

Evolving Fuzzy logic Systems for creative personalized Socially Assistive Robots

Dell'Anna, Davide; Jamshidnejad, Anahita

DOI

[10.1016/j.engappai.2022.105064](https://doi.org/10.1016/j.engappai.2022.105064)

Publication date

2022

Document Version

Final published version

Published in

Engineering Applications of Artificial Intelligence

Citation (APA)

Dell'Anna, D., & Jamshidnejad, A. (2022). Evolving Fuzzy logic Systems for creative personalized Socially Assistive Robots. *Engineering Applications of Artificial Intelligence*, 114, Article 105064.
<https://doi.org/10.1016/j.engappai.2022.105064>

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.



Evolving Fuzzy logic Systems for creative personalized Socially Assistive Robots

Davide Dell'Anna*, Anahita Jamshidnejad

Department of Control and Operations, Delft University of Technology, Delft, The Netherlands

ARTICLE INFO

Keywords:

Evolving Fuzzy logic Systems
Personalized Socially Assistive Robots
Robot creativity

ABSTRACT

Socially Assistive Robots (SARs) are increasingly used in dementia and elderly care. In order to provide effective assistance, SARs need to be personalized to individual patients and account for stimulating their divergent thinking in creative ways. Rule-based fuzzy logic systems provide effective methods for automated decision-making of SARs. However, expanding and modifying the rules of fuzzy logic systems to account for the evolving needs, preferences, and medical conditions of patients can be tedious and costly. In this paper, we introduce EFS4SAR, a novel Evolving Fuzzy logic System for Socially Assistive Robots that supports autonomous evolution of the fuzzy rules that steer the behavior of the SAR. EFS4SAR combines traditional rule-based fuzzy logic systems with evolutionary algorithms, which model the process of evolution in nature and have shown to result in creative behaviors. We evaluate EFS4SAR via computer simulations on both synthetic and real-world data. The results show that the fuzzy rules evolved over time are not only personalized with respect to the personal preferences and therapeutic needs of the patients, but they also meet the following criteria for creativity of SARs: originality and effectiveness of the therapeutic tasks proposed to the patients. Compared to existing evolving fuzzy systems, EFS4SAR achieves similar effectiveness with higher degree of originality.

1. Introduction

Socially Assistive Robots (SARs) (Van Wynsberghe, 2016; Broekens et al., 2009; Syriopoulou-Delli and Gkiolnta, 2022) are growingly used to help implementing non-pharmaceutical therapies in dementia care, with positive effects on the cognitive, psychological, and global functioning of patients (Van Mierlo et al., 2010; Sharif et al., 2018; Rebok et al., 2014).

To ensure the effectiveness of the implemented therapies, the behavior of SARs must be personalized to individual patients (Ascensão and Jamshidnejad, 2022; Abdi et al., 2018; Sabanovic et al., 2013). Personalization can take place either during individual therapeutic sessions (short-term personalization (Tsiakas et al., 2018)) via adaptation of the difficulty and type of tasks proposed to patients based on their capabilities and personality (Syriopoulou-Delli and Gkiolnta, 2022; Abdi et al., 2018; Sabanovic et al., 2013; Ramachandran et al., 2018), or over the course of multiple sessions (long-term personalization), by explicitly accounting for the (evolving) cultural and social background of patients (Van Mierlo et al., 2010; Epp, 2003) and by regularly and creatively stimulating patients to explore new activities so to preserve their divergent thinking and cognitive capacities (Palmiero et al., 2012; Liberati et al., 2012; Hannemann, 2006).

While a large corpus of solutions for short-term personalization of SARs exists in the literature, only few works report on long-term personalization and decision-making of SARs (Moro et al., 2018; Umbrico

et al., 2020; Rossi et al., 2017). Existing works are either tailored to specific tasks (Moro et al., 2018), or they address general user groups (e.g., extrovert or introvert patients (Sabanovic et al., 2013)). Clabaugh et al. (2019), for instance, propose a hierarchical framework for human-robot learning tailored for specific tasks for children with autism spectrum disorder. Tapus et al. (2008), use the degree of extroversion of the patients to adjust the volume and pitch of the voice, the distance from the patients, and the speed of movements of a SAR.

Fuzzy logic and Fuzzy Inference Systems (FISs) (Bai and Wang, 2006) have identified as appropriate tools for encoding and leveraging the necessary cultural, social, and medical knowledge in the automated decision-making of SARs (Bruno et al., 2017; Mobahi and Ansari, 2003; Vitiello et al., 2017). Indeed, in socially assistive contexts, the available and relevant knowledge (e.g., the preference, needs, and background of a patient) is often expressed via (fuzzy and ambiguous) linguistic terms. SARs that operate based on fuzzy logic allow the domain experts (e.g., therapists or caregivers) to express their knowledge via linguistic IF-THEN fuzzy rules for the robot, and to easily explain the resulting behavior of the robot (Garibaldi, 2019). Furthermore, FISs support the representation of social and cultural aspects that are essential for effective culture-aware robots (Bruno et al., 2017).

Despite their expressiveness and adequacy, approaches based on fuzzy logic generally require a tedious process to expand and adapt

* Corresponding author.

E-mail addresses: d.dellanna@tudelft.nl (D. Dell'Anna), a.jamshidnejad@tudelft.nl (A. Jamshidnejad).

the rules to new cases and to the evolving needs of the patients (Moro et al., 2018). For this reason, more dynamic learning-based approaches, e.g., based on Interactive Reinforcement Learning (Tsiakas et al., 2018) or supervised learning (Tapus et al., 2009) have generally been preferred to rule-based ones in the literature for the decision-making of SARs (Moro et al., 2018; Hemminghaus and Kopp, 2017; Liu et al., 2014). Learning-based approaches, however, suffer from a lack of expressiveness when it comes to representing social, verbal, or cultural aspects, which are crucial for long-term personalization and decision-making of SARs (Bruno et al., 2017).

In the last decade, numerous solutions for automated evolution of fuzzy rule-based systems have emerged (Skrjanc et al., 2019; Lughofer, 2015). Evolving Fuzzy Systems (EFSs) are powerful tools capable of self-adapting and self-developing in order to reflect the dynamical changes in the input data. EFSs autonomously evolve the fuzzy rules of a fuzzy system by identifying and refining the antecedent and consequent parts of the rules, and by dynamically adapting the parameters of the membership functions characterizing the linguistic terms in the rules. The evolution of the rules is based, for the antecedents, on the density of the data (Angelov and Filev, 2004, 2005), or the distance, error, and statistical contribution of new data w.r.t. the current model (Lughofer, 2008; Rong et al., 2006; Pratama et al., 2014b), and, for the consequent, on the fuzzily weighted recursive least squares and maximum correntropy criterion (Brüggemann and Bitmead, 2021; Angelov and Filev, 2004; Bao et al., 2018) approaches. The numerous (neuro-) fuzzy evolving systems in the literature include, but are not limited to, DENFIS (Kasabov and Song, 2002), eTS (Angelov and Filev, 2004), SAFIS (Rong et al., 2006), FLEXFIS (Lughofer, 2008), PANFIS (Pratama et al., 2014a), GENEFIS (Pratama et al., 2014b), McIT2FIS (Subramanian et al., 2014), SAFL (Gu and Shen, 2021), MEEFIS (Gu, 2021), CEFNS (Bao et al., 2018), and PENsemble (Pratama et al., 2018).

