior probability that the fixed effect for human feedback is greater than zero based on the guidelines by Chechile [90]. We further report 95% Highest Density Intervals (HDIs).

Delayed effect of human feedback. Besides the direct effect of human feedback on the effort and return likelihood, there might also be a delayed effect. For example, if human feedback increases a person's self-efficacy, then the person may spend a lot of effort on future activities even when not receiving additional human feedback. To determine whether having received human feedback leads to a higher effort and return likelihood, we fit two further statistical models. For both dependent variables (i.e., effort and return likelihood), we fit a model containing a general mean, a random intercept for each participant, and a

fixed effect for whether participants had received human feedback until then. We again fit both models with diffuse priors and used posterior probabilities and 95% HDIs to assess whether the effect of having received human feedback is positive.

Delayed effect of multiple human feedback messages. The delayed effect of human feedback might be stronger for people who have received multiple feedback messages. To determine whether having received more human feedback leads to a higher effort and return likelihood, we created two further statistical models by extending the previous two models with a fixed effect for the number of times participants had received human feedback until then. Again, we fit both models with diffuse priors and used posterior probabilities and 95% HDIs to assess whether the effect of multiple human feedback messages is positive.

RQ2: LONG-TERM EFFECTS OF OPTIMALLY ALLOCATED HUMAN FEEDBACK ON ENGAGEMENT

While our inferential analysis of delayed human feedback effects already looked a few steps into the future, it was based on randomly allocated human feedback. However, in some situations, giving human feedback might also be detrimental in the long run. So now, we want to use simulations to assess the long-term effects of optimally allocated human feedback based on a person's state. With optimally allocated human feedback we mean feedback that is only given in situations a) where it is ultimately more beneficial than not giving feedback, and b) where this benefit outweighs the economic cost of giving human feedback.

RL model. To study the long-term effects of optimally allocated human feedback, we designed and trained an RL model for deciding when to allocate human feedback. Starting with a base model that maximizes the effort people spend on their activities over time, we add the consideration of human feedback costs, and later for RQ3 of other ethical principles for allocating feedback. Figure 6.7 visualizes our final model, whose components we describe subsequently.

We can define our approach as a Markov Decision Process (MDP) $\langle S, A, R, T, \gamma \rangle$. The action space A consisted of two actions (i.e., giving human feedback no/yes), the reward function $R : S \times A \rightarrow [0, 1]$ was determined by the self-reported effort spent on activities, $T : S \times A \times S \rightarrow [0, 1]$ was the transition function, and the discount factor γ was set to 0.85 to favor rewards obtained earlier over rewards obtained later as also done in previous work (e.g., [20, 113]). The finite state space S described the state a person was in and was captured by their perceived importance of and self-efficacy for preparing for quitting smoking/vaping as well as their appreciation of receiving human feedback. The goal of an MDP is to learn an optimal policy $\pi^* : S \rightarrow \Pi(A)$ that describes an action to take in each state that maximizes the expected cumulative discounted reward $E \left[\sum_{t=0}^{\infty} \gamma^t r_t \right]$. The optimal Q-value function $Q^* : S \times A \rightarrow \mathbb{R}$ describes the expected cumulative discounted reward for executing action a in state s and π^* in all subsequent states.

State space We considered six features to describe the state space: 1) the perceived importance, 2) self-efficacy, 3) the difficulty of the assigned activity based on the activity difficulty ratings by Albers et al. [26], 4) energy, 5) human feedback appreciation, and 6)

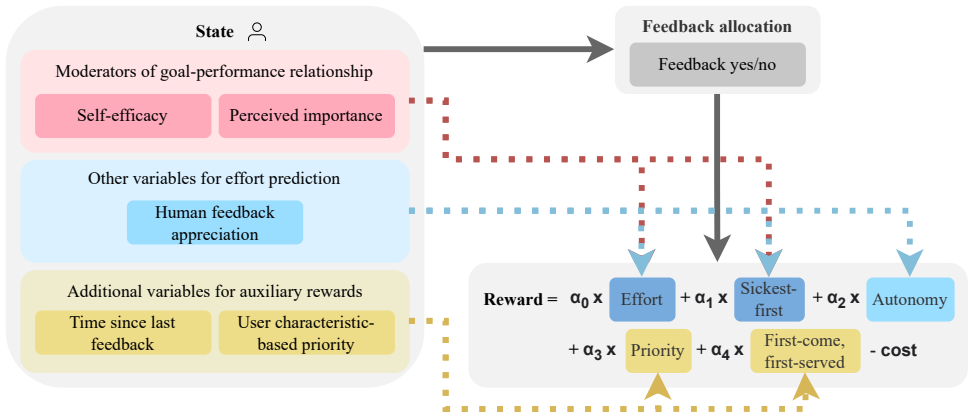


Figure 6.7: Visualization of the final RL model. The arrows indicate which state features are used to predict the different reward functions. The five reward functions can be combined by setting different weights α .

6

the session number. The first three features were considered since goal-setting theory posits that goal commitment, facilitated by importance and self-efficacy, and task difficulty are moderators of the effects that goals have on performance [224]. More precisely, low commitment and high task difficulty might make it harder for people to reach their goals, which may make human feedback more beneficial. We further included energy since it was shown to be an important predictor of the effort people spend on preparatory activities for quitting smoking in a previous study [16]. Moreover, since the novelty of the intervention may influence people’s motivation to do the activities [21], we also captured the session number.

To reduce the size of the state space and thus create a more robust model, we selected three abstracted base state features based on our collected data. Specifically, using the G-algorithm [88] and its adaptation by Albers et al. [20] as inspiration, we iteratively selected the feature for which the Q-values for the abstracted feature values were most different. We thereby specified the first selected feature to have three and the second and third features two abstracted values. Abstract features were computed based on percentiles. For example, to create an abstract feature with two values, we set all values less than or equal to the median to 0 and those greater than the median to 1. Besides reducing the required data, selecting a subset of the state features also has the advantage that the virtual coach would in the future need to ask people fewer questions per session, which is in line with keeping smoker demands to a minimum [243]. The three selected features were 1) perceived importance with three values, 2) self-efficacy with two values, and 3) human feedback appreciation with two values. The base state space thus had size $3 \times 2 \times 2 = 12$. We refer to the resulting base states with three-digit strings such as 201 (here perceived importance is high, self-efficacy is low, and human feedback appreciation is high). Figure C.10 and Figure C.11 show the mean effort and number of samples per combination of values for the three selected features.

Action space The action space was defined by two actions: giving ($a = 1$) and not giving ($a = 0$) human feedback.

Reward Just as in the algorithm by Albers et al. [20], the base reward signal was based on asking people how much effort they spent on their previous activity on a scale from 0 to 10. Based on the sample population mean effort \bar{e} , the reward $r \in [0, 1]$ for an effort response e was computed as follows:

$$r = \begin{cases} \frac{e}{2\bar{e}} & \text{if } e < \bar{e} \\ 1 - \frac{10-e}{2(10-\bar{e})} & \text{if } e > \bar{e} \\ 0.5 & \text{otherwise.} \end{cases}$$

The idea behind this reward signal was that an effort response equal to the mean effort was awarded a reward of 0.5, and that rewards for efforts greater and lower than the mean were each equally spaced.

Reward and transition functions The reward and transition functions were estimated from our data.

Human feedback cost Due to budget constraints, the base reward may cause human feedback to be allocated to more people than can be economically afforded. To be able to reduce the amount of allocated human feedback, we introduce a cost factor c to be included in the reward computation that depends on the action a :

$$r_c = \begin{cases} r & \text{if } a = 0 \\ r - c & \text{if } a = 1 \end{cases}$$

Training the model We computed 0.001-optimal policies and corresponding Q^* with Gauss-Seidel value iteration from the Python MDP Toolbox. We use $\pi^{*,c}$ to refer to an optimal policy for a certain cost c .

Long-term effects of unlimited human feedback. First, we assume we have no economic budget constraints and can allocate as much human feedback as we wish (i.e., $c = 0$). To assess the effectiveness of such unlimited human feedback over time, we ran simulations based on our collected data to compare four different policies concerning the mean reward per activity assignment over time: 1) the optimal policy $\pi^{*,0}$, 2) the policy of always assigning human feedback, 3) a theoretical average policy where each of the two actions is taken $\frac{1}{2}$ times for each person at each time step, and 4) the policy of never assigning human feedback. To obtain a realistic population, the simulated people were initially distributed across the state features following the distribution we observed in the first session of our study (Figure C.12).

Long-term effects of limited human feedback. In practice, budget constraints might make it impossible to always allocate human feedback according to $\pi^{*,0}$. To reduce the amount of allocated human feedback, we added different human feedback costs to the base reward, and assessed the resulting mean reward and amount of allocated human feedback over time. The considered costs were chosen such that the resulting optimal policies $\pi^{*,c}$ all differ in the number of states that are allocated feedback. We again used as starting population the distribution of people across the 12 states we observed in our study’s first session.

RQ3: EFFECT OF DIFFERENT ETHICAL ALLOCATION PRINCIPLES ON HUMAN FEEDBACK RECEIVED BY SMOKER SUBGROUPS

Given that we can only provide limited human feedback, we cannot allocate human feedback to everybody. The RL models we have trained for RQ2 all allocate human feedback to those who will see the largest increase in effort spent on preparatory activities over time because of the feedback. This can be seen as maximizing total benefits according to the allocation principle that Persad et al. [297] call *prognosis*. However, we can also use other ethical principles in our RL model. Here, we now want to assess the effects of incorporating different ethical allocation principles on the subgroups of smokers that receive feedback.

Potential live application. To get a realistic assessment of the effect of incorporating different ethical allocation principles, let us first define a potential live smoking cessation application. Suppose we have an application in which people have up to nine sessions with a virtual coach, after each of which they can get feedback from a human coach. As people sometimes drop out of eHealth applications before completing them [48, 123], we assume a 15% chance people drop out of our application after each session³. The spots of people who have either completed all nine sessions or have dropped out are given to new people. These new people are distributed across the 12 base states as in the first session of our study. Taking about six minutes to write a feedback message, the human coach can give feedback to 58 people every day. Assuming 166 spots in the application, this amounts to 35%. Therefore, the human feedback costs in our analyses were set such that on average about 35% of people receive feedback every day.

Extending the RL model with other ethical allocation principles. To also reward allocating human feedback according to ethical principles other than *prognosis*, we created the four auxiliary (i.e., additional) rewards shown in Table 6.4. We use *first-come, first-served* to illustrate the effect of treating people equally. Note that the ethical principles *youngest first*, *instrumental value*, and *reciprocity* can all be represented by setting an individual characteristic-based priority level. To compute these auxiliary rewards, we extended the state space by two features, each with three values: 1) a random individual characteristic-based priority level that remains fixed for each person, and 2) time since the last human feedback. Both of these state features only influence the auxiliary reward and not the base reward (i.e., *prognosis*). Each auxiliary reward $r_{aux} \in [0, 1]$ is then computed as $r_{aux} = \frac{aux - aux_{min}}{aux_{max} - aux_{min}}$, where *aux* is a person’s value for the measure underlying the auxiliary

³This is the average percentage of negative return likelihood ratings per session of our longitudinal research study.

reward (e.g., the time since the last human feedback) and aux_{min} and aux_{max} are the lowest and highest possible values for the measure.

Table 6.4: Allocation principles by Persad et al. [297] with the addition of autonomy and corresponding rewards. We also show the mean weight assigned to the principles by participants of our post-questionnaire. We use *first-come, first-served* to illustrate the effect of treating people equally.

Allocation principles	Reward	Weight
Prognosis	<i>Base reward</i> : Prioritize people who will see the largest increase in effort because of the feedback	30.82%
Treating people equally (lottery, first-come, first-served, least amount of human feedback so far ¹)	<i>First-come, first-served</i> : Prioritize people who have spent the longest time since the last feedback	22.18%
Sickest first	<i>Sickest first</i> : Prioritize people who would spend the lowest effort without feedback	25.34%
Youngest first, instrumental value, reciprocity	<i>Priority</i> : Prioritize people with a higher individual characteristic-based priority	13.04%
Autonomy	<i>Autonomy</i> : Prioritize people who appreciate feedback the most	8.62%

¹ We supplemented the principles for treating people equally by Persad et al. [297] with the principle of prioritizing people with the least amount of human feedback so far. The reason is that while Persad et al. [297] focus on medical resources that can be allocated to each person only once, human feedback can in our context also be allocated more than once.

Policy comparison. Using the rewards from Table 6.4 and the weights given to them by smokers, we compared six policies based on which states they allocate feedback: 1) the optimal policy based on the base reward, 2–5) the four optimal policies for using the base reward together with either *first-come, first-served*, *sickest first*, *autonomy*, or *priority* with the two rewards weighted based on the smoker-preferred weights, and 6) the optimal policy based on all five rewards weighted according to weights derived from smokers' preferred principles for allocating human feedback (Table 6.4). Due to the relatively large drop in reward between human feedback costs of 0.07 and 0.09 observed for our analysis of the long-term effects of limited feedback (Figure 6.3a), we set the human feedback cost to 0.07 for the base reward-based optimal policy, which means that after each session around 35% of people get feedback (Figure C.13b). Since incorporating auxiliary rewards can change the amount of allocated feedback, we tuned the costs for the other policies such that these policies also allocate feedback to around 35% of people, thus allowing for a fair comparison between policies.

7

CONCLUSION

CONCLUSIONS

The research presented in this thesis studies how Reinforcement Learning (RL) can be used to make the support provided in virtual coach-based smoking cessation interventions more effective. After establishing user needs for such an intervention and providing recommendations for addressing them (Chapter 2), this thesis investigates using RL to adapt 1) *how* people are persuaded (Chapter 3 and 4), 2) *what* they are asked to do (Chapter 5), and 3) *who* they are supported by (Chapter 6). In doing so, the goal is to answer the following overarching research question:

How can reinforcement learning be used to make the support in virtual coach-based smoking cessation interventions more effective?

Below we provide the conclusions we can draw for the resulting four research questions and one hypothesis.

RQ1: What are users' needs for a virtual coach-based smoking cessation intervention?

By understanding users' needs, we can design interventions that satisfy them, increasing the chance of effective user engagement. For our case of a digital smoking cessation intervention, we identified 14 main themes describing user needs. These themes fall into four sets, pertaining to 1) users' behaviors concerning an eHealth application, 2) the users, 3) other parties that may be involved in a behavior, and 4) the environment in which behaviors are performed. Examples of considered behaviors include following a physical activity program while quitting smoking, receiving motivational messages, or consulting a general practitioner. The overall most common theme was whether behaviors were perceived as useful because they provided motivation, encouragement, help, advice, or learning opportunities. Smokers' perceptions of usefulness thereby did not necessarily match those of experts. For example, physical activity can facilitate quitting smoking [152, 305], yet, several smokers did not share this view. The timing and intensity of behaviors also was a concern. In the set relating to users themselves, themes were their motivation for and importance of change, need for autonomy, and further personal characteristics such as their personality. The next set of themes concerned other involved parties such as a general practitioner or virtual coach. Major themes here were companionableness, accountability, and whether the other parties were human or AI. Environmental factors formed the last set of themes. Key themes here included the availability of sufficient time, the helpfulness of support from one's social environment, the usefulness of prompts and triggers, and the diversity of other environmental factors. Overall, these needs are often interconnected; for example, the perceived usefulness of behaviors is linked to the availability of sufficient time and a user's motivation to change. To address these needs, we formulated several recommendations. Rather than just focusing on one need at a time, connected needs should be addressed together.

H1: Subsequently incorporating 1) states, 2) the consideration of future states, and 3) the weighting of samples based on the similarity of people into an algorithm that selects the best persuasive strategy is more effective than not incorporating the respective element.

One of the user needs we identified for our first research question (RQ1) is also related to receiving motivational support. Choosing a motivating way to propose preparatory activities for quitting smoking is one way to address this. Here we specifically focused on different persuasive strategies, such as informing users that peers recommend doing an activity. To examine the benefit of a personalized RL algorithm in choosing persuasive strategies, we analyzed whether a more complex algorithm (i.e., one that contains more algorithm elements) has a more positive effect on the effort people spend on their activities and their perceived motivational impact of the conversations with the virtual coach. Our findings provide some support that perceived motivational impact and effort are positively affected, with the latter becoming apparent only after some time. Looking at the algorithm complexity levels separately, it seems that the level that considers users' current and future states but not their user characteristic-based similarity is the most effective. Support for this is even stronger when we only look at those people who were most involved in their activities. People with low activity involvement, on the other hand, not only provided lower effort and motivation ratings than highly involved people but also did not show higher susceptibility to higher algorithm complexity levels.

Besides the effect of choosing persuasive strategies based on algorithms with different complexities, we also examined the effect of choosing persuasive strategies based on data collected for different activity types. This was possible because half of the preparatory activities addressed quitting smoking directly, while the other half targeted becoming physically active as a way to facilitate quitting smoking [152, 305]. Comparing the best persuasive strategies learned for the two activity types, we find that their similarity is relatively low. Thus, there seems to be limited generalizability of the best persuasive strategies learned for activities for quitting smoking and activities for becoming physically active.

RQ2: How does predicting behavior based on user characteristics compare to doing so based on states?

Compared to predicting behavior based on states, doing so based on user characteristics has the advantage that data on them needs to be collected only once at the start of the intervention rather than before each persuasive attempt. Examples of such user characteristics are users' stage of change for becoming physically active, personality, or smoking status. Yet, we find that we can better predict the effort people spend on preparatory activities after being persuaded with different persuasive strategies based on states than user characteristics measurable before an intervention. If, however, we also consider user characteristics that pertain to how users experienced the intervention, namely, users' overall involvement in their activities, considering user states no longer clearly outperforms considering user characteristics. Yet, such information on involvement is, just like information on states, also not available before an intervention. Overall, our findings show that behavior can best be predicted based on factors that are conceptually closest to the behavior. States are

conceptually closest, followed by overarching measures of users' experience with the whole intervention, and then intervention-independent user characteristics such as personality.

RQ3: How can we create an RL model for building human competencies that combines the views of experts and users?

The most prevalent need we observed for *RQ1* is that users perceive behaviors associated with a smoking cessation intervention as useful. Since the usefulness views of experts and users may differ, creating an RL model that combines both views may be beneficial. To this end, we designed and followed a five-step pipeline. First, we established competency-building activities by determining which activities experts typically propose to prepare for quitting smoking. Next, we determined the competencies that smokers and experts perceive as being built by these activities. Using the repertory grid technique [134], we found nine smoker-identified competencies (e.g., self-efficacy, awareness of positive outcomes, knowledge of how to maintain/achieve mental well-being) and six expert-identified competencies for experts (e.g., clear future identity, motivation, insights into personal strengths). Afterward, we created nine belief-changing activities to convince smokers of the usefulness of the nine smoker-identified competencies. And finally, we designed and trained the RL model. Using simulations to evaluate the RL model, we find that proposing activities based on the model allows smokers to build 91% of expert-identified competencies within five interactions. While all model components contribute to this, users' current state based on their usefulness beliefs, energy, and degrees of having built the expert competencies is the largest contributing factor. The contributions of the transitions to the next states, on the other hand, are small. This confirms the value of considering, if not the *future*, at least the *current* views of smokers and experts.

One possible reason why the contributions of the transitions to next states are small is related to the effects of belief-changing activities. Specifically, while we did find credible evidence that these activities positively affect usefulness beliefs, these effects tend to be small. As such, belief-changing activities were also not included in the learned optimal strategy for building competencies.

RQ4: How effective is adding human support to a virtual coach-based behavior change intervention in the long term?

In *RQ1*, we saw several user needs relating to other parties involved in a behavior, such as virtual coaches or general practitioners. Users were concerned about the accountability toward these parties and their companionableness. As also expressed by users, both accountability and companionableness can be influenced by whether another party is human or AI: including support from a human coach can make people feel more accountable [210, 250] and satisfied [327]. The simulations we conducted for *RQ4* show that providing human feedback more often generally, but not always, leads to a higher effort spent on preparatory activities over time. Specifically, when users' perceived importance of preparing to quit is high and either their self-efficacy or human feedback appreciation is low, it is better in the long run not to provide human feedback. While it is thus optimal to provide human feedback in all but two states, about half of the human feedback messages

can be removed without a large drop in effort, for example when human resources are scarce. In that case, the largest long-term increase in effort can be obtained by giving feedback to people who perceive preparing to quit as not that important and at the same time have high self-efficacy for preparing to quit.

While allocating limited human feedback to those who would benefit most from it seems logical, it is, in fact, only one of several ethical principles for allocating scarce healthcare resources. Other principles include allocating feedback to those who want it most (*autonomy*), those who have waited the longest (*first-come, first-served*), or those who have the lowest chance of successfully preparing to quit without human feedback (*sickest first*). By choosing from these allocation principles, our simulation results show we take a moral stance on who may increase their chance of successful smoking cessation and positive health outcomes. Analyzing the "standard" benefit-maximizing model, we find that it prioritizes people who want feedback and are already doing well. We show that to favor other smoker subgroups, the model can be extended to incorporate other allocation principles. However, interdependencies between principles exist. For example, since we find that people who are not doing well tend not to want feedback and benefit less from it, giving more feedback to them means doing worse in respecting people's autonomy and maximizing total benefit.

LIMITATIONS

There are several limitations to the conclusions this thesis draws. These pertain to the study participants, the setting in which our studies took place, and modeling assumptions.

SELECTION OF PARTICIPANTS

One set of limitations pertaining to all conclusions is that we conducted our studies with participants recruited from the crowdsourcing platform Prolific. Prolific provides access to participants from almost all countries of the Organization for Economic Co-operation and Development (OECD) as well as South Africa [311] and hence a geographically and culturally diverse population. However, participants might be less diverse in other regards, thus limiting generalization. For example, since younger and more highly educated people were represented more in our sample, our findings might generalize less well to older and less highly educated populations. Furthermore, because our participants were enrolled on Prolific, they had at least some experience with digital services. Previous work on recruiting older people on Prolific [370], for instance, saw that most of them used technology frequently, with almost all of them using a computer at least several times a day and most owning a smartphone. Besides high age and a low education level, low eHealth literacy is another barrier to eHealth use [8, 318]. Thus, especially our conclusions regarding user needs for *RQ1* should be interpreted with this limitation in mind: the needs of users with more barriers to using an application like ours are potentially not sufficiently represented. Yet, while our findings might not hold for less technology-savvy people, such people might also be less inclined [261] and suited [8, 318] to use a stand-alone virtual coach-based smoking cessation intervention.

Another limitation concerning the use of Prolific is that participants were paid for completing sessions with the virtual coaches. Even though participants were informed that

their payment was independent of their reported effort and experience with their activities to account for self-interest and loss aversion biases [117], participants might have felt some obligation to complete the activities. Therefore, the effects of persuasive strategies, effort spent on competency-building and belief-changing activities, and the effects of human feedback might be different in a setting in which no payment is provided. These differences might be both *absolute* in the sense that unpaid people might generally spend less effort on activities and *relative* in the sense that what is optimal in which state may change. Neither of these differences would, however, likely change our conclusions regarding the effectiveness of RL in adapting smoking cessation support. Furthermore, it is noteworthy that using a crowdsourcing platform for participant recruitment does not necessarily imply poor data quality. In fact, Douglas et al. [116] saw that the percentage of participants who provided meaningful responses (e.g., passed attention checks, followed instructions, remembered previously presented information) was higher for a sample recruited from Prolific than for an undergraduate sample recruited through the university.

Thirdly, all our participants were contemplating or preparing for quitting smoking (or vaping in the study for RQ4). While people in the preparation stage intend to take action within the next month, people in the contemplation stage only intend to change after that. For the latter, action-focused preparatory activities such as thinking about smoking-relevant routines or learning relaxation exercises might not yet be very relevant [308]. This might have led to lower involvement with our intervention and thus different processing of persuasive messages, preparatory activities, and human feedback [229] than we would see for people in the preparation stage. Our findings from our exploratory subgroup analysis for H1 support this. While our conclusions might thus not fully generalize to a pure preparation stage sample, they are based on a much larger sample of smokers given that 40% of smokers are in the contemplation stage [308]. In light of only 20% of smokers being in the preparation stage, an intervention also suited for people in the contemplation stage can have a much larger impact. Moreover, people who quit smoking can also return to the contemplation stage after events such as a relapse [308]. Notably, we did exclude people in the precontemplation stage who comprise another 40% of smokers [308]. People in this stage do not intend to change in the next six months, a period we deemed too long for action-oriented preparatory activities like the ones proposed in our studies to be relevant. Our findings hence do not necessarily generalize to all smokers.

SETTING

Another set of limitations concerns the setting in which we conducted our studies. As described in the introduction to this thesis, we chose to focus on minimal interventions for preparing for quitting smoking to assess the effects of our algorithms in a meaningful and ethical way. The downside of this is that it is not clear to which degree our findings generalize to a) a full smoking cessation application and b) a different domain. Regarding the former, it could be that our algorithm effects interact with those of other components of a full smoking cessation application such as self-monitoring and goal-setting. Moreover, the effects of our algorithms could also interact with each other. For example, using the persuasive strategy *authority* might have a different effect when a person also receives feedback from a human coach. Before studying such interaction effects, understanding the isolated effects of these algorithms can help to relatively quickly identify algorithms

that are promising to be studied further. This was the aim of this thesis. Regarding generalization to other application domains, it is likely that the relative effectiveness of different persuasive strategies, types of activities, and human support chosen in different states differs between domains. This is supported by our observation that the agreement between optimal persuasive strategies for activities for smoking cessation versus activities for increasing physical activity is rather low. Similarly, Alslaity and Tran [34] showed that the application domain influences which of the persuasive strategies by Cialdini [93] is most effective. However, given the ample literature suggesting that current states influence behavior and behavior, in turn, influences states in various health behavior change domains, we expect our more general conclusions regarding the effectiveness of using RL to adapt the three *how-*, *what-* and *who-*dimensions of support to also hold for other types of health behavior change. Nevertheless, more work is needed to confirm this.

Another limitation is that our conclusions for *H1*, *RQ2*, *RQ3*, and *RQ4* are based on self-reported measures of engagement with activities rather than more objective ones. While more objective measures might be desirable, self-reported measures of engagement are a convenient and efficient way to understand users' perspectives [273]. Focusing on users' perspectives of performance rather than actual performance, self-reported measures are thus not to be seen as a less objective version of behavioral or psychological measures, but rather as measuring a different response process [102]. Such user perspectives on engagement might be much easier to relate to actual engagement than more objective measures. A user spending a lot of time on a task, for instance, might be highly engaged but might also be confused [273]. Moreover, even when we know how to interpret more objective measures, it is not always clear how they can be obtained. Measuring if people watched a video might be possible; measuring if people printed a picture and placed it somewhere they can see it every day seems less feasible. Similarly, we used a self-reported measure of motivation in our analysis for *H1*. Even though more objective measures such as ones relying on psychophysiological reactions are possible in a controlled lab setting (e.g., [349]), collecting them in a real-world setting raises feasibility and privacy concerns, thus reducing the real-world relevance of such measures.

Not only did we use self-reported measures, but we also focused on proximal ones like effort and motivation rather than more distal ones capturing the actual behavior we want to accomplish (i.e., smoking abstinence). Our goal is that people do their preparatory activities more thoroughly, so that they are better prepared to quit smoking, so that they can better quit and stay quit. We measured the effectiveness of our algorithms only in terms of the first element of this chain, which means that our conclusions do not necessarily generalize to smoking abstinence. However, not only does the focus on engagement with activities facilitate faster and less noisy evaluations of algorithm effects [45], but it is also still meaningfully linked to smoking abstinence. In fact, preparation for behavior change is commonly included in smoking cessation interventions (e.g., [243, 266, 278]) to increase the chance of successful change.

Finally, our analyses are based on a limited set of possible state features and user characteristics. While we formed our candidate set based on literature suggesting the relevance of these state features and user characteristics, we might have missed more relevant ones. This especially concerns our conclusions drawn for *RQ2*, where we directly compared the effects of states versus user characteristics in predicting behavior after

persuasive attempts and found states to be more effective. Notably, however, we did already consider more than 30 user characteristics capturing demographics, smoking and physical activity, personality, and need for cognition.

MODELING ASSUMPTIONS

Any predictive model makes assumptions since it is impossible to account for every possible variable affecting the outcome [325]. In building our RL models, we assumed that the Markov property holds. This means we assumed that rewards and transitions to next states depend only on the current state. There are three main reasons why this assumption may not have fully held in our studies. First, there might be a novelty effect [101, 333]. According to this novelty effect, people's initially high curiosity and expectation fade over time as they become aware of a new system's limitations. Only in the RL model built for *RQ4* did we include the session number as a potential state feature to capture such changes over time. The fact that the session number was not selected as one of the three state features suggests that if there was a novelty effect, it was not large. However, larger novelty effects could have been at play in the earlier two studies. This might have especially been the case for the first study which was conducted before the introduction of large language models to the general public where people were less familiar with technology such as chatbots and hence possibly more curious. Studies on older technologies such as computers show that people's attitudes toward them can become more negative over time. For example, students' attitudes toward computers became more pessimistic between 1988 and 2001 [193]. However, sometimes becoming more pessimistic about the negative impacts of a technology can go hand-in-hand with becoming more optimistic about positive impacts [193]. Second, there might have been delayed effects of actions beyond those captured by the reward and next state. For example, in the study conducted for *RQ3*, it might have taken people more time to thoroughly reflect on belief-changing activities and change their usefulness beliefs accordingly. And also the effects of competency-building activities might take longer to fully manifest. For instance, some people mentioned in their activity experience descriptions that they reflected on their future selves several times or thought about a previous activity again after a new one had been assigned. Third, there might be an effect of (not) repeating actions. For instance, repeating a persuasive strategy could make it more [76] or ultimately less [75] effective. Similarly, the first human feedback message could be more effective than later ones beyond what is captured by transitions to different states. Us not accounting for these possible violations of the Markov property potentially limits the conclusions drawn from our simulation-based analyses in two main ways. First, the effects of belief-changing activities studied for *RQ3* might actually be larger if delayed changes beyond the next session are considered. Second, the long-term effects of choosing activities for *RQ3* and adding human feedback for *RQ4* might be different than projected by our simulations.

Besides not accounting for possible violations of the Markov property during our longitudinal studies, there are additional concerns for a full-length behavior change intervention. For example, the cognitive effort and awareness required to do activities might ultimately decrease if people form habits for doing them [138]. Our simulation-based analyses of long-term effects for *RQ3* and *RQ4* did not account for such effects possibly arising after several weeks. Thus, while our findings might hold for preparatory activities that are

intended to be done at most a few times within at most a few weeks, they might be less applicable to activities done repeatedly for a longer time such as taking a certain number of steps every day.

More generally, while human data-based simulations are a common way to assess RL models [384], our conclusions for *RQ2*, *RQ3*, and *RQ4* depend on the datasets we used for the simulations. Even though we collected relatively large, and expensive, datasets comprising between 1,710 and 2,366 interaction samples, we obtained few samples for some state-action combinations due to some states occurring less frequently than others. For those state-action combinations, our simulations are at risk of overfitting to the limited data, thus potentially limiting generalizability.

FUTURE WORK

The first obvious next step is to perform randomized controlled trials similar to the one we performed for *H1* to see how well the algorithm effects we observed in human data-based simulations generalize. To confirm the link between doing preparatory activities more thoroughly and achieving and maintaining smoking abstinence, information on smoking abstinence right after the intervention and at 6- and 12-month follow-ups [169] should also be collected. The best-performing algorithms should then also be implemented and tested in full smoking cessation interventions. A challenge in this context are the possible interaction effects not only between the algorithms and other static intervention components such as self-monitoring and goal-setting but also between multiple algorithms. Simply rolling out the individually trained models would fail to account for such interaction effects. Also learning from the data collected during the intervention could address this. However, this would require trying different actions in each state rather than performing a fixed optimal action per state, thus reducing predictability. Moreover, due to the possibly large size of a single combined model, it might be worthwhile to carefully think about which model components are expected to influence each other. Incorporating such structure can greatly reduce the amount of required data [251].

Incorporating multiple models that use different variables to make decisions raises the question of how the necessary data can be collected without burdening users. If data is collected via self-reported measures, the added benefit of collecting a variable should be greater than the cost in terms of making the application less attractive [243]. One alternative to self-reports is data collected with sensors. However, not only may improvements in sensing technology be needed to obtain higher data quality, lower the burden on users, and increase adherence [395], but switching from self-reports to sensors also requires rethinking the variables used. While it may be possible to estimate variables such as self-efficacy with sensors to some degree in some contexts (e.g., in a rehabilitation setting [155, 280]), the estimation would introduce additional noise and thus increase the amount of required data to make reliable predictions. Alternatively, one could dynamically choose which self-report measures to collect so that users remain engaged and prediction uncertainty is reduced as much as possible (e.g., [120]).

In trading off variance and bias, all models proposed and tested in this thesis combined or "pooled" data from multiple individuals. However, individuals may differ in their response to persuasive strategies, activity proposals, and human feedback beyond what is captured by states and user characteristics. Kaptein and Eckles [182], for example, found a large

heterogeneity in responses to persuasive strategies by Cialdini [93] relative to the effect sizes of these persuasive strategies, with the heterogeneity remaining when incorporating user characteristics. While the heterogeneity observed by Kaptein and Eckles [182] could also be due to people's states, our results for *RQ2* show that large prediction errors remain even when accounting for people's states. Given that it can be worse for an individual to receive the on average best persuasive strategy than no persuasive attempt [182], it might be beneficial to also adapt to individual users. Since data on individual users is often not available before an intervention, such adaptation can also occur online. Examples of adapting support to single users include choosing persuasive strategies by Cialdini [93] [183], selecting reinforcement strategies for a robot teacher [331], and determining when to send tooth brushing reminders [366]. Since pooling data across users can initially lead to a higher reward per user even if effects are heterogeneous [366], more weight could be given to data from single individuals only after some time (e.g., [183]). Obtaining a high reward at the start is welcome given that many users quickly drop out from eHealth applications [123].

Compared to using fixed pre-trained models as we did in our studies, adapting models online based on the new data collected during an intervention might make the models more effective. However, this is not the only concern. Instead, we often also want to be able to use the data, which has been expensive to collect, to perform a range of secondary analyses. For instance, we might want to estimate the overall effect of human feedback across individuals. And we want to do this in a replicable way [401]. Replicability here means that applying the same adaptive algorithm to different individuals drawn from the same population should lead to similar analysis results after the intervention. This is not guaranteed in digital health interventions where responses to types of support can change over time (i.e., are non-stationary) [401]. Algorithms that optimally balance effectiveness during the intervention and power of post-intervention analyses exist for some types of RL algorithms (e.g., [122]), but a better understanding of how this could work in cases where future states are also explicitly considered is welcome. In addition, algorithms may explicitly try to account for (some forms of) non-stationarity. Mintz et al. [249], for example, model the non-stationarity that can arise due to repeating actions within a short time frame.

Lastly, while RL models with a consideration of current and future states have been applied to other problems such as timing running notifications [381], suggesting step goals [140], and choosing messages for diabetes prevention [195], it would be interesting to see how well insights on the effectiveness of different RL model components generalize across domains. This could also further support the behavior change theories from which model components are often derived.

CONTRIBUTIONS

This work makes several scientific and societal contributions, outlined below.

SCIENTIFIC

In answering the overarching research question on how RL can be used to create more effective support in virtual coach-based smoking cessation interventions, this thesis primarily

contributes insights into user needs for such an intervention and a better understanding of the effectiveness of (components of) RL models in adapting the *how*-, *what*-, and *who*-dimensions of support.

User needs and how to address them. First, this thesis contributes 14 themes describing user needs that are informed by 671 people interacting with an intervention prototype for about two weeks. We showed that these needs are often interconnected and pertain to a behavior itself, a user, other parties that may be involved in a behavior, and the environment in which a behavior is performed. The needs and our recommendations for addressing them provide a useful starting point for future developments of virtual smoking cessation coaches. It is especially important that perceived usefulness is taken into account, as this was the most prevalent theme. Furthermore, our publicly available data on 5,074 free-text responses from 671 people together with several user characteristics facilitates future analyses of the needs of specific smoker subgroups (e.g., those with a lower socioeconomic status) regarding specific intervention components (e.g., identity-related activities). Our data on how users experienced their activities from the other two longitudinal studies can further be used to enrich this dataset. Moreover, the more than 800 introductions to human coaches from our third longitudinal study can help to determine what kind of human support users want and to improve human support accordingly.

Empirically evaluated RL models. Additionally, this thesis provides RL models for adapting *how* smokers are persuaded, *what* they are asked to do, and *who* they are supported by. The effectiveness of (components of) these RL models has been empirically evaluated. For example, we showed the benefit of considering current and future user states based on the COM-B model in choosing persuasive strategies, especially when people are highly involved in the intervention. Considering user characteristics, either alone or in addition to current and future states, was not found to be more effective. Given the large number of possible algorithm components, these insights can help narrow down the components considered for new algorithms for adapting the support provided in eHealth applications, in particular ones for smoking cessation. In addition, our findings provide support for several theories from which the components were derived. For example, findings from the *how*-dimension lend support to the COM-B model. Our publicly available datasets with 2,366, 1,710, and 2,326 interaction samples and corresponding user characteristics for the three support dimensions can further be used to perform more targeted analyses. Especially useful are the datasets collected for the *what*- and *who*-dimensions, where actions (i.e., activities for the *what*-dimension and the allocation of human feedback for the *who*-dimension) were selected randomly in all sessions.

Research on smoking cessation interventions. Besides the insights directly gained by answering our research questions, the publicly available material created for our studies is another contribution that can facilitate future research on smoking cessation interventions. This includes the persuasive messages developed for the first study, overall 44 competency-building and 9 belief-changing activities, the code of the virtual coaches Sam [9], Mel [10], and Kai [12], and the 523 human feedback messages written for our third study [29]. In addition, our dataset on 469 action plans for doing preparatory activities formulated by

smokers [25] can be used to help other smokers create such plans, which can be difficult without guidance [109, 304]. The same dataset also contains smokers' reflections on the usefulness of preparatory activities, which may help to convince new smokers of their usefulness. Lastly, the nine smoker-identified competencies for quitting smoking can be used to tailor smoking cessation material to smokers' usefulness beliefs, while the six expert-identified competencies can help track how prepared smokers are to quit smoking.

SOCIETAL

On a broader level, this thesis also contributes to individuals' smoking cessation efforts, in-person smoking cessation support, and the design and development of smoking cessation interventions.

Helping individuals quit smoking. Our insights, into how RL can be used to make smoking cessation support provided by virtual coaches more effective, can be used to develop smoking cessation interventions that help more people successfully quit smoking. Given that it takes many smokers more than 30 quit attempts before being successful [86], more effective support is welcome. Moreover, by helping individuals quit smoking, the large burden smoking places on the healthcare system can be reduced. In light of the current strain on the financial, staffing, and societal sustainability of healthcare [105], this is highly valuable.

Improving in-person smoking cessation support. In addition to *digital* smoking cessation support, our insights can also help improve *in-person* smoking cessation support. For example, our RL model for choosing activities that build competencies for quitting smoking can also inform the activity choices of a human healthcare professional. Similarly, understanding factors influencing whether human feedback given in an eHealth application is effective in the long run may help identify people who would benefit from additional in-person support. Or, framed differently, it may help identify people who may need less support, thus freeing up resources for those who do need more help. Some of the people needing more help from human healthcare professionals may then also be people who are enrolled in a virtual coach-based smoking cessation application. This thesis presents one way such a blended intervention could be implemented. Blended smoking cessation interventions are promising given that adherence to them might be similar to the one to fully face-to-face interventions [341] while reducing the need for healthcare professionals.