EFSs learn from data in a one-pass fashion, i.e., from data samples that arrive as a stream, and quickly react to changes in the patterns in the input data (Lughofer and Angelov, 2011). This permits EFSs to continuously self-improve their performance, and makes them particularly effective when dealing with nonlinear, non-stationary problems (Skrjanc et al., 2019). However, this real-time reactivity can also lead to poor global predictions and to the so-called “unlearning effect” (Gu et al., 2021). To overcome this limitation, the more recent literature employs evolutionary algorithms as a mechanism underlying the evolution of the fuzzy systems (Skrjanc et al., 2019; Hidalgo et al., 2020). Evolutionary algorithms (Mitchell, 1998), which include for example Genetic Algorithms (GA) (Mitchell, 1998) and Particle Swarm Optimization (PSO) algorithms (Kennedy and Eberhart, 1995), are inspired by the process of natural selection, where the fittest individuals in a population are naturally selected for reproduction and produce offspring of the next generation. Existing nature-inspired evolving fuzzy systems include for example PSO-ALMMo (Gu et al., 2021), a PSO-based enhancement of the original autonomous learning multiple model (ALMMo) fuzzy system (Angelov et al., 2018), ANFIS-GA and ANFIS-PSO (Moayedi et al., 2020). Furthermore, Multi-Objective Evolutionary Fuzzy Systems (MOEFSs) (Fazzolari et al., 2013) such as D-MOFARC (Fazzolari et al., 2014) and SKMOEFS-MPAES_RCS (Gallo et al., 2020) are Multi-Objective Evolutionary Algorithms (MOEAs) that concurrently optimize the two conflicting objectives of accuracy (to be maximized) and explainability (to be maximized by minimizing the number of rules learned) of the EFS.

EFSs are successfully applied to a wide variety of real-world applications related to system identification and streaming data processing (e.g., classification and regression) (Giannoglou et al., 2015; Chua and Tan, 2011; Lu, 2015; Sharkawy, 2010; Skrijanc et al., 2019) in contexts such as volatility forecasting (Maciel et al., 2017), actions recognition (Skrjanc et al., 2018), brain signals classification (de Jesús Rubio et al., 2019), etc. However, up until now, EFSs have not been applied for the control and long-term personalization of SARs.

In this paper, we introduce EFS4SAR (Evolving Fuzzy logic System for Socially Assistive Robots), an evolving fuzzy rule-based system specifically tailored for long-term personalization and creative decision-making of SARs. EFS4SAR employs genetic algorithms to autonomously evolve over time the rule base of a fuzzy system used for the decision-making of a SAR during interactions with patients. EFS4SAR introduces autonomous dynamics into traditional rule-based approaches for SARs, which otherwise, despite their expressiveness and adequacy, are usually excluded from dynamic SAR applications due to their lack of dynamicity. Evolutionary algorithms, furthermore, including genetic algorithms, have been shown to stimulate creativity (Cluzel et al., 2012), and to lead to surprising and original effective outcomes (Hornby et al., 2006), which makes them in line with the *standard definition of creativity* (Runco and Jaeger, 2012) – which states that creativity requires both originality (i.e., novelty) and effectiveness (i.e., capability to achieve satisfactory performance) (Lehman et al., 2020). By leveraging genetic algorithms as a long-term evolutionary fuzzy mechanism, therefore, EFS4SAR permits to creatively challenge the patients via non-obvious and non-repetitive suggestions, while at the same time optimizing the rules to be personalized to the particular patient.

Since the rules used for the decision-making of a SAR *actively* influence the feedback of the patients, an effective EFS for SARs needs to explicitly account, in evolving such rules, for the needs and preferences of the patients, for the differences between different therapeutic approaches, and for the need of diversity in the activities proposed to patients. These aspects are not considered in existing EFSs, which mainly focus on the (passive) task of approximating a target system and are optimized to maximize the accuracy of the fuzzy rule base (e.g., for classification and regression tasks) (Skrjanc et al., 2019).

The following aspects distinguish EFS4SAR from existing EFSs and characterize the contributions of this paper.

1. EFS4SAR evolves the fuzzy rules that steer the behavior of SARs so that the resulting therapeutic activities proposed to patients are personalized, effective, and diverse. To evolve the fuzzy rules, in addition to searching for fuzzy rules that accurately characterize a target system (in our case, the patient), as commonly done by EFSs, EFS4SAR also considers the therapeutic activities that have been performed in previous interactions in order to account for diversity of the resulting fuzzy rules.
2. EFS4SAR uses a novel fitness function specifically tailored to the SAR domain. This fitness function efficiently assesses new fuzzy rules based on (i) a second fuzzy logic rule base encoding the indications of the therapists as well as the available knowledge about patients, (ii) the evolving needs and preferences of patients, encoded via a dynamic credit assignment approach, and (iii) a diversity index that is important for cognitive stimulation of patients.
3. In addition to EFS4SAR, this paper also presents a novel approach for integrating EFSs in the long-term decision-making and personalization of SARs.

We report experimental results of EFS4SAR for both synthetic and real-world data (the latter consisting of 15 data sets of human preferences about daily activities, which are released in the supplementary material of this paper (Dell'Anna and Jamshidnejad, 2022)), and we perform comparisons with commonly used state-of-the-art EFSs. In addition to standard metrics (i.e., accuracy and number of rules), we also consider metrics that are specifically related to two characteristic features of creativity (i.e., originality and effectiveness of the therapeutic tasks proposed to patients (Runco and Jaeger, 2012)). In doing so, we also assess for the first time the creativity of several state-of-the-art EFSs methods.

The rest of the paper is organized as follows. Section 2 describes the proposed approaches. Section 3 focuses on the evolution of fuzzy rules via evolutionary algorithms. In Sections 4.1 and 4.2, we evaluate the developed approaches via the results obtained from computer

Table 1
Frequently-used mathematical notations.

H	Set of all possible therapies
a_t^p	Therapeutic activity suggested by the SAR at time t
H_a	Therapeutic category to which activity a corresponds
A	Set of all possible activities
c-FIS	Creative Fuzzy Inference System
a-FIS	Assessor Fuzzy Inference System
V	Set of input variables of the c-FIS
v_x	Number of possible realizations of variable x
L	List of all possible realizations of the output variables of the c-FIS
I_t	Vector of crisp inputs of the c-FIS at time t
O_t	Vector of crisp outputs of the c-FIS at time t
$c_{r,t}$	Credit of c-rule r at time t
w_h	Weight of therapy $h \in H$
$\bar{e}_{a,t}$	Average value of the outputs of the a-FIS at time t , given activity a
R_t	Rule base of the c-FIS at time t
$s_{r,t}$	Firing strength of rule r at time t
R_t^+	Set of the rules in R_t with $s_{r,t} > 0$
f_{t,a,I_t}	Feedback received by the SAR from the patient about the activity a suggested at time t given I_t
q_t	Payoff computed for the feedback f_{t,a,I_t}
β	Learning rate for the credit of the rules
χ_r	Chromosome encoding c-rule r
$F_{r,t}$	Fitness of c-rule r at time t
$R^{\max}$	Maximum value of the rule base of the c-FIS
η	Repetition cost
$z_{a,n}$	Number of times the activity a was suggested in the last n time steps
N	Number of samples required by the fitness function
P	List of populations of chromosomes
p^*	Best population in P
γ	Weight in $[0, 1]$ given to d_p for the evaluation of a population p ; $1 - \gamma$ is the weight given to $\bar{F}_p$
d_p	Diversity index of population p
$\bar{F}_p$	Average fitness of chromosomes in population p , normalized in $[0, 1]$
$g_{p,S}$	Gini-Simpson index of population p w.r.t. a set of species $S \in \{H, A\}$
$n_{p,s}$	Number of individuals in p belonging to a species $s \in S$
κ	Weight in $[0, 1]$ given to $g_{p,A}$ in the calculation of d_p ; $1 - \kappa$ is the weight given to $g_{p,H}$

simulations on synthetic and real-world data. Finally, in Section 5 we conclude the paper and propose topics for future research.

In the rest of the paper, we follow the mathematical notations given in Table 1.

2. An integrated AI-based approach for creative personalized Socially Assistive Robots

In this section, we describe the approach that we propose for creative long-term personalization and decision-making of SARs (see Figs. 1 and 2). We distinguish three phases, described in details below, which are meant to, respectively, (i) provide the SAR with the available knowledge about the personal preferences, physical and mental status, and needs of the patient, and about the therapeutic interventions relevant for the patient, (ii) steer the behavior of the SAR during the interactions with the patients, (iii) use the knowledge acquired during the interactions to evolve the SAR's fuzzy rules via EFS4SAR.

2.1. Human set-up phase

During the *human set-up phase*, the expert (e.g., the therapist or caregiver) performs the following three tasks: First, the expert decides which therapeutic categories should be considered during the interactions with a particular patient, and associates *weights* to the therapeutic categories in order to express the priorities for the SAR. In the example given in Fig. 1, among the five therapeutic categories, *cognitive stimulation* (Jang et al., 2015), *music/video therapy* (Woods et al., 2018) and *multimodal exercise program* (Vaughan et al., 2014; Trautwein et al., 2017) are chosen (indicated by the colored boxes), while the others

(*reminiscence therapy* (Woods et al., 2018), and *CBT* (Anon, 0000)) are excluded.

Secondly, the expert defines an initial set of *c-rules* (from *creative* rules, since these are the rules that will evolve over time as per Section 3) based on the available knowledge from the patients, including the existing norms and guidelines (e.g., based on Anon, 0000) and common practices in accepted therapeutic interventions. C-rules are fuzzy IF-THEN rules that constitute the rule base of the creative fuzzy inference system *c-FIS* which is used to steer the behavior of the SAR during interactions with patients.

The premise of a c-rule is a conjunction of terms, each of them assigning a linguistic value to a linguistic variable from the set of possible input variables V . The consequent is a single term assigning a linguistic value to *one* of the possible outputs of the SAR, i.e., to one of the activities that the SAR can undertake with the patients (e.g., playing a difficult tic-tac-toe game).

The c-rules relate relevant information to activities that should be suggested and performed by the SAR. Relevant information is defined, during the set-up phase, by linguistic variables and membership functions that can be interpreted by the fuzzy inference system, and includes environmental factors (e.g., time of the day, weather conditions, etc.), personal preferences, personal physical conditions, and personal mental conditions of the patients (i.e., four factors relevant for effective personalization of non-pharmaceutical interventions in dementia care (Van Mierlo et al., 2010)). An example of a c-rule is “IF *temperature is moderate* AND *boredom is very high* THEN *take a long walk*”.

A c-rule r is assigned a *credit* c_r , which is a numerical value that expresses the importance and expected effectiveness of r . The credits of the initial rules are tuned by the expert to reflect the available knowledge, and they will be automatically updated by the SAR over time via EFS4SAR based on the experience acquired during the interactions, to provide long-term personalization of the SAR's suggestions to patients.

Finally, the expert defines a set of fuzzy *a-rules* (from *assessor* rules), which characterize the preferences of patients (e.g., *the patient likes to have a long walk when it's sunny*), and encode personalized indications given by the caregivers (e.g., *the patient should not have long walks*) and the available medical knowledge concerning the different therapeutic categories. In addition to the variables that can be used for the inputs of c-rules, the input of a-rules also includes the activities that constitute the *output* of c-rules. The consequent of a-rules assigns linguistic values to *feedback measures*, i.e., variables that represent an evaluation of the interaction with the patient. These measures constitute the feedback that the patients may explicitly provide to the SAR during and after an interaction or that may be detected by the SAR. Examples of these measures that are relevant for dementia care (Tsiakas et al., 2018) include the level of boredom and agitation of the patient during the interactions, any problems for falling asleep, the emotions expressed during interactions, and the level of engagement of the patient (Marti et al., 2006). Such feedback can be both explicit (e.g., a verbal statement, or tests and questionnaires (Anon, 0000)) and implicit (e.g., expressed and measured by means of the tone of the voice, the level of participation, facial expressions, eye contact (Webb et al., 2020)). An example of an a-rule is “IF *boredom is very high* AND *game activity is difficult* THEN *participation is very low*”.

The a-rules compose the rule base of the a-FIS, which is used during the interactions with patients to predict the expected outcome of an activity and help selecting the activities, and to define the fitness function that guides the evolution of the c-rules so that the new rules, while selected in a creative way, are also in line with the patient's preferences, physical and mental health condition, etc.

2.2. Interaction phase

We define an interaction as one continuous session during which patients and the SAR interact by performing joint activities for a given amount of time (Marti et al., 2006; Tapus et al., 2009).




Fig. 1. Overview of the proposed framework. The colored input boxes indicate four of factors relevant in dementia care and that can be used as input for the FISs. Purple boxes indicate possible activities of the robot, and dark blue ones indicate feedback measures.




Fig. 2. Creative personalization phase.

During an interaction, the c-FIS determines the candidate activities that suit the current inputs to the c-FIS. Suppose that for the input vector I_t received by the SAR at time t , two c-rules are activated, and the c-FIS outputs two suggested activities: *to take a 40 min walk* or *to play a tic-tac-toe game with the robot with the difficulty level 3*. The suggested activities¹ go through a selection block (Activity Selector in Fig. 1), which selects the activity a_t^* from the set of possible activities (outputs of the c-FIS) O_t suggested for the current time t that maximizes a combination of three factors, according to the following relationship:

$$a_t^* = \arg \max_{a \in O_t} \left(\bar{e}_{a,t} \cdot w_{H_a} \cdot \frac{\sum_{r_a \in R_t^+} c_{r_a,t}}{|R_t^+|} \right) \quad (1)$$

In (1), $\bar{e}_{a,t}$ is the average value of the outputs of the a-FIS for activity a at time t , which we call *expected feedback* for the input given to the c-FIS at time t and a candidate activity a ; w_{H_a} indicates the *weight* associated to the therapeutic category to which activity a corresponds; and the last term is the mean of the credits of all c-rules that contributed (i.e., with a non-zero firing strength) to the suggestion of activity a at time t . Note that $|\cdot|$ is used for the cardinality of a set. In other words, the activity selector determines the optimal activity to execute by considering the preferences and needs of the patients based on the a-rules, the importance of different therapeutic categories, and the experience (encoded in the credits of the rules) acquired during previous interactions with the patient.

Once an activity is selected, the SAR suggests it to the patient, and gives assistance during the task. The process of suggesting and performing an activity with the patient is called a *step* of an interaction. During the course of an interaction, the SAR receives implicit or explicit feedback (called *realized feedback*) from the patient about the proposed activities. The realized feedback is used during the interaction to give new information to the c-rules and a-rules for the next steps of the same interaction, and to personalize the c-rules for future interactions

¹ Technically, the crisp output of activity-related variables *walk* and *play-tic-tac-toe* will have values 40 and 3, belonging to the respective universes of discourse of the variables. All other outputs of the c-FIS will have a default value outside their universe of discourse (e.g., a value -1, if the universe of discourse of the variable is $[x > -1, \dots]$), indicating that no suggestion is given for a certain activity output of the c-FIS.

(i.e., long-term personalization). For this reason, an *interactions memory* stores tuples $\langle I_t, a_t^*, \bar{e}_{a_t^*}, f_{t,a_t^*}, \{s_{r,t} \mid r \in R_t\} \rangle$, with a_t^* the activity performed by the SAR when the input vector I_t was given to the c-FIS and to the a-FIS at time t ; $\bar{e}_{a_t^*}$ the expected feedback for activity a_t^* at time t ; f_{t,a_t^*} the realized feedback after executing activity a_t^* ; $\{s_{r,t} \mid r \in R_t\}$ the set of firing strengths of the rules in the c-FIS at time t .

2.3. Creative personalization phase

In the creative personalization phase, the SAR leverages the knowledge stored in the *interactions memory* to evolve the c-rules. The evolution of the c-rules is performed in between different interactions (or sessions) by the EFS4SAR algorithm, which personalizes the activities that will be proposed to the patients while maintaining an adequate degree of variability, therefore making the suggestions of the SAR “creative”. At the end of the personalization phase, the fittest c-rules found with the evolutionary algorithm replace the current ones in the c-FIS and are used to steer the behavior of the SAR during the next interactions. In Section 3 we describe in detail EFS4SAR.

3. EFS4SAR: Evolving Fuzzy logic System for Socially Assistive Robots

In this section we discuss EFS4SAR: the proposed method for evolving the fuzzy rules that determine the possible behaviors of the SAR (i.e., the c-rules). EFS4SAR adopts a genetic algorithm (Mitchell, 1998) that relies on the credits that are assigned to the rules and are updated over time.