Designing and developing smoking cessation interventions. Our recommendations for addressing the user needs we identified can aid the design and development of smoking cessation interventions that better meet user needs. In addition, our findings support implementing and testing our three algorithms in a full smoking cessation intervention. The cold-start problem could thereby be avoided by using our data to pretrain the algorithms. And material to help operationalize these algorithms (e.g., persuasive messages, competency-building and belief-changing activities) is publicly available as well. Since the *how-*, *what-*, and *who-*dimensions of support are also relevant to other health behavior change domains, our algorithms and the underlying material and data can further serve

as inspiration for developing eHealth applications for changing other behaviors such as physical inactivity or unhealthy eating.

REFLECTION ON ETHICAL RESEARCH

Inspired by the topology on ethical machine learning research put forward by Ashurst et al. [41], we here take a moment to reflect on the integrity of our research, the way we mitigated research process harms, and possible downstream consequences of future uses of our research. The concerns discussed here are by no means exhaustive. Yet, they should give an idea of the ethical issues at stake.

RESEARCH INTEGRITY

Two primary concerns concerning the integrity of our research are reproducibility and replicability.

Reproducibility. Reproducing results means "obtaining consistent results using the same input data, computational methods, and conditions of analysis" [264]. Reproducibility thus ensures that results are accurate and can be trusted. To facilitate reproducing our quantitative analyses, we have published all underlying data and analysis code in public repositories. Moreover, we provide Docker images to make using the same computational environment easier. Publishing Docker images has also been especially recommended for Bayesian analyses like the ones we conduct for *H1*, *RQ3*, and *RQ4* [372].

Replicability. Replicability of a study's results means that "applying the same methods to the same scientific question produces similar results" [264]. While it is normal that not all findings can be replicated due to, among others, the intrinsic variation and complexity of the phenomena that are studied, there are also avoidable sources of non-replicability [264]. In their book, National Academies of Sciences, Engineering, and Medicine [264] describe six such sources, including publication bias, misaligned incentives, inappropriate statistical inference, poor study design, errors, and incomplete study reporting. Our main attempt to increase the chance that our findings can be replicated was the preregistration of all of our studies in the Open Science Framework (OSF) [14–17]. Not only do the preregistration forms provide detailed information on the way the studies were conducted, but they also help to make clear which parts of our analyses are exploratory versus confirmatory. For example, while we performed a confirmatory analysis of the effectiveness of different algorithm components for *H1*, our analyses of algorithm components for *RQ2*, *RQ3*, and *RQ4* were exploratory. For the qualitative analyses conducted for *RQ1*, we further used both data source triangulation and researcher triangulation to increase the validity of our findings [82]. For the former, we examined quantitative data and literature; for the latter, we included a second researcher with a different background. To give an idea of the reliability, we report agreement coefficients for the coding performed by two researchers. Lastly, our published analysis code helps to determine whether computational errors cause non-replicability.

MITIGATING RESEARCH PROCESS HARMS

Risk to participants. To assess and mitigate potential risks to participants of our studies, we sought and obtained approval from the Human Research Ethics Committee of Delft University of Technology for all of our studies. The main risk for participants was that potentially sensitive information could leak and be connected to them. To mitigate this risk, sensitive data was stored on secure servers at Delft University of Technology and only accessible to the core research team comprising three researchers. Only anonymized data was shared outside this research team. Participants were informed about this risk before participating and only those who provided informed consent could participate.

Compensation of participants. All participants of our crowdsourcing studies were compensated for their time based on the minimum payment rules on the Prolific platform, which ranged from 5 GBP per hour in 2021 to 6 GBP per hour in 2024. While it has been argued that crowdsourcing platforms exploit vulnerable people, Moss et al. [256] showed in the case of MTurk that participants are not worse off financially than the rest of the population and complete studies primarily not for their primary income but as a supplement. Still, our compensation is in line with the minimum wage in some (e.g., Poland [248]) but not all (e.g., Germany [71]) of the OECD countries participants can be from Prolific [311].

DOWNSTREAM CONSEQUENCES

Main concerns regarding the possible impact of future uses of our research relate to security and privacy, health disparities, and application in settings where system and user goals are not aligned and/or participation is mandatory.

7

Security and privacy. While we built small models that require relatively little input from users, they still depend on user data such as motivation, self-efficacy, or effort spent on activities. This data is subject to a variety of security and privacy risks, including unauthorized access, theft, and alteration [175] if the models are implemented in practice. Rather than making it obligatory for everybody to provide this information, users could be given the option to choose from more or less personalized support (e.g., [165]). Since users may not be fully aware of their vulnerability to privacy and security threats [50], they may need to be assisted in making an informed choice.

Reinforcing health disparities. Since we performed our studies primarily with relatively highly educated participants with at least some technical knowledge, our findings might not generalize well to people with lower levels of education and less technical knowledge. Smoking cessation interventions developed based on our findings might thus benefit those people less. Given that a lower socioeconomic status is generally associated with increased smoking prevalence [139], this means that our research may reinforce existing health disparities. Therefore, it is important to test the generalizability of our findings to less highly educated people with less technical knowledge in future work and make adjustments if needed. Notably, it has also been argued for eHealth more generally that even though it has the potential to reduce health disparities by increasing access to healthcare in low-resource settings [315], it often first increases health disparities [364].

Application in settings with misaligned goals and/or mandatory participation.

This thesis proposed three algorithms that try to learn which type of support, in which situation, helps to get smokers to do activities that help them prepare for quitting smoking. While trying to get people to do something might not seem ethical, the context and the way this is done matter. In a smoking cessation intervention, people, presumably, participate voluntarily to quit smoking, and the intervention's purpose is to help them achieve this goal. More generally, Oinas-Kukkonen [275] define behavior change support systems, which include smoking cessation interventions, as "an information system designed to form, alter or reinforce attitudes, behaviors or an act of complying without using deception, coercion or inducements." However, algorithms like ours can also be used in settings in which these criteria are not necessarily met. Kaptein et al. [184], for example, apply their algorithm for choosing persuasive strategies to an online retail setting after earlier applying it to reduce snacking [183]. In such an online retail setting, the algorithm's goal (i.e., maximizing profits) does not necessarily align with people's goals. Moreover, systems differ in how freely people can choose whether to participate in them. For instance, participation in government-led systems such as China's credit scoring system [216] or opt-out procedures for organ donation implemented in countries such as the Netherlands and Switzerland [254] is mandatory. Our research could, in theory, contribute to developing models for such settings. However, given that we used standard RL techniques already publicly available in textbooks [357] and created smoking cessation-specific models that are not directly transferable to other settings, a direct contribution of our research to this seems unlikely for now.

FINAL REMARKS

Unhealthy behavior, and especially smoking, substantially strains the financial, staffing, and societal sustainability of healthcare. While changing such behaviors independently is challenging, eHealth applications, particularly those incorporating conversational agents taking the role of virtual coaches, offer the potential to provide effective support anytime and anywhere. Personalizing the support these applications give might further improve satisfaction, engagement, and health outcomes. A particularly promising approach is to account for users' current and future states, which can be accomplished with RL. Therefore, this thesis investigated how RL can be used to make the support in virtual coach-based smoking cessation interventions more effective. Findings show the benefit of considering current and future states when adapting the three *how-*, *what-*, and *who-*dimensions of smoking cessation support over time. This supports using RL to create behavior change support that is effective in the long run. Deriving model components from the behavior change literature, our results are based on small models that are relatively interpretable and data-efficient, thus fostering transparency and the preservation of privacy. We encourage further research on such psychology-informed RL.

A

Table A.1: Preparatory activities. Titles and formulations for the 24 preparatory activities that were used in the study, 12 each for smoking cessation and physical activity increase. 4 activities had another activity as prerequisite.

Title and formulation	Prereq.
SMOKING CESSATION	
<p>1 Desired future self (writing). Having high aspiration to quit smoking may aid quitting successfully. Thus, before the next session, I advise you to think about the person that you would like to be once you have successfully quit smoking. For example, you might want to be a "grandfather who can play football with his grandchildren" or a "nurse who can walk up the stairs to the fourth floor without getting out of breath." Write down everything that comes to your mind.</p>	
<p>2 Reasons for quitting smoking. Having high aspiration to quit smoking may aid quitting successfully. So, before the next session, I advise you to identify and write down reasons why you want to stop smoking. After writing them down, think about which reasons are most important to you and order them accordingly.</p>	
<p>3 Feared future self (writing). Having high motivation to quit smoking may aid quitting successfully. Thus, before the next session, I advise you to think about who you do NOT want to be in the future but might become if you continue to smoke. For example, you might NOT want to be a "mother who dies early of coronary heart disease like her mother did" or a "husband who is frowned upon by his wife" or a "man who is dependent on a substance." Write down everything that comes to your mind.</p>	

Table A.1: (continued)

	Title and formulation	Prereq.
4	<p>Feared future self (picture). Having high motivation to quit smoking may help to quit successfully. So, before the next session, I advise you to think about who you do NOT want to be in the future but might become if you continue to smoke. For example, you might NOT want to be a "mother who dies early of coronary heart disease like her mother did" or a "husband who is frowned upon by his wife" or a "man who is dependent on a substance." Then, look for or take a picture that best captures your feared future self. Save or print this picture so that you can see it every day.</p>	
5	<p>Fighting match visualization. Focusing on your goal of successfully quitting smoking may help to quit. Thus, before the next session, I advise you to take some time to visualize smoking as a battle. For example, you might see yourself and a cigarette as two boxers in a fighting match. Then imagine yourself winning this battle. Visualize clearly how you win and what it feels like to be the winner. Write down a few words about your winning experience.</p>	
6	<p>Smoking-relevant routines. Getting fewer cravings to smoke may make it easier to successfully quit smoking. Therefore, before the next session, I advise you to think about routines in your daily life that often cause you to get cravings to smoke. For example, you might have experienced that if you go to bed very late and thus sleep less, you smoke more the next day. Or maybe you have noticed that if you skip your breakfast, you always smoke on your way to work but NOT otherwise. How could you change these routines to reduce or even avoid those cravings? Write down everything that comes to your mind.</p>	12
7	<p>High risk situations. Preparing for situations in which avoiding to smoke is difficult may make it easier to successfully quit smoking. Thus, before the next session, I advise you to think about situations in which you might find it difficult to refrain from smoking. For example, this could be during your lunch break at work, when you meet your best friend, or when you watch TV. How could you deal with these situations so that you do NOT smoke? Write down your plans in a few words.</p>	12

Table A.1: (continued)

	Title and formulation	Prereq.
8	<p>Coping with cravings. Planning how to resist urges to smoke may make it easier to successfully quit smoking. Therefore, before the next session, I advise you to think about activities that you could do to keep yourself busy when you feel the urge to smoke so that you do NOT smoke. These urges typically last a few minutes; think about something that you could do in the meantime until the urge has passed. For example, you could water your plants, eat a carrot, do 10 push-ups, or do something for another person in need. Write down everything that comes to your mind.</p>	
9	<p>Relaxation exercise. Tensing and relaxing areas of the body can reduce cravings and withdrawal symptoms, because it is very difficult to feel tense or uptight in a relaxed body. Thus, before the next session, I advise you to watch the following 15-minute video to learn progressive muscle relaxation (which is a way of relaxing your body): https://www.youtube.com/watch?v=ih002wUzgcZab_channel=MarkConnelly. Even if you have already heard of this technique, it might be a good idea to refresh your memory.</p>	
10	<p>Personal rule. Having strong determination to refrain from smoking may help to quit successfully. So, before the next session, I advise you to take some time to create a personal rule that helps you to refrain from smoking. Possible examples include "Not a puff - no matter what," "Say no to smoking, yes to life" or "Smoking is NOT an option." Write down your rule on a piece of paper and repeat it to yourself 3 times. Put the piece of paper with your rule somewhere you can see it every day.</p>	
11	<p>Education on body repair. Having a strong desire to refrain from smoking may aid quitting successfully. Thus, before the next session, I advise you to watch the following 12-minute video on how the body starts repairing itself immediately, as soon as a person stops smoking: https://www.youtube.com/watch?v=ZhTOC0T3P3cZab_channel=RespiratoryTherapyZone. What information from the video is most relevant to you? Take a few notes.</p>	

Table A.1: (continued)

	Title and formulation	Prereq.
12	<p>Recording smoking behavior. Preparing for situations in which you commonly smoke may make it easier to successfully quit smoking. Therefore, I recommend that you record the situations in which you smoke before the next session. Take note of one or two keywords to describe the situation and the number of cigarettes that you smoked. For example, you might note "Lunch break, 2 cigarettes" or "TV, 5 cigarettes." It might be helpful to take these notes on your phone, or you could carry a small piece of paper and pen in your pocket.</p>	
<p>PHYSICAL ACTIVITY INCREASE</p>		
13	<p>Recording physical activity. Becoming more physically active (e.g. exercise, take walks, sit less) may make it easier to successfully quit smoking. One important step for becoming more physically active is to know one's current level. This allows to later set a precise goal and hence to feel more motivated. So, I recommend that you record your current behavior with regards to physical activity before the next session. Try to keep track of how much time you spend 1) sitting, 2) working out and 3) being moderately active (e.g. taking a walk, biking to the grocery store). For this, it might be helpful to keep a piece of paper and pen on your kitchen table, or maybe you have a smart watch that can record these types of behavior for you.</p>	
14	<p>Barriers. Becoming more physically active (e.g. exercise, take walks, sit less) may help you to successfully quit smoking. One important step for becoming more physically active is to remove possible obstacles. Thus, before the next session, I advise you to think about things that make it difficult for you to be physically active. For example, this could be that you do NOT have a raincoat to bike to the grocery store when it is raining, that you do NOT want to work out alone, or that you are at work all day and too exhausted by the time that you come home. What are possible solutions to your barriers? For instance, you could buy a raincoat, join a running group, or take a walk during your lunch break at work. Write down everything that comes to your mind.</p>	
15	<p>Reasons for becoming more physically active. Quitting smoking may be easier if you become more physically active (e.g. exercise, take walks, sit less). One crucial step for this is to have a strong desire to become more physically active. Therefore, before the next session, I advise you to identify and write down reasons why you want to become more physically active. After writing them down, think about which reasons are most important to you and order them accordingly.</p>	

Table A.1: (continued)

	Title and formulation	Prereq.
16	<p>Recommended physical activity. Quitting smoking may be easier if you become more physically active (e.g. exercise, take walks, sit less). One important step for becoming more physically active is to set a specific goal and thus to feel more aspiration. Therefore, before the next session, I advise you to watch the following 2-minute video on how much and which type of physical activity is recommended: https://www.youtube.com/watch?v=AAPhWbG_zLsZab_channel=TREKGroup. Then, compare your physical activity behavior to the recommended amounts for the different types of physical activity. Write down which recommended amounts you meet or exceed, and which ones you do NOT meet.</p>	13
17	<p>Desired future self (writing). Quitting smoking may be easier if you become more physically active (e.g. exercise, take walks, sit less). One important step for this is to have high ambition to become more physically active. Thus, before the next session, I advise you to think about the person that you would like to be once you have become more physically active. For example, you might want to be a "grandfather who can play football with his grandchildren" or a "nurse who can walk up the stairs to the fourth floor without getting out of breath." Write down everything that comes to your mind.</p>	
18	<p>Feared future self (writing). It may be easier to successfully quit smoking if you become more physically active (e.g. exercise, take walks, sit less). One crucial step for this is to have high determination to become more physically active. Therefore, before the next session, I advise you to think about who you do NOT want to be in the future but might become if you fail to become more physically active. For example, you might NOT want to be a "mother who dies early of coronary heart disease like her father did" or a "daughter who is frowned upon by her mother" or a "man who is dependent on his wife in his everyday life." Write down everything that comes to your mind.</p>	

Table A.1: (continued)

	Title and formulation	Prereq.
19	<p>Feared future self (picture). It may be easier to successfully quit smoking if you become more physically active (e.g. exercise, take walks, sit less). One crucial step for this is to have high determination to become more physically active. Therefore, before the next session, I advise you to think about who you do NOT want to be in the future but might become if you fail to become more physically active. For example, you might NOT want to be a "mother who dies early of coronary heart disease like her father did" or a "daughter who is frowned upon by her mother" or a "man who is dependent on his wife in his everyday life." Then, look for or take a picture that best captures your feared future self. Save or print this picture so that you can see it every day.</p>	
20	<p>Fighting match visualization. Becoming more physically active (e.g. exercise, take walks, sit less) may help you to successfully quit smoking. One important step for this is to focus on the goal of becoming more physically active. Thus, before the next session, I advise you to take some time to visualize becoming more physically active as a battle. For example, you might see yourself and non-active version of yourself as two boxers in a fighting match. Then imagine yourself winning this battle. Visualize clearly how you win and what it feels like to be the winner. Write down a few words about your winning experience.</p>	
21	<p>Desired future self (picture). Quitting smoking may be easier if you become more physically active (e.g. exercise, take walks, sit less). One crucial step for this is to have high motivation to become more physically active. Thus, before the next session, I advise you to think about the person that you would like to be once you have become more physically active. For example, you might want to be a "grandfather who can play football with his grandchildren" or a "nurse who can walk up the stairs to the fourth floor without getting out of breath." Then look for or take a picture that best captures your desired future self. Save or print this picture so that you can see it every day.</p>	
22	<p>Plan for becoming more physically active. Becoming more physically active (e.g. exercise, take walks, sit less) may help you to successfully quit smoking. One crucial part for this is to create a plan for becoming more physically active. Therefore, before the next session, I advise you to think about what you could do to become more physically active. For example, you could get up from your desk after every 30 minutes of sitting, bike to the grocery store, do 10 squats every morning, or join a running group. Write down everything that comes to your mind. Which plan do you want to focus on? Highlight this plan.</p>	13

Table A.1: (continued)

	Title and formulation	Prereq.
23	<p>Impact of physical activity on cravings. Becoming more physically active (e.g. exercise, take walks, spend less time sitting) may help you to successfully quit smoking. One crucial step for this is to have high aspiration to become more physically active. So, before the next session, I advise you to watch the following 5-minute video about the possible positive impact of physical activity on dealing with cravings to smoke: https://www.youtube.com/watch?v=StM10jzbt1kZab_channel=TreeHouseRecovery. What do you think about the information in the video? Write down your thoughts in a few words.</p>	
24	<p>Personal rule. Being more physically active (e.g. exercise, take walks, spend less time sitting) may aid you to stop smoking. One important aspect for this is to have strong resolve to become more physically active. So, before the next session, I advise you to take some time to create a personal rule that helps you to become more physically active. Possible examples include "10 squats - no matter what," "Say no to sitting, yes to life" or "Driving to the grocery store is NOT an option." Write down your rule on a piece of paper and repeat it to yourself 3 times. Put the piece of paper with your rule somewhere you can see it every day.</p>	

Abbreviations: Prereq., Prerequisite.

Table A.2: Interaction scenarios. Some scenarios about similar interactions are grouped together to facilitate their analysis.

Scenario	Topic	Scenario group
INTERACTION WITH VIRTUAL COACH		
1	Follow physical activity program while quitting smoking	Follow PA program
2	Plan for smoking HRSs in the mornings	Regular HRS planning/reflection
3	Plan for physical activity HRSs on Sundays	Regular HRS planning/reflection
4	Help button for smoking HRSs	Help button for HRS
5	Help button for physical activity HRSs	Help button for HRS
6	Reflect on smoking HRSs in the evenings	Regular HRS planning/reflection
7	Reflect on physical activity HRSs on Sundays	Regular HRS planning/reflection
8	Discuss repeated failure of reaching physical activity goals	Discuss repeated failure PA
9	Receive motivational messages	Receive motivational messages
INTERACTION WITH SOCIAL ENVIRONMENT		
10	Tell SE about quit attempt	Involve SE
11	Discuss with an SO how they can support the quit attempt	Involve SE
INTERACTION WITH GENERAL PRACTITIONER		
12	Consult GP at start of quit attempt	Involve GP
13	Consult GP in case of smoking relapse	Involve GP

Abbreviations: PA, Physical activity; HRS, High risk situation; SE, Social environment; SO, Significant other; GP, General practitioner.

Table A.3: Links to the videos of the interaction scenarios. We provide links to both the male and the female version.

Scenario	Topic	Link to video
INTERACTION WITH VIRTUAL COACH		
1	Follow physical activity program while quitting smoking	<u>Female:</u> https://youtu.be/_0lmuAuJbfU <u>Male:</u> https://youtu.be/gHZZKX6dQMo
2	Plan for smoking HRSs in the mornings	<u>Female:</u> https://youtu.be/DyZPA781nn4 <u>Male:</u> https://youtu.be/zpWY28vtSZM
3	Plan for physical activity HRSs on Sundays	<u>Female:</u> https://youtu.be/JFz17RAXD0U <u>Male:</u> https://youtu.be/CA78pzMoFuo
4	Help button for smoking HRSs	<u>Female:</u> https://youtu.be/aLrCZyFvs1o <u>Male:</u> https://youtu.be/sLBF2BuMv80
5	Help button for physical activity HRSs	<u>Female:</u> https://youtu.be/5j1a6N1hs8s <u>Male:</u> https://youtu.be/HZLtG4LfvIs
6	Reflect on smoking HRSs in the evenings	<u>Female:</u> https://youtu.be/dkUkzneVDyM <u>Male:</u> https://youtu.be/Raq7Qyvh45A
7	Reflect on physical activity HRSs on Sundays	<u>Female:</u> https://youtu.be/nczb97qrgOg <u>Male:</u> https://youtu.be/FXC7oZLLzvQ
8	Discuss repeated failure of reaching physical activity goals	<u>Female:</u> https://youtu.be/UHds1aEB1GQ <u>Male:</u> https://youtu.be/STO8tEGpWXI
9	Receive motivational messages	<u>Female:</u> https://youtu.be/eV6TLfa-hIs <u>Male:</u> https://youtu.be/IGKRjhK40Cs
INTERACTION WITH SOCIAL ENVIRONMENT		
10	Tell SE about quit attempt	<u>Female:</u> https://youtu.be/Soz1TeF0EiY

Table A.3: (continued)

Scenario	Topic	Link to video
11	Discuss with an SO how they can support the quit attempt	<u>Male:</u> https://youtu.be/wDxMd-vwJgM <u>Female:</u> https://youtu.be/UPxHprogZc0 <u>Male:</u> https://youtu.be/ZbopPvhQJWw
INTERACTION WITH GENERAL PRACTITIONER		
12	Consult GP at start of quit attempt	<u>Female:</u> https://youtu.be/1DTeHgBA71w <u>Male:</u> https://youtu.be/2WJQDn251DM
13	Consult GP in case of smoking relapse	<u>Female:</u> https://youtu.be/R1-xdeRexUw <u>Male:</u> https://youtu.be/w5W1DGhKTn4

Abbreviations: HRS, High risk situation; SE, Social environment; SO, Significant other; GP, General practitioner.

Table A.4: Details on how we measured the activity feedback, and barriers and motivators for doing the preparatory activities.

Measure	Explanation	Specification	Source
✍ ACTIVITY FEEDBACK			
Activity effort	Effort spent on activity from previous session	Scale from 0 ("Nothing") to 10 ("Extremely strong").	Adapted based on Hutchinson and Tenenbaum [170].
Activity experience	Experience with activity from previous session	Free-text response to the question "How did you approach, do, or experience your assigned activity?" If the effort spent on an activity was lower than four, the question was supplemented by the sentence "If you have no feedback on the activity, just type 'None' in the text field." In the post-questionnaire, this additional sentence was provided independent of the activity effort.	
Activity experience modification	Changes and/or additions for the activity experience	After a user had provided an answer to the activity experience question, the virtual coach played the answer back to the user and gave them the option to make changes or additions. If a user indicated that they would like to make a change or addition, they were asked to provide another free-text response.	

Table A.4: (continued)

Measure	Explanation	Specification	Source
Barriers	Barriers for doing the activities	Free-text response to the question "What were barriers for you to do your assigned activities? Barriers are factors that may have restricted, impeded, or blocked your completion of your activities. Please enter "None" if there was nothing and nobody that restricted, impeded, or blocked your completion of your activities."	Definition of barrier in the APA Dictionary of Psychology [36].
Motivators	Motivators for doing the activities	Free-text response to the question "What aspects motivated you to do your assigned activities? Please enter "None" if there was nothing and nobody that motivated you."	

Table A.5: Question and scale endpoints for each interaction scenario.

Scenario	Topic	Question	Scale endpoints
INTERACTION WITH VIRTUAL COACH			
1	Follow physical activity program while quitting smoking	If this were you, would you follow this advice and also follow the program for becoming more physically active?	-5 ("Definitely no"), 5 ("Definitely yes")
2	Plan for smoking HRSs in the mornings	If this were you, would you make such a plan with your virtual coach in the morning?	-5 ("Definitely no"), 5 ("Definitely yes")
3	Plan for physical activity HRSs on Sundays	If this was you, would you make such a plan with your virtual coach on Sundays?	-5 ("Definitely no"), 5 ("Definitely yes")
4	Help button for smoking HRSs	If this was you, would you press the help button when you have a craving?	-5 ("Definitely no"), 5 ("Definitely yes")
5	Help button for physical activity HRSs	If this was you, would you press the help button in this situation?	-5 ("Definitely no"), 5 ("Definitely yes")
6	Reflect on smoking HRSs in the evenings	If this was you, would you follow this advice and consult your virtual coach in the evening?	-5 ("Definitely no"), 5 ("Definitely yes")
7	Reflect on physical activity HRSs on Sundays	If this was you, would you consult your virtual coach on Sundays?	-5 ("Definitely no"), 5 ("Definitely yes")
8	Discuss repeated failure of reaching physical activity goals	If this was you, would you take this advice and consult your virtual coach?	-5 ("Definitely no"), 5 ("Definitely yes")

Table A.5: (continued)

Scenario	Topic	Question	Scale endpoints
9	Receive motivational messages	If this were you, what do you think would be the impact of such a motivational message on you doing the activity?	-5 ("Definitely negative"), 5 ("Definitely positive")
INTERACTION WITH SOCIAL ENVIRONMENT			
10	Tell SE about quit attempt	If this was you, would you follow up on this advice and tell your social environment about your quit attempt?	-5 ("Definitely no"), 5 ("Definitely yes")
11	Discuss with an SO how they can support the quit attempt	If this were you, would you follow up on this advice and reach out to a significant other for additional support?	-5 ("Definitely no"), 5 ("Definitely yes")
INTERACTION WITH GENERAL PRACTITIONER			
12	Consult GP at start of quit attempt	If this was you, would you follow up on this advice and contact your GP?	-5 ("Definitely no"), 5 ("Definitely yes")
13	Consult GP in case of smoking relapse	If this was you, would you contact your GP if your app advises you to do so?	-5 ("Definitely no"), 5 ("Definitely yes")

Abbreviations: HRS, High risk situation; SE, Social environment; SO, Significant other; GP, General practitioner.

Table A.6: Participant characteristics. Characteristics of the 671 participants with at least one valid free-text response.

Characteristic	Value
AGE	
- 18 – 30, n (%)	315 (46.94%)
- 31 – 40, n (%)	168 (25.04%)
- 41 – 50, n (%)	100 (14.90%)
- 51 – 60, n (%)	70 (10.43%)
- 61 – 74, n (%)	18 (2.68%)
GENDER	
- Female, n (%)	349 (52.01%)
- Male, n (%)	310 (46.20%)
- Other, n (%)	11 (1.64%)
- <i>No data</i> , n (%)	1 (0.15%)
HIGHEST COMPLETED EDUCATION LEVEL	
- No formal qualifications, n (%)	5 (0.75%)
- Secondary education (e.g. GED/GCSE), n (%)	76 (11.33%)
- High school diploma/A-levels, n (%)	170 (25.34%)
- Technical/community college, n (%)	103 (15.35%)
- Undergraduate degree (BA/BSc/other), n (%)	211 (31.45%)
- Graduate degree (MA/MSc/MPhil/other), n (%)	95 (14.16%)
- Doctorate degree (PhD/other), n (%)	7 (1.04%)
- <i>No data</i> , n (%)	4 (0.60%)
PREVIOUS QUIT ATTEMPT OF AT LEAST 24 HOURS	
- Yes, n (%)	528 (78.69%)
- No, n (%)	143 (21.31%)
SMOKING FREQUENCY	
- Once a day, n (%)	36 (5.37%)
- 2 – 5 times a day, n (%)	165 (24.59%)
- 6 – 10 times a day, n (%)	213 (31.74%)
- 11 – 19 times a day, n (%)	190 (28.32%)
- More than 20 times a day, n (%)	64 (9.54%)
- <i>No data</i> , n (%)	3 (0.45%)
TTM-STAGE FOR BECOMING PHYSICALLY ACTIVE	
- Precontemplation, n (%)	44 (6.56%)
- Contemplation, n (%)	188 (28.02%)
- Preparation, n (%)	159 (23.70%)
- Action, n (%)	83 (12.37%)
- Maintenance, n (%)	197 (29.36%)

Table A.6: (continued)

Characteristic	Value
TTM-STAGE FOR QUITTING SMOKING	
- Contemplation, n (%)	577 (85.99%)
- Preparation, n (%)	94 (14.01%)
WEEKLY EXERCISE AMOUNT	
- Never (0 – 60 minutes per week), n (%)	174 (25.93%)
- Sometimes (60 – 150 minutes per week), n (%)	322 (47.99%)
- Often (more than 150 minutes per week), n (%)	170 (25.34%)
- <i>No data</i> , n (%)	5 (0.75%)

Abbreviations: GED, General educational development; GCSE, General certificate of secondary education; BA, Bachelor of Arts; BSc, Bachelor of Science; MA, Master of Arts; MSc, Master of Science; MPhil, Master of Philosophy; PhD, Doctor of Philosophy; TTM, Transtheoretical model.




Figure A.7: Coding scheme. The coding scheme consists of 4 codes at the highest level, 15 codes at the second level, and 86 codes at the third level. We show the reliability for each coding level.

A

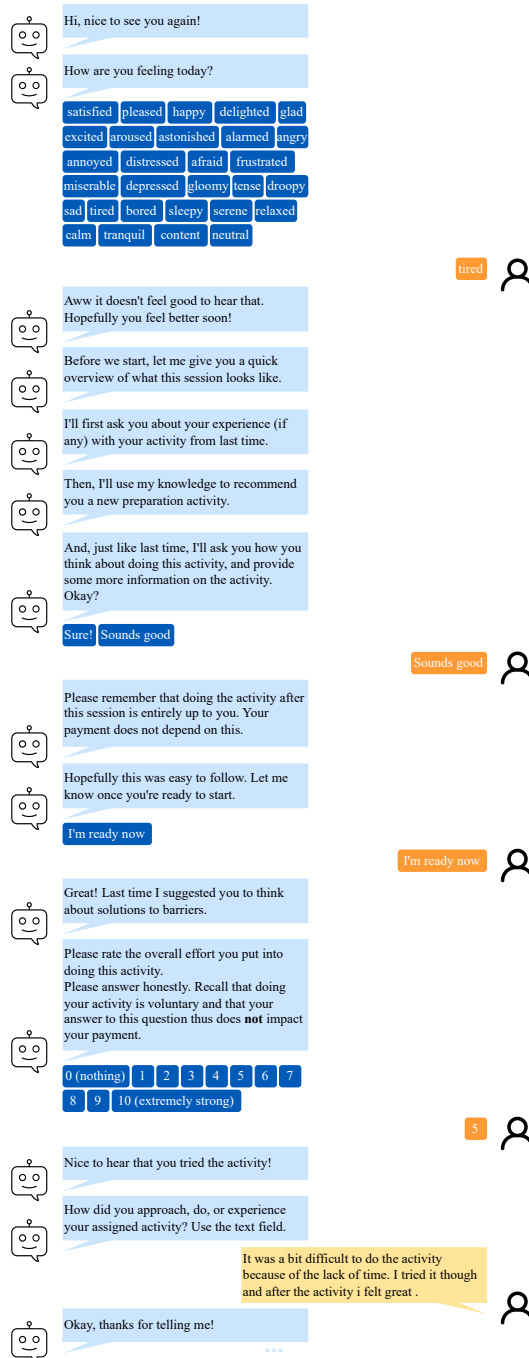


Figure A.8: Conversation example. The example shows the start of the second conversational session with Sam, including responses given by a participant.




Figure A.9: Percentage of times each code from the coding scheme appears in each response type as well as across all response types together. The response types are the activity experiences, barriers, motivators, and the groups of interaction scenarios. N denotes the number of responses.

B

B

Table B.1: Examples of persuasive message templates and resulting messages.

Persuasion type	Template example	Message example
Action planning	Please think about the time after this and before the next session. When, where and how do you plan to do your recommended preparation activity? Please formulate a rule for completing your activity that has the form "If <situation>, then I will <do activity>." For instance, "If I get out of the shower tonight, then I will <do activity>." Type your rule into this chat. The more precise, concrete and personally you formulate your rule, the more it can help you.	Please think about the time after this and before the next session. When, where and how do you plan to do your recommended preparation activity? Please formulate a rule for completing your activity that has the form "If <situation>, then I will <do activity>." For instance, "If I get out of the shower tonight, then I will <i>identify reasons why I want to stop smoking</i> ." Type your rule into this chat. The more precise, concrete and personally you formulate your rule, the more it can help you.
Authority	Experts recommend <doing activity> to <positive impact of activity>.	Experts recommend <i>identifying reasons why you want to stop smoking to increase your aspiration to stop smoking</i> .
Commitment	You've committed to become somebody who has quit smoking. <doing activity> may help you to become this person.	You've committed to become somebody who has quit smoking. <i>Identifying reasons why you want to stop smoking</i> may help you to become this person.
Consensus	Most people think that <doing activity> may help to <positive impact of activity>.	Most people think that <i>identifying reasons why you want to stop smoking</i> may help to <i>increase your aspiration to stop smoking</i> .




Figure B.2: Participant flow through the study components in the experiment. The numbers next to the downward arrows denote how many people started the study components. We show the distribution across the four algorithm complexity levels for the participants who did not respond to the invitation to a study component after the randomization. Note that participants can return their submission on Prolific to withdraw from a study.

Table B.3: Participant characteristics for each algorithm complexity level.

Characteristic	Algorithm complexity level			
	1	2	3	4
NUMBER				
- n	162	167	163	174
GENDER				
- Female, n (%)	84 (51.85%)	87 (52.10%)	83 (50.92%)	91 (52.30%)
- Male, n (%)	76 (46.91%)	76 (45.51%)	78 (47.85%)	79 (45.40%)
AGE				
- 18–30, n (%)	73 (45.06%)	84 (50.30%)	81 (49.69%)	74 (42.53%)
- 31–40, n (%)	47 (29.01%)	31 (18.56%)	38 (23.31%)	50 (28.74%)
- 41–50, n (%)	21 (12.96%)	29 (17.37%)	19 (11.66%)	31 (17.82%)
- 51–60, n (%)	15 (9.26%)	20 (11.98%)	20 (12.27%)	15 (8.62%)
- 61–74, n (%)	6 (3.70%)	3 (1.80%)	5 (3.07%)	4 (2.30%)
BIG-5 PERSONALITY				
- Agr., Mean (SD)	4.83 (1.17)	4.90 (1.11)	4.85 (1.15)	4.86 (1.14)
- Con., Mean (SD)	4.89 (1.41)	4.88 (1.42)	4.93 (1.45)	4.94 (1.38)
- ES, Mean (SD)	4.23 (1.60)	4.21 (1.57)	4.15 (1.55)	4.16 (1.60)
- Ext., Mean (SD)	3.90 (1.61)	3.89 (1.59)	3.98 (1.63)	3.93 (1.59)
- OE, Mean (SD)	5.20 (1.15)	5.22 (1.13)	5.23 (1.15)	5.19 (1.14)
DROPOUT AFTER SESSION 3*				
- No response, n (%)	18 (11.11%)	15 (8.98%)	20 (12.27%)	18 (10.34%)
EFFORT FOR FIRST ACTIVITY				
- Mean (SD)	5.55 (2.52)	5.54 (2.54)	5.48 (2.56)	5.45 (2.47)
PREVIOUS QUIT ATTEMPT OF AT LEAST 24 HOURS				
- Yes, n (%)	128 (79.01%)	127 (76.05%)	135 (82.82%)	133 (76.44%)
- No, n (%)	34 (20.99%)	40 (23.95%)	28 (17.18%)	41 (23.56%)
TTM-STAGE FOR BECOMING PHYSICALLY ACTIVE				
- Precontemplation, n (%)	9 (5.56%)	14 (8.38%)	10 (6.13%)	10 (5.75%)
- Contemplation, n (%)	44 (27.16%)	43 (25.75%)	49 (30.06%)	50 (28.74%)
- Preparation, n (%)	43 (26.54%)	41 (24.55%)	38 (23.31%)	37 (21.26%)
- Action, n (%)	19 (11.73%)	19 (11.38%)	19 (11.66%)	25 (14.37%)
- Maintenance, n (%)	47 (29.01%)	50 (29.94%)	47 (28.83%)	52 (29.89%)
TTM-STAGE FOR QUITTING SMOKING				
- Contemplation, n (%)	133 (82.10%)	147 (88.02%)	140 (85.89%)	152 (87.36%)

Table B.3: (continued)

Characteristic	Algorithm complexity level			
	1	2	3	4
- Preparation, n (%)	29 (17.90%)	20 (11.98%)	23 (14.11%)	22 (12.64%)

* Session 3 was the first session in which participants were treated differently based on their condition (i.e., persuaded based on the four algorithm complexity levels).

Abbreviations: SD, Standard deviation; Agr., Agreeableness; Con., Conscientiousness; ES, Emotional stability; Ext., Extraversion; OE, Openness to experience.