Next, we first provide details about the update procedure of the credits of the rules, then after briefly illustrating how we encode a c-rule with a binary chromosome, we discuss our definition of fitness of a chromosome representing a c-rule, and finally the procedure for the evolution of the c-rules via genetic algorithms.

3.1. Update of the credits of C-rules

The update rule (2) is applied, before evolving the c-rules, for all c-rules that fired in the last interaction, as are encoded in the interactions memory. Let R_t^+ be the subset of c-rules of the c-FIS with a non-zero firing strength at time t , and q_t the payoff (i.e., a negative value if the activity proposed at time t by the SAR led to poor realized feedback, a positive value otherwise). The update of the credits of the rules is inspired by the Holland Classifier Systems (Geyer-Schulz, 1995), and is given by

$$c_{r,t+1} \leftarrow c_{r,t} + \beta \cdot q_t \cdot \frac{s_{r,t}}{\sum_{r' \in R_t^+} s_{r',t}} \quad (2)$$

In (2), $s_{r,t}$ denotes the firing strength of rule r at time t , and $\beta \in [0, 1]$ is a constant value representing the learning rate. The credit of a rule, therefore, is updated to reflect the response of the patients to

the proposed activity, and the update is proportional to the contribution that the rule had (via its relative firing strength) to the decision taken by the SAR at time t .

3.2. Fuzzy C-rules as binary chromosomes

We follow the *Michigan approach* (Valenzuela-Rendón, 1991) for the fuzzy rules encoding: each individual in the population of the genetic algorithm represents a c-rule and the population of one generation represents a candidate rule base of the c-FIS. Instead of the alternative Pittsburgh approaches (Thrifty, 1991), where one individual represents the entire rule base, we choose this approach because (i) computationally efficient (Koshiyama et al., 2019), and (ii) it allows to easily refine over time the fitness of individual rules based on the feedback received when applying those rules during the interactions with patients.

Specifically, chromosomes in our approach are binary strings that represent c-rules with a variable number of inputs and one output. We consider the case where the membership functions of the linguistic variables used for the inputs and outputs of the fuzzy rules are time-invariant. We follow a standard encoding: a chromosome χ_r encoding c-rule r is a binary string made of $|\chi_r| = (\sum_{i \in V} [\log_2(v_i + 1)]) + [\log_2(\sum_{a \in A} v_a)]$ bits (genes). The first $\sum_{i \in V} [\log_2(v_i + 1)]$ genes characterize the premise of the rule, which can be composed by at most $|V|$ terms (corresponding to the number of possible input variables), and v_x is the number of possible realizations (represented by fuzzy membership functions) of a linguistic variable x . We add an additional *disabled value* to be used when the input variable should not be part of the premise of the rule encoded by the chromosome, thereby supporting rules of different size. The remaining $[\log_2(\sum_{a \in A} v_a)]$ genes characterize the consequent of the rule. Here, since we consider rules with only one output (activity), we encode the *index* of a linguistic value from the ordered list L of all possible realizations of all output variables (activities) for the c-rules, so that, given the order of L , the encoded index indicates both the linguistic variable and its membership function.

3.3. Fitness of a Fuzzy C-rule

Let χ_r be a chromosome encoding a fuzzy rule r as per Section 3.2 and let V be the set of all possible input variables for the c-FIS and $a \in A$ the activity corresponding to the output of rule r . We compute the fitness $F_{r,t}$ of a rule r at time t by means of four different factors, using:

$$F_{r,t} = c_{r,t} \cdot w_{H_a} \cdot \omega_r \cdot \frac{1}{\eta^{z_{a,n}}} \quad (3)$$

In (3), $c_{r,t}$ is the credit of c-rule r at time t (if the rule was never considered before, by default $c_{r,t} = 1$), w_{H_a} is the weight of the therapy to which activity a corresponds, $z_{a,n}$ is the number of times the activity a has been already suggested to the patients in the last n suggestions, $\eta \geq 1$ is a parameter used to determine the cost to associate to multiple repetitions of the same suggestion, and finally ω_r is an assessment of rule r obtained by means of the a-FIS, which will be described below. The first three terms in (3) contribute to the definition of the rule's fitness based on the criteria of satisfactory performance for creativity (i.e., the expected effectiveness of the activity proposed by the rule, based on the acquired experience with the patient and the therapeutic indications provided by the therapist), while the last term ensures that rules that are repeated too many times in recent interactions get a lower fitness (to satisfy the criteria of originality for creativity).

Algorithm 1 details the complete procedure to compute the fitness of a c-rule: First, the input chromosome χ_r is decoded (via `DECODE`) into a c-rule r . If, due to the encoding bit overhead (Mitchell, 1998; Kumar, 2013), the rule encodes values out of the possible bounds (assessed via `INVALID`), the algorithm returns a big negative value M for the fitness of the rule. Otherwise, it determines the output activity a and the values of $z_{a,n}$, $c_{r,t}$, w_{H_a} (via `GETACT`, `GETREP`, `GETCRED` and `GETWEIGHT`, respectively), and synthesizes a fuzzy inference system s-FIS with a rule

base populated only by r . The algorithm then assesses the rule by means of the a-FIS. In doing so, it adopts a Monte Carlo strategy where it repeats N times the assessment. In particular, every time a set $c - in$ of crisp inputs for all possible input variables V is sampled (via `SAMPLE`) from their universe of discourse. Values in $c - in$ are given in input to the synthesized s-FIS so that a value $c - out$ is obtained for the activity output of rule r (via `s-FIS(c - in)`). Given $c - in$ and $c - out$, the a-FIS is then used to assess the new c-rule (via `A-FIS(c - in ∪ {c - out})`), obtaining a set E of crisp expected feedback values per output of the a-FIS. The algorithm then computes a value ω_r , representing the assessment of the rule encoded by chromosome χ_r via the a-FIS for different inputs, as the average of the crisp expected feedback values within the set E .

Applying a Monte Carlo approach permits to evaluate c-rules by considering the entire universe of discourse of the input variables (the larger the value of N the better the evaluation of the rules in general; this however increases the computational cost of computing the fitness of the evolved c-rules). Since we synthesize the rules during the personalization phase, and not online during the interaction phase, we opt for a Monte Carlo approach and synthesize rules that are expected to be effective considering different possible crisp values of the input variables, so to cover different possible cases that can happen during the next interactions.

Algorithm 1 Computing the Fitness Function

```

1: Input: chromosome  $\chi_r$ , assessor fuzzy inference system a-FIS, re-
   required number of samples  $N$ , current time  $t$ , set  $C_t$  of credits of all
   rules at time  $t$ , set of weights  $W$  of all therapies, repetition cost  $\eta$ ,
   interactions memory intmem,  $n$  number of previous suggestions to
   consider, a large number  $M$ 
2: Output: fitness value for chromosome  $\chi_r$  encoding a fuzzy rule  $r$ 
3:  $r \leftarrow \text{DECODE}(\chi_r)$ 
4: if INVALID( $r$ ) then return  $-M$ 
5: end if
6:  $z_{a,n} \leftarrow \text{GETREP}(\text{intmem}, n)$ ;  $a \leftarrow \text{GETACT}(r)$ 
7:  $c_{r,t} \leftarrow \text{GETCRED}(r, C_t)$ ;  $w_{H_a} \leftarrow \text{GETWEIGHT}(H_a, W)$ 
8: s-FIS  $\leftarrow \text{SYNTHESIZE}(\{r\})$ 
9: sampled_fitness  $\leftarrow []$ 
10: repeat
11:    $c - in \leftarrow \text{SAMPLE}(V)$ 
12:    $c - out \leftarrow \text{s-FIS}(c - in)$ 
13:    $E \leftarrow \text{A-FIS}(c - in \cup \{c - out\})$ 
14:   sampled_fitness.APPEND(MEAN( $E$ ))
15: until  $N$  samples are obtained
16:  $\omega_r \leftarrow \text{MEAN}(\text{sampled\_fitness})$ 
17: return  $c_{r,t} \cdot w_{H_a} \cdot \omega_r \cdot \frac{1}{\eta^{z_{a,n}}}$ 

```

3.4. Evolving the C-rules

Given a chromosome χ_r encoding a c-rule and a fitness function that provides an assessment of the encoded rule, we can straightforwardly apply a genetic algorithm to evolve a population of c-rules. After executing the genetic algorithm, we retrieve the best population according to their fitness values, excluding duplicate and invalid rules. The best population of c-rules is the population $p^* \in P$ for which a weighted sum of the fitness of the c-rules and the diversity of the entire c-rule base is maximum, i.e.,

$$p^* = \arg \max_{p \in P} \left(\bar{F}_p \cdot (1 - \gamma) + d_p \cdot \gamma \right) \quad (4)$$

In (4), $\bar{F}_p$ is the average fitness of the rules in population p normalized to the range $[0, 1]$ to be commensurable with d_p , d_p is a *diversity index* of the population p , and $\gamma \in [0, 1]$ is the relative importance given to the diversity index w.r.t. the fitness.

To compute the diversity index of a population, the Gini-Simpson index (Jost, 2006) (a measure that reflects how many different types

(or species) there are in a population) is used. In particular, the diversity index d_p for population p is defined as a weighted sum of two Gini-Simpson indices $g_{p,H}$ and $g_{p,A}$ defined according to the therapeutic categories (e.g., *cognitive stimulation* or *CBT*) and the activities (e.g., *have a walk* or *play tic-tac-toe*) that are covered by the population. We have:

$$d_p = g_{p,H} \cdot (1 - \kappa) + g_{p,A} \cdot \kappa \quad (5)$$

In (5), $\kappa \in [0, 1]$ determines the relative importance of the two Gini-Simpson indices, H is the set of therapeutic categories, and A is the set of possible activities corresponding to these therapeutic categories. In general, the Gini-Simpson index $g_{p,S}$ for population p w.r.t. the set S (where in (5) S is either H or A) is computed by:

$$g_{p,S} = 1 - \frac{\sum_{s \in S} n_{p,s}(n_{p,s} - 1)}{|p|(|p| - 1)} \quad (6)$$

In (6), $n_{p,s}$ is the number of c-rules in population p that correspond to element s of set S . In particular, for the Gini-Simpson index $g_{p,H}$, a c-rule corresponds to an element s in H whenever the output activity a of r belongs to the type of therapeutic category s . For the Gini-Simpson index $g_{p,A}$, a c-rule corresponds to an element s in A whenever the output activity of the c-rule is the same as s .

With (5) and (6), therefore, we characterize two measures of diversity of a rule base that describe how many different therapeutic categories and activities are considered in the rule base. Both aspects are important for our purposes, because one rule base may have high variety in therapeutic categories but rules that always suggest the same activity for a given category. Vice-versa, the rule base may include rules that are very variegated in terms of activities but they may all belong to the same therapeutic category. Our aim is to avoid either of these two cases.

4. Case study

In this section we present our experiments based on computer simulations on both synthetic and real-world data. In Section 4.1, we assess the performance of the proposed approaches for creative decision-making and personalization of SARs by simulating interactions between a SAR and patients via a number of agent-based discrete-event simulations.² In Section 4.2, we compare EFS4SAR, our proposed evolutionary approach, with state-of-the-art evolving fuzzy systems on real-world data obtained based on a survey with human participants.

4.1. Synthetic simulated patients

We run simulations involving two agents: the SAR, which implements the decision-making and personalization approaches proposed in this paper, and the *Patient*, which reacts to the activity suggested by SAR according to its own preferences, which are initially unknown to the SAR. Note that in the rest of this section, SAR and Patient are used interchangeably for SAR agent and Patient agent. A *simulation step* defines one interaction step between SAR and Patient during which one activity (we consider interactions composed by one step) is suggested by SAR, and the corresponding feedback is received from Patient.

We give a weight of unity to all the therapeutic categories in set H (see Section 2), and leave the rule base of the a-FIS empty. Therefore, via our simulations we assess whether SAR can personalize its decisions to Patient in a finite number of interactions based on Patient's immediate feedback only, and with no prior knowledge about Patient (i.e., an empty rule base for the a-FIS).

² Our implementation (see the supplementary material for this paper in Dell'Anna and Jamshidnejad, 2022) relies on the MESA framework (Kazil et al., 2020) for the agent-based simulation, on the PyGAD library for the genetic algorithm (Gad, 2021), and on the skfuzzy library (Warner et al., 2019) for the fuzzy inference system.




Fig. 3. Membership functions describing the linguistic values used in the experiments for (a) the input of the c-FIS (a singleton variable), and for (b) the output activities of the c-FIS, each of them with three possible realizations.

Figs. 3a and 3b illustrate the membership functions for, respectively, the inputs and outputs of the c-FIS of SAR. As it is illustrated in Fig. 3(a), we represent the input of the c-FIS with one crisp value³ within the range $[0, 10]$. The output variables of the c-FIS are considered to describe the therapeutic activities suggested by SAR via three terms (c.f. the three membership functions in Fig. 3(b)) within the range $[0, 10]$. These terms for, e.g., activity “walk” may be “short”, “regular”, and “long”.

In our experiments, we use standard fuzzy operators for fuzzy inference. In particular, we use the *minimum* operator to determine the certainty of the premises and to compute the implied membership functions of the consequent of the rules in the inference engine of the c-FIS, and we use the center of gravity method for defuzzification. Our choice is motivated by the fact that we consider a case where membership functions are time-invariant, so that using different operators would only minimally affect the results of our experiments, which focus on evaluating the personalization and originality of the activities suggested by SAR to Patients.

At the beginning of its simulation step, SAR receives randomly sampled values for its input variables and suggests an activity to Patient using the c-FIS. Every Patient returns a value $f \in [0, 10]$ (where 0 indicates an extremely negative feedback, and 10 indicates an extremely positive one) in response to the activity suggested by SAR. Since we focus on the personalization of SAR's decisions, we focus on a single numerical feedback which provides information about both the patient's preferences and the effectiveness of a certain therapeutic activity in an aggregate way. As explained in Section 2.1, however, multiple *feedback measures* are supported.

Initially, f is randomly assigned by Patient to the proposed activities s.t. in 60% of cases $f < 5$ and in the remaining 40% of cases $f \geq 5$. In doing so, we model the fact that some therapeutic interventions are more effective for a patient (either because of the preferences of patients, or because of the therapeutic output), while others are less effective. Whenever the same activity is suggested by SAR, Patient adjusts the previous feedback according to its *type*, a characteristic defined per Patient that determines how it reacts to repeated suggestions of the same activity over time. In particular, we consider Patients that have a *memory* of size $m = 14$ (i.e., they remember the previous 14 suggestions made by SAR in the previous 14 simulation steps). Given an activity a , every Patient returns a feedback f determined by:

$$f \leftarrow f + \Delta f(n_{a,m}) \quad (7)$$

³ In this case study, having more than one input for the c-FIS does not affect the internal states or preferences of Patient, but only results in an increase in the number of preferences of Patient that SAR should learn. With a single input, however, the number of samples required by the fitness function remains 1 and the number of required rules and generations remains relatively low, which makes repeated simulations with a multitude of different types of Patients computationally affordable. In Section 4.2, which focuses on a more realistic case study, we consider multiple inputs with multiple values.