B

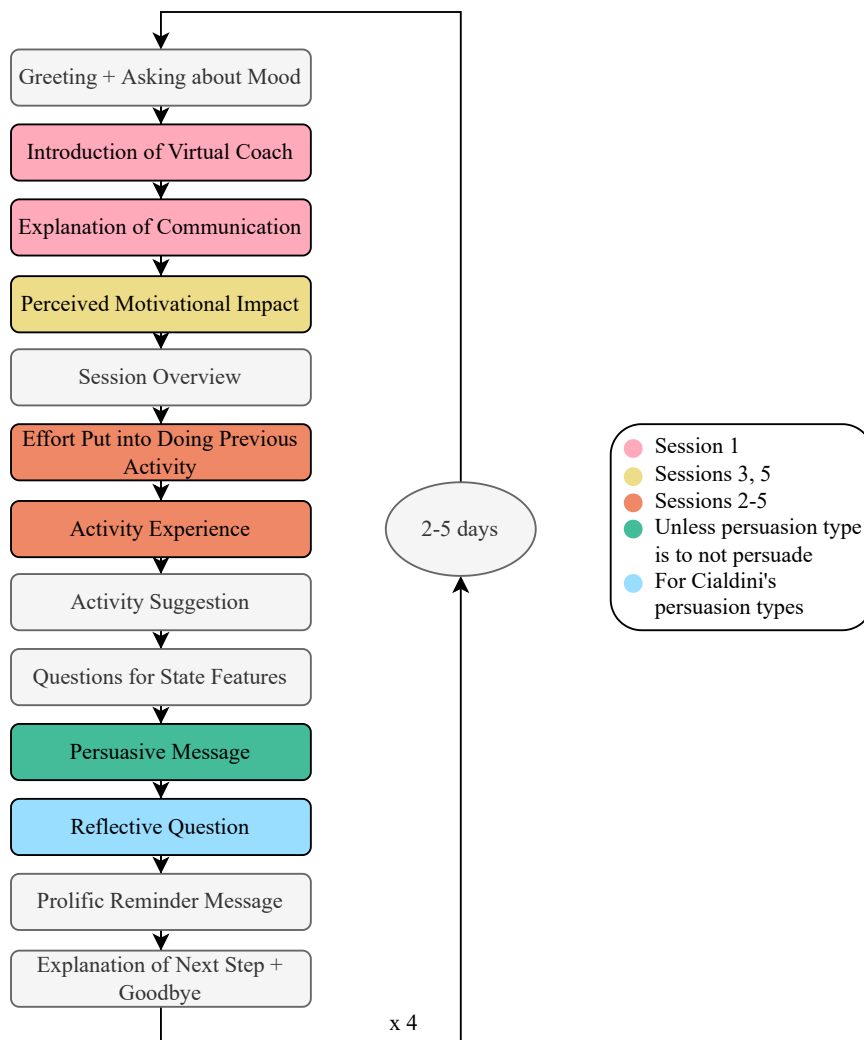
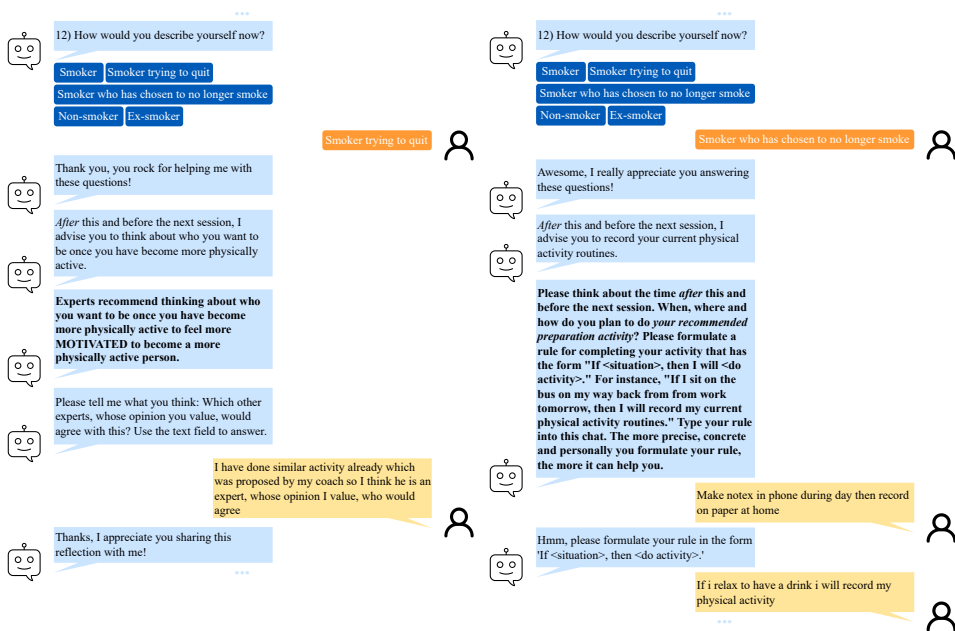


Figure B.4: Structure of the five conversational sessions with the virtual coach Sam.



(a) Authority.

(b) Action planning.

Figure B.5: Two excerpts of actual conversations with the virtual coach. The excerpts include the last state question and persuasion based on the persuasion types authority (a) and action planning (b).

BTable B.6: Interpretation guidelines for Cohen's κ based on Landis and Koch [209].

Cohen's κ	Evaluation
≤ 0	No agreement
0.01-0.20	No to slight agreement
0.21-0.40	Fair agreement
0.41-0.60	Moderate agreement
0.61-0.80	Substantial agreement
0.81-1.00	Almost perfect agreement

Table B.7: Number of times participants reported the effort they spent on each activity and the mean effort that was reported. We include only the samples used in our analysis for *H1* (i.e., not the five samples of participants that were never assigned to a condition due to not completing session 2 and which were only used in the analysis for *H2*).

	Activity	Number	Effort mean (SD)
SMOKING CESSATION			
1	Desired future self (writing)	114	5.91 (2.28)
2	Reasons for quitting smoking	131	5.65 (2.88)
3	Feared future self (writing)	82	5.15 (2.94)
4	Feared future self (picture)	96	5.75 (2.62)
5	Fighting match visualization	118	5.34 (2.70)
6	Smoking-relevant routines*	7	5.86 (3.24)
7	High risk situations*	7	7.71 (1.50)
8	Coping with cravings	127	5.87 (2.52)
9	Relaxation exercise	136	5.18 (3.16)
10	Personal rule	113	4.88 (3.09)
11	Education on body repair	103	5.22 (3.07)
12	Recording smoking behavior	140	5.69 (2.96)
	Total	1174	5.49 (2.85)
PHYSICAL ACTIVITY INCREASE			
13	Recording physical activity	128	5.39 (2.73)
14	Barriers	136	5.60 (2.66)
15	Reasons for becoming more physically active	125	5.80 (2.56)
16	Recommended physical activity*	3	4.33 (4.04)
17	Desired future self (writing)	124	5.27 (2.47)
18	Feared future self (writing)	95	5.84 (2.76)
19	Feared future self (picture)	95	5.35 (2.81)
20	Fighting match visualization	127	5.00 (2.75)
21	Desired future self (picture)	119	4.76 (2.75)
22	Plan for becoming more physically active*	5	5.00 (2.65)
23	Impact of physical activity on cravings	121	5.18 (3.24)
24	Personal rule	109	5.21 (3.13)
	Total	1187	5.33 (2.80)

* Activity had another activity as a prerequisite.

Abbreviations: SD, Standard deviation.

Table B.8: Formulations for the 24 preparatory activities that were used in the study. In case an activity involved watching a video, there were two different formulations for the session and the reminder message, with only the one for the latter containing the link to the video. This was to prevent participants from directly clicking on the video link when reading their activity during the session.

B

	Activity formulation	Prereq.	Type
1	Having high aspiration to quit smoking may aid quitting successfully. Thus, before the next session, I advise you to think about the person that you would like to be once you have successfully quit smoking. For example, you might want to be a "grandfather who can play football with his grandchildren" or a "nurse who can walk up the stairs to the fourth floor without getting out of breath." Write down everything that comes to your mind.		Smoking cessation
2	Having high aspiration to quit smoking may aid quitting successfully. So, before the next session, I advise you to identify and write down reasons why you want to stop smoking. After writing them down, think about which reasons are most important to you and order them accordingly.		Smoking cessation
3	Having high motivation to quit smoking may aid quitting successfully. Thus, before the next session, I advise you to think about who you do NOT want to be in the future but might become if you continue to smoke. For example, you might NOT want to be a "mother who dies early of coronary heart disease like her mother did" or a "husband who is frowned upon by his wife" or a "man who is dependent on a substance." Write down everything that comes to your mind.		Smoking cessation
4	Having high motivation to quit smoking may help to quit successfully. So, before the next session, I advise you to think about who you do NOT want to be in the future but might become if you continue to smoke. For example, you might NOT want to be a "mother who dies early of coronary heart disease like her mother did" or a "husband who is frowned upon by his wife" or a "man who is dependent on a substance." Then, look for or take a picture that best captures your feared future self. Save or print this picture so that you can see it every day.		Smoking cessation

Table B.8: (continued)

	Activity formulation	Prereq.	Type
5	Focusing on your goal of successfully quitting smoking may help to quit. Thus, before the next session, I advise you to take some time to visualize smoking as a battle. For example, you might see yourself and a cigarette as two boxers in a fighting match. Then imagine yourself winning this battle. Visualize clearly how you win and what it feels like to be the winner. Write down a few words about your winning experience.		Smoking cessation
6	Getting fewer cravings to smoke may make it easier to successfully quit smoking. Therefore, before the next session, I advise you to think about routines in your daily life that often cause you to get cravings to smoke. For example, you might have experienced that if you go to bed very late and thus sleep less, you smoke more the next day. Or maybe you have noticed that if you skip your breakfast, you always smoke on your way to work but NOT otherwise. How could you change these routines to reduce or even avoid those cravings? Write down everything that comes to your mind.	12	Smoking cessation
7	Preparing for situations in which avoiding to smoke is difficult may make it easier to successfully quit smoking. Thus, before the next session, I advise you to think about situations in which you might find it difficult to refrain from smoking. For example, this could be during your lunch break at work, when you meet your best friend, or when you watch TV. How could you deal with these situations so that you do NOT smoke? Write down your plans in a few words.	12	Smoking cessation
8	Planning how to resist urges to smoke may make it easier to successfully quit smoking. Therefore, before the next session, I advise you to think about activities that you could do to keep yourself busy when you feel the urge to smoke so that you do NOT smoke. These urges typically last a few minutes; think about something that you could do in the meantime until the urge has passed. For example, you could water your plants, eat a carrot, do 10 push-ups, or do something for another person in need. Write down everything that comes to your mind.		Smoking cessation

Table B.8: (continued)

	Activity formulation	Prereq.	Type
9	Tensing and relaxing areas of the body can reduce cravings and withdrawal symptoms, because it is very difficult to feel tense or uptight in a relaxed body. Thus, before the next session, I advise you to watch the following 15-minute video to learn progressive muscle relaxation (which is a way of relaxing your body): https://www.youtube.com/watch?v=ih002wUzgkcZab_channel=MarkConnelly . Even if you have already heard of this technique, it might be a good idea to refresh your memory.		Smoking cessation
10	Having strong determination to refrain from smoking may help to quit successfully. So, before the next session, I advise you to take some time to create a personal rule that helps you to refrain from smoking. Possible examples include "Not a puff - no matter what," "Say no to smoking, yes to life" or "Smoking is NOT an option." Write down your rule on a piece of paper and repeat it to yourself 3 times. Put the piece of paper with your rule somewhere you can see it every day.		Smoking cessation
11	Having a strong desire to refrain from smoking may aid quitting successfully. Thus, before the next session, I advise you to watch the following 12-minute video on how the body starts repairing itself immediately, as soon as a person stops smoking: https://www.youtube.com/watch?v=ZhTOC0T3P3cZab_channel=RespiratoryTherapyZone . What information from the video is most relevant to you? Take a few notes.		Smoking cessation
12	Preparing for situations in which you commonly smoke may make it easier to successfully quit smoking. Therefore, I recommend that you record the situations in which you smoke before the next session. Take note of one or two keywords to describe the situation and the number of cigarettes that you smoked. For example, you might note "Lunch break, 2 cigarettes" or "TV, 5 cigarettes." It might be helpful to take these notes on your phone, or you could carry a small piece of paper and pen in your pocket.		Smoking cessation

Table B.8: (continued)

	Activity formulation	Prereq.	Type
13	Becoming more physically active (e.g. exercise, take walks, sit less) may make it easier to successfully quit smoking. One important step for becoming more physically active is to know one's current level. This allows to later set a precise goal and hence to feel more motivated. So, I recommend that you record your current behavior with regards to physical activity before the next session. Try to keep track of how much time you spend 1) sitting, 2) working out and 3) being moderately active (e.g. taking a walk, biking to the grocery store). For this, it might be helpful to keep a piece of paper and pen on your kitchen table, or maybe you have a smart watch that can record these types of behavior for you.		Physical activity increase
14	Becoming more physically active (e.g. exercise, take walks, sit less) may help you to successfully quit smoking. One important step for becoming more physically active is to remove possible obstacles. Thus, before the next session, I advise you to think about things that make it difficult for you to be physically active. For example, this could be that you do NOT have a raincoat to bike to the grocery store when it is raining, that you do NOT want to work out alone, or that you are at work all day and too exhausted by the time that you come home. What are possible solutions to your barriers? For instance, you could buy a raincoat, join a running group, or take a walk during your lunch break at work. Write down everything that comes to your mind.		Physical activity increase
15	Quitting smoking may be easier if you become more physically active (e.g. exercise, take walks, sit less). One crucial step for this is to have a strong desire to become more physically active. Therefore, before the next session, I advise you to identify and write down reasons why you want to become more physically active. After writing them down, think about which reasons are most important to you and order them accordingly.		Physical activity increase

Table B.8: (continued)

	Activity formulation	Prereq.	Type
16	<p>Quitting smoking may be easier if you become more physically active (e.g. exercise, take walks, sit less). One important step for becoming more physically active is to set a specific goal and thus to feel more aspiration. Therefore, before the next session, I advise you to watch the following 2-minute video on how much and which type of physical activity is recommended: https://www.youtube.com/watch?v=AAPhWbG_zLsZab_channel=TREKGroup. Then, compare your physical activity behavior to the recommended amounts for the different types of physical activity. Write down which recommended amounts you meet or exceed, and which ones you do NOT meet.</p>	13	Physical activity increase
17	<p>Quitting smoking may be easier if you become more physically active (e.g. exercise, take walks, sit less). One important step for this is to have high ambition to become more physically active. Thus, before the next session, I advise you to think about the person that you would like to be once you have become more physically active. For example, you might want to be a "grandfather who can play football with his grandchildren" or a "nurse who can walk up the stairs to the fourth floor without getting out of breath." Write down everything that comes to your mind.</p>		Physical activity increase
18	<p>It may be easier to successfully quit smoking if you become more physically active (e.g. exercise, take walks, sit less). One crucial step for this is to have high determination to become more physically active. Therefore, before the next session, I advise you to think about who you do NOT want to be in the future but might become if you fail to become more physically active. For example, you might NOT want to be a "mother who dies early of coronary heart disease like her father did" or a "daughter who is frowned upon by her mother" or a "man who is dependent on his wife in his everyday life." Write down everything that comes to your mind.</p>		Physical activity increase

Table B.8: (continued)

	Activity formulation	Prereq.	Type
19	It may be easier to successfully quit smoking if you become more physically active (e.g. exercise, take walks, sit less). One crucial step for this is to have high determination to become more physically active. Therefore, before the next session, I advise you to think about who you do NOT want to be in the future but might become if you fail to become more physically active. For example, you might NOT want to be a "mother who dies early of coronary heart disease like her father did" or a "daughter who is frowned upon by her mother" or a "man who is dependent on his wife in his everyday life." Then, look for or take a picture that best captures your feared future self. Save or print this picture so that you can see it every day.		Physical activity increase
20	Becoming more physically active (e.g. exercise, take walks, sit less) may help you to successfully quit smoking. One important step for this is to focus on the goal of becoming more physically active. Thus, before the next session, I advise you to take some time to visualize becoming more physically active as a battle. For example, you might see yourself and non-active version of yourself as two boxers in a fighting match. Then imagine yourself winning this battle. Visualize clearly how you win and what it feels like to be the winner. Write down a few words about your winning experience.		Physical activity increase
21	Quitting smoking may be easier if you become more physically active (e.g. exercise, take walks, sit less). One crucial step for this is to have high motivation to become more physically active. Thus, before the next session, I advise you to think about the person that you would like to be once you have become more physically active. For example, you might want to be a "grandfather who can play football with his grandchildren" or a "nurse who can walk up the stairs to the fourth floor without getting out of breath." Then look for or take a picture that best captures your desired future self. Save or print this picture so that you can see it every day.		Physical activity increase

Table B.8: (continued)

	Activity formulation	Prereq.	Type
22	Becoming more physically active (e.g. exercise, take walks, sit less) may help you to successfully quit smoking. One crucial part for this is to create a plan for becoming more physically active. Therefore, before the next session, I advise you to think about what you could do to become more physically active. For example, you could get up from your desk after every 30 minutes of sitting, bike to the grocery store, do 10 squats every morning, or join a running group. Write down everything that comes to your mind. Which plan do you want to focus on? Highlight this plan.	13	Physical activity increase
23	Becoming more physically active (e.g. exercise, take walks, spend less time sitting) may help you to successfully quit smoking. One crucial step for this is to have high aspiration to become more physically active. So, before the next session, I advise you to watch the following 5-minute video about the possible positive impact of physical activity on dealing with cravings to smoke: https://www.youtube.com/watch?v=StM10jzbt1kZab_channel=TreeHouseRecovery . What do you think about the information in the video? Write down your thoughts in a few words.		Physical activity increase
24	Being more physically active (e.g. exercise, take walks, spend less time sitting) may aid you to stop smoking. One important aspect for this is to have strong resolve to become more physically active. So, before the next session, I advise you to take some time to create a personal rule that helps you to become more physically active. Possible examples include "10 squats - no matter what," "Say no to sitting, yes to life" or "Driving to the grocery store is NOT an option." Write down your rule on a piece of paper and repeat it to yourself 3 times. Put the piece of paper with your rule somewhere you can see it every day.		Physical activity increase

Abbreviations: Prereq., Prerequisite.




Figure B.9: Mean L_1 -errors based on 100 repetitions when drawing different numbers of samples from the 2,366 samples we gathered. We provide mean L_1 -errors for comparing the estimated and true Q-values for all state-action combinations (yellow) and comparing the true Q-values of the estimated and true optimal actions for all states (blue). True Q-values and optimal actions are those that are computed based on all 2,366 samples. The horizontal lines indicate percentages of the mean L_1 -error for the lowest number of samples compared to the highest number of samples for comparing the estimated and true Q-values for all state-action combinations.

C

C

Table C.1: Mean effort and number of samples for each of the 37 preparatory activities based on the 2,326 collected interaction samples.

	Preparatory activity	Effort mean (SD)	Number
1	Creating motivational slogans/quotes for quitting smoking	6.04 (2.91)	71
2	Creating motivational slogans/quotes for becoming more physically active	5.45 (3.34)	78
3	Testimonial on becoming more physically active	5.41 (3.30)	76
4	Desired future self after quitting smoking - Writing	6.06 (2.76)	70
5	Desired future self after becoming more physically active - Writing	5.25 (2.74)	72
6	Reasons for quitting smoking	6.00 (2.76)	72
7	Reasons for becoming more physically active	5.48 (2.83)	60
8	Personal rule for not smoking	5.27 (3.28)	70
9	Personal rule for becoming more physically active	5.70 (2.56)	74
10	How friends and/or family will receive one's desired future self after quitting smoking*	6.62 (2.42)	21
11	How friends and/or family will receive one's desired future self after becoming more physically active*	5.94 (2.90)	17
12	Focusing on past successes for quitting smoking	5.51 (2.73)	65
13	Focusing on past successes for becoming more physically active	5.40 (2.36)	73
14	Role model for others by quitting smoking	4.80 (2.72)	71
15	Role model for others by becoming more physically active	6.00 (2.68)	72

Table C.1: (continued)

	Preparatory activity	Effort mean (SD)	Number
16	Tracking smoking behavior	6.05 (2.52)	80
17	Tracking physical activity behavior	5.12 (2.91)	74
18	Feared future self when not quitting smoking - Writing	6.21 (2.58)	70
19	Feared future self when not becoming more physically active - Writing	5.89 (2.61)	75
20	Feared future self when not quitting smoking - Picture	5.55 (2.68)	75
21	Feared future self when not becoming more physically active - Picture	5.90 (2.67)	73
22	Visualizing smoking as a battle	5.20 (2.45)	80
23	Visualizing becoming more physically active as a battle	5.73 (2.38)	60
24	Desired future self after quitting smoking - Picture	5.86 (2.54)	77
25	Desired future self after becoming more physically active - Picture	6.34 (2.45)	71
26	Education on sleep	5.58 (3.07)	79
27	Routines that cause cravings*	6.11 (3.10)	18
28	Thinking of high-risk situations and how to cope with them*	5.58 (2.07)	12
29	Alternative behaviors for cravings	6.60 (2.29)	57
30	Progressive muscle relaxation	5.32 (3.17)	75
31	Breathing exercise	6.05 (2.39)	64
32	Exchanging a passive activity for an active one	6.13 (2.45)	82
33	Thinking of solutions to barriers to becoming physically active	5.71 (2.70)	83
34	Education on recommended physical activity*	5.67 (2.90)	12
35	Plan for becoming more physically active*	5.77 (2.89)	13
36	Positive diary	5.86 (3.04)	71
37	Focusing on past success in general	6.25 (2.68)	63
	Total	5.74 (2.75)	2326

* Activity had another activity as prerequisite.

Abbreviations: SD, Standard deviation.


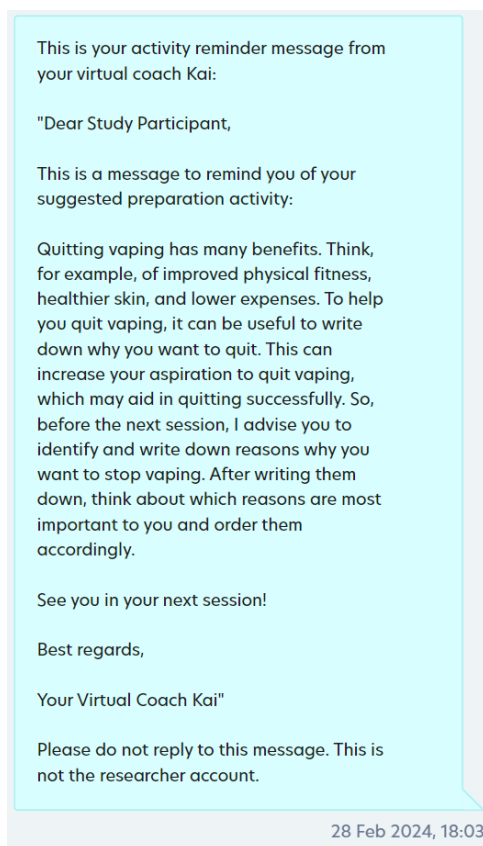


Figure C.2: Transition probabilities between the 12 states without and with human feedback. Only transitions with a probability of at least $\frac{1}{12}$ are shown. We distinguish transitions to a state with a higher or highest $R(s,0)$ (blue), lower or lowest $R(s,0)$ (red), and the same $R(s,0)$ (black), where $R(s,0)$ denotes the immediate reward when not giving feedback in state s . A thicker line denotes a higher probability.



C

Figure C.4: Screenshot of how the activity reminder messages were sent to participants on the crowdsourcing platform Prolific Academic.

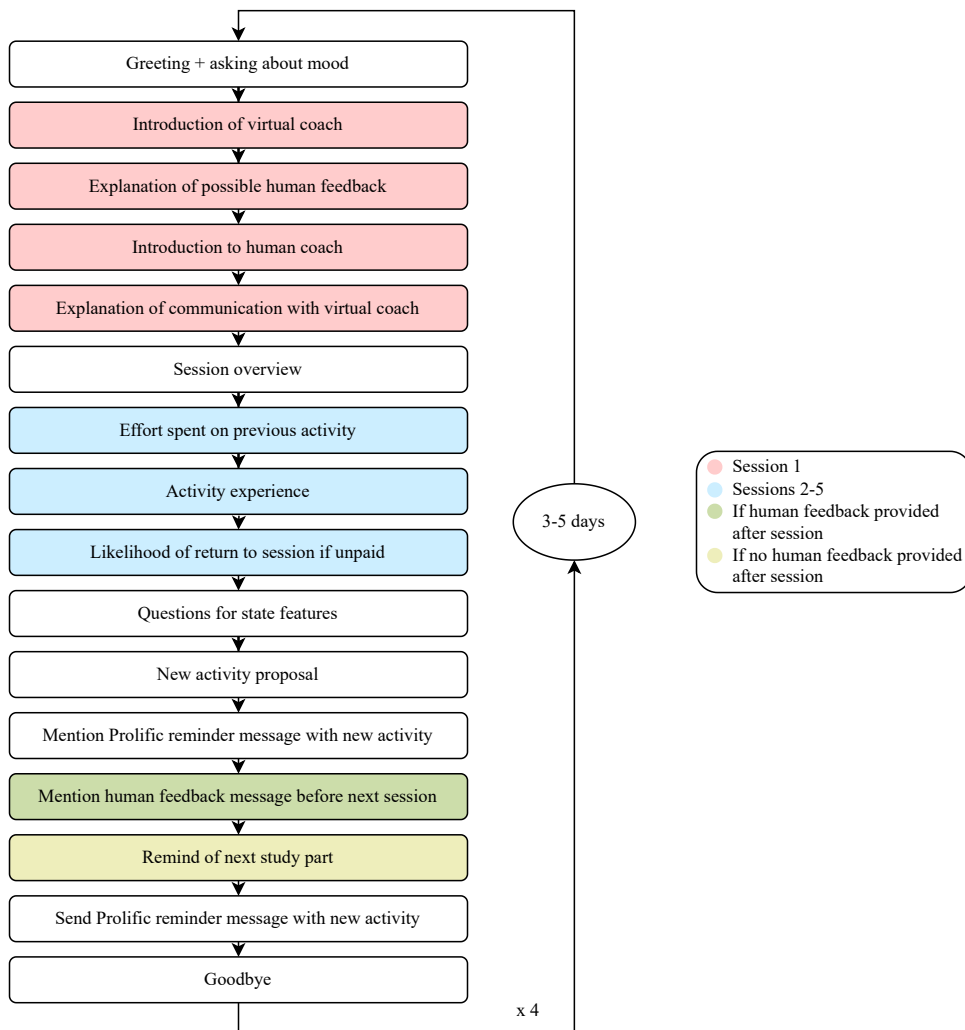


Figure C.5: Structure of the five conversational sessions with Kai. Participants received reminder messages with the formulations of their new activities on the online crowdsourcing platform Prolific Academic, which is where participants were recruited.

Table C.6: Preparatory activities for quitting smoking. Titles and formulations from the reminder messages for the 37 preparatory activities for quitting smoking that were used in the study. We also provide the prerequisite activities where applicable as well as refer to sources the activities are based on. For vapers, the formulations were adapted by, for example, replacing “smoking” with “vaping.” For an example of how the reminder messages were sent to participants, refer to Figure C.4.

	Title and formulation	Prereq.	Sources
1	Creating motivational slogans/quotes for quitting smoking. Having strong motivation to quit smoking helps to quit successfully. Before the next session, I thus suggest you take some time to look for a motivational quote or write down something that motivates you to quit smoking. Place this somewhere you can see it every day, such as your fridge or closet door.		Michie et al. [241]
2	Creating motivational slogans/quotes for becoming more physically active. Having strong motivation to become more physically active helps to succeed. Before the next session, I thus suggest you take some time to look for a motivational quote or write down something that motivates you to become more physically active. Place this somewhere you can see it every day, such as your fridge or closet door.		Michie et al. [241]
3	Testimonial on becoming more physically active. When preparing for becoming more physically active, it can be useful to learn from other people who have succeeded in becoming more physically active. What goal did they set for themselves? And how did they reach it? Before the next session, I thus recommend you watch this short video in which 5 people describe how they reached their physical activity goals: https://youtu.be/m1MHo9fCTG8 . What can you take away from the 5 examples for yourself? Take a few notes on a piece of paper or your phone.		Albers et al. [22], Hizli et al. [163]

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
4	Desired future self after quitting smoking - Writing. Having high aspiration to quit smoking may aid in quitting successfully. Thus, before the next session, I advise you to think about the person that you would like to be once you have successfully quit smoking. For example, you might want to be a "grandfather who can play football with his grandchildren" or a "nurse who can walk up the stairs to the fourth floor without getting out of breath." Write down everything that comes to your mind.		Albers et al. [21], Meijer et al. [236], Penformis et al. [295]
5	Desired future self after becoming more physically active - Writing. Quitting smoking may be easier if you become more physically active (e.g., take walks, swim, or go running). One important step for this is to have a high ambition to become more physically active. Thus, before the next session, I advise you to think about the person that you would like to be once you have become more physically active. For example, you might want to be a "grandfather who can play football with his grandchildren" or a "nurse who can walk up the stairs to the fourth floor without getting out of breath." Write down everything that comes to your mind.		Albers et al. [21], Meijer et al. [236], Penformis et al. [295]
6	Reasons for quitting smoking. Quitting smoking has many benefits. Think, for example, of improved physical fitness, healthier skin, and lower expenses. To help you quit smoking, it can be useful to write down why you want to quit. This can increase your aspiration to quit smoking, which may aid in quitting successfully. So, before the next session, I advise you to identify and write down reasons why you want to stop smoking. After writing them down, think about which reasons are most important to you and order them accordingly.		Albers et al. [21], Michie et al. [243], Trimbos Instituut [369]

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
7	Reasons for becoming more physically active. Quitting smoking may be easier if you become more physically active (e.g., take walks, swim, or go running). One crucial step for this is to have a strong desire to become more physically active. Therefore, before the next session, I advise you to identify and write down reasons why you want to become more physically active. After writing them down, think about which reasons are most important to you and order them accordingly.		Albers et al. [21], Michie et al. [243]
8	Personal rule for not smoking. Having strong determination to refrain from smoking may help to quit successfully. So, before the next session, I advise you to take some time to create a personal rule that helps you to refrain from smoking. Possible examples include "Not a puff - no matter what," "Say no to smoking, yes to life" or "Smoking is NOT an option." Write down your rule on a piece of paper and repeat it to yourself 3 times. Put the piece of paper with your rule somewhere you can see it every day.		Albers et al. [21], Michie et al. [243]
9	Personal rule for becoming more physically active. Being more physically active (e.g., taking walks, swimming, or going running) may aid you to stop smoking. One important aspect for this is to have strong resolve to become more physically active. So, before the next session, I advise you to take some time to create a personal rule that helps you to become more physically active. Possible examples include "10 squats - no matter what," "Say no to sitting, yes to life" or "Driving to the grocery store is NOT an option." Write down your rule on a piece of paper and repeat it to yourself 3 times. Put the piece of paper with your rule somewhere you can see it every day.		Albers et al. [21], Michie et al. [243]

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
10	<p>How friends and/or family will receive one's desired future self after quitting smoking. People are social creatures. This means that what we do is noticed by others around us, and they can react to it. To help you quit smoking, it can be useful to imagine how people who are important to you will receive your non-smoker future self. This can boost your confidence. The reactions of others, such as friends, colleagues, and family, to your future self can be positive, but they could also be negative. It is good to be prepared for both possibilities. Before the next session, I suggest you grab a pen and paper and answer these 4 questions: 1) How would the people who are important to you react to the new you, who has quit smoking? 2) How would you feel about their reactions? 3) How would you react if the people who are important to you react positively to the new you? 4) How would you react if the people who are important to you react negatively to the new you?</p>	4 or 24	Mercken et al. [239], Perfect Fit project
11	<p>How friends and/or family will receive one's desired future self after becoming more physically active. As social beings, our actions are observed by those around us, who may react in various ways. To boost your confidence when preparing for becoming more physically active, it can be beneficial to envision how those who are important to you will respond to your future physically active self. This can be your colleagues, friends, family, or neighbors, for example. While their reactions may be positive, they could also be negative. So it is wise to anticipate and prepare for both possible outcomes. Before the next session, I suggest you grab a pen and paper and answer these 4 questions: 1) How would the people who are important to you react to the new you, who has become more physically active? 2) How would you feel about their reactions? 3) How would you react if the people who are important to you react positively to the new you? 4) How would you react if the people who are important to you react negatively to the new you?</p>	5 or 25	Mercken et al. [239], Perfect Fit project

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
12	<p>Focusing on past successes for quitting smoking. To increase your confidence that you will succeed in quitting smoking, it can help to think back to previous successes. Have you ever quit smoking before or reduced the number of times you smoked? Maybe you once only had one instead of two cigarettes after dinner. Every success in quitting smoking counts! Before the next session, I suggest you take some time to think about your previous successes and make a list on a piece of paper. Take a few minutes to make your list before moving on. Then take a closer look at your list. Try to think about what strengths you have that helped you to achieve these things. Write these strengths down on your list so you do not forget them. You can also hang or place your list somewhere in your home so that you are reminded of your successes and strengths more often. The list shows that you can be proud of yourself.</p>		<p>Michie et al. [241], Perfect Fit project, Stichting Stop Bewust [352]</p>
13	<p>Focusing on past successes for becoming more physically active. Reflecting on past successes can boost your confidence in successfully becoming more physically active. Before the next session, I suggest you think about when you succeeded in being more physically active. Have you ever taken the bike instead of the car, or taken the stairs instead of the escalator? Every small victory counts! Take a few minutes to make your list before moving on. Then take a closer look at your list. Try to think about what you did that helped you to achieve these things. Write it down on your list so you do not forget it. You can also hang or place your list somewhere in your home so you are reminded of your successes more often. These successes show that you can be proud of yourself.</p>		<p>Michie et al. [241], Perfect Fit project, Stichting Stop Bewust [352]</p>

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
14	Role model for others by quitting smoking. Many people want to quit smoking for other people, such as children or friends. Thinking about how quitting smoking makes you a role model by doing something good for others can motivate you during difficult moments in your quitting journey. I, therefore, recommend you think about how quitting smoking makes you a role model for others before the next session. Grab a pen and a piece of paper and write down your thoughts. Or maybe a picture can help you capture your thoughts.		Perfect Fit project, Trim-bos Instituut [369]
15	Role model for others by becoming more physically active. Many people want to become more physically active for other people, such as children or friends. Thinking about how becoming more physically active makes you a role model by doing something good for others can motivate you during difficult moments in your behavior change journey. I, therefore, recommend you think about how becoming more physically active makes you a role model for others before the next session. Grab a pen and a piece of paper and write down your thoughts. Or maybe a picture can help you capture your thoughts.		Perfect Fit project, Trim-bos Instituut [369]
16	Tracking smoking behavior. Preparing for situations in which you commonly smoke may make it easier to successfully quit smoking. Therefore, I recommend that you record the situations in which you smoke before the next session. Take note of one or two keywords to describe the situation and the number of cigarettes that you smoked. For example, you might note "Lunch break, 2 cigarettes" or "TV, 5 cigarettes." It might be helpful to take these notes on your phone, or you could carry a small piece of paper and pen in your pocket.		Albers et al. [21], Michie et al. [243]

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
17	Tracking physical activity behavior. Becoming more physically active (e.g., taking walks, swimming, or boxing) may make it easier to successfully quit smoking. One important step for becoming more physically active is to know one's current level. This allows one to later set a precise goal and hence to feel more motivated. So, I recommend that you record your current behavior with regard to physical activity before the next session. Try to keep track of how much time you spend 1) sitting, 2) working out and 3) being moderately active (e.g., taking a walk, biking to the grocery store). For this, it might be helpful to keep a piece of paper and pen on your kitchen table, or maybe you have a smartwatch that can record these types of behavior for you.		Albers et al. [21], Michie et al. [243]
18	Fearful future self when not quitting smoking - Writing. Having high motivation to quit smoking may aid in quitting successfully. Thus, before the next session, I advise you to think about who you do NOT want to be in the future but might become if you continue to smoke. For example, you might NOT want to be a "mother who dies early of coronary heart disease as her mother did," a "husband who is frowned upon by his wife" or a "man who is dependent on a substance." Write down everything that comes to your mind.		Albers et al. [21], Michie et al. [243]
19	Fearful future self when not becoming more physically active - Writing. It may be easier to successfully quit smoking if you become more physically active (e.g., swim, take walks, or dance). One crucial step for this is to have high determination to become more physically active. Therefore, before the next session, I advise you to think about who you do NOT want to be in the future but might become if you fail to become more physically active. For example, you might NOT want to be a "mother who dies early of coronary heart disease as her father did," a "daughter who is frowned upon by her mother" or a "man who is dependent on his wife in his everyday life." Write down everything that comes to your mind.		Albers et al. [21], Meijer et al. [236], Penfornis et al. [295]

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
20	Feared future self when not quitting smoking - Picture. Having high motivation to quit smoking may help to quit successfully. So, before the next session, I advise you to think about who you do NOT want to be in the future but might become if you continue to smoke. For example, you might NOT want to be a "mother who dies early of coronary heart disease as her mother did," a "husband who is frowned upon by his wife" or a "man who is dependent on a substance." Then, look for or take a picture that best captures your feared future self. Save or print this picture so that you can see it every day.		Albers et al. [21], Meijer et al. [236], Penformis et al. [295]
21	Feared future self when not becoming more physically active - Picture. It may be easier to successfully quit smoking if you become more physically active (e.g. exercise, take walks, sit less). One crucial step for this is to have high determination to become more physically active. Therefore, before the next session, I advise you to think about who you do NOT want to be in the future but might become if you fail to become more physically active. For example, you might NOT want to be a "mother who dies early of coronary heart disease as her father did," a "daughter who is frowned upon by her mother" or a "man who is dependent on his wife in his everyday life." Then, look for or take a picture that best captures your feared future self. Save or print this picture so that you can see it every day.		Albers et al. [21], Meijer et al. [236], Penformis et al. [295]
22	Visualizing smoking as a battle. Focusing on your goal of successfully quitting smoking may help you to quit. Thus, before the next session, I advise you to take some time to visualize smoking as a battle. For example, you might see yourself and a cigarette as two boxers in a fighting match. Then imagine yourself winning this battle. Visualize clearly how you win and what it feels like to be the winner. Write down a few words about your winning experience.		Albers et al. [21], Michie et al. [243]

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
23	Visualizing becoming more physically active as a battle. Becoming more physically active (e.g., swimming, taking walks, or dancing) may help you to successfully quit smoking. One important step for this is to focus on the goal of becoming more physically active. Thus, before the next session, I advise you to take some time to visualize becoming more physically active as a battle. For example, you might see yourself and a non-active version of yourself as two boxers in a fighting match. Then imagine yourself winning this battle. Visualize clearly how you win and what it feels like to be the winner. Write down a few words about your winning experience.		Albers et al. [21], Michie et al. [243]
24	Desired future self after quitting smoking - Picture. Having high aspiration to quit smoking may aid in quitting successfully. Thus, before the next session, I advise you to think about the person that you would like to be once you have successfully quit smoking. For example, you might want to be a "strong woman who lives a healthy life" or a "father who is a good role model for his children." Then look for or take a picture that best captures your desired future self. Save or print this picture so that you can see it every day.		Albers et al. [21], Meijer et al. [236], Penformis et al. [295]
25	Desired future self after becoming more physically active - Picture. Quitting smoking may be easier if you become more physically active (e.g., take walks, dance, or swim). One crucial step for this is to have high motivation to become more physically active. Thus, before the next session, I advise you to think about the person that you would like to be once you have become more physically active. For example, you might want to be a "grandfather who can play football with his grandchildren" or a "nurse who can walk up the stairs to the fourth floor without getting out of breath." Then look for or take a picture that best captures your desired future self. Save or print this picture so that you can see it every day.		Albers et al. [21], Meijer et al. [236], Penformis et al. [295]

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
26	Education on sleep. If you do not smoke, you sometimes sleep worse. This can make it more difficult to remain quit. You can sleep worse, for example, if you have the same caffeine intake (e.g., coffee, tea, energy drinks, chocolate) as before quitting because caffeine is metabolized less quickly once you quit. Before the next session, I thus recommend you watch this short video for a few tips for better sleep: https://www.youtube.com/watch?v=nysjq8VIwI8Zab_channel=EveryMindMatters . How do you plan to use these tips to improve your sleep after quitting smoking? Take a few notes on your phone or a piece of paper.		Perfect Fit project, Trim-bos Instituut [369]
27	Routines that cause cravings. Getting fewer cravings to smoke may make it easier to successfully quit smoking. Therefore, before the next session, I advise you to think about routines in your daily life that often cause you to get cravings to smoke. For example, you might have experienced that if you go to bed very late and thus sleep less, you smoke more the next day. Or maybe you have noticed that if you skip your breakfast, you always smoke on your way to work but NOT otherwise. How could you change these routines to reduce or even avoid those cravings? Write down everything that comes to your mind.	16	Albers et al. [21], Michie et al. [243]
28	Thinking of high-risk situations and how to cope with them. Preparing for situations in which avoiding smoking is difficult may make it easier to successfully quit smoking. Thus, before the next session, I advise you to think about situations in which you might find it difficult to refrain from smoking. For example, this could be during your lunch break at work, when you meet your best friend, or when you watch TV. How could you deal with these situations so that you do NOT smoke? Write down your plans in a few words.	16	Albers et al. [21], Michie et al. [243]

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
29	Alternative behaviors for cravings. Planning how to resist urges to smoke may make it easier to successfully quit smoking. Therefore, before the next session, I advise you to think about activities that you could do to keep yourself busy when you feel the urge to smoke so that you do NOT smoke. These urges typically last a few minutes; think about something that you could do in the meantime until the urge has passed. For example, you could water your plants, eat a carrot, do 10 push-ups, or do something for another person in need. Write down everything that comes to your mind.		Albers et al. [21], Michie et al. [243]
30	Progressive muscle relaxation. Tensing and relaxing areas of the body can reduce cravings and withdrawal symptoms because it is very difficult to feel tense or uptight in a relaxed body. Thus, before the next session, I advise you to watch the following 15-minute video to learn progressive muscle relaxation (which is a way of relaxing your body): https://www.youtube.com/watch?v=ih002wUzgkC&channel=MarkConnelly . Even if you have already heard of this technique, it might be a good idea to refresh your memory.		Albers et al. [21], Michie et al. [243]
31	Breathing exercise. When you quit smoking, you may feel restless or irritable during the first days or weeks. To help you quit and stay quit, it can help to learn how to manage these nicotine withdrawal symptoms. One way to manage them is through breathing exercises. So, before the next session, I suggest you to watch this 3-minute video to learn how to do box breathing: https://www.youtube.com/watch?v=tEmt1Znux58&channel=SunnybrookHospital . Even if you have already heard of this technique, it might be a good idea to refresh your memory.		National Cancer Institute [265]

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
32	<p>Exchanging a passive activity for an active one. Becoming more physically active (e.g., taking walks, going running, swimming) may help you to successfully quit smoking. One crucial part for this is to think about ways you can incorporate physical activity into your daily life. One way to do this is to exchange a passive activity for an active one. Therefore, before the next session, I advise you to think about ways you could exchange a passive activity for an active one. For example, you could take the stairs instead of the escalator, bike to work instead of taking the bus, or work at a standing desk. Grab a piece of paper or your phone and write down everything that comes to your mind. Which exchange do you want to focus on? Highlight this exchange.</p>		Voedingscentrum [380]
33	<p>Thinking of solutions to barriers to becoming physically active. Becoming more physically active (e.g., taking walks, boxing, dancing) may help you to successfully quit smoking. One important step for becoming more physically active is to remove possible obstacles. Thus, before the next session, I advise you to think about things that make it difficult for you to be physically active. For example, this could be that you do NOT have a raincoat to bike to the grocery store when it is raining, that you do NOT want to work out alone, or that you are at work all day and too exhausted by the time that you come home. What are possible solutions to your barriers? For instance, you could buy a raincoat, join a running group, or take a walk during your lunch break at work. Write down everything that comes to your mind.</p>		Albers et al. [21], Alfaifi et al. [31]

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
34	<p>Education on recommended physical activity. Quitting smoking may be easier if you become more physically active (e.g., swim, take walks, go running). One important step for becoming more physically active is to set a specific goal and thus to feel more aspiration. Therefore, before the next session, I advise you to watch the following 2-minute video on how much and which type of physical activity is recommended: https://www.youtube.com/watch?v=AAPhWbG_zLsZab_channel=TREKGroup. Then, compare your physical activity behavior to the recommended amounts for the different types of physical activity. Write down which recommended amounts you meet or exceed, and which ones you do NOT meet.</p>	17	Albers et al. [21]
35	<p>Plan for becoming more physically active. Becoming more physically active (e.g., taking walks, dancing, swimming) may help you to successfully quit smoking. One crucial part for this is to create a plan for becoming more physically active. Therefore, before the next session, I advise you to think about what you could do to become more physically active. For example, you could get up from your desk after every 30 minutes of sitting, bike to the grocery store, do 10 squats every morning, or join a running group. Write down everything that comes to your mind. Which plan do you want to focus on? Highlight this plan.</p>	17	Albers et al. [21]

Table C.6: (continued)

	Title and formulation	Prereq.	Sources
36	Positive diary. This activity is called "Positive Diary" and helps you think positively and feel good. This can help you quit smoking and become more physically active. In the evening before going to bed, think about the day you had. Write down 2 or 3 things that happened that you are grateful for, happy about, or that went well. For example, "Someone smiled at me in the supermarket," "I did not smoke today," or "I took a nice walk with a friend." Writing down these positive moments can help you feel better, about yourself and about your day. You can write down anything! It does not have to be about being more physically active or quitting smoking, but can be any enjoyable moment from the day, big or small. This will help you to focus on positive things.		Perfect Fit project, Sutton [356]
37	Focusing on past success in general. To increase your confidence that you will succeed in changing your behavior (e.g., quitting smoking), it can help to think back to previous successes. Before the next session, I suggest you take a moment to reflect on times when you succeeded in something or felt satisfied. Have you ever won a competition? Did you cook something delicious recently? Or maybe you learned a new language? Grab a pen and paper and make a list of your success moments. No success is too small to write down! Take a few minutes to make your list before moving on. Then take a closer look at your list. Try to think about what you did that helped you to achieve your successes. Write it down on your list so you do not forget it. You can also hang or place your list somewhere in your home so that you are reminded of your successes more often. The list shows that you can be proud of yourself.		Michie et al. [241], Perfect Fit project, Stichting Stop Bewust [352]

Abbreviations: Prereq., Prerequisite.

Table C.7: Allocation principles in our post-questionnaire and corresponding principles by Persad et al. [297] together with the mean weights assigned to the principles by participants. For each principle, we also quote the allocation preferences initially expressed in free text by a participant who *afterward* assigned a relatively high weight to the principle. This weight afterward assigned to the principle is given in parentheses. Note that since participants expressed their initial allocation preferences in free text *before* they saw and weighted the principles, they might also have assigned a high weight to a principle they had not initially considered. As such, there is not always a clear match between the example quotes and the principles. We provide direct, uncorrected quotes.

	Allocation principle	Principle by Persad et al. [297]	Weight	Example quote of initial preferences
TREATING PEOPLE EQUALLY				
1	Random	Lottery	9.69%	"Randomly to make it fair I suppose" (100%)
2	Longest time since last human feedback	First-come, first-served	6.45%	"Randomness" (25%)
3	Least amount of human feedback so far	/	6.04%	"Help those first who need the most help but also be equal like give feedback at least once to each person" (20%)
FAVORING THE WORST-OFF: PRIORITARIANISM				
4	Least likely to successfully prepare for quitting [smoking/vaping] without human feedback	Sickest first	13.51%	"When the person seems to be struggling with the program" (70%)
5	Most likely to experience negative consequences of [smoking/vaping] in the future without human feedback	Sickest first	11.82%	"help those who seem struggle more or it is more urgent to quit as soon as possible" (30%)
6	Youngest first	Youngest first	5.31%	"The most urgent cases should take priority over the milder ones ..." (35%)
MAXIMIZING TOTAL BENEFITS: UTILITARIANISM				
7	Largest increase in chance of successfully preparing for quitting [smoking/vaping] because of the feedback	Prognosis	16.42%	"The people who are struggling the most should receive the feedback" (100%)

Table C.7: (continued)

	Allocation principle	Principle by Persad et al. [297]	Weight	Example quote of initial preferences
8	Largest reduction in negative consequences of [smoking/vaping] in the future because of the feedback	Prognosis	14.40%	"I think the virtual coach can use an algorithm to decide who is struggling and in the most need of the human coaches time" (40%)
PROMOTING AND REWARDING SOCIAL USEFULNESS				
9	Largest value to society in the future (e.g., health-care staff, workers producing influenza vaccine, people who agree to improve their health and thus use fewer resources in the future)	Instrumental value	3.83%	"how much they want to quit smoking and how much effort they put into interacting with kia" (30%)
10	Past usefulness or sacrifice (e.g., past organ donors, people who participated in vaccine research, people who made healthy lifestyle choices that reduced their need for resources in the past)	Reciprocity	3.90%	"virtual coach should prioritize individuals who are making significant progress or facing specific challenges in their journey to quit smoking" (25%)
RESPECTING AUTONOMY				
11	Highest appreciation of human feedback	/	8.62%	"People who indicated that they wanted human feedback after a session with Kai" (100%)

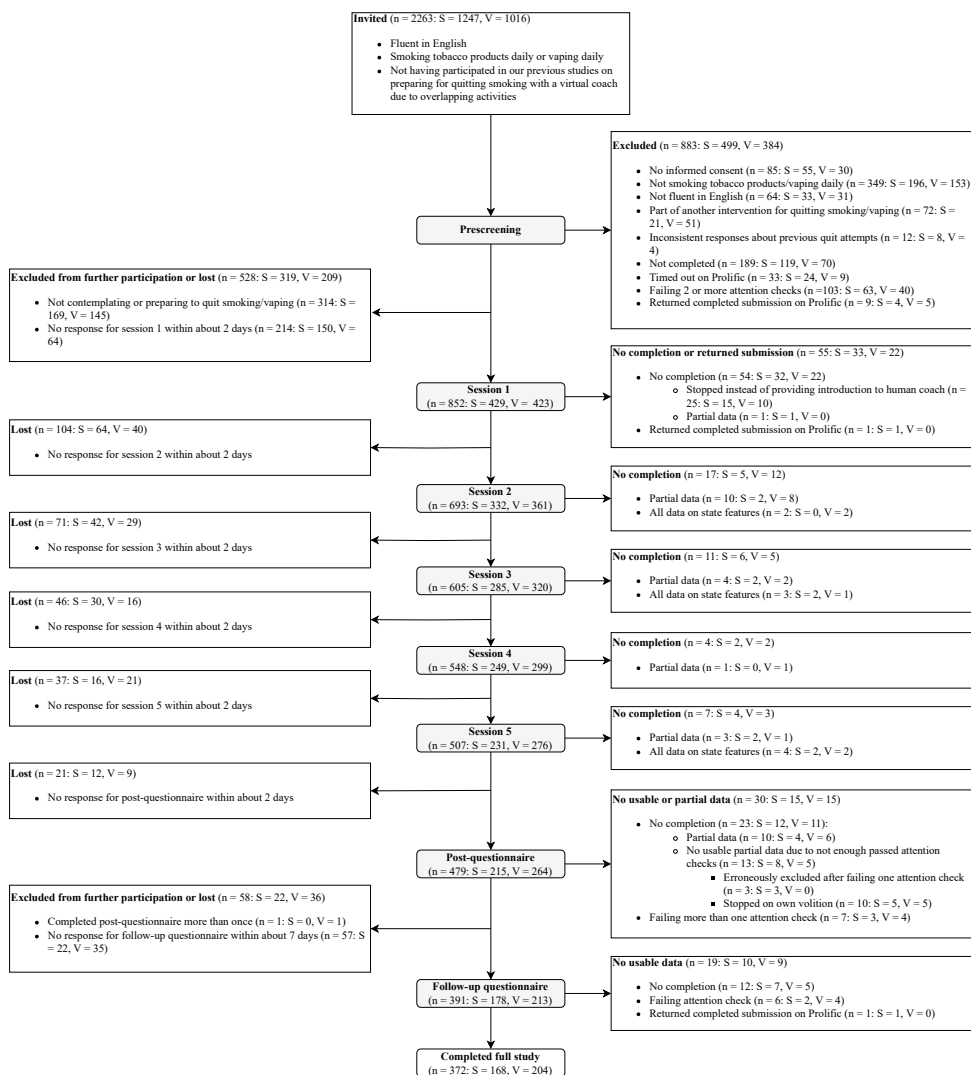


Figure C.8: Participant flow. Participants on Prolific can return their submissions. Data from the follow-up questionnaire is not included in this thesis.

Table C.9: Characteristics of the 679 participants with at least one interaction sample. People participated in the study as either smokers or vapers.

Characteristic	Value
AGE (IN YEARS)	
- Mean (SD)	36.30 (11.21)
- Range	19 – 71
GENDER	
- Female, n (%)	330 (48.60%)
- Male, n (%)	335 (49.34%)
- Other, n (%)	14 (2.06%)
HIGHEST COMPLETED EDUCATION LEVEL	
- No formal qualifications, n (%)	5 (0.74%)
- Secondary education (e.g. GED/GCSE), n (%)	61 (8.98%)
- High school diploma/A-levels, n (%)	139 (20.47%)
- Technical/community college, n (%)	90 (13.25%)
- Undergraduate degree (BA/BSc/other), n (%)	263 (38.73%)
- Graduate degree (MA/MSc/MPhil/other), n (%)	107 (15.76%)
- Doctorate degree (PhD/other), n (%)	9 (1.33%)
- Don't know/not applicable, n (%)	5 (0.74%)
PARTICIPATION AS SMOKER VS. VAPER	
- Smoker, n (%)	352 (51.84%)
- Vaper, n (%)	327 (48.16%)
SMOKING/VAPING FREQUENCY	
- Once a day, n (%)	30 (4.42%)
- 2 – 5 times a day, n (%)	109 (16.05%)
- 6 – 10 times a day, n (%)	136 (20.03%)
- 11 – 19 times a day, n (%)	160 (23.56%)
- More than 20 times a day, n (%)	244 (35.94%)
TTM-STAGE FOR BECOMING PHYSICALLY ACTIVE	
- Precontemplation, n (%)	24 (3.53%)
- Contemplation, n (%)	182 (26.80%)
- Preparation, n (%)	146 (21.50%)
- Action, n (%)	98 (14.43%)
- Maintenance, n (%)	228 (33.58%)
WEEKLY EXERCISE AMOUNT	
- Never (0 – 60 minutes per week), n (%)	190 (27.98%)
- Sometimes (60 – 150 minutes per week), n (%)	301 (44.33%)

Table C.9: (continued)

Characteristic	Value
- Often (more than 150 minutes per week), n (%)	187 (27.54%)

Abbreviations: SD, Standard deviation; GED, General educational development; GCSE, General certificate of secondary education; BA, Bachelor of Arts; BSc, Bachelor of Science; MA, Master of Arts; MSc, Master of Science; MPhil, Master of Philosophy; PhD, Doctor of Philosophy; TTM, Transtheoretical model.

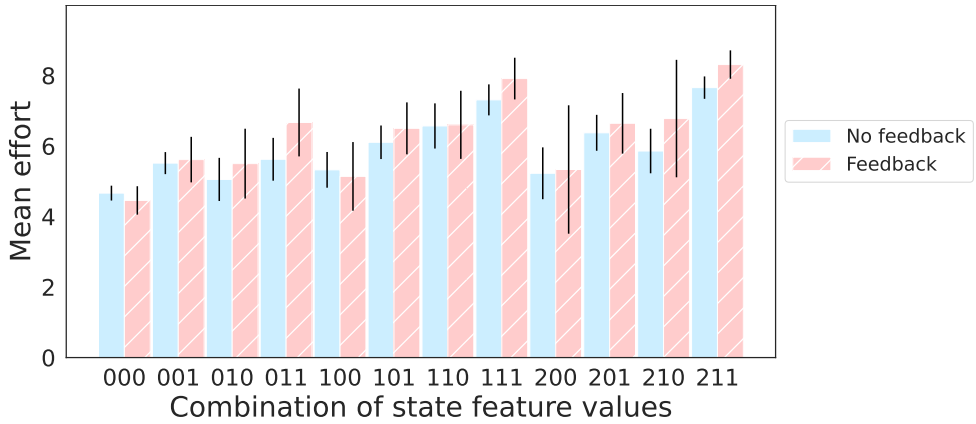


Figure C.10: Mean effort with 95% credible intervals per action and combination of values for the three selected state features. We denote the values of the three features using this order: 1) perceived importance, 2) self-efficacy, and 3) human feedback appreciation.




Figure C.11: Number of samples per action and combination of values for the three selected state features. We denote the values of the three features using this order: 1) perceived importance, 2) self-efficacy, and 3) human feedback appreciation. When we have less than 10 samples (dashed horizontal line) for a feature value combination and action, we impute with the mean effort spent on preparatory activities for the effort prediction in our RL-based analyses.

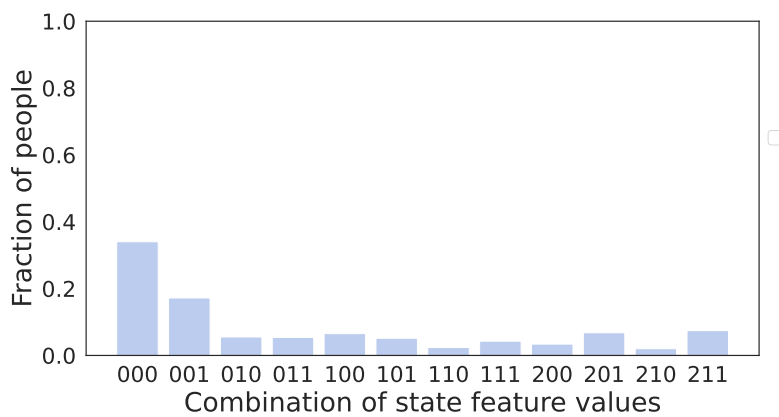


Figure C.12: Initial distribution of people across the 12 base states observed in the first session of our longitudinal study. We denote the values of the three features using this order: 1) perceived importance, 2) self-efficacy, and 3) human feedback appreciation.




Figure C.13: Mean reward (a) and percentage of people receiving feedback (b) per time step for the optimal policy and different costs for providing human feedback in our potential live application.

SUMMARY FOR LAY AUDIENCE

Motivated by previous work indicating that including human support can make people feel more accountable and satisfied with an eHealth intervention, we examined in this work how effective it is to add support from a human coach to a virtual coach-based intervention in which people do preparatory activities (e.g., envision desired future self, track smoking behavior, learn a breathing exercise) to prepare for successful smoking cessation. Adding extensive human support would undermine one of the goals of eHealth interventions, namely, the reliance on scarce and expensive healthcare professionals. We thus focused on relatively low-cost human support, in our case feedback messages written by Master's students in Psychology. Using data from a crowdsourcing study in which 679 daily smokers and vapers interacted with a text-based virtual coach in up to five sessions spread over about two weeks, we analyzed psychological, economic, and ethical factors that play a role when allocating this support. We find that having received human feedback increases both the retention and the engagement with activities assigned by the virtual coach. However, there are also situations where not giving feedback is better in the long run. When providing human feedback is expensive, it is most effective in the long run to give this feedback to people who perceive preparing to quit as not that important and at the same time have high self-efficacy for preparing to quit. Notably, while it may seem intuitive to allocate limited human feedback to those who would benefit most from it, it is, in fact, only one of several ethical principles for allocating scarce medical resources. Alternatives include allocating feedback to those who want it most or to individuals who have the lowest likelihood of successfully quitting without such support. Our analysis of a "standard" model for allocating human feedback that optimizes the overall population-level benefit over time shows that such a model prioritizes people who want feedback and are already doing well. While we show that we can include other allocation principles in the model to favor other smoker subgroups, we find that these principles are often interdependent. For example, since we observed that people who are not doing well tend to not want human feedback and benefit less from it, giving more feedback to them means doing worse in respecting people's autonomy and maximizing the overall benefit. As different smoker subgroups benefit depending on the chosen allocation principles, our findings show that moral decisions are unavoidable when human support is limited.

BIBLIOGRAPHY

- [1] Rob Aalbers and Anne-Fleur Roos. Zorguitgaven, ons een zorg? <https://www.cpb.nl/sites/default/files/omnidownload/CPB-Publicatie-Zorguitgaven-ons-een-zorg.pdf>, 2022.
- [2] Amal Abdulrahman, Deborah Richards, and Ayse Aysin Bilgin. Reason explanation for encouraging behaviour change intention. In Frank Dignum, Alessio Lomuscio, Ulle Endriss, and Ann Nowé, editors, *AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021*, pages 68–77. ACM, 2021. doi: 10.5555/3463952.3463967.
- [3] Inger L Abma, Gert J Olthuis, and Anke JM Oerlemans. Comparing the dutch icu triage guidelines to guidelines from abroad: differences and similarities. *Netherlands Journal of Critical Care*, 29(6):238–45, 2021.
- [4] Inger L Abma, Gert J Olthuis, Irma THM Maassen, Marjan L Knippenberg, Miriam Moviat, Annie J Hasker, AG Buenen, Bernard G Fikkers, and Anke JM Oerlemans. Putting icu triage guidelines into practice: A simulation study using observations and interviews. *PLoS ONE*, 18(8):e0286978, 2023. doi: 10.1371/journal.pone.0286978.
- [5] Marc A Adams, James F Sallis, Gregory J Norman, Melbourne F Hovell, Eric B Hekler, and Elyse Perata. An adaptive physical activity intervention for overweight adults: a randomized controlled trial. *PLoS ONE*, 8(12):e82901, 2013. doi: 10.1371/journal.pone.0082901.
- [6] Icek Ajzen. The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2):179–211, 1991. doi: 10.1016/0749-5978(91)90020-T.
- [7] Ulrika Akenine, Mariagnese Barbera, Cathrien RL Beishuizen, Mandana Fallah Pour, Juliette Guillemont, Anna Rosenberg, Nicola Coley, Francesca Mangialasche, Lotta Salo, Stephanie Savy, A Jeannette Pols, Sandrine Andrieu, Edo Richard, Eric Moll van Charante, and Miia Kivipelto. Attitudes of at-risk older adults about prevention of cardiovascular disease and dementia using ehealth: a qualitative study in a european context. *BMJ Open*, 10:e037050, 2020. doi: 10.1136/bmjopen-2020-037050.
- [8] Isra Al-Dhahir, Thomas Reijnders, Jasper S Faber, Rita J van den Berg-Emons, Veronica R Janssen, Roderik A Kraaijenhagen, Valentijn T Visch, Niels H Chavannes, and Andrea WM Evers. The barriers and facilitators of ehealth-based lifestyle intervention programs for people with a low socioeconomic status: scoping review. *Journal of Medical Internet Research*, 24(8):e34229, 2022. doi: 10.2196/34229.

- [9] Nele Albers. Reinforcement Learning-based Persuasion for a Conversational Agent to Support Behavior Change: Code. <https://doi.org/10.5281/zenodo.6319356>, 2022. URL https://github.com/PerfectFit-project/virtual_coach_rl_persuasion_algorithm.
- [10] Nele Albers. Virtual Coach Mel for Proposing Useful Preparatory Activities for Quitting Smoking. <https://doi.org/10.5281/zenodo.8302492>, 2023. URL https://github.com/PerfectFit-project/virtual_coach_useful_activities.
- [11] Nele Albers. Example conversational session with the virtual coach mel for preparing to quit smoking. https://youtu.be/RywcZSHOs_g, 2023.
- [12] Nele Albers. Virtual Coach Kai for Preparing for Quitting Smoking with Human Support. <https://doi.org/10.5281/zenodo.11102861>, 2024. URL https://github.com/PerfectFit-project/virtual_coach_human_inv.
- [13] Nele Albers. Example dialog with the virtual coach kai for preparing to quit smoking with human support. <https://youtu.be/gLiy5AnolYE>, 2024.
- [14] Nele Albers and Willem-Paul Brinkman. Perfect fit - experiment to gather data for and test a reinforcement learning-approach for motivating people. <https://doi.org/10.17605/osf.io/k2uac>, 2021.
- [15] Nele Albers and Willem-Paul Brinkman. Perfect fit - beliefs about and competencies built by preparatory activities for quitting smoking and becoming more physically active. <https://doi.org/10.17605/osf.io/cax6f>, 2022.
- [16] Nele Albers and Willem-Paul Brinkman. Perfect fit - learning to propose useful preparatory activities for quitting smoking and becoming more physically active. <https://doi.org/10.17605/osf.io/nuy4w>, 2023.
- [17] Nele Albers and Willem-Paul Brinkman. Perfect fit - learning when to involve a human coach in an ehealth application for preparing for quitting smoking or vaping. <https://doi.org/10.17605/osf.io/78cnr>, 2024.
- [18] Nele Albers, Mark A Neerinx, and Willem-Paul Brinkman. Acceptance of a virtual coach for quitting smoking and becoming physically active: Dataset. <https://doi.org/10.4121/19934783>, 2022.
- [19] Nele Albers, Mark A Neerinx, and Willem-Paul Brinkman. Addressing people's current and future states in a reinforcement learning algorithm for persuading to quit smoking and to be physically active: Data and analysis code. <https://doi.org/10.4121/21533055>, 2022.
- [20] Nele Albers, Mark A Neerinx, and Willem-Paul Brinkman. Addressing people's current and future states in a reinforcement learning algorithm for persuading to quit smoking and to be physically active. *PLoS ONE*, 17(12):e0277295, 2022. doi: 10.1371/journal.pone.0277295.

-
- [21] Nele Albers, Mark A Neerincx, Kristell M Penforis, and Willem-Paul Brinkman. Users' needs for a digital smoking cessation application and how to address them: A mixed-methods study. *PeerJ*, 10:e13824, 2022. doi: 10.7717/peerj.13824.
- [22] Nele Albers, Beyza Hizli, Bouke L Scheltinga, Eline Meijer, and Willem-Paul Brinkman. Setting physical activity goals with a virtual coach: Vicarious experiences, personalization and acceptance. *Journal of Medical Systems*, 47:15, 2023. doi: 10.1007/s10916-022-01899-9.
- [23] Nele Albers, Mark A Neerincx, Nadyne L Aretz, Mahira Ali, Arsen Ekinci, and Willem-Paul Brinkman. Attitudes toward a virtual smoking cessation coach: Relationship and willingness to continue. In Alexander Meschtscherjakov, Cees Midden, and Jaap Ham, editors, *Persuasive Technology: 18th International Conference, PERSUASIVE 2023, Eindhoven, The Netherlands, April 19–21, 2023, Proceedings*, pages 265–274, 2023. doi: 10.1007/978-3-031-30933-5_17.
- [24] Nele Albers, Mark A Neerincx, and Willem-Paul Brinkman. Persuading to prepare for quitting smoking with a virtual coach: Using states and user characteristics to predict behavior - Data, analysis code and appendix. <https://doi.org/10.4121/22153898>, 2023.
- [25] Nele Albers, Mark A Neerincx, and Willem-Paul Brinkman. Preparing for quitting smoking and becoming more physically active with a virtual coach: Reflections for persuasive messages and action plans. <https://doi.org/10.4121/21905271>, 2023.
- [26] Nele Albers, Mark A Neerincx, and Willem-Paul Brinkman. Difficulty and time perceptions of preparatory activities for quitting smoking: Dataset. <https://doi.org/10.4121/5198F299-9C7A-40F8-8206-C18DF93EE2A0>, 2023.
- [27] Nele Albers, Mark A Neerincx, and Willem-Paul Brinkman. Persuading to prepare for quitting smoking with a virtual coach: Using states and user characteristics to predict behavior. In Noa Agmon, Bo An, Alessandro Ricci, and William Yeoh, editors, *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 - 2 June 2023*, pages 717–726. ACM, 2023. doi: 10.5555/3545946.3598704.
- [28] Nele Albers, Mark Neerincx, and Willem-Paul Brinkman. Reinforcement learning for proposing smoking cessation activities that build competencies: Combining two worldviews in a virtual coach - data, analysis code, and appendix for the phd thesis chapter. <https://doi.org/10.4121/9C4D9C35-3330-4536-AB8D-D5BB237C277D>, 2024.
- [29] Nele Albers, Mark Neerincx, and Willem-Paul Brinkman. Human feedback messages for preparing for quitting smoking: Dataset. <https://doi.org/10.4121/7E88CA88-50E9-4E8D-A049-6266315A2ECE>, 2024.

- [30] Nele Albers, Francisco Melo, Mark Neerincx, Olya Kudina, and Willem-Paul Brinkman. The impact of human feedback in a chatbot-based smoking cessation intervention: An empirical study into psychological, economic, and ethical factors - data and analysis code for the phd thesis chapter. <https://doi.org/10.4121/1d9aa8eb-9e63-4bf5-98a3-f359dbc932a4>, 2025.
- [31] Yousef Alfaifi, Floriana Grasso, and Valentina Tamma. An ontology of psychological barriers to support behaviour change. In Patty Kostkova, Floriana Grasso, Carlos Castillo, Yelena Mejova, Arnold Bosman, and Michael Edelstein, editors, *Proceedings of the 2018 International Conference on Digital Health, DH 2018, Lyon, France, April 23-26, 2018*, pages 11–15. ACM, 2018. doi: 10.1145/3194658.3194680.
- [32] Nurcan Alkış and Tuğba Taşkaya Temizel. The impact of individual differences on influence strategies. *Personality and Individual Differences*, 87:147–152, 2015. ISSN 0191-8869. doi: 10.1016/j.paid.2015.07.037.
- [33] Alice Alphonse, Kezia Stewart, Jamie Brown, and Olga Perski. Exploring users' experiences with a quick-response chatbot within a popular smoking cessation smartphone app: semistructured interview study. *JMIR Formative Research*, 6(7): e36869, 2022. doi: 10.2196/36869.
- [34] Alaa Alsilaity and Thomas Tran. On the impact of the application domain on users' susceptibility to the six weapons of influence. In Sandra Burri Gram-Hansen, Tanja Svarre Jonassen, and Cees J H Midden, editors, *Persuasive Technology. Designing for Future Change - 15th International Conference on Persuasive Technology, PERSUASIVE 2020, Aalborg, Denmark, April 20-23, 2020, Proceedings*, volume 12064 of *Lecture Notes in Computer Science*, pages 3–15, 2020. doi: 10.1007/978-3-030-45712-9_1.
- [35] Teresa Amabile and Steven Kramer. *The progress principle: Using small wins to ignite joy, engagement, and creativity at work*. Harvard Business Review Press, 2011.
- [36] American Psychological Association. Apa dictionary of psychology, 2018. URL <https://dictionary.apa.org/barrier>.
- [37] American Psychological Association. Apa dictionary of psychology, 2018. URL <https://dictionary.apa.org/state>. Accessed on 01.07.2024.
- [38] Dean F Anderson and Charles M Cychosz. Development of an exercise identity scale. *Perceptual and Motor Skills*, 78(3):747–751, 1994. doi: 10.1177/003151259407800313.
- [39] Sandra Andraszewicz, Benjamin Scheibehenne, Jörg Rieskamp, Raoul Grasman, Josine Verhagen, and Eric-Jan Wagenmakers. An introduction to bayesian hypothesis testing for management research. *Journal of Management*, 41(2):521–543, 2015. doi: 10.1177/0149206314560412.
- [40] Rikke Aune Asbjørnsen, Jobke Wentzel, Mirjam Lien Smedsrød, Jøran Hjelmsæth, Matthew M Clark, Lise Solberg Nes, and Julia EWC Van Gemert-Pijnen. Identifying persuasive design principles and behavior change techniques supporting end user values and needs in ehealth interventions for long-term weight loss maintenance:

- qualitative study. *Journal of Medical Internet Research*, 22(11):e22598, 2020. doi: 10.2196/22598.
- [41] Carolyn Ashurst, Solon Barocas, Rosie Campbell, and Deborah Raji. Disentangling the components of ethical research in machine learning. In *FACCT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency, Seoul, Republic of Korea, June 21 - 24, 2022*, pages 2057–2068. ACM, 2022. doi: 10.1145/3531146.3533781.
- [42] Turgay Ayer, Oguzhan Alagoz, Natasha K Stout, and Elizabeth S Burnside. Heterogeneity in women’s adherence and its role in optimal breast cancer screening policies. *Management Science*, 62(5):1339–1362, 2016. doi: 10.1287/mnsc.2015.2180.
- [43] Rasmus Bååth. Bayesian first aid: A package that implements bayesian alternatives to the classical *.test functions in r. In *UseR! 2014 - the International R User Conference*, 2014.
- [44] Stephen Babb, Ann Malarcher, Gillian Schauer, Kat Asman, and Ahmed Jamal. Quitting smoking among adults—united states, 2000–2015. *Morbidity and Mortality Weekly Report*, 65(52):1457–1464, 2017. doi: 10.15585/mmwr.mm6552a1.
- [45] Timothy B Baker, David H Gustafson, and Dhavan Shah. How can research keep up with ehealth? ten strategies for increasing the timeliness and usefulness of ehealth research. *Journal of Medical Internet Research*, 16(2):e36, 2014. doi: 10.2196/jmir.2925.
- [46] Albert Bandura. Self-efficacy: toward a unifying theory of behavioral change. *Psychological Review*, 84(2):191–215, 1977. doi: 10.1037/0033-295X.84.2.191.
- [47] Harald Baumeister, Lars Reichler, Marie Munzinger, and Jiayi Lin. The impact of guidance on internet-based mental health interventions—a systematic review. *Internet Interventions*, 1(4):205–215, 2014. doi: 10.1016/j.invent.2014.08.003.
- [48] Amit Baumel, Frederick Muench, Stav Edan, and John M Kane. Objective user engagement with mental health apps: systematic search and panel-based usage analysis. *Journal of Medical Internet Research*, 21(9):e14567, 2019. doi: 10.2196/14567.
- [49] Tom L Beauchamp and James F Childress. *Principles of Biomedical Ethics*. Oxford University Press, USA, 5 edition, 2001.
- [50] Xavier Bellekens, Andrew Hamilton, Preetila Seeam, Kamila Nieradzinska, Quentin Franssen, and Amar Seeam. Pervasive ehealth services a security and privacy risk awareness survey. In *2016 International Conference On Cyber Situational Awareness, Data Analytics And Assessment (CyberSA)*, pages 1–4, 2016. doi: 10.1109/CyberSA.2016.7503293.
- [51] Dror Ben-Zeev, Susan M Kaiser, and Izabela Krzos. Remote “hovering” with individuals with psychotic disorders and substance use: feasibility, engagement, and therapeutic alliance with a text-messaging mobile interventionist. *Journal of Dual Diagnosis*, 10(4):197–203, 2014. doi: 10.1080/15504263.2014.962336.

- [52] Thomas Berger, Katja Hämmerli, Nina Gubser, Gerhard Andersson, and Franz Caspar. Internet-based treatment of depression: a randomized controlled trial comparing guided with unguided self-help. *Cognitive Behaviour Therapy*, 40(4):251–266, 2011. doi: 10.1080/16506073.2011.616531.
- [53] Shlomo Berkovsky, Jill Freyne, and Harri Oinas-Kukkonen. Influencing individually: Fusing personalization and persuasion. *ACM Transactions on Interactive Intelligent Systems*, 2(2):1–8, 2012. doi: 10.1145/2209310.2209312.
- [54] Mauro Bertolotti, Valentina Carfora, and Patrizia Catellani. Different frames to reduce red meat intake: the moderating role of self-efficacy. *Health Communication*, 35(4):475–482, 2019. doi: 10.1080/10410236.2019.1567444.
- [55] David Best, Melinda Beckwith, Catherine Haslam, S Alexander Haslam, Jolanda Jetten, Emily Mawson, and Dan I Lubman. Overcoming alcohol and other drug addiction as a process of social identity transition: The social identity model of recovery (simor). *Addiction Research & Theory*, 24(2):111–123, 2016. doi: 10.3109/16066359.2015.1075980.
- [56] Robbert-Jan Beun, Willem-Paul Brinkman, Siska Fitriani, Fiemke Griffioen-Both, Corine H G Horsch, Jaap Lancee, and Sandor Spruit. Improving adherence in automated e-coaching - A case from insomnia therapy. In Alexander Meschtscherjakov, Boris E R de Ruyter, Verena Fuchsberger, Martin Murer, and Manfred Tscheligi, editors, *Persuasive Technology - 11th International Conference, PERSUASIVE 2016, Salzburg, Austria, April 5-7, 2016, Proceedings*, volume 9638 of *Lecture Notes in Computer Science*, pages 276–287, 2016. doi: 10.1007/978-3-319-31510-2_24.
- [57] Timothy W Bickmore, Lisa Caruso, Kerri Clough-Gorr, and Tim Heeren. ‘it’s just like you talk to a friend’ relational agents for older adults. *Interacting with Computers*, 17(6):711–735, 2005. doi: 10.1016/j.intcom.2005.09.002.
- [58] Lucas Blair. What video games can teach us about badges and pathways. In *Digital Badges in Education*, pages 62–70. Routledge, 2016.
- [59] Olivier A Blanson Henkemans, Paul J M van der Boog, Jasper Lindenberg, Charles A P G van der Mast, Mark A Neerincx, and Bertie J H M Zwetsloot-Schonk. An online lifestyle diary with a persuasive computer assistant providing feedback on self-management. *Technology and Health Care*, 17(3):253–267, 2009. doi: 10.3233/THC-2009-0545.
- [60] Herbert Bless, Gerd Bohner, Norbert Schwarz, and Fritz Strack. Mood and persuasion: A cognitive response analysis. *Personality and Social Psychology Bulletin*, 16(2):331–345, 1990. doi: 10.1177/0146167290162013.
- [61] Tom Bocklisch, Joey Faulkner, Nick Pawlowski, and Alan Nichol. Rasa: Open source language understanding and dialogue management. *arXiv preprint arXiv:1712.05181*, 2017.

-
- [62] Waransanang Boontarig, Wichian Chutimaskul, Vithida Chongsuphajaisiddhi, and Borworn Papisratorn. Factors influencing the thai elderly intention to use smart-phone for e-health services. In *2012 IEEE Symposium on Humanities, Science and Engineering Research*, pages 479–483, 2012. doi: 10.1109/SHUSER.2012.6268881.
- [63] Robert Bowman, Orla Cooney, Anja Thieme, Benjamin R Cowan, and Gavin Doherty. Pervasive therapy: Designing conversation-based interfaces for ecological momentary intervention. *IEEE Pervasive Computing*, 21(4):55–63, 2022. doi: 10.1109/MPRV.2022.3209068.
- [64] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101, 2006. doi: 10.1191/1478088706qp063oa.
- [65] Virginia Braun and Victoria Clarke. Thematic analysis. In *APA handbook of research methods in psychology*, volume 2, pages 57–71. American Psychological Association, 2012. doi: 10.1037/13620-004.
- [66] Linda D Breeman, Mike Keesman, Douwe E Atsma, Niels H Chavannes, Veronica Janssen, Lisette van Gemert-Pijnen, Hareld Kempers, Wessel Kraaij, Fabienne Rauwers, Thomas Reijnders, Wilma Scholte op Reimer, Jobke Wentzel, Roderik A Kraaijenhagen, and Andrea W M Evers. A multi-stakeholder approach to ehealth development: Promoting sustained healthy living among cardiovascular patients. *International Journal of Medical Informatics*, 147:104364, 2021. doi: 10.1016/j.ijmedinf.2020.104364.
- [67] Robert L Brennan and Dale J Prediger. Coefficient kappa: Some uses, misuses, and alternatives. *Educational and Psychological Measurement*, 41(3):687–699, 1981. doi: 10.1177/001316448104100307.
- [68] Willem-Paul Brinkman. Guest editorial - cognitive engineering in mental health computing. *Journal of CyberTherapy & Rehabilitation (JCR)*, 4(1):9–13, 2011.
- [69] Willem-Paul Brinkman, Reinder Haakma, and DG Bouwhuis. The theoretical foundation and validity of a component-based usability questionnaire. *Behaviour & Information Technology*, 28(2):121–137, 2009. doi: 10.1080/01449290701306510.
- [70] Susan A Brown, Alan R Dennis, and Viswanath Venkatesh. Predicting collaboration technology use: Integrating technology adoption and collaboration research. *Journal of Management Information Systems*, 27(2):9–54, 2010. doi: 10.2753/MIS0742-1222270201.
- [71] Bundesministerium für Arbeit und Soziales. Einführung und anpassungen des mindestlohns. <https://www.bmas.de/DE/Arbeit/Arbeitsrecht/Mindestlohn/Einfuehrung-und-Anpassungen-Mindestlohn/einfuehrung-und-anpassung-mindestlohn.html>, 2024. Accessed on 27.08.2024.
- [72] Franziska Burger, Mark A Neerinx, and Willem-Paul Brinkman. Natural language processing for cognitive therapy: Extracting schemas from thought records. *PLoS ONE*, 16(10):e0257832, 2021. doi: 10.1371/journal.pone.0257832.