Fig. 4. (a) Change $\Delta f(n_{a,m})$ applied to the previous feedback f given to an activity by different types of Patients based on their memory of the last m simulation steps, and (b) examples of four types of Patients resulting from using (8) with different parameters.

where

$$\Delta f(n_{a,m}) = \begin{cases} \lambda n_{a,m}, & \text{if } 0 \leq n_{a,m} \leq \tau_1. \\ \lambda \tau_1, & \text{if } \tau_1 < n_{a,m} \leq \tau_2. \\ \lambda \tau_1 - \frac{\lambda \tau_1}{\tau_3 - \tau_2} (n_{a,m} - \tau_2), & \text{if } \tau_2 < n_{a,m} \leq \tau_3. \\ -\lambda (n_{a,m} - \tau_3), & \text{if } n_{a,m} > \tau_3. \end{cases} \quad (8)$$

In (8), $n_{a,m}$ is the number of times that activity a has been suggested in the last m simulation steps. Every Patient is characterized by four parameters: τ_1, τ_2, τ_3 (with $\tau_1 \leq \tau_2 \leq \tau_3 \leq m$) and λ . If an activity was not suggested in the last m simulation steps (i.e., $n_{a,m} = 0$), Patient returns the same feedback returned the most recent time for the same activity (or, the first time, a random feedback). If the activity was already suggested but for not more than τ_1 times, i.e., $0 < n_{a,m} \leq \tau_1$, then Patient increases the previously given feedback by $\Delta f(n_{a,m}) = \lambda n_{a,m}$, with λ characterizing how quickly the feedback of Patient changes. Whenever $\tau_1 < n_{a,m} \leq \tau_2$, Patient keeps increasing steadily the feedback, indicating that Patient continues enjoying more and more the activity. When $\tau_2 < n_{a,m} \leq \tau_3$, the increase of feedback is reduced. Whenever $n_{a,m} > \tau_3$, $\Delta f(n_{a,m})$ changes its sign. Different combinations of $\tau_1, \tau_2, \tau_3, \lambda$ allow to characterize a variety of Patients with very different behaviors and dynamics. Fig. 4(a) illustrates function f given by (8) for τ_1, τ_2, τ_3 being non-zero and non-equal. Fig. 4(b) provides examples of four types of Patient obtained via different combinations of the four parameters: the dashed red curve describes a patient that never changes its feedback, the dash-dotted orange curve characterizes a patient that enjoys repeating over and over the same activity, the solid black curve characterizes a patient with an almost same behavioral trend but more moderate, while the dotted blue curve characterizes a patient that enjoys to repeat the same activity a few times, but is annoyed by too many repetitions.

Metrics for creative personalization. We are interested in providing, through a creative process, a balanced trade-off between the comfort/satisfaction of the patients and the therapeutic value (in terms of cognitive stimulation) of proposing a variety of activities. We quantify these aspects via two metrics, called *feedback* and *repetitions*.

Feedback is the feedback provided as per (7) by Patient for SAR. We analyze both the average and the trend of the feedback throughout the simulation. A SAR that can effectively learn the preferences of Patient and the effectiveness of different therapeutic interventions, and can personalize its decision-making accordingly, will result in feedback values that tend to 10.

Repetitions measures the number of activities in the last m simulation steps that are identical to the activity suggested at the current simulation step i . We analyze the average number ρ of repetitions in the simulation, i.e.,

$$\rho = \frac{\sum_{i \in [1,s]} n_{a_i,m}}{s} \quad (9)$$

In (9), s is the number of simulation steps, and $n_{a_i,m}$ is the number of times that activity a suggested at simulation step i was suggested in the previous m simulation steps.

The two metrics *feedback* and *repetitions* are in line with the two criteria (effectiveness, i.e., satisfactory performance, and originality, respectively) that characterize creativity as per the standard definition proposed by Runco and Jaeger (2012). *Feedback* is a measure of *effectiveness* of the personalization, for it determines the usefulness, fit and appropriateness of the suggested activities. *Repetitions* is a measure of *originality* of the personalization, for it quantifies the uniqueness and diversity of the proposed activities over time. A creative SAR that performs as intended will result in high feedback and low repetitions.

Experimental settings. We conduct the following experiments, whose details are summarized in Table 2, and where we analyze the performance of a system S that steers SAR according to the approaches proposed in this paper for decision-making and personalization. In particular,

- We study the importance of the repetition cost in the calculation of the fitness of a c-rule by comparing the decision-making systems S and S_R for which EFS4SAR is implemented with and without a repetition cost, i.e., with respectively $\eta = 10$ and $\eta = 1$ (see (3)).
- We study the importance of the *diversity index* in the choice of the best population by comparing the decision-making systems S and S_D for which EFS4SAR is implemented with and without a diversity index, i.e., with respectively $\gamma = 0.5$ and $\gamma = 0$ (see (4)).
- We study the robustness of the decision-making system S when the feedback received by Patient is affected by noise. This may occur due to the errors of the sensors (Webb et al., 2020), impreciseness of the tests conducted with the patients, or due to the inability of patients to clearly communicate their feedback (e.g., as a result of severe cognitive impairments).

For the *Noise* experiments, we consider three types of noise that we apply to the feedback of Patient: a *gaussian* noise (i.e., a noise that applies a change to the feedback sampled from a normal distribution with mean $\mu = 0$ and standard deviation $\sigma = 2$, cropped to a minimum $g^{\min} = -10$ and a maximum $g^{\max} = 10$); a more disruptive noise, which we call *inverse gaussian*, obtained by first generating a temporary gaussian noise n_t and then producing the final noise n by applying the rule: $n = g^{\min} - n_t$ if $n_t < 0$, $n = g^{\max} - n_t$ if $n_t > 0$, and, when $n = 0$, in 50% of the cases $n = g^{\max}$ and in 50% $n = g^{\min}$. Finally we consider a *reversed feedback* noise, obtained as $10 - f$, with $f \in [0, 10]$ the feedback of Patient. For each type of noise we test two scenarios: in scenario *Noise 0.2*, the noise is generated with probability 0.2 (we chose this value to reflect the accuracy reported in Webb et al., 2020), so that it is applied on average on 20% of the received feedback; in scenario *Noise 1*, the noise is applied to all received feedback.

Moreover, as a *baseline*, we consider a decision-making system S_B that randomly shuffles its c-rules per simulation step. In line with the definition of creativity, such *merely original* decision-making system S_B is expected to lead to very few repetitions but also a low feedback.

We consider, therefore, 10 different decision-making systems, i.e., $S, S_B, S_R, S_D, S_{g0.2}, S_{g1}, S_{ig0.2}, S_{ig1}, S_{r0.2}$, and S_{r1} . We randomly synthesize 200 Patients by uniformly sampling for each of them their parameters τ_1, τ_2, τ_3 from $[0, 14]$ (more specifically, we uniformly sample τ_3 from $[0, 14]$, then we sample τ_2 from $[0, \tau_3]$, finally τ_1 from $[0, \tau_2]$) and λ from $[0.01, 0.05]$. For each decision-making system, we run 200 different simulations, each of them simulating a sequence of

Table 2

Details of the parameters of (a) the tested decision-making systems, i.e., our proposal S , the baseline S_B , a variation of S (S_R) that ignores the repetition cost, and a variation of S (S_D) that ignores the diversity index, and (b) the algorithms, common for all experiments. Values g , ig and r stand for *gaussian*, *inverse gaussian* and *reversed feedback* noise, respectively.

Decision-making system	S	S_B	S_R	S_D	Noise
β	0.04	0	0.04	0.04	0.04
η	10	1	1	10	10
$z_{a,n}$	7	—	—	7	7
γ	0.5	0	0.5	0	0.5
Noise type	—	—	—	—	{ g, ig, r }
Noise prob.	0	0	0	0	{0.2, 1}

(a)

Parameter	Val	Parameter	Val
Mating top %	80%	Generations	50
Selection type	sss	Samples N	1
Crossover type	uniform	Inputs c-FIS	1
Crossover prob.	1	Outputs c-FIS	117
Mutation prob.	0.25	Possible activities	117
κ	0.75	Possible therapies	7
Rules	20	Memory size m	14
Default credit rule	1	Simulations	200 (5x each)
w_h	$1 \forall h \in H$	Interactions/Sim	4000
Big M	100	Steps/Interactions	1
Domain feedback f	[0, 10]	Payoff q_i for f	$\alpha \cdot f^{1.8}$, with $\alpha = 1$ if $f \geq 5$, $\alpha = -1$ else

(b)

4000 interaction steps between the SAR and one of the 200 randomly synthesized Patients. Since in every simulation a different patient is considered, every decision-making system is tested in 200 different conditions (patients with different preferences and behaviors). In order to account for the randomness of the evolutionary algorithm in interaction of the SAR with every patient, we repeat every simulation 5 times to obtain multiple data points for every simulation. In doing so, the results (of the total $10 \times 200 \times 5$ simulations) provide statistically relevant information over a variety of conditions.