- [73] Thomas A Burling, Gary D Marshall, and Andrea L Seidner. Smoking cessation for substance abuse inpatients. *Journal of Substance Abuse*, 3(3):269–276, 1991. doi: 10.1016/S0899-3289(10/80011-2).
- [74] Robert E Burnkrant and H Rao Unnava. Effects of self-referencing on persuasion. *Journal of Consumer Research*, 22(1):17–26, 1995. doi: 10.1086/209432.
- [75] John T Cacioppo and Richard E Petty. Effects of message repetition and position on cognitive response, recall, and persuasion. *Journal of Personality and Social Psychology*, 37(1):97–109, 1979. doi: 10.1037/0022-3514.37.1.97.
- [76] John T Cacioppo and Richard E Petty. Central and peripheral routes to persuasion: The role of message repetition. In *Psychological Processes and Advertising Effects: Theory, Research, and Applications*. Lawrence Erlbaum Associates, 1985.
- [77] John T Cacioppo, Richard E Petty, and Chuan Feng Kao. The efficient assessment of need for cognition. *Journal of Personality Assessment*, 48(3):306–307, 1984. doi: 10.1207/s15327752jpa4803_13.
- [78] Laura Callaghan, Hua-Hie Yong, Ron Borland, K Michael Cummings, Sara C Hitchman, and Geoffrey T Fong. What kind of smoking identity following quitting would elevate smokers relapse risk? *Addictive Behaviors*, 112:106654, 2020. ISSN 0306-4603. doi: 10.1016/j.addbeh.2020.106654.
- [79] Canada Health Infoway. The emerging benefits of electronic medical record use in community-based care. *Pricewaterhouse Coopers LLP*, 2013.
- [80] Valentina Carfora, Mauro Bertolotti, and Patrizia Catellani. Informational and emotional daily messages to reduce red and processed meat consumption. *Appetite*, 141:104331, 2019. doi: 10.1016/j.appet.2019.104331.
- [81] Valentina Carfora, Francesca Di Massimo, Rebecca Rastelli, Patrizia Catellani, and Marco Piastra. Dialogue management in conversational agents through psychology of persuasion and machine learning. *Multimedia Tools and Applications*, 79(47):35949–35971, 2020. doi: 10.1007/s11042-020-09178-w.
- [82] Nancy Carter, Denise Bryant-Lukosius, Alba DiCenso, Jennifer Blythe, and Alan J Neville. The use of triangulation in qualitative research. *Oncology Nursing Forum*, 41(5):545–547, 2014. doi: 10.1188/14.ONF.545-547.
- [83] Patrizia Catellani, Valentina Carfora, and Marco Piastra. Connecting social psychology and deep reinforcement learning: A probabilistic predictor on the intention to do home-based physical activity after message exposure. *Frontiers in Psychology*, 12:696770, 2021. doi: 10.3389/fpsyg.2021.696770.
- [84] CBS. Health spending per capita second highest in the eu. <https://www.cbs.nl/en-gb/news/2022/49/health-spending-per-capita-second-highest-in-the-eu>, 2022. Accessed on 11.07.2024.

-
- [85] Joseph Cesario, E Tory Higgins, and Abigail A Scholer. Regulatory fit and persuasion: Basic principles and remaining questions. *Social and Personality Psychology Compass*, 2(1):444–463, 2008. doi: 10.1111/j.1751-9004.2007.00055.x.
- [86] Michael Chaiton, Lori Diemert, Joanna E Cohen, Susan J Bondy, Peter Selby, Anne Philipneri, and Robert Schwartz. Estimating the number of quit attempts it takes to quit smoking successfully in a longitudinal cohort of smokers. *BMJ Open*, 6(6): e011045, 2016. doi: 10.1136/bmjopen-2016-011045.
- [87] Michael O Chaiton, Joanna E Cohen, Paul W McDonald, and Susan J Bondy. The heaviness of smoking index as a predictor of smoking cessation in canada. *Addictive Behaviors*, 32(5):1031–1042, 2007. doi: 10.1016/j.addbeh.2006.07.008.
- [88] David Chapman and Leslie Pack Kaelbling. Input generalization in delayed reinforcement learning: An algorithm and performance comparisons. In John Mylopoulos and Raymond Reiter, editors, *Proceedings of the 12th International Joint Conference on Artificial Intelligence. Sydney, Australia, August 24-30, 1991*, pages 726–731, San Francisco, CA, USA, 1991. Morgan Kaufmann Publishers.
- [89] Janine Chapman, Christopher J Armitage, and Paul Norman. Comparing implementation intention interventions in relation to young adults’ intake of fruit and vegetables. *Psychology & Health*, 24(3):317–332, 2009. doi: 10.1080/08870440701864538.
- [90] Richard A Chechile. *Bayesian Statistics for Experimental Scientists: A General Introduction Using Distribution-Free Methods*. MIT Press, 2020.
- [91] Nan-Chen Chen, Margaret Drouhard, Rafal Kocielnik, Jina Suh, and Cecilia R Aragon. Using machine learning to support qualitative coding in social science: Shifting the focus to ambiguity. *ACM Transactions on Interactive Intelligent Systems*, 8(2):1–20, 2018. doi: 10.1145/3185515.
- [92] Paul F Christiano, Jan Leike, Tom B Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M Wallach, Rob Fergus, S V N Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 4302–4310, Red Hook, NY, USA, 2017. Curran Associates Inc.
- [93] Robert B Cialdini. *Influence: The Psychology of Persuasion*. Harper Business, New York, USA, 2006.
- [94] Alexandra Cimpean, Catholijn M Jonker, Pieter Jules Karel Libin, and Ann Nowe. A group and individual aware framework for fair reinforcement learning. In *Proceedings of the Adaptive and Learning Agents Workshop (ALA 2024)*, 2024.
- [95] Benjamin Clément, Didier Roy, Pierre-Yves Oudeyer, and Manuel Lopes. Multi-armed bandits for intelligent tutoring systems. *Journal of Educational Data Mining*, 7(2): 20–48, 2015. doi: 10.5281/zenodo.3554667.

- [96] Jacob Cohen. *Statistical power analysis for the behavioral sciences*. Routledge, 2 edition, 1988.
- [97] Jacob Cohen. A power primer. *Psychological Bulletin*, 112(1):155–159, 1992. doi: 10.1037//0033-2909.112.1.155.
- [98] Sunny Consolvo, David W McDonald, and James A Landay. Theory-driven design strategies for technologies that support behavior change in everyday life. In Dan R Olsen Jr., Richard B Arthur, Ken Hinckley, Meredith Ringel Morris, Scott E Hudson, and Saul Greenberg, editors, *Proceedings of the 27th International Conference on Human Factors in Computing Systems, CHI 2009, Boston, MA, USA, April 4-9, 2009*, pages 405–414, New York, NY, USA, 2009. ACM. doi: 10.1145/1518701.1518766.
- [99] Jae Cooper, Ron Borland, Hua-Hie Yong, Ann McNeill, Rachael L Murray, Richard J O’Connor, and K Michael Cummings. To what extent do smokers make spontaneous quit attempts and what are the implications for smoking cessation maintenance? findings from the international tobacco control four country survey. *Nicotine & Tobacco Research*, 12(suppl_1):S51–S57, 2010. doi: 10.1093/ntr/ntq052.
- [100] Ângelo Costa, Stella Heras, Javier Palanca, Jaume Jordán, Paulo Novais, and Vicente Julian. Using argumentation schemes for a persuasive cognitive assistant system. In Natalia Criado Pacheco, Carlos Carrascosa, Nardine Osman, and Vicente Julián Inglada, editors, *Multi-Agent Systems and Agreement Technologies - 14th European Conference, EUMAS 2016, and 4th International Conference, AT 2016, Valencia, Spain, December 15-16, 2016, Revised Selected Papers*, volume 10207 of *Lecture Notes in Computer Science*, pages 538–546, 2016. doi: 10.1007/978-3-319-59294-7_43.
- [101] Emmelyn AJ Croes and Marjolijn L Antheunis. Can we be friends with mitsuku? a longitudinal study on the process of relationship formation between humans and a social chatbot. *Journal of Social and Personal Relationships*, 38(1):279–300, 2021. doi: 10.1177/026540752095946.
- [102] Junhua Dang, Kevin M King, and Michael Inzlicht. Why are self-report and behavioral measures weakly correlated? *Trends in Cognitive Sciences*, 24(4):267–269, 2020. doi: 10.1016/j.tics.2020.01.007.
- [103] Nathan Davis, Kristina Shiroma, Bo Xie, Tom Yeh, Xu Han, and Atami De Main. Designing ehealth tutorials with and for older adults. In *Information: Equity, Diversity, Inclusion, Justice, and Relevance - Proceedings of the 84th ASIS&T Annual Meeting, ASIST 2021, Salt Lake City, UT, USA, October 30 - November 2, 2021*, volume 58 of *Proceedings of the Association for Information Science and Technology*, pages 92–103, 2021. doi: 10.1002/pr2.439.
- [104] Maartje M A de Graaf, Somaya Ben Allouch, and J A G M van Dijk. What makes robots social?: A user’s perspective on characteristics for social human-robot interaction. In Adriana Tapus, Elisabeth André, Jean-Claude Martin, François Ferland, and Mehdi Ammi, editors, *Social Robotics - 7th International Conference, ICSR 2015, Paris, France, October 26-30, 2015, Proceedings*, volume 9388 of *Lecture Notes in Computer Science*, pages 184–193, 2015. doi: 10.1007/978-3-319-25554-5_19.

-
- [105] Marianne de Visser, Arnoud Boot, Gijsbert Werner, Arthur van Riel, and Mérove Gijsberts. Kiezen voor houdbare zorg. mensen, middelen en maatschappelijk draagvlak. <https://www.wrr.nl/publicaties/rapporten/2021/09/15/kiezen-voor-houdbare-zorg>, 2021.
- [106] Roelof Anne Jelle de Vries. *Theory-Based and Tailor-Made: Motivational Messages for Behavior Change Technology*. PhD thesis, University of Twente, Netherlands, 2018.
- [107] Roelof Anne Jelle de Vries, Khiet P Truong, Cristina Zaga, Jamy Li, and Vanessa Evers. A word of advice: how to tailor motivational text messages based on behavior change theory to personality and gender. *Personal and Ubiquitous Computing*, 21(4): 675–687, 2017. doi: 10.1007/S00779-017-1025-1.
- [108] Edward L Deci and Richard M Ryan. *Intrinsic motivation and self-determination in human behavior*. Perspectives in Social Psychology. Springer New York, 1985.
- [109] Laurent Degroote, Delfien Van Dyck, Ilse De Bourdeaudhuij, Annick De Paepe, and Geert Crombez. Acceptability and feasibility of the mhealth intervention ‘myday-plan’ to increase physical activity in a general adult population. *BMC Public Health*, 20, 2020. doi: 10.1186/s12889-020-09148-9.
- [110] Shuang Di, Jeremy Petch, Hertz C Gerstein, Ruoqing Zhu, and Diana Sherifali. Optimizing health coaching for patients with type 2 diabetes using machine learning: Model development and validation study. *JMIR Formative Research*, 6(9):e37838, 2022. doi: 10.2196/37838.
- [111] Francesca Di Massimo, Valentina Carfora, Patrizia Catellani, and Marco Piastra. Applying psychology of persuasion to conversational agents through reinforcement learning: an exploratory study. In Raffaella Bernardi, Roberto Navigli, and Giovanni Semeraro, editors, *Proceedings of the Sixth Italian Conference on Computational Linguistics, Bari, Italy, November 13-15, 2019*, volume 2481 of *CEUR Workshop Proceedings*, 2019.
- [112] Carlo C DiClemente, James O Prochaska, Scott K Fairhurst, Wayne F Velicer, Mary M Velasquez, and Joseph S Rossi. The process of smoking cessation: an analysis of precontemplation, contemplation, and preparation stages of change. *Journal of Consulting and Clinical Psychology*, 59(2):295–304, 1991. doi: 10.1037//0022-006x.59.2.295.
- [113] Martin Dierikx, Nele Albers, Bouke L Scheltinga, and Willem-Paul Brinkman. Collaboratively setting daily step goals with a virtual coach: Using reinforcement learning to personalize initial proposals. In Nilufar Baghaei, Raian Ali, Khin Than Win, and Kiemute Oyibo, editors, *Persuasive Technology - 19th International Conference, PERSUASIVE 2024, Wollongong, NSW, Australia, April 10-12, 2024, Proceedings*, volume 14636 of *Lecture Notes in Computer Science*, pages 100–115, 2024. doi: 10.1007/978-3-031-58226-4_9.
- [114] Arie Dijkstra. Personalization/computer-tailoring in persuasive technology: Tailoring ingredients target psychological processes. In Rita Orji, Michaela Reisinger, Marc Busch, Arie Dijkstra, Agnis Stibe, and Manfred Tscheligi, editors, *Proceedings*

- of the International Workshop on Personalization in Persuasive Technology co-located with the 11th International Conference on Persuasive Technology (PT 2016), Salzburg, Austria, April 5th, 2016*, volume 1582 of *CEUR Workshop Proceedings*, pages 6–12, 2016.
- [115] Gavin Doherty, David Coyle, and John Sharry. Engagement with online mental health interventions: an exploratory clinical study of a treatment for depression. In Joseph A Konstan, Ed H Chi, and Kristina Höök, editors, *CHI Conference on Human Factors in Computing Systems, CHI '12, Austin, TX, USA - May 05 - 10, 2012*, pages 1421–1430. ACM, 2012. doi: 10.1145/2207676.2208602.
- [116] Benjamin D Douglas, Patrick J Ewell, and Markus Brauer. Data quality in online human-subjects research: Comparisons between mturk, prolific, cloudresearch, qualtrics, and sona. *PLoS ONE*, 18(3):e0279720, 2023. doi: 10.1371/journal.pone.0279720.
- [117] Tim Draws, Alisa Rieger, Oana Inel, Ujwal Gadiraju, and Nava Tintarev. A checklist to combat cognitive biases in crowdsourcing. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 9, pages 48–59, 2021. doi: 10.1609/hcomp.v9i1.18939.
- [118] David Dubois. *Competency-Based HR Management*. Black Well Publishing, Hoboken, NJ, USA, 1998.
- [119] David Dunning. *Self-insight: Roadblocks and detours on the path to knowing thyself*. Psychology Press, 1 edition, 2012.
- [120] Kirstin Early, Jennifer Mankoff, and Stephen E Fienberg. Dynamic question ordering in online surveys. *Journal of Official Statistics*, 33(3):625–657, 2017. doi: 10.1515/jos-2017-0030.
- [121] Tracy Epton, Sinead Currie, and Christopher J Armitage. Unique effects of setting goals on behavior change: Systematic review and meta-analysis. *Journal of Consulting and Clinical Psychology*, 85(12):1182–1198, 2017. doi: 10.1037/ccp0000260.
- [122] Akram Erraqabi, Alessandro Lazaric, Michal Valko, Emma Brunskill, and Yun-En Liu. Trading off rewards and errors in multi-armed bandits. In Aarti Singh and Xiaojin (Jerry) Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA*, volume 54 of *Proceedings of Machine Learning Research*, pages 709–717, 2017.
- [123] Gunther Eysenbach. The law of attrition. *Journal of Medical Internet Research*, 7(1), 2005. doi: 10.2196/jmir.7.1.e11.
- [124] Jasper S Faber, Isra Al-Dhahir, Thomas Reijnders, Niels H Chavannes, Andrea WM Evers, Jos J Kraal, HJG van den Berg-Emons, and Valentijn T Visch. Attitudes toward health, healthcare, and ehealth of people with a low socioeconomic status: a community-based participatory approach. *Frontiers in Digital Health*, 3:690182, 2021. doi: 10.3389/fdgth.2021.690182.

-
- [125] Ahmed Fadhil and Silvia Gabrielli. Addressing challenges in promoting healthy lifestyles: the al-chatbot approach. In *Proceedings of the 11th EAI International Conference on Pervasive Computing Technologies for Healthcare, PervasiveHealth 2017, Barcelona, Spain, May 23-26, 2017*, pages 261–265, New York, NY, USA, 2017. ACM. doi: 10.1145/3154862.3154914.
- [126] Yichen E Fang, Zhixian Zhang, Ray Wang, Bolu Yang, Chen Chen, Claudia Nisa, Xin Tong, and Lijing L Yan. Effectiveness of ehealth smoking cessation interventions: systematic review and meta-analysis. *Journal of Medical Internet Research*, 25:e45111, 2023. doi: 10.2196/45111.
- [127] Esteve Fernández, Anna Schiaffino, Carme Borrell, Joan Benach, Carles Ariza, Josep Maria Ramon, Jorge Twose, Manel Nebot, and Anton Kunst. Social class, education, and smoking cessation: long-term follow-up of patients treated at a smoking cessation unit. *Nicotine & Tobacco Research*, 8(1):29–36, 2006. doi: 10.1080/14622200500264432.
- [128] Michela Ferron and Paolo Massa. Transtheoretical model for designing technologies supporting an active lifestyle. In Franca Garzotto, Massimo Zancanaro, Antonella De Angeli, and Fabio Paternò, editors, *Biannual conference of the Italian chapter of SIGCHI, CHIItaly '13, Trento, Italy - September 16 - 20, 2013*, pages 1–8. ACM, 2013. doi: 10.1145/2499149.2499158.
- [129] Siska Fitriani, Merijn Bruijnes, Fengxiang Li, and Willem-Paul Brinkman. Questionnaire items for evaluating artificial social agents-expert generated, content validated and reliability analysed. In *IVA '21: ACM International Conference on Intelligent Virtual Agents, Virtual Event, Japan, September 14-17, 2021*, pages 84–86. ACM, 2021. doi: 10.1145/3472306.3478341.
- [130] Siska Fitriani, Corine Horsch, Robbert Jan Beun, Fiemke Griffioen-Both, and Willem-Paul Brinkman. Factors affecting user’s behavioral intention and use of a mobile-phone-delivered cognitive behavioral therapy for insomnia: A small-scale utaut analysis. *Journal of Medical Systems*, 45(12):110, 2021. doi: 10.1007/s10916-021-01785-w.
- [131] Brian J Fogg. *Persuasive technology: using computers to change what we think and do*. Morgan Kaufmann, USA, 1 edition, 2002.
- [132] Evan M Forman, Stephanie G Kerrigan, Meghan L Butryn, Adrienne S Juarascio, Stephanie M Manasse, Santiago Ontañón, Diane H Dallal, Rebecca J Crochiere, and Danielle Moskow. Can the artificial intelligence technique of reinforcement learning use continuously-monitored digital data to optimize treatment for weight loss? *Journal of Behavioral Medicine*, 42(2):276–290, 2019. doi: 10.1007/s10865-018-9964-1.
- [133] Evan M Forman, Michael P Berry, Meghan L Butryn, Charlotte J Hagerman, Zhuoran Huang, Adrienne S Juarascio, Erica M LaFata, Santiago Ontañón, J Mick Tilford, and Fengqing Zhang. Using artificial intelligence to optimize delivery of weight loss treatment: protocol for an efficacy and cost-effectiveness trial. *Contemporary Clinical Trials*, 124:107029, 2023. doi: 10.1016/j.cct.2022.107029.

- [134] Fay Fransella, Richard Bell, and Don Bannister. *A manual for repertory grid technique*. John Wiley & Sons, 2 edition, 2004.
- [135] Caroline Free, R Whittaker, R Knight, T Abramsky, A Rodgers, and IG Roberts. Txt2stop: a pilot randomised controlled trial of mobile phone-based smoking cessation support. *Tobacco Control*, 18(2):88–91, 2009. doi: 10.1136/tc.2008.026146.
- [136] Marina Fridin and Mark Belokopytov. Acceptance of socially assistive humanoid robot by preschool and elementary school teachers. *Computers in Human Behavior*, 33:23–31, 2014. doi: 10.1016/j.chb.2013.12.016.
- [137] Kristel M Gallagher, John A Updegraff, Alexander J Rothman, and Linda Sims. Perceived susceptibility to breast cancer moderates the effect of gain-and loss-framed messages on use of screening mammography. *Health Psychology*, 30(2):145–152, 2011. doi: 10.1037/a0022264.
- [138] Benjamin Gardner and Amanda L Rebar. Habit formation and behavior change. In *Oxford Research Encyclopedia of Psychology*. Oxford University Press, 04 2019. doi: 10.1093/acrefore/9780190236557.013.129.
- [139] Bridgette E Garrett, Brandi N Martell, Ralph S Caraballo, and Brian A King. Socioeconomic differences in cigarette smoking among sociodemographic groups. *Preventing Chronic Disease*, 16, 2019. doi: 10.5888/pcd16.180553.
- [140] Fabio Gasparetti, Luca Maria Aiello, and Daniele Quercia. Personalized weight loss strategies by mining activity tracker data. *User Modeling and User-Adapted Interaction*, 30(3):447–476, 2020. doi: 10.1007/S11257-019-09242-7.
- [141] Mohammad Ghavamzadeh, Shie Mannor, Joelle Pineau, and Aviv Tamar. Bayesian reinforcement learning: A survey. *Foundations and Trends® in Machine Learning*, 8 (5-6):359–483, 2015. doi: 10.1561/22000000049.
- [142] Koel Ghorai and Pradeep Ray. Factors affecting the acceptance of mobile based multi-feature service for smoking cessation using utaut. *European Journal of Biomedical Informatics*, 15(1), 2019.
- [143] Marcello A Gómez-Maureira, Isabelle Kniestedt, Max Van Duijn, Carolien Rieffe, and Aske Plaat. Level design patterns that invoke curiosity-driven exploration: An empirical study across multiple conditions. *Proceedings of the ACM on Human-Computer Interaction*, 5(CHI PLAY):1–32, 2021. doi: 10.1145/3474698.
- [144] Colin Goodman and Timothy Houk. Misapplying autonomy: why patient wishes cannot settle treatment decisions. *Theoretical Medicine and Bioethics*, 43(5-6):289–305, 2022. doi: 10.1007/s11017-022-09593-0.
- [145] Goren Gordon, Samuel Spaulding, Jacqueline Kory Westlund, Jin Joo Lee, Luke Plummer, Marayna Martinez, Madhurima Das, and Cynthia Breazeal. Affective personalization of a social robot tutor for children’s second language skills. In Dale Schuurmans and Michael P Wellman, editors, *Proceedings of the Thirtieth AAAI*

Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA, pages 3951–3957. AAAI Press, 2016. doi: 10.1609/AAAI.V30I1.9914.

- [146] Samuel D Gosling, Peter J Rentfrow, and William B Swann. A very brief measure of the big-five personality domains. *Journal of Research in Personality*, 37(6):504–528, 2003. doi: 10.1016/S0092-6566(03/00046-1.
- [147] Trisha Greenhalgh, Joseph Wherton, Chrysanthi Papoutsis, Jennifer Lynch, Gemma Hughes, Susan Hinder, Nick Fahy, Rob Procter, and Sara Shaw. Beyond adoption: a new framework for theorizing and evaluating nonadoption, abandonment, and challenges to the scale-up, spread, and sustainability of health and care technologies. *Journal of Medical Internet Research*, 19(11), 2017. doi: 10.2196/jmir.8775.
- [148] Aishwarya Gupta, Divya Hathwar, and A Vijayakumar. Introduction to AI chatbots. *International Journal of Engineering Research & Technology*, 9(7):255–258, 2020.
- [149] David H Gustafson, Michael G Boyle, Bret R Shaw, Andrew Isham, Fiona McTavish, Stephanie Richards, Christopher Schubert, Michael Levy, and Kim Johnson. An e-health solution for people with alcohol problems. *Alcohol Research & Health*, 33(4): 327–337, 2011.
- [150] Kilem L Gwet. *Handbook of inter-rater reliability: The definitive guide to measuring the extent of agreement among raters*. Advanced Analytics, LLC, 4 edition, 2014.
- [151] Kilem L Gwet. *Handbook of inter-rater reliability: The definitive guide to measuring the extent of agreement among raters*, volume 1: Analysis of Categorical Ratings. AgreeStat Analytics, 5 edition, 2021.
- [152] Marcela Haasova, Fiona C Warren, Michael Ussher, Kate Janse Van Rensburg, Guy Faulkner, Mark Cropley, James Byron-Daniel, Emma S Everson-Hock, Hwajung Oh, and Adrian H Taylor. The acute effects of physical activity on cigarette cravings: systematic review and meta-analysis with individual participant data. *Addiction*, 108(1):26–37, 2013. doi: 10.1111/j.1360-0443.2012.04034.x.
- [153] Martin S Hagger and Aleksandra Luszczynska. Implementation intention and action planning interventions in health contexts: State of the research and proposals for the way forward. *Applied Psychology: Health and Well-Being*, 6(1):1–47, 2014. doi: 10.1111/aphw.12017.
- [154] Sajanee Halko and Julie A Kientz. Personality and persuasive technology: an exploratory study on health-promoting mobile applications. In Thomas Ploug, Per F V Hasle, and Harri Oinas-Kukkonen, editors, *Persuasive Technology, 5th International Conference, PERSUASIVE 2010, Copenhagen, Denmark, June 7-10, 2010. Proceedings*, volume 6137 of *Lecture Notes in Computer Science*, pages 150–161, 2010. doi: 10.1007/978-3-642-13226-1_16.
- [155] Yuka Hayakawa, Kaori Fujinami, Junichi Yamamoto, and Airi Tsuji. Self-efficacy measurement method using regression models with anticipatory gaze for supporting rehabilitation. In Klaus Miesenberger, Petr Penáz, and Makato Kobayashi, editors,

- Computers Helping People with Special Needs - 19th International Conference, ICCHP 2024, Linz, Austria, July 8-12, 2024, Proceedings, Part II*, volume 14751 of *Lecture Notes in Computer Science*, pages 311–319, 2024. doi: 10.1007/978-3-031-62849-8_38.
- [156] Conor F Hayes, Roxana Rădulescu, Eugenio Bargiacchi, Johan Källström, Matthew Macfarlane, Mathieu Reymond, Timothy Verstraeten, Luisa M Zintgraf, Richard Dazeley, Fredrik Heintz, Enda Howley, Athirai A Irissappane, Patrick Mannion, Ann Nowé, Gabriel Ramos, Marcello Restelli, Peter Vamplew, and Diederik M Roijers. A practical guide to multi-objective reinforcement learning and planning. *Autonomous Agents and Multi-Agent Systems*, 36:26, 2022. doi: 10.1007/s10458-022-09552-y.
- [157] Richard L Hayes and Ramona Oppenheim. Constructivism: Reality is what you make it. In Thomas L Sexton and Barbara L Griffin, editors, *Constructivist thinking in counseling practice, research, and training*, pages 19–40. Teachers College Press, 1997.
- [158] Linwei He, Divyaa Balaji, Reinout W Wiers, Marjolijn L Antheunis, and Emiel Kraemer. Effectiveness and acceptability of conversational agents for smoking cessation: A systematic review and meta-analysis. *Nicotine & Tobacco Research*, 25(7):1241–1250, 12 2022. ISSN 1469-994X. doi: 10.1093/ntr/ntac281.
- [159] Olivier A Blanson Henkemans, Bert PB Bierman, Joris Janssen, Rosemarijn Looije, Mark A Neerinx, Marierose MM van Dooren, Jitske LE de Vries, Gert Jan van der Burg, and Sasja D Huisman. Design and evaluation of a personal robot playing a self-management education game with children with diabetes type 1. *International Journal of Human-Computer Studies*, 106:63–76, 2017. doi: 10.1016/j.ijhcs.2017.06.001.
- [160] Charlie Hewitt, Ioannis Politis, Theocharis Amanatidis, and Advait Sarkar. Assessing public perception of self-driving cars: the autonomous vehicle acceptance model. In Wai-Tat Fu, Shimei Pan, Oliver Brdiczka, Polo Chau, and Gaelle Calvary, editors, *Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI 2019, Marina del Rey, CA, USA, March 17-20, 2019*, pages 518–527. ACM, 2019. doi: 10.1145/3301275.3302268.
- [161] Takuya Hiraoka, Graham Neubig, Sakriani Sakti, Tomoki Toda, and Satoshi Nakamura. Reinforcement learning of cooperative persuasive dialogue policies using framing. In Jan Hajic and Junichi Tsujii, editors, *COLING 2014, 25th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, August 23-29, 2014, Dublin, Ireland*, pages 1706–1717. Association for Computational Linguistics, 2014.
- [162] Rosemary Hiscock, Linda Bauld, Amanda Amos, Jennifer A Fidler, and Marcus Munafò. Socioeconomic status and smoking: a review. *Annals of the New York Academy of Sciences*, 1248(1):107–123, 2012. doi: 10.1111/j.1749-6632.2011.06202.x.
- [163] Beyza Hizli, Nele Albers, and Willem-Paul Brinkman. Data and code underlying the master thesis: Goal-setting dialogue for physical activity with a virtual coach. <https://doi.org/10.4121/20047328>, 2022.

-
- [164] Rink Hoekstra, Richard D Morey, Jeffrey N Rouder, and Eric-Jan Wagenmakers. Robust misinterpretation of confidence intervals. *Psychonomic Bulletin & Review*, 21(5):1157–1164, 2014. doi: 10.3758/s13423-013-0572-3.
- [165] Santiago Hors-Fraile, Shwetambara Malwade, Francisco Luna-Perejon, Claudio Amaya, Antón Civit, Francine Schneider, Panagiotis Bamidis, Shabbir Syed-Abdul, Yu-Chuan Li, and Hein De Vries. Opening the black box: Explaining the process of basing a health recommender system on the i-change behavioral change model. *IEEE Access*, 7:176525–176540, 2019. doi: 10.1109/ACCESS.2019.2957696.
- [166] Corine Horsch, Jaap Lancee, Robbert Jan Beun, Mark A Neerincx, and Willem-Paul Brinkman. Adherence to technology-mediated insomnia treatment: a meta-analysis, interviews, and focus groups. *Journal of Medical Internet Research*, 17(9):e214, 2015. doi: 10.2196/jmir.4115.
- [167] Corine Horsch, Sandor Spruit, Jaap Lancee, Rogier van Eijk, Robbert Jan Beun, Mark Neerincx, and Willem-Paul Brinkman. Reminders make people adhere better to a self-help sleep intervention. *Health and Technology*, 7:173–188, 2017. doi: 10.1007/s12553-016-0167-x.
- [168] Neil Howlett, Joerg Schulz, Daksha Trivedi, Nicholas Troop, and Angel Chater. Determinants of weekly sitting time: construct validation of an initial com-b model and comparison of its predictive validity with the theory of planned behaviour. *Psychology & Health*, 36(1):96–114, 2021. doi: 10.1080/08870446.2020.1763994.
- [169] John R Hughes, Josue P Keely, Ray S Niaura, Deborah J Ossip-Klein, Robyn L Richmond, and Gary E Swan. Measures of abstinence in clinical trials: issues and recommendations. *Nicotine & Tobacco Research*, 5(1):13–25, 2003.
- [170] Jasmin C Hutchinson and Gershon Tenenbaum. Perceived effort—can it be considered gestalt? *Psychology of Sport and Exercise*, 7(5):463–476, 2006. doi: 10.1016/j.psychsport.2006.01.007.
- [171] Institute for Health Metrics and Evaluation (IHME). Global burden of disease 2021: Findings from the gbd 2021 study. https://www.healthdata.org/sites/default/files/2024-05/GBD_2021_Booklet_FINAL_2024.05.16.pdf, 2024.
- [172] Iris ten Klooster, Hanneke Kip, Lisette van Gemert-Pijnen, Rik Crutzen, and Saskia Kelders. A systematic review on ehealth technology personalization approaches. *iScience*, 27:110771, 2024. ISSN 2589-0042. doi: 10.1016/j.isci.2024.110771.
- [173] David-Zacharie Issom, Marie-Dominique Hardy-Dessources, Marc Romana, Gunnar Hartvigsen, and Christian Lovis. Toward a conversational agent to support the self-management of adults and young adults with sickle cell disease: Usability and usefulness study. *Frontiers in Digital Health*, 3:600333, 2021. doi: 10.3389/fgth.2021.600333.

- [174] Isabel Izquierdo, Julio Olea, and Francisco José Abad. Exploratory factor analysis in validation studies: Uses and recommendations. *Psicothema*, 26(3):395–400, 2014. doi: 10.7334/psicothema2013.349.
- [175] Sharmin Jahan, Mozammel Chowdhury, Rafiqul Islam, and Junbin Gao. Security and privacy protection for ehealth data. In Robin Doss, Selwyn Piramuthu, and Wei Zhou, editors, *Future Network Systems and Security*, pages 197–205, Cham, 2018. Springer International Publishing. doi: 10.1007/978-3-319-94421-0_16.
- [176] Rusi Jaspal, Kevin R Carriere, and Fathali M Moghaddam. Bridging micro, meso, and macro processes in social psychology. In Jaan Valsiner, Giuseppina Marsico, Nandita Chaudhary, Tatsuya Sato, and Virginia Dazzani, editors, *Psychology as the Science of Human Being*, pages 265–276. Springer International Publishing, Cham, 2016. doi: 10.1007/978-3-319-21094-0_15.
- [177] Raouf Jaziri and Saleh Alnahdi. Choosing which covid-19 patient to save? the ethical triage and rationing dilemma. *Ethics, Medicine and Public Health*, 15:100570, 2020. ISSN 2352-5525. doi: 10.1016/j.jemep.2020.100570.
- [178] Oliver P John and Sanjay Srivastava. The big-five trait taxonomy: History, measurement, and theoretical perspectives. In L A Pervin and O P John, editors, *Handbook of personality: Theory and research*, pages 102–138. Guilford Press, New York, 1999.
- [179] Yilin Kang, Ah-Hwee Tan, and Chunyan Miao. An adaptive computational model for personalized persuasion. In Qiang Yang and Michael J Wooldridge, editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 61–67. AAAI Press, 2015.
- [180] Atreyi Kankanhalli, Qihui Xia, Peilin Ai, and Xi Zhao. Understanding personalization for health behavior change applications: A review and future directions. *AIS Transactions on Human-Computer Interaction*, 13(3):316–349, 2021. doi: 10.17705/1thci.00152.
- [181] Maurits Kaptein. Customizing persuasive messages; the value of operative measures. *Journal of Consumer Marketing*, 35(2):208–217, 2018. doi: 10.1108/JCM-11-2016-1996.
- [182] Maurits Kaptein and Dean Eckles. Heterogeneity in the effects of online persuasion. *Journal of Interactive Marketing*, 26(3):176–188, 2012. doi: 10.1016/j.intmar.2012.02.002.
- [183] Maurits Kaptein, Panos Markopoulos, Boris De Ruyter, and Emile Aarts. Personalizing persuasive technologies: Explicit and implicit personalization using persuasion profiles. *International Journal of Human-Computer Studies*, 77:38–51, 2015. doi: 10.1016/j.ijhcs.2015.01.004.
- [184] Maurits Kaptein, Richard McFarland, and Petri Parvinen. Automated adaptive selling. *European Journal of Marketing*, 56(5/6):1037–1059, 2018. doi: 10.1108/EJM-08-2016-0485.

-
- [185] Maria Karekla, Stella Nicoleta Savvides, and Andrew Gloster. An avatar-led intervention promotes smoking cessation in young adults: a pilot randomized clinical trial. *Annals of Behavioral Medicine*, 54(10):747–760, 2020. doi: 10.1093/abm/kaaa013.
- [186] Ravi Karkar, Jasmine Zia, Roger Vilardaga, Sonali R Mishra, James Fogarty, Sean A Munson, and Julie A Kientz. A framework for self-experimentation in personalized health. *Journal of the American Medical Informatics Association*, 23(3):440–448, 2016. doi: 10.1093/jamia/ocv150.
- [187] George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6): 422–440, 2021. doi: 10.1038/s42254-021-00314-5.
- [188] Arian Kashefi, Kyana Bosschaerts, Stephen Murphy, Sofie Van Hoecke, Mariek Vanden Abeele, Lieven De Marez, Femke Ongenaes, and Peter Conradie. Unlocking the potential of mhealth for smoking cessation: An expert viewpoint. In June Wei and George Margetis, editors, *Human-Centered Design, Operation and Evaluation of Mobile Communications - 5th International Conference, MOBILE 2024, Held as Part of the 26th HCI International Conference, HCII 2024, Washington, DC, USA, June 29 - July 4, 2024, Proceedings, Part I*, volume 14737 of *Lecture Notes in Computer Science*, pages 59–79, 2024. doi: 10.1007/978-3-031-60458-4_5.
- [189] David J Kavanagh and Gordon H Bower. Mood and self-efficacy: Impact of joy and sadness on perceived capabilities. *Cognitive Therapy and Research*, 9(5):507–525, 1985. doi: 10.1007/BF01173005.
- [190] Frances Kay-Lambkin, Amanda Baker, Terry Lewin, and Vaughan Carr. Acceptability of a clinician-assisted computerized psychological intervention for comorbid mental health and substance use problems: treatment adherence data from a randomized controlled trial. *Journal of Medical Internet Research*, 13(1), 2011. doi: 10.2196/jmir.1522.
- [191] Saskia M Kelders, Llewellyn Ellardus Van Zyl, and Geke DS Ludden. The concept and components of engagement in different domains applied to ehealth: a systematic scoping review. *Frontiers in Psychology*, 11:926, 2020. doi: 10.3389/fpsyg.2020.00926.
- [192] George A Kelly. *The Psychology of Personal Constructs, Volumes 1 and 2*. W. W. Norton, 1955.
- [193] Charles E Thomas Kent A Walstrom and Ann Weber. Changes in student computer technology attitudes over 20 years: 1988 to 2009. *Journal of Computer Information Systems*, 51(2):81–86, 2010. doi: 10.1080/08874417.2010.11645471.
- [194] Hager Khechine and Marc Augier. Adoption of a social learning platform in higher education: An extended UTAUT model implementation. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*, pages 53–62, 2019.
- [195] Sarah Eve Kinsey, Jack Wolf, Nalini Saligram, Varun Ramesan, Meeta Walavalkar, Nidhi Jaswal, Sandhya Ramalingam, Arunesh Sinha, and Thanh Hong Nguyen.

- Building a personalized messaging system for health intervention in underprivileged regions using reinforcement learning. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*, pages 6022–6030, 2023. doi: 10.24963/IJCAI.2023/668.
- [196] Predrag Klasnja, Eric B Hekler, Saul Shiffman, Audrey Boruvka, Daniel Almirall, Ambuj Tewari, and Susan A Murphy. Microrandomized trials: An experimental design for developing just-in-time adaptive interventions. *Health Psychology*, 34S(0): 1220–1228, 2015. doi: 10.1037/hea0000305.
- [197] Britt Klein, David Austin, Ciaran Pier, Litza Kiropoulos, Kerrie Shandley, Joanna Mitchell, Kathryn Gilson, and Lisa Ciechomski. Internet-based treatment for panic disorder: Does frequency of therapist contact make a difference? *Cognitive Behaviour Therapy*, 38(2):100–113, 2009. doi: 10.1080/16506070802561132.
- [198] Michel Klein, Nataliya Mogles, and Arlette van Wissen. An intelligent coaching system for therapy adherence. *IEEE Pervasive Computing*, 12(3):22–30, 2013. doi: 10.1109/MPRV.2013.41.
- [199] Ahmet Baki Kocaballi, Shlomo Berkovsky, Juan C Quiroz, Liliana Laranjo, Huong Ly Tong, Dana Rezazadegan, Agustina Briatore, and Enrico Coiera. The personalization of conversational agents in health care: systematic review. *Journal of Medical Internet Research*, 21(11):e15360, 2019. doi: 10.2196/15360.
- [200] Felicitas Kraemer, Kees Van Overveld, and Martin Peterson. Is there an ethics of algorithms? *Ethics and Information Technology*, 13:251–260, 2011. doi: 10.1007/s10676-010-9233-7.
- [201] Justin Kruger and David Dunning. Unskilled and unaware of it: how difficulties in recognizing one’s own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, 77(6):1121–1134, 1999. doi: 10.1037//0022-3514.77.6.1121.
- [202] John Kruschke. *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan*. Academic Press, 2 edition, 2014.
- [203] John K Kruschke. Rejecting or accepting parameter values in bayesian estimation. *Advances in Methods and Practices in Psychological Science*, 1(2):270–280, 2018. doi: 10.1177/2515245918771304.
- [204] A Kulhánek, R Gabrhelík, D Novák, V Burda, and H Brendryen. ehealth intervention for smoking cessation for czech tobacco smokers: pilot study of user acceptance. *Adiktologie*, 18(2):81–85, 2018.
- [205] Frederick J Kviz, Melissa A Clark, Kathleen S Crittenden, Sally Freels, and Richard B Warnecke. Age and readiness to quit smoking. *Preventive Medicine*, 23(2):211–222, 1994.

-
- [206] Sawsen Lakhal, Hager Khechine, and Daniel Pascot. Student behavioural intentions to use desktop video conferencing in a distance course: integration of autonomy to the utaut model. *Journal of Computing in Higher Education*, 25:93–121, 2013. doi: 10.1007/s12528-013-9069-3.
- [207] Timo A Lakka, Kirsikka Aittola, Elina Järvelä-Reijonen, Tanja Tilles-Tirkkonen, Reija Männikkö, Niina Lintu, Leila Karhunen, Marjukka Kolehmainen, Marja Harjumaa, Elina Mattila, Riia Järvenpää, Miikka Ermes, Santtu Mikkonen, Janne Martikainen, Kaisa Poutanen, Ursula Schwab, Pilvikki Absetz, Jaana Lindström, and Jussi Pihlajamäki. Real-world effectiveness of digital and group-based lifestyle interventions as compared with usual care to reduce type 2 diabetes risk—a stop diabetes pragmatic randomised trial. *The Lancet Regional Health—Europe*, 24:100527, 2023. doi: 10.1016/j.lanepe.2022.100527.
- [208] George Lakoff. Why it matters how we frame the environment. *Environmental Communication*, 4(1):70–81, 2010. doi: 10.1080/17524030903529749.
- [209] J Richard Landis and Gary G Koch. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174, 1977.
- [210] Yi-Chieh Lee, Naomi Yamashita, and Yun Huang. Exploring the effects of incorporating human experts to deliver journaling guidance through a chatbot. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1):1–27, 2021. doi: 10.1145/3449196.
- [211] William Leeson, Adam Resnick, Daniel Alexander, and John Rovers. Natural language processing (nlp) in qualitative public health research: a proof of concept study. *International Journal of Qualitative Methods*, 18:1609406919887021, 2019. doi: 10.1177/160940691988702.
- [212] Sietske Lems. De huisarts-patiëntrelatie anno 2006. *Amsterdam/Utrecht: TNS NIPO/NHG*, 2006.
- [213] Aniek J Lentferink, Hilbrand KE Oldenhuis, Martijn de Groot, Louis Polstra, Hugo Velthuisen, and Julia EWC van Gemert-Pijnen. Key components in ehealth interventions combining self-tracking and persuasive ecoaching to promote a healthier lifestyle: a scoping review. *Journal of Medical Internet Research*, 19(8):e277, 2017. doi: 10.2196/jmir.7288.
- [214] Qiang Li, Jason Hsia, and Gonghuan Yang. Prevalence of smoking in china in 2010. *New England Journal of Medicine*, 364(25):2469–2470, 2011. doi: 10.1056/NEJMc1102459.
- [215] Shen Li, Zhan Qu, Yiyang Li, and Xuelei Ma. Efficacy of e-health interventions for smoking cessation management in smokers: a systematic review and meta-analysis. *eClinicalMedicine*, 68, 2024. doi: 10.1016/j.eclinm.2023.102412.
- [216] Fan Liang and Yuchen Chen. The making of “good” citizens: China’s social credit systems and infrastructures of social quantification. *Policy & Internet*, 14(1):114–135, 2022. doi: 10.1002/poi3.291.

- [217] Peng Liao, Kristjan Greenewald, Predrag Klasnja, and Susan Murphy. Personalized heartsteps: A reinforcement learning algorithm for optimizing physical activity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(1), 2020. doi: 10.1145/3381007.
- [218] Yanhui Liao, Qiuxia Wu, Jinsong Tang, Fengyu Zhang, Xuyi Wang, Chang Qi, Haoyu He, Jiang Long, Brian C Kelly, and Joanna Cohen. The efficacy of mobile phone-based text message interventions (‘happy quit’) for smoking cessation in china. *BMC Public Health*, 16:833, 2016. doi: 10.1186/s12889-016-3528-5.
- [219] Silje Stangeland Lie, Bjørg Karlsen, Ellen Renate Oord, Marit Graue, and Bjørg Oftedal. Dropout from an ehealth intervention for adults with type 2 diabetes: a qualitative study. *Journal of Medical Internet Research*, 19(5):e187, 2017. doi: 10.2196/jmir.7479.
- [220] Jasy Suet Yan Liew, Nancy McCracken, Shichun Zhou, and Kevin Crowston. Optimizing features in active machine learning for complex qualitative content analysis. In Cristian Danescu-Niculescu-Mizil, Jacob Eisenstein, Kathleen R McKeown, and Noah A Smith, editors, *Proceedings of the Workshop on Language Technologies and Computational Social Science@ACL 2014, Baltimore, MD, USA, June 26, 2014*, pages 44–48. Association for Computational Linguistics, 2014. doi: 10.3115/V1/W14-2513.
- [221] Hui-Fei Lin and Fuyuan Shen. Regulatory focus and attribute framing: Evidence of compatibility effects in advertising. *International Journal of Advertising*, 31(1): 169–188, 2012. doi: 10.2501/IJA-31-1-169-188.
- [222] Jianchang Lin, Li-An Lin, and Serap Sankoh. A general overview of adaptive randomization design for clinical trials. *Journal of Biometrics & Biostatistics*, 7(2):294, 2016. doi: 10.4172/2155-6180.1000294.
- [223] Yu-li Liu, Wenjia Yan, Bo Hu, Zhuoyang Li, and Yik Ling Lai. Effects of personalization and source expertise on users’ health beliefs and usage intention toward health chatbots: Evidence from an online experiment. *Digital Health*, 8:20552076221129718, 2022. doi: 10.1177/20552076221129718.
- [224] Edwin A Locke and Gary P Latham. Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *American Psychologist*, 57(9):705–717, 2002. doi: 10.1037/0003-066X.57.9.705.
- [225] Lan Luo, Cara Arizmendi, and Kathleen M Gates. Exploratory factor analysis (efa) programs in r. *Structural Equation Modeling: A Multidisciplinary Journal*, 26(5): 819–826, 2019. doi: 10.1080/10705511.2019.1615835.
- [226] Francisco Lupiáñez-Villanueva, Miquel Àngel Mayer, and Joan Torrent. Opportunities and challenges of web 2.0 within the health care systems: an empirical exploration. *Informatics for Health and Social Care*, 34(3):117–126, 2009. doi: 10.1080/17538150903102265.
- [227] Kien Hoa Ly, Ann-Marie Ly, and Gerhard Andersson. A fully automated conversational agent for promoting mental well-being: a pilot rct using mixed methods. *Internet Interventions*, 10:39–46, 2017. doi: 10.1016/j.invent.2017.10.002.

-
- [228] Tambiama Madiega. Eu guidelines on ethics in artificial intelligence: Context and implementation. *European Parliamentary Research Service*, pages 1–13, 2019.
- [229] Durairaj Maheswaran and Joan Meyers-Levy. The influence of message framing and issue involvement. *Journal of Marketing Research*, 27(3):361–367, 1990. doi: 10.2307/3172593.
- [230] Milena Soriano Marcolino, João Antonio Queiroz Oliveira, Marcelo D’Agostino, Antonio Luiz Ribeiro, Maria Beatriz Moreira Alkmim, and David Novillo-Ortiz. The impact of mhealth interventions: systematic review of systematic reviews. *JMIR mHealth and uHealth*, 6(1), 2018. doi: 10.2196/mhealth.8873.
- [231] Judith Masthoff and Julita Vassileva. 9 personalized persuasion for behavior change. In Mirjam Augstein, Eelco Herder, and Wolfgang Wörndl, editors, *Personalized Human-Computer Interaction*, page 205. De Gruyter Oldenbourg, 2023. doi: 10.1515/9783110988567-009.
- [232] Roger C Mayer, James H Davis, and F David Schoorman. An integrative model of organizational trust. *The Academy of Management Review*, 20(3):709–734, 1995. doi: 10.2307/258792.
- [233] Edward McAuley. Self-efficacy and the maintenance of exercise participation in older adults. *Journal of Behavioral Medicine*, 16(1):103–113, 1993. doi: 10.1007/BF00844757.
- [234] Richard McElreath. *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. Chapman and Hall/CRC, 2020.
- [235] Eline Meijer, Winifred A Gebhardt, Colette Van Laar, Ramin Kawous, and Sarah CAM Beijik. Socio-economic status in relation to smoking: The role of (expected and desired) social support and quitter identity. *Social Science & Medicine*, 162:41–49, 2016. doi: 10.1016/j.socscimed.2016.06.022.
- [236] Eline Meijer, Winifred A Gebhardt, Colette van Laar, Bas van den Putte, and Andrea WM Evers. Strengthening quitter self-identity: An experimental study. *Psychology & Health*, 33(10):1229–1250, 2018. doi: 10.1080/08870446.2018.1478976.
- [237] Eline Meijer, Janneke S Korst, Kristiene G Oosting, Eline Heemskerk, Sander Hermsen, Marc C Willemsen, Bas van den Putte, Niels H Chavannes, and Jamie Brown. “at least someone thinks i’m doing well”: a real-world evaluation of the quit-smoking app stopcoach for lower socio-economic status smokers. *Addiction Science & Clinical Practice*, 16(1):1–14, 2021. doi: 10.1186/s13722-021-00255-5.
- [238] Eline Meijer, Kristell Penfornis, Nele Albers, Bouke Scheltinga, Douwe Atsma, Niels Chavannes, Sven van der Burg, Willem-Paul Brinkman, and Winnie Gebhardt. Perfect fit: een identiteitsgebaseerde virtuele coach die ondersteunt bij stoppen met roken en meer bewegen. *Nederlands Netwerk voor Tabaksonderzoek (NNvT) Congres*, 2021.

- [239] Liesbeth Mercken, M Candel, L Van Osch, and Hein de Vries. No smoke without fire: The impact of future friends on adolescent smoking behaviour. *British Journal of Health Psychology*, 16(1):170–188, 2011. doi: 10.1348/135910710X531608.
- [240] S Michie, R West, R Campbell, J Brown, and H Gainforth. *ABC of behaviour change theories: an essential resource for researchers, policy makers and practitioners*. Silverback Publishing, 2014.
- [241] Susan Michie, Stefanie Ashford, Falko F Sniehotta, Stephan U Dombrowski, Alex Bishop, and David P French. A refined taxonomy of behaviour change techniques to help people change their physical activity and healthy eating behaviours: the calo-re taxonomy. *Psychology & Health*, 26(11):1479–1498, 2011. doi: 10.1080/08870446.2010.540664.
- [242] Susan Michie, Maartje M Van Stralen, and Robert West. The behaviour change wheel: a new method for characterising and designing behaviour change interventions. *Implementation Science*, 6(1):42, 2011. doi: 10.1186/1748-5908-6-42.
- [243] Susan Michie, Jamie Brown, Adam WA Geraghty, Sascha Miller, Lucy Yardley, Benjamin Gardner, Lion Shahab, Andy McEwen, John A Stapleton, and Robert West. Development of stopadvisor: a theory-based interactive internet-based smoking cessation intervention. *Translational Behavioral Medicine*, 2(3):263–275, 2012. doi: 10.1007/s13142-012-0135-6.
- [244] Susan Michie, Michelle Richardson, Marie Johnston, Charles Abraham, Jill Francis, Wendy Hardeman, Martin P Eccles, James Cane, and Caroline E Wood. The behavior change technique taxonomy (v1) of 93 hierarchically clustered techniques: building an international consensus for the reporting of behavior change interventions. *Annals of Behavioral Medicine*, 46(1):81–95, 2013. doi: 10.1007/s12160-013-9486-6.
- [245] Susan Michie, Lou Atkins, and Robert West. *The behaviour change wheel: A guide to designing interventions*. Silverback Publishing, 1 edition, 2014.
- [246] Susan Michie, Lucy Yardley, Robert West, Kevin Patrick, and Felix Greaves. Developing and evaluating digital interventions to promote behavior change in health and health care: recommendations resulting from an international workshop. *Journal of Medical Internet Research*, 19(6):e232, 2017. doi: 10.2196/jmir.7126.
- [247] Katarzyna Milcarz, Kinga Polańska, Łukasz Balwicki, Teresa Makowiec-Dąbrowska, Wojciech Hanke, Leokadia Bąk-Romaniszyn, and Dorota Kaleta. Perceived barriers and motivators to smoking cessation among socially-disadvantaged populations in poland. *International Journal of Occupational Medicine and Environmental Health*, 32(3):363–377, 2019. doi: 10.13075/ijom.1896.01377.
- [248] Ministry of Family, Labour and Social Policy, Republic of Poland. Minimum wage. <https://www.gov.pl/web/family/minimum-wage>, 2023. Accessed on 27.08.2024.

-
- [249] Yonatan Mintz, Anil Aswani, Philip Kaminsky, Elena Flowers, and Yoshimi Fukuoka. Nonstationary bandits with habituation and recovery dynamics. *Operations Research*, 68(5):1493–1516, 2020. doi: 10.1287/opre.2019.1918.
- [250] Elliot G Mitchell, Rosa Maimone, Andrea Cassells, Jonathan N Tobin, Patricia Davidson, Arlene M Smaldone, and Lena Mamykina. Automated vs. human health coaching: exploring participant and practitioner experiences. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), 2021. doi: 10.1145/3449173.
- [251] Aditya Mohan, Amy Zhang, and Marius Lindauer. Structure in deep reinforcement learning: A survey and open problems. *Journal of Artificial Intelligence Research*, 79: 1167–1236, 2024. doi: 10.1613/JAIR.1.15703.
- [252] David C Mohr, Pim Cuijpers, and Kenneth Lehman. Supportive accountability: a model for providing human support to enhance adherence to ehealth interventions. *Journal of Medical Internet Research*, 13(1):e30, 2011. doi: 10.2196/jmir.1602.
- [253] Ali H Mokdad, James S Marks, Donna F Stroup, and Julie L Gerberding. Actual causes of death in the united states, 2000. *JAMA*, 291(10):1238–1245, 2004. doi: 10.1001/jama.291.10.1238.
- [254] Alberto Molina-Pérez, David Rodríguez-Arias, and Janet Delgado. Differential impact of opt-in, opt-out policies on deceased organ donation rates: a mixed conceptual and empirical study. *BMJ Open*, 12(9):e057107, 2022. doi: 10.1136/bmjopen-2021-057107.
- [255] Joao Luis Zeni Montenegro, Cristiano André da Costa, and Rodrigo da Rosa Righi. Survey of conversational agents in health. *Expert Systems with Applications*, 129: 56–67, 2019. doi: 10.1016/j.eswa.2019.03.054.
- [256] Aaron J Moss, Cheskie Rosenzweig, Jonathan Robinson, Shalom N Jaffe, and Leib Litman. Is it ethical to use mechanical turk for behavioral research? relevant data from a representative survey of mturk participants and wages. *Behavior Research Methods*, 55(8):4048–4067, 2023. doi: 10.3758/s13428-022-02005-0.
- [257] Aisha Muhammad Abdullahi, Rita Orji, and Kiemute Oyibo. Personalizing persuasive technologies: Do gender and age affect susceptibility to persuasive strategies? In *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization, UMAP 2018, Singapore, July 08-11, 2018*, pages 329–334, New York, NY, USA, 2018. ACM. doi: 10.1145/3213586.3225246.
- [258] Marcus R Munafo, Jessica I Zetteler, and Taane G Clark. Personality and smoking status: A meta-analysis. *Nicotine & Tobacco Research*, 9(3):405–413, 2007. doi: 10.1080/14622200701188851.
- [259] Tamara Munzner. *Visualization analysis and design*. CRC Press, 1 edition, 2014.
- [260] P Karen Murphy and Lucia Mason. Changing knowledge and beliefs. In Patricia A Alexander and Philip H Winne, editors, *Handbook of Educational Psychology*, pages 305–324. Lawrence Erlbaum Associates, New York, 2 edition, 2006.

- [261] Tom Nadarzynski, Oliver Miles, Aimee Cowie, and Damien Ridge. Acceptability of artificial intelligence (ai)-led chatbot services in healthcare: A mixed-methods study. *Digital Health*, 5, 2019. doi: 10.1177/2055207619871808.
- [262] Nadia Nahar, Shurui Zhou, Grace Lewis, and Christian Kästner. More engineering, no silos: Rethinking processes and interfaces in collaboration between interdisciplinary teams for machine learning projects. *arXiv preprint arXiv:2110.10234*, 2021.
- [263] Clifford Nass and Youngme Moon. Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56(1):81–103, 2000. doi: 10.1111/0022-4537.00153.
- [264] National Academies of Sciences, Engineering, and Medicine. *Reproducibility and Replicability in Science*. The National Academies Press, Washington, DC, 2019. ISBN 978-0-309-48616-3. doi: 10.17226/25303. URL <https://nap.nationalacademies.org/catalog/25303/reproducibility-and-replicability-in-science>.
- [265] National Cancer Institute. Handling nicotine withdrawal and triggers when you decide to quit tobacco, 2022. URL <https://www.cancer.gov/about-cancer/causes-prevention/risk/tobacco/withdrawal-fact-sheet>. Accessed in 2023.
- [266] Felix Naughton, Aimie Hope, Chloë Siegele-Brown, Kelly Grant, Garry Barton, Caitlin Notley, Cecilia Mascolo, Tim Coleman, Lee Shepstone, Stephen Sutton, A Toby Prevost, David Crane, Felix Greaves, and Juliet High. An automated, online feasibility randomized controlled trial of a just-in-time adaptive intervention for smoking cessation (quit sense). *Nicotine & Tobacco Research*, 25(7):1319–1329, 2023. doi: 10.1093/ntr/ntad032.
- [267] Rachel K Nelligan, Rana S Hinman, Pek Ling Teo, and Kim L Bennell. Exploring attitudes and experiences of people with knee osteoarthritis toward a self-directed ehealth intervention to support exercise: Qualitative study. *JMIR Rehabilitation and Assistive Technologies*, 7(2):e18860, 2020. doi: 10.2196/18860.
- [268] Allen Newell and Herbert Alexander Simon. *Human problem solving*. Prentice-Hall, Englewood Cliffs, NJ, 1972.
- [269] Ritesh Noothigattu, Djallel Bouneffouf, Nicholas Mattei, Rachita Chandra, Piyush Madan, Kush R Varshney, Murray Campbell, Moninder Singh, and Francesca Rossi. Teaching AI agents ethical values using reinforcement learning and policy orchestration. *IBM Journal of Research and Development*, 63(4/5):2:1–2:9, 2019. doi: 10.1147/JRD.2019.2940428.
- [270] Donald A. Norman. *The design of everyday things*. MIT Press, Cambridge, MA, 2013.
- [271] GJ Norman, SV Benisovich, CR Nigg, and JS Rossi. Examining three exercise staging algorithms in two samples. In *19th Annual Meeting of the Society of Behavioral Medicine*, 1998.

-
- [272] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453, 2019. doi: 10.1126/science.aax2342.
- [273] Heather L O’Brien and Elaine G Toms. The development and evaluation of a survey to measure user engagement. *Journal of the American Society for Information Science and Technology*, 61(1):50–69, 2010. doi: 10.1002/asi.21229.
- [274] Daan Odijk, Björn Burscher, Rens Vliegenthart, and Maarten de Rijke. Automatic thematic content analysis: Finding frames in news. In Adam Jatowt, Ee-Peng Lim, Ying Ding, Asako Miura, Taro Tezuka, Gaël Dias, Katsumi Tanaka, Andrew J Flanagan, and Bing Tian Dai, editors, *Social Informatics - 5th International Conference, SocInfo 2013, Kyoto, Japan, November 25-27, 2013, Proceedings*, volume 8238 of *Lecture Notes in Computer Science*, pages 333–345, Cham, 2013. Springer International Publishing. doi: 10.1007/978-3-319-03260-3_29.
- [275] Harri Oinas-Kukkonen. Behavior change support systems: A research model and agenda. In Thomas Ploug, Per Hasle, and Harri Oinas-Kukkonen, editors, *Persuasive Technology, 5th International Conference, PERSUASIVE 2010, Copenhagen, Denmark, June 7-10, 2010. Proceedings*, volume 6137 of *Lecture Notes in Computer Science*, pages 4–14, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. ISBN 978-3-642-13226-1.
- [276] Harri Oinas-Kukkonen and Marja Harjumaa. A systematic framework for designing and evaluating persuasive systems. In Harri Oinas-Kukkonen, Per F V Hasle, Marja Harjumaa, Katarina Segerståhl, and Peter Øhrstrøm, editors, *Persuasive Technology, Third International Conference, PERSUASIVE 2008, Oulu, Finland, June 4-6, 2008. Proceedings*, *Lecture Notes in Computer Science*, pages 164–176, 2008. doi: 10.1007/978-3-540-68504-3_15.
- [277] Harri Oinas-Kukkonen and Marja Harjumaa. Persuasive systems design: Key issues, process model, and system features. *Communications of the Association for Information Systems*, 24:28, 2009. doi: 10.17705/1cais.02428.
- [278] Eduardo Olano-Espinosa, Jose Francisco Avila-Tomas, Cesar Minue-Lorenzo, Blanca Matilla-Pardo, María Encarnación Serrano Serrano, F Javier Martinez-Suberviola, Mario Gil-Conesa, Isabel Del Cura-González, and Dejal@ Group. Effectiveness of a conversational chatbot (dejal@ bot) for the adult population to quit smoking: pragmatic, multicenter, controlled, randomized clinical trial in primary care. *JMIR mHealth and uHealth*, 10(6):e34273, 2022. doi: 10.2196/34273.
- [279] Travis E Oliphant. A bayesian perspective on estimating mean, variance, and standard-deviation from data. *Faculty Publications*, 2006.
- [280] Temitayo A Olugbade, Nadia Bianchi-Berthouze, Nicolai Marquardt, and Amanda C de C Williams. Human observer and automatic assessment of movement related self-efficacy in chronic pain: from exercise to functional activity. *IEEE Transactions on Affective Computing*, 11(2):214–229, 2018. doi: 10.1109/TAFFC.2018.2798576.

- [281] Annet Onnes. Monitoring AI systems: a problem analysis, framework and outlook. In Stefan Schlobach, María Pérez-Ortiz, and Myrthe Tielman, editors, *HAI 2022: Augmenting Human Intellect - Proceedings of the First International Conference on Hybrid Human-Artificial Intelligence, Amsterdam, The Netherlands, 13-17 June 2022*, volume 354 of *Frontiers in Artificial Intelligence and Applications*, pages 238–240, 2022. doi: 10.3233/FAIA220202.
- [282] Annet Onnes, Mehdi Dastani, and Silja Renooij. Bayesian network conflict detection for normative monitoring of black-box systems. In Michael Franklin and Soon Ae Chun, editors, *Proceedings of the Thirty-Sixth International Florida Artificial Intelligence Research Society Conference, FLAIRS 2023, Clearwater Beach, FL, USA, May 14-17, 2023*, 2023. doi: 10.32473/FLAIRS.36.133240.
- [283] Santiago Ontañón. An overview of distance and similarity functions for structured data. *Artificial Intelligence Review*, 53(7):5309–5351, 2020. doi: 10.1007/s10462-020-09821-w.
- [284] Harm op den Akker, Miriam Cabrita, Riëks op den Akker, Valerie M Jones, and Hermie J Hermens. Tailored motivational message generation: A model and practical framework for real-time physical activity coaching. *Journal of Biomedical Informatics*, 55:104–115, 2015. ISSN 1532-0464. doi: 10.1016/j.jbi.2015.03.005.
- [285] Rita Orji, Julita Vassileva, and Regan L Mandryk. Modeling the efficacy of persuasive strategies for different gamer types in serious games for health. *User Modeling and User-Adapted Interaction*, 24(5):453–498, 2014. doi: 10.1007/s11257-014-9149-8.
- [286] Oladapo Oyebo, Chinenye Ndulue, Dinesh Mulchandani, Ashfaq A Zamil Adib, Mona Alhasani, and Rita Orji. Tailoring persuasive and behaviour change systems based on stages of change and motivation. In Yoshifumi Kitamura, Aaron Quigley, Katherine Isbister, Takeo Igarashi, Pernille Bjørn, and Steven Mark Drucker, editors, *CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8-13, 2021*. ACM, 2021. doi: 10.1145/3411764.3445619.
- [287] Oladapo Oyebo, Darren Steeves, and Rita Orji. Persuasive strategies and emotional states: towards designing personalized and emotion-adaptive persuasive systems. *User Modeling and User-Adapted Interaction*, 2024. doi: 10.1007/s11257-023-09390-x.
- [288] Kiemute Oyibo, Rita Orji, and Julita Vassileva. Effects of personality on cialdini’s persuasive strategies. In Peter w de Vries and Thomas Van Rompay, editors, *12th International Conference on Persuasive Technology, PERSUASIVE 2017, Amsterdam, The Netherlands, April 4-6, 2017, Adjunct Proceedings*, pages 54–55, 2017.
- [289] Kiemute Oyibo, Ifeoma Adaji, Rita Orji, Babatunde Olabenjo, and Julita Vassileva. Susceptibility to persuasive strategies: A comparative analysis of nigerians vs. canadians. In Tanja Mitrovic, Jie Zhang, Li Chen, and David Chin, editors, *Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization, UMAP 2018, Singapore, July 08-11, 2018*, pages 229–238. ACM, 2018. doi: 10.1145/3209219.3209239.

-
- [290] Claudia Pagliari, David Sloan, Peter Gregor, Frank Sullivan, Don Detmer, James P Kahan, Wija Oortwijn, and Steve MacGillivray. What is ehealth (4): a scoping exercise to map the field. *Journal of Medical Internet Research*, 7(1), 2005. doi: 10.2196/jmir.7.1.e9.
- [291] George Papathanasiou, Anastasia Mamali, Spyridon Papafloratos, and Efthimia Zerva. Effects of smoking on cardiovascular function: the role of nicotine and carbon monoxide. *Health Science Journal*, 8(2):274–290, 2014.
- [292] Pablo Paredes, Ran Gilad-Bachrach, Mary Czerwinski, Asta Roseway, Kael Rowan, and Javier Hernandez. Poptherapy: Coping with stress through pop-culture. In Andreas Hein, Susanne Boll, and Friedrich Köhler, editors, *Proceedings of the 8th International Conference on Pervasive Computing Technologies for Healthcare, PervasiveHealth 2014, Oldenburg, Germany, May 20-23, 2014*, pages 109–117, 2014. doi: 10.4108/ICST.PERVASIVEHEALTH.2014.255070.
- [293] Moonkyoung Park, Baram Kang, Ahyun Ryu, YueLin Li, and Rhayun Song. Motivational factors for smoking behaviors in individuals with metabolic syndrome. *Patient Preference and Adherence*, 15:2847–2854, 2021. doi: 10.2147/PPA.S343874.
- [294] Kristell Penfornis, Winifred Gebhardt, and Eline Meijer. Mijn toekomstige zelf is (niet) gestopt met roken: Een experimentele studie naar de effecten van een toekomstige-zelf interventie op de zelfidentiteit van rokers. *Nederlands Netwerk voor Tabaksonderzoek (NNvT) Congres*, 2021.
- [295] Kristell M Penfornis, Winifred A Gebhardt, Ralph CA Rippe, Colette Van Laar, Bas van den Putte, and Eline Meijer. My future-self has (not) quit smoking: An experimental study into the effect of a future-self intervention on smoking-related self-identity constructs. *Social Science & Medicine*, 320:115667, 2023. doi: 10.1016/j.socscimed.2023.115667.
- [296] Evan K. Perrault and Samantha A. Nazione. Informed consent—uninformed participants: Shortcomings of online social science consent forms and recommendations for improvement. *Journal of Empirical Research on Human Research Ethics*, 11(3): 274–280, 2016. doi: 10.1177/1556264616654610.
- [297] Govind Persad, Alan Wertheimer, and Ezekiel J Emanuel. Principles for allocation of scarce medical interventions. *The Lancet*, 373(9661):423–431, 2009. doi: 10.1016/S0140-6736(09)60137-9.
- [298] Richard E Petty and John T Cacioppo. The elaboration likelihood model of persuasion. In *Communication and Persuasion: Central and Peripheral Routes to Attitude Change*, pages 1–24. Springer New York, New York, NY, 1986. doi: 10.1007/978-1-4612-4964-1_1.
- [299] John D Piette, Karen B Farris, Sean Newman, Larry An, Jeremy Sussman, and Satinder Singh. The potential impact of intelligent systems for mobile health self-management support: Monte carlo simulations of text message support for medication adherence. *Annals of Behavioral Medicine*, 49(1):84–94, 2015. doi: 10.1007/s12160-014-9634-7.

- [300] John D Piette, Sean Newman, Sarah L Krein, Nicolle Marinec, Jenny Chen, David A Williams, Sara N Edmond, Mary Driscoll, Kathryn M LaChappelle, Robert D Kerns, Marianna Maly, H Myra Kim, Karen B Farris, Diana M Higgins, Eugenia Buta, and Alicia A Heapy. Patient-centered pain care using artificial intelligence and mobile health tools: a randomized comparative effectiveness trial. *JAMA Internal Medicine*, 182(9):975–983, 2022. doi: 10.1001/jamainternmed.2022.3178.
- [301] John D Piette, Laura Thomas, Sean Newman, Nicolle Marinec, Joel Krauss, Jenny Chen, Zhenke Wu, and Amy SB Bohnert. An automatically adaptive digital health intervention to decrease opioid-related risk while conserving counselor time: Quantitative analysis of treatment decisions based on artificial intelligence and patient-reported risk measures. *Journal of Medical Internet Research*, 25:e44165, 2023. doi: 10.2196/44165.
- [302] Alina Pommeranz, Joost Broekens, Pascal Wiggers, Willem-Paul Brinkman, and Catholijn M Jonker. Designing interfaces for explicit preference elicitation: a user-centered investigation of preference representation and elicitation process. *User Modeling and User-Adapted Interaction*, 22:357–397, 2012. doi: 10.1007/S11257-011-9116-6.
- [303] Louise Poppe, Celien Van der Mispel, Ilse De Bourdeaudhuij, Maité Verloigne, Samyah Shadid, and Geert Crombez. Users’ thoughts and opinions about a self-regulation-based ehealth intervention targeting physical activity and the intake of fruit and vegetables: a qualitative study. *PLoS ONE*, 12(12):e0190020, 2017. doi: 10.1371/journal.pone.0190020.
- [304] Louise Poppe, Celien Van der Mispel, Geert Crombez, Ilse De Bourdeaudhuij, Helene Schroé, and Maité Verloigne. How users experience and use an ehealth intervention based on self-regulation: mixed-methods study. *Journal of Medical Internet Research*, 20(10):e10412, 2018. doi: 10.2196/10412.
- [305] Carly S Priebe, John Atkinson, and Guy Faulkner. Run to quit: An evaluation of a scalable physical activity-based smoking cessation intervention. *Mental Health and Physical Activity*, 13:15–21, 2017. doi: 10.1016/j.mhpa.2017.08.001.
- [306] Carly S Priebe, Mark Beauchamp, Kelly Wunderlich, and Guy Faulkner. “i’m a runner not a smoker”: Changes in identity as predictors of smoking cessation and physical activity. *Psychology of Sport and Exercise*, 49:101702, 2020. doi: 10.1016/j.psychsport.2020.101702.
- [307] James O Prochaska and Carlo C DiClemente. Stages and processes of self-change of smoking: toward an integrative model of change. *Journal of Consulting and Clinical Psychology*, 51(3):390–395, 1983. doi: 10.1037/0022-006X.51.3.390.
- [308] James O Prochaska and Wayne F Velicer. The transtheoretical model of health behavior change. *American Journal of Health Promotion*, 12(1):38–48, 1997. doi: 10.4278/0890-1171-12.1.38.

-
- [309] James O Prochaska, Colleen A Redding, and Kerry E Evers. The transtheoretical model and stages of change. In *Health behavior: Theory, Research, and Practice*, pages 125–148. Jossey-Bass/Wiley, 5 edition, 2015.
- [310] Prolific. Who are the participants on prolific? <https://researcher-help.prolific.co/hc/en-gb/articles/360009220833-Who-are-the-participants-on-Prolific>, 2022. Accessed on 20.06.2022.
- [311] Prolific. Who are the participants on prolific? <https://researcher-help.prolific.com/hc/en-gb/articles/360009220833-Who-are-the-participants-on-Prolific>, 2023. Accessed on 30.11.2023.
- [312] GM Azmal Ali Quaosar, Md Rakibul Hoque, and Yukun Bao. Investigating factors affecting elderly’s intention to use m-health services: an empirical study. *Telemedicine and e-Health*, 24(4):309–314, 2018. doi: 10.1089/tmj.2017.0111.
- [313] Mashfiqui Rabbi, Angela Pfammatter, Mi Zhang, Bonnie Spring, and Tanzeem Choudhury. Automated personalized feedback for physical activity and dietary behavior change with mobile phones: a randomized controlled trial on adults. *JMIR mHealth and uHealth*, 3(2):e42, 2015. doi: 10.2196/mhealth.4160.
- [314] Nikita B Rajani, Nikolaos Mastellos, and Filippos T Filippidis. Self-efficacy and motivation to quit of smokers seeking to quit: quantitative assessment of smoking cessation mobile apps. *JMIR mHealth and uHealth*, 9(4):e25030, 2021. doi: 10.2196/25030.
- [315] Margot Rakers, Steven van de Vijver, Paz Bossio, Nic Moens, Michiel Rauws, Millicent Orera, Hongxia Shen, Cynthia Hallensleben, Evelyn Brakema, Nick Guldemon, Niels H Chavannes, and María Villalobos-Quesada. Series: ehealth in primary care. part 6: Global perspectives: Learning from ehealth for low-resource primary care settings and across high-, middle-and low-income countries. *European Journal of General Practice*, 29(1):2241987, 2023. doi: 10.1080/13814788.2023.2241987.
- [316] Simona Ramanauskaitė and Asta Slotkienė. Hierarchy-based competency structure and its application in e-evaluation. *Applied Sciences*, 9(17):3478, 2019. doi: 10.3390/app9173478.
- [317] Hedieh Ranjbartabar, Deborah Richards, Ayse Aysin Bilgin, and Cat Kutay. Do you mind if i ask? addressing the cold start problem in personalised relational agent conversation. In *IVA ’21: ACM International Conference on Intelligent Virtual Agents, Virtual Event, Japan, September 14-17, 2021*, pages 167–174. ACM, 2021. doi: 10.1145/3472306.3478357.
- [318] Fabienne Reiners, Janienke Sturm, Lisette JW Bouw, and Eveline JM Wouters. Sociodemographic factors influencing the use of ehealth in people with chronic diseases. *International Journal of Environmental Research and Public Health*, 16(4):645, 2019. doi: 10.3390/ijerph16040645.

- [319] Melanie Elise Renfrew, Darren Peter Morton, Jason Kyle Morton, Jason Scott Hinze, Geraldine Przybylko, and Bevan Adrian Craig. The influence of three modes of human support on attrition and adherence to a web-and mobile app-based mental health promotion intervention in a nonclinical cohort: Randomized comparative study. *Journal of Medical Internet Research*, 22(9):e19945, 2020. doi: 10.2196/19945.
- [320] Research2Guidance. mhealth economics 2017 – current status and future trends in mobile health. <https://research2guidance.com/product/mhealth-economics-2017-current-status-and-future-trends-in-mobile-health/>, 2017.
- [321] Julia Rheker, Gerhard Andersson, and Cornelia Weise. The role of “on demand” therapist guidance vs. no support in the treatment of tinnitus via the internet: a randomized controlled trial. *Internet Interventions*, 2(2):189–199, 2015. doi: 10.1016/j.invent.2015.03.007.
- [322] Ryan Rhodes and Patrick Boudreau. Physical activity and personality traits. In *Oxford Research Encyclopedia of Psychology*. Oxford University Press, 2017. doi: 10.1093/acrefore/9780190236557.013.210.
- [323] Ryan E Rhodes and NEI Smith. Personality correlates of physical activity: a review and meta-analysis. *British Journal of Sports Medicine*, 40(12):958–965, 2006. doi: 10.1136/bjism.2006.028860.
- [324] Rijksinstituut voor Volksgezondheid en Milieu (RIVM). Volksgezondheid toekomst verkenning 2018: Een gezond vooruitzicht. synthese. <https://www.rivm.nl/publicaties/volksgezondheid-toekomst-verkenning-2018-gezond-vooruitzicht-synthese>, 2018.
- [325] Kara Rogers. Scientific modeling. <https://www.britannica.com/science/scientific-modeling>, 2023. Accessed on 23.08.2024.
- [326] Ronald W Rogers. A protection motivation theory of fear appeals and attitude change. *The Journal of Psychology*, 91(1):93–114, 1975. doi: 10.1080/00223980.1975.9915803.
- [327] Pieter J Rohrbach, Alexandra E Dingemans, Philip Spinhoven, Joost R Van Ginkel, Marjolein Fokkema, Tom F Wilderjans, Stephanie Bauer, and Eric F Van Furth. Effectiveness of an online self-help program, expert-patient support, and their combination for eating disorders: Results from a randomized controlled trial. *International Journal of Eating Disorders*, 55(10):1361–1373, 2022. doi: 10.1002/eat.23785.
- [328] Jeremy Rose and Oskar MacGregor. The architecture of algorithm-driven persuasion. *Journal of Information Architecture*, 6(1), 2021.
- [329] Ariel Rosenfeld and Sarit Kraus. Strategical argumentative agent for human persuasion. In Gal A Kaminka, Maria Fox, Paolo Bouquet, Eyke Hüllermeier, Virginia Dignum, Frank Dignum, and Frank van Harmelen, editors, *ECAI 2016 - 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, The Netherlands - Including Prestigious Applications of Artificial Intelligence (PAIS 2016)*,

- volume 285 of *Frontiers in Artificial Intelligence and Applications*, pages 320–328, 2016. doi: 10.3233/978-1-61499-672-9-320.
- [330] Stéphane Ross and Joelle Pineau. Model-based bayesian reinforcement learning in large structured domains. In David A McAllester and Petri Myllymäki, editors, *UAI 2008, Proceedings of the 24th Conference in Uncertainty in Artificial Intelligence, Helsinki, Finland, July 9-12, 2008*, volume 2008, pages 476–483, Arlington, Virginia, USA, 2008. AUAI Press.
- [331] Sayanti Roy, Christopher Crick, Emily Kieson, and Charles Abramson. A reinforcement learning model for robots as teachers. In *27th IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2018, Nanjing, China, August 27-31, 2018*, pages 294–299. IEEE, 2018. doi: 10.1109/ROMAN.2018.8525563.
- [332] Richard M Ryan and Edward L Deci. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, 55(1):68–78, 2000. doi: 10.1037/0003-066X.55.1.68.
- [333] Soheil Sadeghi, Somit Gupta, Stefan Gramatovici, Jiannan Lu, Hao Ai, and Ruhan Zhang. Novelty and primacy: a long-term estimator for online experiments. *Technometrics*, 64(4):524–534, 2022. doi: 10.1080/00401706.2022.2124309.
- [334] Shaun Scholes and Alison Neave. Health survey for england 2016: Physical activity in adults. <https://discovery.ucl.ac.uk/id/eprint/10040937/1/HSE16-Adult-phy-act.pdf>, 2017.
- [335] Dylan GM Schouten, Rosie T Paulissen, Marieke Hanekamp, Annemarie Groot, Mark A Neerinx, and Anita HM Cremers. Low-literates’ support needs for societal participation learning: empirical grounding of theory-and model-based design. *Cognitive Systems Research*, 45:30–47, 2017. doi: 10.1016/j.cogsys.2017.04.007.
- [336] Steven A Schroeder. We can do better—improving the health of the american people. *New England Journal of Medicine*, 357(12):1221–1228, 2007. doi: 10.1056/NEJMsa073350.
- [337] Stephen M Schueller, Kathryn Noth Tomasino, and David C Mohr. Integrating human support into behavioral intervention technologies: The efficiency model of support. *Clinical Psychology: Science and Practice*, 24(1):27–45, 2017. doi: 10.1111/cpsp.12173.
- [338] Andreas Richard Schwerdtfeger, Catalina Schmitz, and Matthias Warken. Using text messages to bridge the intention-behavior gap? a pilot study on the use of text message reminders to increase objectively assessed physical activity in daily life. *Frontiers in Psychology*, 3:270, 2012. doi: 10.3389/fpsyg.2012.00270.
- [339] Muzafer Sherif and Carl I Hovland. *Social judgement: Assimilation and contrast effects in communication and attitude change*. Yale University Press, 1961.
- [340] Donghee Shin, Bu Zhong, and Frank A Biocca. Beyond user experience: What constitutes algorithmic experiences? *International Journal of Information Management*, 52:102061, 2020. ISSN 0268-4012. doi: 10.1016/j.ijinfomgt.2019.102061.