The number (4000) of steps performed in every simulation was determined during a preliminary experimentation to guarantee the SAR reaches a steady-state behavior. Similarly, the values of the other parameters reported in Table 2 have been tuned through extensive experimentation. The tuning was firstly driven by a time-resource balancing need, so that an adequate number of experiments could be performed with our available computation resources. More specifically, first we determined the parameters that most heavily affected the computation time (i.e., the number of inputs of the c-FIS, the number of samples for the fitness function, and the number of c-rules). Once fixed these values respectively to 1, 1, and 20, we followed a grid-search method to explore possible combinations of values for the other parameters,⁴ and identified values that led to the best results in terms of *feedback* and *repetitions*. For example, in our experiments, running the genetic algorithm for more than 50 generations led to an increase in the number of repetitions in the resulting suggestions. With more than 50 generations, rules with high credits (i.e., rule that were previously successfully employed by SAR) had more chances to be selected. Conversely, less than 50 generations had the opposite effect, leading to “too original” rules that did not leverage the knowledge learned during previous interactions.

⁴ Specifically, we considered the following values: {10,20,50,70,100} for the number of generations, {20,50,80} for the % of mating parents, {sss, random} for the selection type (with sss indicating steady-state selection), {0.05,0.15,0.25} for the mutation probability, and {0.25,0.5,0.75} for κ . Crossover type was set to *uniform* since we are interested in altering every gene of the chromosomes, and memory size to 14 to indicate an hypothetical period of 2 weeks (assuming one interaction per day).

Table 3

Average and standard deviation of *feedback* and *repetitions* for different systems. “Effect on *feedback* S ” and “Effect on *repetitions* S ” report the interpretation of the effect size (Cohen’s d) w.r.t. decision-making system S . The star symbol (*) shows that a Wilcoxon Signed-Rank test indicates a significant difference from decision-making system S with p -value $p < 0.05$. The double star symbol (**) indicates a significant difference with p -value $p < 0.01$. The absence of a star symbol indicates the lack of significant difference from decision-making system S .

System	Feedback	Effect on feedback S	Repetitions	Effect on repetitions S
S	8.17 ± 1.00		1.32 ± 0.24	
S_B	4.50 ± 0.18	Large (5.103)**	0.13 ± 0.01	Large (7.103)**
S_R	6.47 ± 1.53	Large (1.314)**	5.34 ± 2.28	Large (2.477)**
S_D	8.17 ± 1.04	None (0.000)	1.38 ± 0.27	Small (0.216)**
$S_{g0.2}$	8.12 ± 1.02	None (0.041)**	1.34 ± 0.26	None (0.057)
S_{g1}	8.03 ± 1.03	None (0.131)**	1.28 ± 0.22	None (0.188)
$S_{ig0.2}$	7.99 ± 1.01	None (0.177)**	1.28 ± 0.25	None (0.179)*
S_{ig1}	5.83 ± 1.09	Large (2.234)**	1.03 ± 0.14	Large (1.496)**
$S_{r0.2}$	7.95 ± 1.01	Small (0.219)**	1.26 ± 0.24	Small (0.287)**
S_{r1}	2.88 ± 1.35	Large (4.451)**	1.22 ± 0.23	Small (0.440)**

Results. Table 3 reports the average and standard deviation of the *feedback* and *repetitions* metrics obtained in the 200 simulations (each repeated 5 times) with each of the 10 considered decision-making systems described above, and the results of Wilcoxon Signed-Rank tests⁵ (Rey and Neuhauser, 2011) comparing the results with those obtained with the decision-making system S , and the effect size. The same results are also visualized via box plots in Figs. 5 and 6.

Proposed decision-making system S compared to the baseline decision-making system S_B . Based on Figs. 5(a) and 5(b), for the baseline decision-making system S_B , the average feedback is 4.5 and, on average, an activity is repeatedly suggested 0.13 times every 14 simulation steps. This is in line with the expectations: randomly re-combining and mutating the c-rules per simulation step without following any particular selection criteria can lead to low numbers of repetition, but also to suggestions that are not personalized and are thus ineffective for Patient. When we consider the proposed decision-making system S , instead, from Figs. 5(a) and 5(b) we see a significant difference: the average measured feedback is significantly higher (on average 8.17), while, on average, an activity is repeatedly suggested only 1.32 times every 14 simulation steps. These results confirm that the proposed system behaves as intended with a variety of Patients, providing a good balance between personalization and variety in the provided suggestions.

Fig. 5(c) reports the average feedback per simulation step for the 200×5 simulations for the considered decision-making systems. While we see that the feedback obtained with the baseline decision-making system S_B (in red) oscillates around 4.5, the decision-making system S (in black) requires 532 simulation steps to reach for the first time an average feedback of 8.17 over all simulations. Moreover, in about 1500 simulation steps the feedback converges to a steady range of values. Note that the feedback stabilizes at around 8 for S and does not converge to its maximum, i.e., 10. This result, which is expected and intended, illustrates that (i) the credit assignment mechanism effectively encodes the acquired knowledge about Patient (i.e., it learns personalized suggestions), (ii) despite the personalized suggestions, the proposed decision-making system S also provides a variegated set of suggestions (observe also the continuous oscillations of the feedback in Fig. 5(c)), regularly incentivizing Patients to try also sub-optimal (less preferred) activities with the intent of challenging their routine and activating their creativity and divergent thinking.

⁵ Wilcoxon Signed-Rank (non-parametric) tests do not rely on assumptions on distribution and variance of data for the analysis of matched-pair data. The null hypothesis is that the differences between matched-pair data have a distribution with center zero.




Fig. 5. Box plots illustrating (a) the average feedback of Patients, (b) the average number of identical activities suggested per 14 simulation step, and (c) the trend of the average feedback (cropped for readability at simulation step 1500, after which the trend only showed slight improvements not visible from the plot), obtained in 200 simulations (each repeated 5 times) using the different decision-making systems.




Fig. 6. Box plots illustrating the (a) average feedback of Patient, and (b) average number of identical activities suggested to Patient per 14 simulation step. These results correspond to using the proposed decision-making system S in 200 simulations each time with different noise (gaussian noise for $S_{g,x}$, inverted gaussian for $S_{ig,x}$, and reverse feedback for $S_{r,x}$, with $x \in (0.2, 1)$ denoting the probability that the noise affects the feedback per simulation step).

Importance of the repetition cost. Decision-making system S_R leads to significantly higher numbers of repetitions according to Fig. 5(b) (an activity is repeatedly suggested, on average, 5.34 times every 14 simulation steps), with a significant reduction of the feedback given by Patient (see Fig. 5(b)). From Fig. 5(b) (blue box), we see that ignoring the repetition cost can lead to edge cases where the same activity is suggested by SAR in almost every single simulation step (average repetition close to 14). The decision-making system S_R tends to prioritize exploitation (i.e., suggesting over and over the activities that are known to be better than others according to Patient's feedback) and to almost ignore exploration (i.e., trying different activities). Note that, as a consequence of the definition of Patient, repeating over and over the same activities does not necessarily lead to the maximum feedback, because when too many repetitions occur, Patient may start decreasing the feedback. This is also illustrated by Fig. 5(c) where the decision-making system S_R quickly reaches a high feedback value (note the initial peak of the dark blue line), exploiting the acquired

knowledge about the preferred activities of Patient, but the received feedback drops to values lower than those of the proposed decision-making system S , stabilizing at an average feedback around 6.5. From these results, we conclude that considering the repetition cost appears to be fundamental not only to provide variegated suggestions (low repetitions), but also effective ones (i.e., high feedback).

Importance of the diversity index. Compared to the proposed decision-making system S , ignoring the diversity index via the decision-making system S_D , leads to an increase in the number of repetitions, and to an almost identical feedback. Considering the diversity index via the proposed decision-making system S leads, on average, to a significant improvement of the (already low number of) repetitions of about 4.5% (small effect size), and to no decrease of the feedback. These results from the fact that, in computing the fitness of the c-rules, the evolutionary process already accounts for both their effectiveness and their diversity (via considering the repetition cost). By tuning parameter γ in (4), which in our experiments was set to 0.5, it is possible to adjust the required diversity of the c-rules, and thus the diversity of the activities suggested by SAR to Patient, without significantly affecting the effectiveness of the proposed therapies.