- [341] Lutz Siemer, Marjolein GJ Brusse-Keizer, Marloes G Postel, Somaya Ben Allouch, Robbert Sanderman, and Marcel E Pieterse. Adherence to blended or face-to-face smoking cessation treatment and predictors of adherence: randomized controlled trial. *Journal of Medical Internet Research*, 22(7):e17207, 2020. doi: 10.2196/17207.
- [342] Luuk PA Simons, Wouter AC van den Heuvel, and Catholijn M Jonker. ehealth whatsapp group for social support: Preliminary results. In Andreja Pucihar, Mirjana Kljajic Borstnar, Pascal Ravesteijn, Juergen Seitz, and Roger W H Bons, editors, *31st Bled eConference: Digital Transformation: Meeting the Challenges, Bled, Slovenia, June 17-20, 2018*, page 28, 2018.
- [343] Ben Singh, Mavra Ahmed, Amanda E Staiano, Claire Gough, Jasmine Petersen, Corneel Vandelanotte, Chelsea Kracht, Christopher Huong, Zenong Yin, Maria F Vasiloglou, Chen-Chia Pan, Camille E Short, Matthew Mclaughlin, Lauren von Klinggraeff, Christopher D Pfledderer, Lisa J Moran, Alyssa M Button, and Carol A Maher. A systematic umbrella review and meta-meta-analysis of ehealth and mhealth interventions for improving lifestyle behaviours. *npj Digital Medicine*, 7(1):179, 2024. doi: 10.1038/s41746-024-01172-y.
- [344] Ronal Singh, Tim Miller, Henrietta Lyons, Liz Sonenberg, Eduardo Velloso, Frank Vetere, Piers Howe, and Paul Dourish. Directive explanations for actionable explainability in machine learning applications. *ACM Transactions on Interactive Intelligent Systems*, 13(4):1–26, 2023. doi: 10.1145/3579363.
- [345] Claudia Sinoo, Sylvia van Der Pal, Olivier A Blanson Henkemans, Anouk Keizer, Bert PB Bierman, Rosemarijn Looije, and Mark A Neerinx. Friendship with a robot: Children’s perception of similarity between a robot’s physical and virtual embodiment that supports diabetes self-management. *Patient Education and Counseling*, 101(7):1248–1255, 2018. doi: 10.1016/j.pec.2018.02.008.
- [346] Falko F Sniehotta, Urte Scholz, Ralf Schwarzer, Bärbel Fuhrmann, Ulrich Kiwus, and Heinz Völler. Long-term effects of two psychological interventions on physical exercise and self-regulation following coronary rehabilitation. *International Journal of Behavioral Medicine*, 12(4):244–255, 2005. doi: 10.1207/s15327558ijbm1204_5.
- [347] Falko F Sniehotta, Ralf Schwarzer, Urte Scholz, and Benjamin Schüz. Action planning and coping planning for long-term lifestyle change: theory and assessment. *European Journal of Social Psychology*, 35(4):565–576, 2005. doi: 10.1002/ejsp.258.
- [348] Ingrid Konstane Ledel Solem, Cecilie Varsi, Hilde Eide, Olöf Birna Kristjansdottir, Jelena Mirkovic, Elin Børøsund, Mette Haaland-Øverby, Karina Heldal, Karlein Mg Schreurs, Lori B Waxenberg, Karen Elizabeth Weiss, Eleshia J Morrison, and Lise Solberg Nes. Patients’ needs and requirements for ehealth pain management interventions: Qualitative study. *Journal of Medical Internet Research*, 21(4):e13205, 2019. doi: 10.2196/13205.
- [349] Hanne AA Spelt, Chao Zhang, Joyce HDM Westerink, Jaap Ham, and Wijnand A IJsselsteijn. Persuasion-induced physiology as predictor of persuasion effectiveness.

-
- IEEE Transactions on Affective Computing*, 13(3):1593–1604, 2020. doi: 10.1109/TAFFC.2020.3022109.
- [350] James P Stevens. *Applied Multivariate Statistics for the Social Sciences*. Routledge, 5 edition, 2009.
- [351] Wayne T Steward, Tamera R Schneider, Judith Pizarro, and Peter Salovey. Need for cognition moderates responses to framed smoking-cessation messages. *Journal of Applied Social Psychology*, 33(12):2439–2464, 2003. doi: 10.1111/j.1559-1816.2003.tb02775.x.
- [352] Stichting Stop Bewust. URL <https://stichtingstopbewust.nl/>. Accessed in 2023.
- [353] Victor J Strecher, Brenda McEvoy DeVellis, Marshall H Becker, and Irwin M Rosenstock. The role of self-efficacy in achieving health behavior change. *Health Education Quarterly*, 13(1):73–92, 1986. doi: 10.1177/109019818601300108.
- [354] Laura L Struik, Joan L Botorff, Neill Bruce Baskerville, and John L Oliffe. The crush the crave quit smoking app and young adult smokers: qualitative case study of affordances. *JMIR mHealth and uHealth*, 6(6):e134, 2018. doi: 10.2196/mhealth.9489.
- [355] Yongjun Sung and Sejung Marina Choi. Increasing power and preventing pain: The moderating role of self-construal in advertising message framing. *Journal of Advertising*, 40(1):71–86, 2011. doi: 10.2753/JOA0091-3367400105.
- [356] Jeremy Sutton. What is a thought diary in cbt? 5 templates and examples, 2021. URL <https://positivepsychology.com/thought-diary/>. Accessed in 2023.
- [357] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. Bradford Books, 2018.
- [358] Antonio Terracciano and Paul T Costa Jr. Smoking and the five-factor model of personality. *Addiction*, 99(4):472–481, 2004. doi: 10.1111/j.1360-0443.2004.00687.x.
- [359] Rosemary Josekutty Thomas, Judith Masthoff, and Nir Oren. Personalising healthy eating messages to age, gender and personality: using cialdini’s principles and framing. In George A Papadopoulos, Tsvi Kuflik, Fang Chen, Carlos Duarte, and Wai-Tat Fu, editors, *Companion Publication of the 22nd International Conference on Intelligent User Interfaces, IUI 2017, Limassol, Cyprus, March 13-16, 2017*, pages 81–84. ACM, 2017. doi: 10.1145/3030024.3040986.
- [360] Rosemary Josekutty Thomas, Judith Masthoff, and Nir Oren. Adapting healthy eating messages to personality. In Peter W de Vries, Harri Oinas-Kukkonen, Liseth Siemons, Nienke Beerlage-de Jong, and Lisette van Gemert-Pijnen, editors, *Persuasive Technology: Development and Implementation of Personalized Technologies to Change Attitudes and Behaviors - 12th International Conference, PERSUASIVE 2017, Amsterdam, The Netherlands, April 4-6, 2017, Proceedings*, volume 10171 of *Lecture Notes in*

- Computer Science*, pages 119–132, Cham, 2017. Springer International Publishing. doi: 10.1007/978-3-319-55134-0_10.
- [361] Emma M Thompson, Louise Destree, Lucy Albertella, and Leonardo F Fontenelle. Internet-based acceptance and commitment therapy: a transdiagnostic systematic review and meta-analysis for mental health outcomes. *Behavior Therapy*, 52(2): 492–507, 2021. doi: 10.1016/j.beth.2020.07.002.
- [362] Myrthe L Tielman, Mark A Neerinx, and Willem-Paul Brinkman. Design and evaluation of personalized motivational messages by a virtual agent that assists in post-traumatic stress disorder therapy. *Journal of Medical Internet Research*, 21(3): e9240, 2019. doi: 10.2196/jmir.9240.
- [363] Myrthe L Tielman, Mark A Neerinx, Claudia Pagliari, Albert Rizzo, and Willem-Paul Brinkman. Considering patient safety in autonomous e-mental health systems—detecting risk situations and referring patients back to human care. *BMC Medical Informatics and Decision Making*, 19(1):47, 2019. doi: 10.1186/s12911-019-0796-x.
- [364] Stefan Timmermans and Rebecca Kaufman. Technologies and health inequities. *Annual Review of Sociology*, 46:583–602, 2020. doi: 10.1146/annurev-soc-121919-054802.
- [365] Nickolai Titov, Gavin Andrews, Matthew Davies, Karen McIntyre, Emma Robinson, and Karen Solley. Internet treatment for depression: a randomized controlled trial comparing clinician vs. technician assistance. *PLoS ONE*, 5(6):e10939, 2010. doi: 10.1371/journal.pone.0010939.
- [366] Anna L Trella, Kelly W Zhang, Inbal Nahum-Shani, Vivek Shetty, Finale Doshi-Velez, and Susan A Murphy. Designing reinforcement learning algorithms for digital interventions: pre-implementation guidelines. *Algorithms*, 15:255, 2022. doi: 10.3390/a15080255.
- [367] Anna L Trella, Kelly W Zhang, Inbal Nahum-Shani, Vivek Shetty, Finale Doshi-Velez, and Susan A Murphy. Reward design for an online reinforcement learning algorithm supporting oral self-care. In Brian Williams, Yiling Chen, and Jennifer Neville, editors, *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, pages 15724–15730. AAAI Press, 2023. doi: 10.1609/AAAI.V37I13.26866.
- [368] Trimbos Instituut. Richtlijn behandeling van tabaksverslaving en stoppen met roken ondersteuning: Herziening 2016. <https://www.trimbos.nl/docs/7435d8f9-c8f4-4302-aec6-f4fddade81f2.pdf>, 2016.
- [369] Trimbos Instituut. De stopcoach, 2023. URL <https://www.trimbos.nl/aanbod/interventies/de-stopcoach/>.

-
- [370] Anne M Turner, Thomas Engelsma, Jean O Taylor, Rashmi K Sharma, and George Demiris. Recruiting older adult participants through crowdsourcing platforms: Mechanical turk versus prolific academic. In *AMIA Annual Symposium Proceedings*, volume 2020, pages 1230–1238, 2020.
- [371] John A Updegraff and Alexander J Rothman. Health message framing: Moderators, mediators, and mysteries. *Social and Personality Psychology Compass*, 7(9):668–679, 2013. doi: 10.1111/spc3.12056.
- [372] Rens van de Schoot, Sarah Depaoli, Ruth King, Bianca Kramer, Kaspar Märtens, Mahlet G Tadesse, Marina Vannucci, Andrew Gelman, Duco Veen, Joukje Willemsen, and Christopher Yau. Bayesian statistics and modelling. *Nature Reviews Methods Primers*, 1(1):1, 2021. doi: 10.1038/s43586-020-00001-2.
- [373] Julia EWC van Gemert-Pijnen, Nicol Nijland, Maarten van Limburg, Hans C Ossebaard, Saskia M Kelders, Gunther Eysenbach, and Erwin R Seydel. A holistic framework to improve the uptake and impact of ehealth technologies. *Journal of Medical Internet Research*, 13(4), 2011. doi: 10.2196/jmir.1672.
- [374] Milon HM van Vliet, Anke Versluis, Niels H Chavannes, Bouke L Scheltinga, Nele Albers, Kristell M Penfornis, Walter Baccinelli, Eline Meijer, and on behalf of the Perfect Fit consortium. Protocol of a mixed-methods evaluation of perfect fit: A personalized mhealth intervention with a virtual coach to promote smoking cessation and physical activity in adults. *Digital Health*, 10:20552076241300020, 2024. doi: 10.1177/20552076241300020.
- [375] Simine Vazire and Erika N Carlson. Self-knowledge of personality: Do people know themselves? *Social and Personality Psychology Compass*, 4(8):605–620, 2010. doi: 10.1111/j.1751-9004.2010.00280.x.
- [376] Viswanath Venkatesh, Michael G Morris, Gordon B Davis, and Fred D Davis. User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3): 425–478, 2003.
- [377] Viswanath Venkatesh, James YL Thong, and Xin Xu. Consumer acceptance and use of information technology: extending the unified theory of acceptance and use of technology. *MIS Quarterly*, 36(1):157–178, 2012. doi: 10.2307/41410412.
- [378] Johannes William Vergouw, Hanneke Smits-Pelser, Marijke C Kars, Thijs van Houwelingen, Harmieke van Os-Medendorp, Helianthe Kort, and Nienke Bleijenberg. Needs, barriers and facilitators of older adults towards ehealth in general practice: a qualitative study. *Primary Health Care Research & Development*, 21:e54, 2020. doi: 10.1017/S1463423620000547.
- [379] Damon J Vidrine, Faith E Fletcher, Heather E Danysh, Salma Marani, Jennifer Irvin Vidrine, Scott B Cantor, and Alexander V Prokhorov. A randomized controlled trial to assess the efficacy of an interactive mobile messaging intervention for underserved smokers: Project action. *BMC Public Health*, 12(1):696, 2012. doi: 10.1186/1471-2458-12-696.

- [380] Voedingscentrum. Zo maak je een eigen eetwissel. URL <https://www.voedingscentrum.nl/nl/thema/eetwissel/zo-maak-je-een-eigen-eetwissel.aspx>. Accessed in 2023.
- [381] Shihan Wang, Chao Zhang, Ben Kröse, and Herke van Hoof. Optimizing adaptive notifications in mobile health interventions systems: reinforcement learning from a data-driven behavioral simulator. *Journal of Medical Systems*, 45(12):102, 2021. doi: 10.1007/s10916-021-01773-0.
- [382] Wenxin Wang, Céline L van Lint, Willem-Paul Brinkman, Ton JM Rövekamp, Sandra van Dijk, Paul van der Boog, and Mark A Neerincx. Guided or factual computer support for kidney patients with different experience levels and medical health situations: preferences and usage. *Health and Technology*, 9(3):329–342, 2019. doi: 10.1007/s12553-019-00295-7.
- [383] Klaus Weber, Niklas Rach, Wolfgang Minker, and Elisabeth André. How to win arguments. *Datenbank-Spektrum*, 20(2):161–169, 2020. doi: 10.1007/s13222-020-00345-9.
- [384] Thure Georg Weimann and Carola Gißke. Unleashing the potential of reinforcement learning for personalizing behavioral transformations with digital therapeutics: A systematic literature review. In Hannes Schlieter, Ana L N Fred, and Hugo Gamboa, editors, *Proceedings of the 17th International Joint Conference on Biomedical Engineering Systems and Technologies, BIOSTEC 2024, Volume 2, Rome, Italy, February 21-23, 2024*, pages 230–245, 2024. doi: 10.5220/0012474700003657.
- [385] Neil D Weinstein. Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, 39(5):806–820, 1980. doi: 10.1037/0022-3514.39.5.806.
- [386] Robert West and Jamie Brown. *Theory of Addiction*. John Wiley & Sons, 2013.
- [387] Robert West and Susan Michie. A brief introduction to the com-b model of behaviour and the prime theory of motivation [v1]. *Qeios*, 2020. doi: 10.32388/WW04E6.2.
- [388] Robyn Whittaker, Enid Dorey, Dale Bramley, Chris Bullen, Simon Denny, C Raina Elley, Ralph Maddison, Hayden McRobbie, Varsha Parag, Anthony Rodgers, and Penny Salmon. A theory-based video messaging mobile phone intervention for smoking cessation: randomized controlled trial. *Journal of Medical Internet Research*, 13(1):e10, 2011. doi: 10.2196/jmir.1553.
- [389] Marc C Willemsen, Hein de Vries, Gerard van Breukelen, and Brian Oldenburg. Determinants of intention to quit smoking among dutch employees: the influence of the social environment. *Preventive Medicine*, 25(2):195–202, 1996. doi: 10.1006/pmed.1996.0046.
- [390] Geoffrey G Williams, Marylene Gagné, Richard M Ryan, and Edward L Deci. Facilitating autonomous motivation for smoking cessation. *Health Psychology*, 21(1):40–50, 2002.

-
- [391] Jessica Wilson, Milena Heinsch, David Betts, Debbie Booth, and Frances Kay-Lambkin. Barriers and facilitators to the use of e-health by older adults: a scoping review. *BMC Public Health*, 21(1):1556, 2021. doi: 10.1186/s12889-021-11623-w.
- [392] Charlotte E Wittekind, Ansgar Feist, Brooke C Schneider, Steffen Moritz, and Anja Fritzsche. The approach-avoidance task as an online intervention in cigarette smoking: a pilot study. *Journal of Behavior Therapy and Experimental Psychiatry*, 46: 115–120, 2015. ISSN 0005-7916. doi: 10.1016/j.jbtep.2014.08.006.
- [393] Emily Wolstenholme, Valentina Carfora, Patrizia Catellani, Wouter Poortinga, and Lorraine Whitmarsh. Explaining intention to reduce red and processed meat in the uk and italy using the theory of planned behaviour, meat-eater identity, and the transtheoretical model. *Appetite*, 166:105467, 2021. doi: 10.1016/j.appet.2021.105467.
- [394] Lan Xiao, Veronica Yank, and Jun Ma. Algorithm for balancing both continuous and categorical covariates in randomized controlled trials. *Computer Methods and Programs in Biomedicine*, 108(3):1185–1190, 2012. doi: 10.1016/j.cmpb.2012.06.001.
- [395] Min-Jeong Yang, Steven K Sutton, Laura M Hernandez, Sarah R Jones, David W Wetter, Santosh Kumar, and Christine Vinci. A just-in-time adaptive intervention (jitai) for smoking cessation: Feasibility and acceptability findings. *Addictive Behaviors*, 136:107467, 2023. doi: 10.1016/j.addbeh.2022.107467.
- [396] Lucy Yardley, Bonnie J Spring, Heleen Riper, Leanne G Morrison, David H Crane, Kristina Curtis, Gina C Merchant, Felix Naughton, and Ann Blandford. Understanding and promoting effective engagement with digital behavior change interventions. *American Journal of Preventive Medicine*, 51(5):833–842, 2016. doi: 10.1016/j.amepre.2016.06.015.
- [397] Yuan Yuan, Fengli Xu, Hancheng Cao, Guozhen Zhang, Pan Hui, Yong Li, and Depeng Jin. Persuade to click: Context-aware persuasion model for online textual advertisement. *IEEE Transactions on Knowledge and Data Engineering*, 35(2):1938–1951, 2021. doi: 10.1109/TKDE.2021.3110724.
- [398] Mohan Zalake, Alexandre Gomes de Siqueira, Krishna Vaddiparti, and Benjamin Lok. The effects of virtual human’s verbal persuasion strategies on user intention and behavior. *International Journal of Human-Computer Studies*, 156:102708, 2021. doi: 10.1016/J.IJHCS.2021.102708.
- [399] Chao Zhang, Joaquin Vanschoren, Arlette van Wissen, Daniël Lakens, Boris de Ruyter, and Wijnand A IJsselstein. Theory-based habit modeling for enhancing behavior prediction in behavior change support systems. *User Modeling and User-Adapted Interaction*, 32(3):389–415, 2022. doi: 10.1007/s11257-022-09326-x.
- [400] Jingwen Zhang, Yoo Jung Oh, Patrick Lange, Zhou Yu, and Yoshimi Fukuoka. Artificial intelligence chatbot behavior change model for designing artificial intelligence chatbots to promote physical activity and a healthy diet. *Journal of Medical Internet Research*, 22(9):e22845, 2020. doi: 10.2196/22845.

-
- [401] Kelly W Zhang, Nowell Closser, Anna L Trella, and Susan A Murphy. Replicable bandits for digital health interventions. *arXiv preprint arXiv:2407.15377*, 2024.
- [402] Mo Zhou, Yonatan Mintz, Yoshimi Fukuoka, Ken Goldberg, Elena Flowers, Philip Kaminsky, Alejandro Castillejo, and Anil Aswani. Personalizing mobile fitness apps using reinforcement learning. *CEUR workshop proceedings*, 2068, 2018.
- [403] ZonMw. E-health. <https://www.zonmw.nl/en/e-health>. Accessed on 22.07.2024.

SIKS DISSERTATIONS

- 2016 01 Syed Saiden Abbas (RUN), Recognition of Shapes by Humans and Machines
- 02 Michiel Christiaan Meulendijk (UU), Optimizing medication reviews through decision support: prescribing a better pill to swallow
- 03 Maya Sappelli (RUN), Knowledge Work in Context: User Centered Knowledge Worker Support
- 04 Laurens Rietveld (VUA), Publishing and Consuming Linked Data
- 05 Evgeny Sherkhonov (UvA), Expanded Acyclic Queries: Containment and an Application in Explaining Missing Answers
- 06 Michel Wilson (TUD), Robust scheduling in an uncertain environment
- 07 Jeroen de Man (VUA), Measuring and modeling negative emotions for virtual training
- 08 Matje van de Camp (TiU), A Link to the Past: Constructing Historical Social Networks from Unstructured Data
- 09 Archana Nottamkandath (VUA), Trusting Crowdsourced Information on Cultural Artefacts
- 10 George Karafotias (VUA), Parameter Control for Evolutionary Algorithms
- 11 Anne Schuth (UvA), Search Engines that Learn from Their Users
- 12 Max Knobbout (UU), Logics for Modelling and Verifying Normative Multi-Agent Systems
- 13 Nana Baah Gyan (VUA), The Web, Speech Technologies and Rural Development in West Africa - An ICT4D Approach
- 14 Ravi Khadka (UU), Revisiting Legacy Software System Modernization
- 15 Steffen Michels (RUN), Hybrid Probabilistic Logics - Theoretical Aspects, Algorithms and Experiments
- 16 Guangliang Li (UvA), Socially Intelligent Autonomous Agents that Learn from Human Reward
- 17 Berend Weel (VUA), Towards Embodied Evolution of Robot Organisms
- 18 Albert Meroño Peñuela (VUA), Refining Statistical Data on the Web
- 19 Julia Efremova (TU/e), Mining Social Structures from Genealogical Data
- 20 Daan Odijk (UvA), Context & Semantics in News & Web Search
- 21 Alejandro Moreno Celleri (UT), From Traditional to Interactive Playspaces: Automatic Analysis of Player Behavior in the Interactive Tag Playground
- 22 Grace Lewis (VUA), Software Architecture Strategies for Cyber-Foraging Systems
- 23 Fei Cai (UvA), Query Auto Completion in Information Retrieval
- 24 Brend Wanders (UT), Repurposing and Probabilistic Integration of Data; An Iterative and data model independent approach
- 25 Julia Kiseleva (TU/e), Using Contextual Information to Understand Searching and Browsing Behavior
- 26 Dilhan Thilakarathne (VUA), In or Out of Control: Exploring Computational Models to Study the Role of Human Awareness and Control in Behavioural Choices, with Applications in Aviation and Energy Management Domains
- 27 Wen Li (TUD), Understanding Geo-spatial Information on Social Media
- 28 Mingxin Zhang (TUD), Large-scale Agent-based Social Simulation - A study on epidemic prediction and control
- 29 Nicolas Höning (TUD), Peak reduction in decentralised electricity systems - Markets and prices for flexible planning
- 30 Ruud Mattheij (TiU), The Eyes Have It
- 31 Mohammad Khelghati (UT), Deep web content monitoring
- 32 Eelco Vriezolkolk (UT), Assessing Telecommunication Service Availability Risks for Crisis Organisations
- 33 Peter Bloem (UvA), Single Sample Statistics, exercises in learning from just one example
- 34 Dennis Schunselaar (TU/e), Configurable Process Trees: Elicitation, Analysis, and Enactment
- 35 Zhaochun Ren (UvA), Monitoring Social Media: Summarization, Classification and Recommendation
- 36 Daphne Karreman (UT), Beyond R2D2: The design of nonverbal interaction behavior optimized for robot-specific morphologies

-
- 37 Giovanni Sileno (UvA), Aligning Law and Action - a conceptual and computational inquiry
 38 Andrea Minuto (UT), Materials that Matter - Smart Materials meet Art & Interaction Design
 39 Merijn Bruijnes (UT), Believable Suspect Agents; Response and Interpersonal Style Selection for an Artificial Suspect
 40 Christian Detweiler (TUD), Accounting for Values in Design
 41 Thomas King (TUD), Governing Governance: A Formal Framework for Analysing Institutional Design and Enactment Governance
 42 Spyros Martzoukos (UvA), Combinatorial and Compositional Aspects of Bilingual Aligned Corpora
 43 Saskia Koldijk (RUN), Context-Aware Support for Stress Self-Management: From Theory to Practice
 44 Thibault Sellam (UvA), Automatic Assistants for Database Exploration
 45 Bram van de Laar (UT), Experiencing Brain-Computer Interface Control
 46 Jorge Gallego Perez (UT), Robots to Make you Happy
 47 Christina Weber (UL), Real-time foresight - Preparedness for dynamic innovation networks
 48 Tanja Buttler (TUD), Collecting Lessons Learned
 49 Gleb Polevoy (TUD), Participation and Interaction in Projects. A Game-Theoretic Analysis
 50 Yan Wang (TiU), The Bridge of Dreams: Towards a Method for Operational Performance Alignment in IT-enabled Service Supply Chains
-
- 2017 01 Jan-Jaap Oerlemans (UL), Investigating Cybercrime
 02 Sjoerd Timmer (UU), Designing and Understanding Forensic Bayesian Networks using Argumentation
 03 Daniël Harold Telgen (UU), Grid Manufacturing; A Cyber-Physical Approach with Autonomous Products and Reconfigurable Manufacturing Machines
 04 Mrunal Gawade (CWI), Multi-core Parallelism in a Column-store
 05 Mahdieh Shadi (UvA), Collaboration Behavior
 06 Damir Vandic (EUR), Intelligent Information Systems for Web Product Search
 07 Roel Bertens (UU), Insight in Information: from Abstract to Anomaly
 08 Rob Konijn (VUA), Detecting Interesting Differences: Data Mining in Health Insurance Data using Outlier Detection and Subgroup Discovery
 09 Dong Nguyen (UT), Text as Social and Cultural Data: A Computational Perspective on Variation in Text
 10 Robby van Delden (UT), (Steering) Interactive Play Behavior
 11 Florian Kunneman (RUN), Modelling patterns of time and emotion in Twitter #anticipointment
 12 Sander Leemans (TU/e), Robust Process Mining with Guarantees
 13 Gijs Huisman (UT), Social Touch Technology - Extending the reach of social touch through haptic technology
 14 Shoshannah Tekofsky (TiU), You Are Who You Play You Are: Modelling Player Traits from Video Game Behavior
 15 Peter Berck (RUN), Memory-Based Text Correction
 16 Aleksandr Chuklin (UvA), Understanding and Modeling Users of Modern Search Engines
 17 Daniel Dimov (UL), Crowdsourced Online Dispute Resolution
 18 Ridho Reinanda (UvA), Entity Associations for Search
 19 Jeroen Vuurens (UT), Proximity of Terms, Texts and Semantic Vectors in Information Retrieval
 20 Mohammadbashir Sedighi (TUD), Fostering Engagement in Knowledge Sharing: The Role of Perceived Benefits, Costs and Visibility
 21 Jeroen Linssen (UT), Meta Matters in Interactive Storytelling and Serious Gaming (A Play on Worlds)
 22 Sara Magliacane (VUA), Logics for causal inference under uncertainty
 23 David Graus (UvA), Entities of Interest – Discovery in Digital Traces
 24 Chang Wang (TUD), Use of Affordances for Efficient Robot Learning
 25 Veruska Zamborlini (VUA), Knowledge Representation for Clinical Guidelines, with applications to Multimorbidity Analysis and Literature Search
 26 Merel Jung (UT), Socially intelligent robots that understand and respond to human touch
 27 Michiel Joosse (UT), Investigating Positioning and Gaze Behaviors of Social Robots: People's Preferences, Perceptions and Behaviors
 28 John Klein (VUA), Architecture Practices for Complex Contexts
 29 Adel Alhuraibi (TiU), From IT-Business Strategic Alignment to Performance: A Moderated Mediation Model of Social Innovation, and Enterprise Governance of IT"
 30 Wilma Latuny (TiU), The Power of Facial Expressions

-
- 31 Ben Ruijl (UL), Advances in computational methods for QFT calculations
- 32 Thajer Samar (RUN), Access to and Retrieval of Content in Web Archives
- 33 Brigit van Loggem (OU), Towards a Design Rationale for Software Documentation: A Model of Computer-Mediated Activity
- 34 Maren Scheffel (OU), The Evaluation Framework for Learning Analytics
- 35 Martine de Vos (VUA), Interpreting natural science spreadsheets
- 36 Yuanhao Guo (UL), Shape Analysis for Phenotype Characterisation from High-throughput Imaging
- 37 Alejandro Montes Garcia (TU/e), WiBAF: A Within Browser Adaptation Framework that Enables Control over Privacy
- 38 Alex Kayal (TUD), Normative Social Applications
- 39 Sara Ahmadi (RUN), Exploiting properties of the human auditory system and compressive sensing methods to increase noise robustness in ASR
- 40 Altaf Hussain Abro (VUA), Steer your Mind: Computational Exploration of Human Control in Relation to Emotions, Desires and Social Support For applications in human-aware support systems
- 41 Adnan Manzoor (VUA), Minding a Healthy Lifestyle: An Exploration of Mental Processes and a Smart Environment to Provide Support for a Healthy Lifestyle
- 42 Elena Sokolova (RUN), Causal discovery from mixed and missing data with applications on ADHD datasets
- 43 Maaïke de Boer (RUN), Semantic Mapping in Video Retrieval
- 44 Garm Lucassen (UU), Understanding User Stories - Computational Linguistics in Agile Requirements Engineering
- 45 Bas Testerink (UU), Decentralized Runtime Norm Enforcement
- 46 Jan Schneider (OU), Sensor-based Learning Support
- 47 Jie Yang (TUD), Crowd Knowledge Creation Acceleration
- 48 Angel Suarez (OU), Collaborative inquiry-based learning
-
- 2018 01 Han van der Aa (VUA), Comparing and Aligning Process Representations
- 02 Felix Mannhardt (TU/e), Multi-perspective Process Mining
- 03 Steven Bosems (UT), Causal Models For Well-Being: Knowledge Modeling, Model-Driven Development of Context-Aware Applications, and Behavior Prediction
- 04 Jordan Janeiro (TUD), Flexible Coordination Support for Diagnosis Teams in Data-Centric Engineering Tasks
- 05 Hugo Huurdeman (UvA), Supporting the Complex Dynamics of the Information Seeking Process
- 06 Dan Ionita (UT), Model-Driven Information Security Risk Assessment of Socio-Technical Systems
- 07 Jieting Luo (UU), A formal account of opportunism in multi-agent systems
- 08 Rick Smetsers (RUN), Advances in Model Learning for Software Systems
- 09 Xu Xie (TUD), Data Assimilation in Discrete Event Simulations
- 10 Julienka Mollee (VUA), Moving forward: supporting physical activity behavior change through intelligent technology
- 11 Mahdi Sargolzaei (UvA), Enabling Framework for Service-oriented Collaborative Networks
- 12 Xixi Lu (TU/e), Using behavioral context in process mining
- 13 Seyed Amin Tabatabaei (VUA), Computing a Sustainable Future
- 14 Bart Joosten (TiU), Detecting Social Signals with Spatiotemporal Gabor Filters
- 15 Naser Davarzani (UM), Biomarker discovery in heart failure
- 16 Jaebok Kim (UT), Automatic recognition of engagement and emotion in a group of children
- 17 Jianpeng Zhang (TU/e), On Graph Sample Clustering
- 18 Henriette Nakad (UL), De Notaris en Private Rechtspraak
- 19 Minh Duc Pham (VUA), Emergent relational schemas for RDF
- 20 Manxia Liu (RUN), Time and Bayesian Networks
- 21 Aad Slotmaker (OU), EMERGO: a generic platform for authoring and playing scenario-based serious games
- 22 Eric Fernandes de Mello Araújo (VUA), Contagious: Modeling the Spread of Behaviours, Perceptions and Emotions in Social Networks
- 23 Kim Schouten (EUR), Semantics-driven Aspect-Based Sentiment Analysis
- 24 Jered Vroon (UT), Responsive Social Positioning Behaviour for Semi-Autonomous Telepresence Robots
- 25 Riste Gligorov (VUA), Serious Games in Audio-Visual Collections
- 26 Roelof Anne Jelle de Vries (UT), Theory-Based and Tailor-Made: Motivational Messages for Behavior Change Technology

-
- 27 Maikel Leemans (TU/e), Hierarchical Process Mining for Scalable Software Analysis
- 28 Christian Willems (UT), Social Touch Technologies: How they feel and how they make you feel
- 29 Yu Gu (TiU), Emotion Recognition from Mandarin Speech
- 30 Wouter Beek (VUA), The "K" in "semantic web" stands for "knowledge": scaling semantics to the web
-
- 2019 01 Rob van Eijk (UL), Web privacy measurement in real-time bidding systems. A graph-based approach to RTB system classification
- 02 Emmanuelle Beauxis Aussalet (CWI, UU), Statistics and Visualizations for Assessing Class Size Uncertainty
- 03 Eduardo Gonzalez Lopez de Murillas (TU/e), Process Mining on Databases: Extracting Event Data from Real Life Data Sources
- 04 Ridho Rahmadi (RUN), Finding stable causal structures from clinical data
- 05 Sebastiaan van Zelst (TU/e), Process Mining with Streaming Data
- 06 Chris Dijkshoorn (VUA), Nichesourcing for Improving Access to Linked Cultural Heritage Datasets
- 07 Soude Fazeli (TUD), Recommender Systems in Social Learning Platforms
- 08 Frits de Nijs (TUD), Resource-constrained Multi-agent Markov Decision Processes
- 09 Fahimeh Alizadeh Moghaddam (UvA), Self-adaptation for energy efficiency in software systems
- 10 Qing Chuan Ye (EUR), Multi-objective Optimization Methods for Allocation and Prediction
- 11 Yue Zhao (TUD), Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs
- 12 Jacqueline Heinerman (VUA), Better Together
- 13 Guanliang Chen (TUD), MOOC Analytics: Learner Modeling and Content Generation
- 14 Daniel Davis (TUD), Large-Scale Learning Analytics: Modeling Learner Behavior & Improving Learning Outcomes in Massive Open Online Courses
- 15 Erwin Walraven (TUD), Planning under Uncertainty in Constrained and Partially Observable Environments
- 16 Guangming Li (TU/e), Process Mining based on Object-Centric Behavioral Constraint (OCBC) Models
- 17 Ali Hurriyetoglu (RUN), Extracting actionable information from microtexts
- 18 Gerard Wagenaar (UU), Artefacts in Agile Team Communication
- 19 Vincent Koeman (TUD), Tools for Developing Cognitive Agents
- 20 Chide Groenouwe (UU), Fostering technically augmented human collective intelligence
- 21 Cong Liu (TU/e), Software Data Analytics: Architectural Model Discovery and Design Pattern Detection
- 22 Martin van den Berg (VUA), Improving IT Decisions with Enterprise Architecture
- 23 Qin Liu (TUD), Intelligent Control Systems: Learning, Interpreting, Verification
- 24 Anca Dumitrache (VUA), Truth in Disagreement - Crowdsourcing Labeled Data for Natural Language Processing
- 25 Emiel van Miltenburg (VUA), Pragmatic factors in (automatic) image description
- 26 Prince Singh (UT), An Integration Platform for Synchromodal Transport
- 27 Alessandra Antonaci (OU), The Gamification Design Process applied to (Massive) Open Online Courses
- 28 Esther Kuindersma (UL), Cleared for take-off: Game-based learning to prepare airline pilots for critical situations
- 29 Daniel Formolo (VUA), Using virtual agents for simulation and training of social skills in safety-critical circumstances
- 30 Vahid Yazdanpanah (UT), Multiagent Industrial Symbiosis Systems
- 31 Milan Jelisavcic (VUA), Alive and Kicking: Baby Steps in Robotics
- 32 Chiara Sironi (UM), Monte-Carlo Tree Search for Artificial General Intelligence in Games
- 33 Anil Yaman (TU/e), Evolution of Biologically Inspired Learning in Artificial Neural Networks
- 34 Negar Ahmadi (TU/e), EEG Microstate and Functional Brain Network Features for Classification of Epilepsy and PNES
- 35 Lisa Facey-Shaw (OU), Gamification with digital badges in learning programming
- 36 Kevin Ackermans (OU), Designing Video-Enhanced Rubrics to Master Complex Skills
- 37 Jian Fang (TUD), Database Acceleration on FPGAs
- 38 Akos Kadar (OU), Learning visually grounded and multilingual representations
-
- 2020 01 Armon Toubman (UL), Calculated Moves: Generating Air Combat Behaviour
- 02 Marcos de Paula Bueno (UL), Unraveling Temporal Processes using Probabilistic Graphical Models
- 03 Mostafa Deghani (UvA), Learning with Imperfect Supervision for Language Understanding

-
- 04 Maarten van Gompel (RUN), Context as Linguistic Bridges
- 05 Yulong Pei (TU/e), On local and global structure mining
- 06 Preethu Rose Anish (UT), Stimulation Architectural Thinking during Requirements Elicitation - An Approach and Tool Support
- 07 Wim van der Vegt (OU), Towards a software architecture for reusable game components
- 08 Ali Mirsoleimani (UL), Structured Parallel Programming for Monte Carlo Tree Search
- 09 Myriam Traub (UU), Measuring Tool Bias and Improving Data Quality for Digital Humanities Research
- 10 Alifah Syamsiyah (TU/e), In-database Preprocessing for Process Mining
- 11 Sepideh Mesbah (TUD), Semantic-Enhanced Training Data Augmentation Methods for Long-Tail Entity Recognition Models
- 12 Ward van Breda (VUA), Predictive Modeling in E-Mental Health: Exploring Applicability in Personalised Depression Treatment
- 13 Marco Virgolin (CWI), Design and Application of Gene-pool Optimal Mixing Evolutionary Algorithms for Genetic Programming
- 14 Mark Raasveldt (CWI/UL), Integrating Analytics with Relational Databases
- 15 Konstantinos Georgiadis (OU), Smart CAT: Machine Learning for Configurable Assessments in Serious Games
- 16 Ilona Wilmont (RUN), Cognitive Aspects of Conceptual Modelling
- 17 Daniele Di Mitri (OU), The Multimodal Tutor: Adaptive Feedback from Multimodal Experiences
- 18 Georgios Methenitis (TUD), Agent Interactions & Mechanisms in Markets with Uncertainties: Electricity Markets in Renewable Energy Systems
- 19 Guido van Capelleveen (UT), Industrial Symbiosis Recommender Systems
- 20 Albert Hankel (VUA), Embedding Green ICT Maturity in Organisations
- 21 Karine da Silva Miras de Araujo (VUA), Where is the robot?: Life as it could be
- 22 Maryam Masoud Khamis (RUN), Understanding complex systems implementation through a modeling approach: the case of e-government in Zanzibar
- 23 Rianne Conijn (UT), The Keys to Writing: A writing analytics approach to studying writing processes using keystroke logging
- 24 Lenin da Nóbrega Medeiros (VUA/RUN), How are you feeling, human? Towards emotionally supportive chatbots
- 25 Xin Du (TU/e), The Uncertainty in Exceptional Model Mining
- 26 Krzysztof Leszek Sadowski (UU), GAMBIT: Genetic Algorithm for Model-Based mixed-Integer optimization
- 27 Ekaterina Muravyeva (TUD), Personal data and informed consent in an educational context
- 28 Bibeg Limbu (TUD), Multimodal interaction for deliberate practice: Training complex skills with augmented reality
- 29 Ioan Gabriel Bucur (RUN), Being Bayesian about Causal Inference
- 30 Bob Zadok Blok (UL), Creatief, Creatiever, Creatiefst
- 31 Gongjin Lan (VUA), Learning better - From Baby to Better
- 32 Jason Rhuggenaath (TU/e), Revenue management in online markets: pricing and online advertising
- 33 Rick Gilsing (TU/e), Supporting service-dominant business model evaluation in the context of business model innovation
- 34 Anna Bon (UM), Intervention or Collaboration? Redesigning Information and Communication Technologies for Development
- 35 Siamak Farshidi (UU), Multi-Criteria Decision-Making in Software Production
-
- 2021 01 Francisco Xavier Dos Santos Fonseca (TUD), Location-based Games for Social Interaction in Public Space
- 02 Rijk Mercur (TUD), Simulating Human Routines: Integrating Social Practice Theory in Agent-Based Models
- 03 Seyyed Hadi Hashemi (UvA), Modeling Users Interacting with Smart Devices
- 04 Ioana Jivet (OU), The Dashboard That Loved Me: Designing adaptive learning analytics for self-regulated learning
- 05 Davide Dell'Anna (UU), Data-Driven Supervision of Autonomous Systems
- 06 Daniel Davison (UT), "Hey robot, what do you think?" How children learn with a social robot
- 07 Armel Lefebvre (UU), Research data management for open science
- 08 Nardie Fanchamps (OU), The Influence of Sense-Reason-Act Programming on Computational Thinking

-
- 09 Cristina Zaga (UT), The Design of Robothings. Non-Anthropomorphic and Non-Verbal Robots to Promote Children's Collaboration Through Play
- 10 Quinten Meertens (UvA), Misclassification Bias in Statistical Learning
- 11 Anne van Rossum (UL), Nonparametric Bayesian Methods in Robotic Vision
- 12 Lei Pi (UL), External Knowledge Absorption in Chinese SMEs
- 13 Bob R. Schadenberg (UT), Robots for Autistic Children: Understanding and Facilitating Predictability for Engagement in Learning
- 14 Negin Samaeemofrad (UL), Business Incubators: The Impact of Their Support
- 15 Onat Ege Adali (TU/e), Transformation of Value Propositions into Resource Re-Configurations through the Business Services Paradigm
- 16 Esam A. H. Ghaleb (UM), Bimodal emotion recognition from audio-visual cues
- 17 Dario Dotti (UM), Human Behavior Understanding from motion and bodily cues using deep neural networks
- 18 Remi Wieten (UU), Bridging the Gap Between Informal Sense-Making Tools and Formal Systems - Facilitating the Construction of Bayesian Networks and Argumentation Frameworks
- 19 Roberto Verdecchia (VUA), Architectural Technical Debt: Identification and Management
- 20 Masoud Mansoury (TU/e), Understanding and Mitigating Multi-Sided Exposure Bias in Recommender Systems
- 21 Pedro Thiago Timbó Holanda (CWI), Progressive Indexes
- 22 Sihang Qiu (TUD), Conversational Crowdsourcing
- 23 Hugo Manuel Proença (UL), Robust rules for prediction and description
- 24 Kaijie Zhu (TU/e), On Efficient Temporal Subgraph Query Processing
- 25 Eoin Martino Grua (VUA), The Future of E-Health is Mobile: Combining AI and Self-Adaptation to Create Adaptive E-Health Mobile Applications
- 26 Benno Kruit (CWI/VUA), Reading the Grid: Extending Knowledge Bases from Human-readable Tables
- 27 Jelte van Waterschoot (UT), Personalized and Personal Conversations: Designing Agents Who Want to Connect With You
- 28 Christoph Selig (UL), Understanding the Heterogeneity of Corporate Entrepreneurship Programs
-
- 2022 01 Judith van Stegeren (UT), Flavor text generation for role-playing video games
- 02 Paulo da Costa (TU/e), Data-driven Prognostics and Logistics Optimisation: A Deep Learning Journey
- 03 Ali el Hassouni (VUA), A Model A Day Keeps The Doctor Away: Reinforcement Learning For Personalized Healthcare
- 04 Ünal Aksu (UU), A Cross-Organizational Process Mining Framework
- 05 Shiwei Liu (TU/e), Sparse Neural Network Training with In-Time Over-Parameterization
- 06 Reza Refaei Afshar (TU/e), Machine Learning for Ad Publishers in Real Time Bidding
- 07 Sambit Praharaj (OU), Measuring the Unmeasurable? Towards Automatic Co-located Collaboration Analytics
- 08 Maikel L. van Eck (TU/e), Process Mining for Smart Product Design
- 09 Oana Andreea Inel (VUA), Understanding Events: A Diversity-driven Human-Machine Approach
- 10 Felipe Moraes Gomes (TUD), Examining the Effectiveness of Collaborative Search Engines
- 11 Mirjam de Haas (UT), Staying engaged in child-robot interaction, a quantitative approach to studying preschoolers' engagement with robots and tasks during second-language tutoring
- 12 Guanyi Chen (UU), Computational Generation of Chinese Noun Phrases
- 13 Xander Wilcke (VUA), Machine Learning on Multimodal Knowledge Graphs: Opportunities, Challenges, and Methods for Learning on Real-World Heterogeneous and Spatially-Oriented Knowledge
- 14 Michiel Overeem (UU), Evolution of Low-Code Platforms
- 15 Jelmer Jan Koorn (UU), Work in Process: Unearthing Meaning using Process Mining
- 16 Pieter Gijbbers (TU/e), Systems for AutoML Research
- 17 Laura van der Lubbe (VUA), Empowering vulnerable people with serious games and gamification
- 18 Paris Mavromoustakos Blom (TiU), Player Affect Modelling and Video Game Personalisation
- 19 Bilge Yigit Ozkan (UU), Cybersecurity Maturity Assessment and Standardisation
- 20 Fakhra Jabeen (VUA), Dark Side of the Digital Media - Computational Analysis of Negative Human Behaviors on Social Media
- 21 Seethu Mariyam Christopher (UM), Intelligent Toys for Physical and Cognitive Assessments
- 22 Alexandra Sierra Rativa (TiU), Virtual Character Design and its potential to foster Empathy, Immersion, and Collaboration Skills in Video Games and Virtual Reality Simulations

-
- 23 Ilir Kola (TUD), Enabling Social Situation Awareness in Support Agents
- 24 Samaneh Heidari (UU), Agents with Social Norms and Values - A framework for agent based social simulations with social norms and personal values
- 25 Anna L.D. Latour (UL), Optimal decision-making under constraints and uncertainty
- 26 Anne Dirkson (UL), Knowledge Discovery from Patient Forums: Gaining novel medical insights from patient experiences
- 27 Christos Athanasiadis (UM), Emotion-aware cross-modal domain adaptation in video sequences
- 28 Onuralp Ulusoy (UU), Privacy in Collaborative Systems
- 29 Jan Kolkmeier (UT), From Head Transform to Mind Transplant: Social Interactions in Mixed Reality
- 30 Dean De Leo (CWI), Analysis of Dynamic Graphs on Sparse Arrays
- 31 Konstantinos Traganos (TU/e), Tackling Complexity in Smart Manufacturing with Advanced Manufacturing Process Management
- 32 Cezara Pastrav (UU), Social simulation for socio-ecological systems
- 33 Brinn Hekkelman (CWI/TUD), Fair Mechanisms for Smart Grid Congestion Management
- 34 Nimat Ullah (VUA), Mind Your Behaviour: Computational Modelling of Emotion & Desire Regulation for Behaviour Change
- 35 Mike E.U. Ligthart (VUA), Shaping the Child-Robot Relationship: Interaction Design Patterns for a Sustainable Interaction
-
- 2023 01 Bojan Simoski (VUA), Untangling the Puzzle of Digital Health Interventions
- 02 Mariana Rachel Dias da Silva (TiU), Grounded or in flight? What our bodies can tell us about the whereabouts of our thoughts
- 03 Shabnam Najafian (TUD), User Modeling for Privacy-preserving Explanations in Group Recommendations
- 04 Gineke Wiggers (UL), The Relevance of Impact: bibliometric-enhanced legal information retrieval
- 05 Anton Bouter (CWI), Optimal Mixing Evolutionary Algorithms for Large-Scale Real-Valued Optimization, Including Real-World Medical Applications
- 06 António Pereira Barata (UL), Reliable and Fair Machine Learning for Risk Assessment
- 07 Tianjin Huang (TU/e), The Roles of Adversarial Examples on Trustworthiness of Deep Learning
- 08 Lu Yin (TU/e), Knowledge Elicitation using Psychometric Learning
- 09 Xu Wang (VUA), Scientific Dataset Recommendation with Semantic Techniques
- 10 Dennis J.N.J. Soemers (UM), Learning State-Action Features for General Game Playing
- 11 Fawad Taj (VUA), Towards Motivating Machines: Computational Modeling of the Mechanism of Actions for Effective Digital Health Behavior Change Applications
- 12 Tessel Bogaard (VUA), Using Metadata to Understand Search Behavior in Digital Libraries
- 13 Injy Sarhan (UU), Open Information Extraction for Knowledge Representation
- 14 Selma Čaušević (TUD), Energy resilience through self-organization
- 15 Alvaro Henrique Chaim Correia (TU/e), Insights on Learning Tractable Probabilistic Graphical Models
- 16 Peter Blomsma (TiU), Building Embodied Conversational Agents: Observations on human nonverbal behaviour as a resource for the development of artificial characters
- 17 Meike Nauta (UT), Explainable AI and Interpretable Computer Vision – From Oversight to Insight
- 18 Gustavo Penha (TUD), Designing and Diagnosing Models for Conversational Search and Recommendation
- 19 George Aalbers (TiU), Digital Traces of the Mind: Using Smartphones to Capture Signals of Well-Being in Individuals
- 20 Arkadiy Dushatskiy (TUD), Expensive Optimization with Model-Based Evolutionary Algorithms applied to Medical Image Segmentation using Deep Learning
- 21 Gerrit Jan de Bruin (UL), Network Analysis Methods for Smart Inspection in the Transport Domain
- 22 Alireza Shojaifar (UU), Volitional Cybersecurity
- 23 Theo Theunissen (UU), Documentation in Continuous Software Development
- 24 Agathe Balayn (TUD), Practices Towards Hazardous Failure Diagnosis in Machine Learning
- 25 Jurian Baas (UU), Entity Resolution on Historical Knowledge Graphs
- 26 Loek Tonnaer (TU/e), Linearly Symmetry-Based Disentangled Representations and their Out-of-Distribution Behaviour
- 27 Ghada Sokar (TU/e), Learning Continually Under Changing Data Distributions
- 28 Floris den Hengst (VUA), Learning to Behave: Reinforcement Learning in Human Contexts
- 29 Tim Draws (TUD), Understanding Viewpoint Biases in Web Search Results
-

- 2024 01 Daphne Miedema (TU/e), On Learning SQL: Disentangling concepts in data systems education
 02 Emile van Krieken (VUA), Optimisation in Neurosymbolic Learning Systems
 03 Feri Wijayanto (RUN), Automated Model Selection for Rasch and Mediation Analysis
 04 Mike Huisman (UL), Understanding Deep Meta-Learning
 05 Yiyong Gou (UM), Aerial Robotic Operations: Multi-environment Cooperative Inspection & Construction Crack Autonomous Repair
 06 Azqa Nadeem (TUD), Understanding Adversary Behavior via XAI: Leveraging Sequence Clustering to Extract Threat Intelligence
 07 Parisa Shayan (TiU), Modeling User Behavior in Learning Management Systems
 08 Xin Zhou (UvA), From Empowering to Motivating: Enhancing Policy Enforcement through Process Design and Incentive Implementation
 09 Giso Dal (UT), Probabilistic Inference Using Partitioned Bayesian Networks
 10 Cristina-Iulia Bucur (VUA), Linkflows: Towards Genuine Semantic Publishing in Science
 11 withdrawn
 12 Peide Zhu (TUD), Towards Robust Automatic Question Generation For Learning
 13 Enrico Liscio (TUD), Context-Specific Value Inference via Hybrid Intelligence
 14 Larissa Capobianco Shimomura (TU/e), On Graph Generating Dependencies and their Applications in Data Profiling
 15 Ting Liu (VUA), A Gut Feeling: Biomedical Knowledge Graphs for Interrelating the Gut Microbiome and Mental Health
 16 Arthur Barbosa Câmara (TUD), Designing Search-as-Learning Systems
 17 Razieh Alidoosti (VUA), Ethics-aware Software Architecture Design
 18 Laurens Stoop (UU), Data Driven Understanding of Energy-Meteorological Variability and its Impact on Energy System Operations
 19 Azadeh Mozafari Mehr (TU/e), Multi-perspective Conformance Checking: Identifying and Understanding Patterns of Anomalous Behavior
 20 Ritsart Anne Plantenga (UL), Omgang met Regels
 21 Federica Vinella (UU), Crowdsourcing User-Centered Teams
 22 Zeynep Ozturk Yurt (TU/e), Beyond Routine: Extending BPM for Knowledge-Intensive Processes with Controllable Dynamic Contexts
 23 Jie Luo (VUA), Lamarck's Revenge: Inheritance of Learned Traits Improves Robot Evolution
 24 Nirmal Roy (TUD), Exploring the effects of interactive interfaces on user search behaviour
 25 Alisa Rieger (TUD), Striving for Responsible Opinion Formation in Web Search on Debated Topics
 26 Tim Gubner (CWI), Adaptively Generating Heterogeneous Execution Strategies using the VOILA Framework
 27 Lincen Yang (UL), Information-theoretic Partition-based Models for Interpretable Machine Learning
 28 Leon Helwerda (UL), Grip on Software: Understanding development progress of Scrum sprints and backlogs
 29 David Wilson Romero Guzman (VUA), The Good, the Efficient and the Inductive Biases: Exploring Efficiency in Deep Learning Through the Use of Inductive Biases
 30 Vijanti Ramautar (UU), Model-Driven Sustainability Accounting
 31 Ziyu Li (TUD), On the Utility of Metadata to Optimize Machine Learning Workflows
 32 Vinicius Stein Dani (UU), The Alpha and Omega of Process Mining
 33 Siddharth Mehrotra (TUD), Designing for Appropriate Trust in Human-AI interaction
 34 Robert Deckers (VUA), From Smallest Software Particle to System Specification - MuDForM: Multi-Domain Formalization Method
 35 Sicui Zhang (TU/e), Methods of Detecting Clinical Deviations with Process Mining: a fuzzy set approach
 36 Thomas Mulder (TU/e), Optimization of Recursive Queries on Graphs
 37 James Graham Nevin (UvA), The Ramifications of Data Handling for Computational Models
 38 Christos Koutras (TUD), Tabular Schema Matching for Modern Settings
 39 Paola Lara Machado (TU/e), The Nexus between Business Models and Operating Models: From Conceptual Understanding to Actionable Guidance
 40 Montijn van de Ven (TU/e), Guiding the Definition of Key Performance Indicators for Business Models
 41 Georgios Siachamis (TUD), Adaptivity for Streaming Dataflow Engines
 42 Emmeke Veltmeijer (VUA), Small Groups, Big Insights: Understanding the Crowd through Expressive Subgroup Analysis

-
- 43 Cedric Waterschoot (KNAW Meertens Instituut), The Constructive Conundrum: Computational Approaches to Facilitate Constructive Commenting on Online News Platforms
 - 44 Marcel Schmitz (OU), Towards learning analytics-supported learning design
 - 45 Sara Salimzadeh (TUD), Living in the Age of AI: Understanding Contextual Factors that Shape Human-AI Decision-Making
 - 46 Georgios Stathis (Leiden University), Preventing Disputes: Preventive Logic, Law & Technology
 - 47 Daniel Daza (VUA), Exploiting Subgraphs and Attributes for Representation Learning on Knowledge Graphs
 - 48 Ioannis Petros Samiotis (TUD), Crowd-Assisted Annotation of Classical Music Compositions
-
- 2025 01 Max van Haastrecht (UL), Transdisciplinary Perspectives on Validity: Bridging the Gap Between Design and Implementation for Technology-Enhanced Learning Systems
 - 02 Jurgen van den Hoogen (JADS), Time Series Analysis Using Convolutional Neural Networks
 - 03 Andra-Denis Ionescu (TUD), Feature Discovery for Data-Centric AI
 - 04 Rianne Schouten (TU/e), Exceptional Model Mining for Hierarchical Data
 - 05 Nele Albers (TUD), Psychology-Informed Reinforcement Learning for Situated Virtual Coaching in Smoking Cessation

ACKNOWLEDGMENTS

I have been fortunate to have had wonderful companions during my PhD journey, each supporting, guiding, and teaching me in their own way, making this journey not just possible, but one I enjoyed so much that I am sad to see it end.

I want to begin by thanking my two supervisors, Willem-Paul and Mark. When I started my PhD, I was surprised at the idea of seeing my daily supervisor Willem-Paul for 1.5 hours every week. But, as it turned out, those 1.5 hours (and often even more) flew by. Willem-Paul, thank you for spending all this time brainstorming, explaining, giving feedback, providing life and career advice, giving humorous anecdotes, and, of course, forcing me to practice Dutch. From you, I have learned so much not only about topics directly related to my research but also about supervising, collaborating, and research ethics. Working on my PhD never felt like something I was doing alone. Instead, I had a true collaborator who was at least as engaged and enthusiastic as I was. Mark, while we met much less frequently, you always impressed me with the connections you made, your ability to simplify my most convoluted sentences, and your skill in capturing and visualizing the main message. Thank you for always again familiarizing yourself with what I was working on and for skillfully thinking along right away. Your calm, steady confidence in my work was reassuring.

I also extend my gratitude to the independent members of my doctoral committee: Tibor Bosse, Catholijn Jonker, Loes Keijsers, Susan Murphy, and Matthijs Spaan. Thank you for your interest in my work and for taking the time to read my thesis, give valuable feedback, and attend the defense.

Heartfelt thanks go to Frans and Miguel as those who got me interested in embarking on this journey in the first place. You not only kindled my interest in reinforcement learning and research but also gave me the confidence to pursue it. I still struggle with writing concise sentences, but I hope I have gotten a bit better.

To my fellow Perfect Fit project members, thank you for making our collaboration so much fun and successful. Bouke, Kristell, Milon, I truly enjoyed our weekly meetings, both when we were rushing to finish the app and when we had time to discuss vacation and career plans. It was great to be companions to each other's PhD journey and I am excited to see where your journeys take you. Sven, Walter, I do not know what we would have done without you taking charge of the development. I certainly would not have had the time to focus on my research as much as I did. And just on the side, you taught me a lot about software development. A huge thanks also to Bendik, Djura, Robin, and the rest of the technical team - it was a pleasure working together during development sprints, especially in person at the eScience Center. Anke, Eline, thanks for managing the project, and our development paper¹, with so much bravado. I am also glad that I could always count on many of you as psychology and physical activity experts - my work would not

¹I still cannot believe that the paper actually got published now.

have been possible without your expertise. And I am happy that I could, in turn, also help some of you with your work.

I also want to express my gratitude to my other group of colleagues, the PhD students, postdocs, and staff at the Interactive Intelligence group. Special thanks go to Enrico for being my favorite person to work across from, Mo for being the most generous tax collector and distraction collaborator, Mani for always again impressing me with deep empathy and in-depth knowledge on so many different topics, Edgar and Laxmi for completing the best office, Pei-Yu for being my active travel companion in Australia, Carolina for the talent to bring any conversation to life, Amir and Zuzanna for short but sweet office partnerships on the 6th and our current 4th floor, Emma for many engaging lunch conversations, Sid for making me feel useful by asking questions about statistics, Rolf for being a rare but excellent climbing companion, Ruben for helping out with outstanding Dutch skills, Pradeep and Antonio for, respectively, reliably accepting and persistently declining my lunch break invitations, Myrthe for sharing her HREC expertise, Mert for helpful brainstorming and career advice sessions, Bram for companionship at conferences and workshops, and Tae for showcasing impressive German skills. Thanks also to Wouter for creating student project prototypes, Ruud for quickly helping every time my laptop's hard disk died, and Anita for answering every single one of the many diverse questions I had over the years. Siska, thank you for spending your free time working with me.

My PhD journey would not have been the same without the weekly meetings with fellow postdocs, PhD students, Master's students, and research assistants. Especially when I started my PhD during the pandemic were they a welcome way to get in touch with and learn from other human beings. And when we could finally meet in person, they continued to be a valued part of the week by providing a glance at adjacent research topics, the opportunity to share experiences, and, of course, a chance to look at preregistration forms. Thanks to everyone who was part of this group, whether for a short time or longer. I especially want to mention Mitchell who got me started with Rasa, Esra who contributed with a lot of positive energy and feedback (and recently also with many ideas for running competitions), Amal who not only joined me in working on reinforcement learning but also provided so much support for our workshop and travels in Australia, and Eva who agreed to take on the fight with Rasa to help develop persuasive dialogs. And Mo, Mani, you were, of course, also invaluable here. I was about to mention you, too, Jennie, but then I remembered that I never saw that cake.

Another integral part of the weekly meetings were the Master's students whom I was fortunate to supervise. Alkis, Beyza, Ramya, you kept me on my toes during the first year of my PhD and made sure I quickly learned about setting up, conducting, and analyzing empirical studies so I could support you in your projects. Moreover, you contributed greatly to the deliverables for the Perfect Fit project. Your personalized testimonials, Beyza, have even helped people quit smoking in the final Perfect Fit app. Martin, Andrei, thanks for agreeing to work on reinforcement learning - it was both insightful and a pleasure to supervise you. Diederik, thanks for joining for your literature survey and for letting me present your initial findings at our workshop. Besides these Master's students, I was also lucky to learn a lot from the Bachelor's students that I had the opportunity to supervise. I am particularly grateful to Nadyne, Mahira, and Arsen for helping me write a paper on smokers' attitudes toward a virtual smoking cessation coach, and to Emma, Bolek,

Johan, and Kriss for helping in creating the Dutch and German translations of the Artificial Social Agent (ASA) Questionnaire. Thanks also to the student assistants who supported my research in various ways: Shruthi and Ioana for patiently double-coding free-text responses, Tom for helping with containerizing Rasa, Sophie and Nora for creating beautiful videos for smoking cessation activities, and Goda and Karina for writing so many feedback messages for smokers and vapers preparing to quit.

For a few months, I was heartily welcomed into another research group, namely, GAIPS in Lisbon. Francisco, Ana, thank you for inviting me, I really enjoyed working with you. And thanks also to all the lovely members of that group as well as the other two visitors from the Netherlands, Anouk and Pieter, who made my time there such a pleasant one. The study I worked on during my time at GAIPS further benefited greatly from the expertise of Olya - thank you for sharing it with me.

In the same cordial way was I invited to join the ProtectMe consortium. Thanks to all the enthusiastic people who made contributing to this project so much fun. It has really been a place to connect with a lot of people doing very interesting and related research, and the great social events organized by the young board around Rick and Esra certainly helped with that. It is not just because I ended up working on an ESM project for mental health myself that I hope we stay in touch.

Last but not least, I want to thank my friends and family in Germany, the Netherlands, and everywhere else. You know who you are. I do want to highlight a few of you who have seen most of my past four years. Vivi, Anna, thanks for sticking around for so long - I am glad we visit each other everywhere we go. Priya, Reka, Hana, Claudia, Bea, Tjasa, you are the best past and present roommates one could ask for. Manolis, Chris, Lazaros, Micha², thanks for being my favorite sports companions (and Claudia, you are that too, of course). Yuan, I really appreciate that we stay in touch despite you now living at the other end of this huge country. Grace, thanks for your unwavering support since pretty much my first day in Delft. Masha, I am so glad that I forced you to be my friend by inviting you over just a few days after you arrived. Jay, thanks for letting me come by so often, and, of course, for your fantastic cover design help. Jose, Gabriel, I am happy we regularly happen to stop by the same apartment. Anurag, I know I can count on you no matter when or what. Jana, Mama, Paps, I would not be where I am if it were not for your support.

²I promise I also enjoy your company outside of sports.

CURRICULUM VITAE

Nele ALBERS

EDUCATION

- 2020 – 2025 **Ph.D. in Computer Science**
Delft University of Technology, Delft, the Netherlands
Promotors: Dr. ir. W.-P. Brinkman and Prof. dr. M.A. Neerincx
- 2018 – 2020 **Master of Science in Computer Science**
Delft University of Technology, Delft, the Netherlands
Supervisor: Dr. F.A. Oliehoek
- 2014 – 2018 **Bachelor of Science in Computer Science**
Fordham University, New York, NY, USA
Marshall University, Huntington, WV, USA

EXPERIENCE

- 2024 – ... **Tilburg University**, Tilburg, the Netherlands
Postdoctoral Researcher
- 2023 **Instituto Superior Técnico**, Lisbon, Portugal
Visiting Researcher
- 2020 – 2024 **Delft University of Technology**, Delft, the Netherlands
Ph.D. Candidate
- 2020 **ShaleProfile B.V.**, Rotterdam, the Netherlands
Business Intelligence Developer
- 2018 **Porsche Digital GmbH**, Ludwigsburg, Germany
IT Intern
- 2017 **zeb.information.technology GmbH & Co.KG**, Münster, Germany
IT Development Intern
- 2016 **GWS Gesellschaft für Warenwirtschafts-Systeme mbH**, Münster,
Germany
IT Development Intern

LIST OF PUBLICATIONS

UNDER REVIEW

- 1. **Nele Albers**, Francisco S Melo, Mark A Neerincx, Olya Kudina, and Willem-Paul Brinkman. The impact of human feedback in a chatbot-based smoking cessation intervention: An empirical study into psychological, economic, and ethical factors. *Under review*.
- 2. **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Reinforcement learning for proposing smoking cessation activities that build competencies: Combining two worldviews in a virtual coach. *Under review*.

2024

1. **Nele Albers**^{*}, Andrea Bönsch^{*}, Jonathan Ehret, Boleslav A Khodakov, and Willem-Paul Brinkman. German and dutch translations of the artificial-social-agent questionnaire instrument for evaluating human-agent interactions. In *Proceedings of the 24th ACM International Conference on Intelligent Virtual Agents, IVA '24*, New York, NY, USA, 2024. ACM. ISBN 9798400706257. doi: 10.1145/3652988.3673928.
2. Anke Versluis, Kristell M Penformis^{*}, Sven A van der Burg^{*}, Bouke L Scheltinga^{*}, Milon HM van Vliet^{*}, **Nele Albers**^{*}, and Eline Meijer. Targeting key risk factors for cardiovascular disease in at-risk individuals: developing a digital, personalized and real-time intervention to facilitate smoking cessation and physical activity. *JMIR Cardio*, 8:e47730, 2024. ISSN 2561-1011. doi: 10.2196/47730.
3. Milon HM van Vliet, Anke Versluis, Niels H Chavannes, Bouke L Scheltinga, **Nele Albers**, Kristell M Penformis, Walter Baccinelli, Eline Meijer, and on behalf of the Perfect Fit consortium. Protocol of a mixed-methods evaluation of Perfect Fit: A personalized mHealth intervention with a virtual coach to promote smoking cessation and physical activity in adults. *Digital Health*, 10:20552076241300020, 2024. doi: 10.1177/20552076241300020.
4. Milon HM van Vliet, Anke Versluis, Niels H Chavannes, Bouke L Scheltinga, **Nele Albers**, Kristell M Penformis, Walter Baccinelli, Eline Meijer, and on behalf of the Perfect Fit consortium. A mixed-methods evaluation of a virtual coach for smoking cessation and physical activity in adults. *European Journal of Public Health*, 34(Supplement_3):ckae144.1191, 2024. ISSN 1464-360X. doi: 10.1093/eurpub/ckae144.1191.
5. Diederik Heijbroek, **Nele Albers**, and Willem-Paul Brinkman. Algorithmic support for health behavior change: A scoping review protocol. In Kiemute Oyibo, Wenzhen Xu, and Elena Vlahu-Gjorgievska, editors, *19th International Conference on Persuasive Technology, Adjunct Proceedings co-located with PERSUASIVE 2024, Wollongong, Australia, April 10th – 12th, 2024*, volume 3728 of *CEUR Workshop Proceedings*, 2024.
6. **Nele Albers**, Amal Abdulrahman, Deborah Richards, Caroline Figueroa, Bibhas Chakraborty, Ananya Bhattacharjee, Linwei He, Mark A Neerincx, Joseph Jay Williams, Nezhil Younsi, Tibor

Bosse, Annemiek Linn, Crystal Smit, and Willem-Paul Brinkman. Workshop on algorithmic behavior change support. In Kiemute Oyibo, Wenzhen Xu, and Elena Vlahu-Gjorgievska, editors, *19th International Conference on Persuasive Technology, Adjunct Proceedings co-located with PERSUASIVE 2024, Wollongong, Australia, April 10th – 12th, 2024*, volume 3728 of *CEUR Workshop Proceedings*, 2024.


7. Martin Dierikx, **Nele Albers**, Bouke L Scheltinga, and Willem-Paul Brinkman. Collaboratively setting daily step goals with a virtual coach: Using reinforcement learning to personalize initial proposals. In Nilufar Baghaei, Raian Ali, Khin Than Win, and Kiemute Oyibo, editors, *Persuasive Technology - 19th International Conference, PERSUASIVE 2024, Wollongong, NSW, Australia, April 10–12, 2024, Proceedings*, volume 14636 of *Lecture Notes in Computer Science*, pages 100–115, 2024. doi: 10.1007/978-3-031-58226-4_9.

2023

1. **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Effectiveness of user states, demographics and traits in persuading to quit smoking. In *35th Benelux Conference on Artificial Intelligence and 32nd Belgian-Dutch Conference on Machine Learning*, 2023.
2. Ramya P Ghantasala, **Nele Albers**, Kristell M Penfornis, Milon HM van Vliet, and Willem-Paul Brinkman. Feasibility of generating structured motivational messages for tailored physical activity coaching. *Frontiers in Digital Health*, 5:1215187, 2023. doi: 10.3389/fdgth.2023.1215187.
- 📄 3. **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Persuading to prepare for quitting smoking with a virtual coach: Using states and user characteristics to predict behavior. In Noa Agmon, Bo An, Alessandro Ricci, and William Yeoh, editors, *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 – 2 June 2023*, pages 717–726. ACM, 2023. doi: 10.5555/3545946.3598704.
4. **Nele Albers**, Mark A Neerincx, Nadyne L Aretz, Mahira Ali, Arsen Ekinci, and Willem-Paul Brinkman. Attitudes toward a virtual smoking cessation coach: Relationship and willingness to continue. In Alexander Meschtscherjakov, Cees Midden, and Jaap Ham, editors, *Persuasive Technology: 18th International Conference, PERSUASIVE 2023, Eindhoven, The Netherlands, April 19–21, 2023, Proceedings*, pages 265–274, 2023. doi: 10.1007/978-3-031-30933-5_17. [**Nominated for best student paper award**].
5. **Nele Albers**, Beyza Hizli, Bouke L Scheltinga, Eline Meijer, and Willem-Paul Brinkman. Setting physical activity goals with a virtual coach: Vicarious experiences, personalization and acceptance. *Journal of Medical Systems*, 47:15, 2023. doi: 10.1007/s10916-022-01899-9.

2022

- 📄 1. **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Addressing people’s current and future states in a reinforcement learning algorithm for persuading to quit smoking and to be physically active. *PLoS ONE*, 17(12):e0277295, 2022. doi: 10.1371/journal.pone.0277295.
2. **Nele Albers**, Mark A Neerincx, Kristell M Penfornis, and Willem-Paul Brinkman. Using a virtual coach to quit smoking: 14 themes for user needs. In *34th Benelux Conference on Artificial Intelligence and 31st Belgian-Dutch Conference on Machine Learning*, 2022.


3. Walter Baccinelli, Sven van der Burg, Robin Richardson, Djura Smits, Cunliang Geng, Lars Ridder, Bouke Scheltinga, **Nele Albers**, Willem-Paul Brinkman, Eline Meijer, and Jasper Reenalda. Reusable virtual coach for smoking cessation and physical activity coaching. In *Proceedings of the 22nd ACM International Conference on Intelligent Virtual Agents, IVA '22*, pages 1–3, New York, NY, USA, 2022. ACM. ISBN 9781450392488. doi: 10.1145/3514197.3551252.
-  4. **Nele Albers**, Mark A Neerincx, Kristell M Penforis, and Willem-Paul Brinkman. Users' needs for a digital smoking cessation application and how to address them: A mixed-methods study. *PeerJ*, 10:e13824, 2022. doi: 10.7717/peerj.13824.

2021

1. **Nele Albers**, Miguel Suau, and Frans A Oliehoek. Using bisimulation metrics to analyze and evaluate latent state representations. In *33rd Benelux Conference on Artificial Intelligence and 30th Belgian-Dutch Conference on Machine Learning*, pages 320–334, 2021.
2. **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Reinforcement learning-based persuasion by a conversational agent for behavior change. In *33rd Benelux Conference on Artificial Intelligence and 30th Belgian-Dutch Conference on Machine Learning*, pages 729–732, 2021.
3. **Nele Albers**, Willem-Paul Brinkman, and Mark A Neerincx. Adaptive data-driven persuasive communication for a conversational agent to support behavior change. In *Doctoral Consortium of the ACM International Conference on Intelligent Virtual Agents (IVA'21)*, 2021.

2020

1. **Nele Albers**, Miguel Suau de Castro, and Frans A Oliehoek. Learning what to attend to: Using bisimulation metrics to explore and improve upon what a deep reinforcement learning agent learns. In *32nd Benelux Conference on Artificial Intelligence and 29th Belgian-Dutch Conference on Machine Learning*, 2020.

 Included in this Thesis.

* Equal contribution.

LIST OF RELEVANT DATASETS AND ANALYSIS CODE

1. **Nele Albers**, Francisco Melo, Mark Neerincx, Olya Kudina, and Willem-Paul Brinkman. The impact of human feedback in a chatbot-based smoking cessation intervention: An empirical study into psychological, economic, and ethical factors - Data and analysis code for the PhD thesis chapter. <https://doi.org/10.4121/1d9aa8eb-9e63-4bf5-98a3-f359dbc932a4>, 2025.
2. **Nele Albers**, Mark Neerincx, and Willem-Paul Brinkman. Reinforcement learning for proposing smoking cessation activities that build competencies: Combining two worldviews in a virtual coach - Data, analysis code, and appendix for the PhD thesis chapter. <https://doi.org/10.4121/9C4D9C35-3330-4536-AB8D-D5BB237C277D>, 2024.
3. **Nele Albers**, Mark Neerincx, and Willem-Paul Brinkman. Human feedback messages for preparing for quitting smoking: Dataset. <https://doi.org/10.4121/7e88ca88-50e9-4e8d-a049-6266315a2ece>, 2024.
4. **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Difficulty and time perceptions of preparatory activities for quitting smoking: Dataset. <https://doi.org/10.4121/5198F299-9C7A-40F8-8206-C18DF93EE2A0>, 2023.
5. **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Persuading to prepare for quitting smoking with a virtual coach: Using states and user characteristics to predict behavior - Data, analysis code and appendix. <https://doi.org/10.4121/22153898>, 2023.
6. **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Preparing for quitting smoking and becoming more physically active with a virtual coach: Reflections for persuasive messages and action plans. <https://doi.org/10.4121/21905271>, 2023.
7. **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Addressing people's current and future states in a reinforcement learning algorithm for persuading to quit smoking and to be physically active: Data and analysis code. <https://doi.org/10.4121/21533055>, 2022.
8. **Nele Albers**, Mark A Neerincx, Kristell M Penformis, and Willem-Paul Brinkman. Users' needs for a digital smoking cessation application and how to address them: Data and analysis code. <https://doi.org/10.4121/20284131>, 2022.
9. **Nele Albers**, Mark A Neerincx, and Willem-Paul Brinkman. Acceptance of a virtual coach for quitting smoking and becoming physically active: Dataset. <https://doi.org/10.4121/19934783>, 2022.

LIST OF RELEVANT PREREGISTRATIONS

1. **Nele Albers** and Willem-Paul Brinkman. Perfect fit - learning when to involve a human coach in an ehealth application for preparing for quitting smoking or vaping. <https://doi.org/10.17605/osf.io/78cnr>, 2024.
2. **Nele Albers** and Willem-Paul Brinkman. Perfect fit - learning to propose useful preparatory activities for quitting smoking and becoming more physically active. <https://doi.org/10.17605/osf.io/nuy4w>, 2023.
3. **Nele Albers** and Willem-Paul Brinkman. Perfect fit - beliefs about and competencies built by preparatory activities for quitting smoking and becoming more physically active. <https://doi.org/10.17605/osf.io/cax6f>, 2022.
4. **Nele Albers** and Willem-Paul Brinkman. Perfect fit - experiment to gather data for and test a reinforcement learning-approach for motivating people. <https://doi.org/10.17605/osf.io/k2uac>, 2021.

LIST OF RELEVANT VIRTUAL COACH IMPLEMENTATIONS

1. **Nele Albers.** Virtual Coach Kai for Preparing for Quitting Smoking with Human Support. <https://doi.org/10.5281/zenodo.11102861>, 2024. URL https://github.com/PerfectFit-project/virtual_coach_human_inv.
2. **Nele Albers.** Virtual Coach Mel for Proposing Useful Preparatory Activities for Quitting Smoking. <https://doi.org/10.5281/zenodo.8302492>, 2023. URL https://github.com/PerfectFit-project/virtual_coach_useful_activities.
3. **Nele Albers.** Reinforcement Learning-Based Persuasion for a Conversational Agent to Support Behavior Change: Code. <https://doi.org/10.5281/zenodo.6319356>, 2022. URL https://github.com/PerfectFit-project/virtual_coach_rl_persuasion_algorithm.