Robustness of decision-making system S to noise. Table 3 reports the results for the noisy feedback cases, and Fig. 6 shows the corresponding box plots. Note that the reported values concern the actual feedback of Patient, and not the measured (noisy) one received by SAR. Considering the gaussian noise, the measured effect size on both feedback and repetition is negligible both when applied in 20% of the cases and when applied consistently at every simulation step. Similarly, in scenario Noise 0.2, an inverted gaussian noise leads to a negligible effect size, and a reversed feedback noise leads to a small effect size (reducing the feedback but also the repetition). The only cases where the noise shows a large effect size on the feedback and repetitions concerned scenario Noise 1 with inverted gaussian and reversed feedback noises, where the average feedback has dropped to 5.83 with inverse gaussian noise, and to 2.88 with reversed feedback. This decrease in feedback is accompanied by a corresponding decrease in repetitions, indicating that the noise significantly affects the credits given to the c-rules, thus the decisions of SAR. We note, however, that both the inverse gaussian and the reversed feedback noises characterize very disruptive types of noises that, if actually occurring in reality at every single interaction step (like in scenario Noise 1), would indicate extremely poor sensors or measurements of the feedback. Applying a reversed feedback noise at every simulation step, for example, corresponds to always providing opposite information from the actual patient's feedback for SAR (e.g., whenever a patient is happy, SAR believes the patient is unhappy, and vice-versa). We can see that this leads SAR to personalizing its behaviors to the opposite behavior that is preferred by a particular patient: observe that the feedback of decision-making system S_{r1} (i.e., a decision-making system that is exposed to the reverse feedback noise in 100% of the cases) has an average of about 2.9. From the results

Table 4

Overview of the Evolving Fuzzy Systems (EFSs) used for comparison with EFS4SAR. Columns *EA* and *MOEA* indicate if the EFS is based, respectively, on evolutionary algorithms and multi-objective evolutionary algorithms. Column *# rules* indicates the approached used to determine the number of rules (one rule per identified cluster, one rule per class, one rule per layer in the case of ensemble models like MEEFIS, or a fixed number).

EFS	Year	EA	MOEA	# rules
SAFLS (Gu and Shen, 2021)	2021	No	No	min. # clusters
FWAadaBoostSOFIES (Gu and Angelov, 2021)	2021	No	No	# classes
MEEFIS (Gu, 2021)	2020	Yes	No	avg nr. layers
PSO-ALMMo (Gu et al., 2021)	2020	Yes	No	# classes
SKMOEFS-MPAES_RCS (Gallo et al., 2020; Antonelli et al., 2014)	2020	Yes	Yes	minimized
ANFIS-GA (Moayedi et al., 2020)	2018	Yes	No	fixed
ANFIS-PSO (Moayedi et al., 2020)	2018	Yes	No	fixed
ALMMo-0 (Angelov et al., 2018)	2017	No	No	# classes
eTS-LS-SVM (Komijani et al., 2012)	2012	No	No	min. # clusters

reported above, we conclude that the proposed decision-making system *S* appears to be robust to the three types of noise that we tested, when occurring with a realistic frequency.

4.2. Fuzzy patients based on human preferences

In this section, we compare EFS4SAR with state-of-the-art evolving fuzzy systems (EFSs). Since most of state-of-the-art EFSs are intended for data streams classification or regression, for comparison purposes we reduce the task of the SAR (i.e., suggesting an activity to the patient) to a data stream classification task. In other words, given a stream of values for the input variables of the SAR, the task of the SAR is to determine an activity (class) to suggest to the patient for each input.

We selected benchmark EFSs among the available implementations of evolving fuzzy classifiers to cover a variety of approaches, including EFSs based on traditional evolving strategies, EFSs based on evolutionary algorithms, and EFSs based on MOEAs. Table 4 reports the selected state-of-the-art EFSs and the corresponding references that provide details on these approaches.

In order to evaluate the EFSs on data that is relevant and realistic for SARs, we use 15 data sets characterizing real-world human preferences about daily activities, which we elicited through a survey with human participants.

Next, we first describe the survey that we conducted to elicit the human preferences, then we illustrate how we used such preferences to model the participants of the survey via fuzzy inference systems, which we used as generators of streams of data to obtain the 15 data sets used in our experiments. Finally, we discuss the results obtained via EFS4SAR and the different EFSs on such data sets.

Human preferences elicitation. We conducted a survey with 15 participants. The participants were asked to indicate which activities they *like*, *do not like*, *love*, or *hate* to do in different circumstances (namely in the *morning*, *afternoon*, and *evening*, during *free*, *working*, or *busy* days, and when it's *sunny*, *cloudy*, or *rainy*). Participants were given a list of 106 examples of activities for their inspiration (divided in 9 categories: *work-related*, *TV Music and Movies*, *Entertainment*, *At home*, *Outdoor*, *Sport*, *Going places*, *Friends and Family*, *Individual*), and they were instructed to indicate new (more personal) ones, if they wanted. Additional activities were considered as part of a 10th category *User defined*. Participants belonged to 4 different nationalities, their age ranged from 20 to 40, and the level of education ranged from undergraduate degree to doctorate. The collected data is also made available in anonymous form in our online supplementary material (Dell'Anna and Jamshidnejad, 2022). Although participants did not suffer from dementia, the corresponding data collected with the survey is relevant and realistic for our purposes, since the main goal of this paper is to evaluate the extent to which EFS4SAR can provide suggestions about daily activities that are both personalized and original.

Data generation. We used the preferences of the participants to construct a model of each participant in the form of a fuzzy inference system whose rule base is constituted by rules encoding the preferences of the participant. Two examples of possible preferences and their

Table 5

Translation of human preferences into fuzzy rules (note that the *love*, *like*, *do not like*, and *hate* translate into, respectively, *very high*, *high*, *low*, and *very low* feedback).

Preference	Fuzzy Rule
when it's sunny and it's a free day I love to exercise	IF <i>weather</i> is <i>sunny</i> AND <i>day-type</i> is <i>free</i> AND <i>to-exercise</i> is <i>performed</i> THEN <i>feedback</i> is <i>very high</i>
when it's a rainy morning during a working day I do not like to work from the office	IF <i>weather</i> is <i>rainy</i> AND <i>day-type</i> is <i>working</i> AND <i>day-time</i> is <i>morning</i> AND <i>to-work-from-the-office</i> is <i>performed</i> THEN <i>feedback</i> is <i>low</i>

Table 6

Number of classes (preferred daily activities) in the 15 real-world data sets. Every data set is composed by 5000 data points, each characterized by 3 dimensions (input variables with numerical values in $[0, 10]$ representing a normalized crisp value for the 3 environmental factors time, weather, and type of the day) and labeled with the activity (class) preferred in that circumstance.

Data set	D0	D1	D2	D3	D4	D5	D6	D7
Classes	13	16	20	21	21	24	24	26
Data set	D8	D9	D10	D11	D12	D13	D14	
Classes	26	26	31	31	36	42	69	

corresponding rules are reported in Table 5. Each resulting FIS is characterized by three input “environmental” variables (*weather*, *day-type*, and *day-time*), each with three possible linguistic realizations explained before, and with a number of input variables based on the indicated preferred daily activities. Each activity is treated as a singleton variable with one possible *default* “performed” value. Each FIS has one output variable, called *feedback*, with 5 possible linguistic values: *very low*, *low*, *average*, *high*, *very high*. The *average* value is used as a default neutral feedback value for cases where no preference is specified. We model the membership function of each value of the input and output variables as triangular functions covering the domain of the variable, ranging from 0 to 10. We use every FIS as a generator of streams of data. Data is constituted of tuples $\langle i_1, i_2, i_3, a \rangle$, where i_1, i_2, i_3 are crisp values randomly sampled from the domain of the three input environmental variables of the FIS, and a is one of the activities preferred by the participant according to the preferences expressed in the survey. More specifically, given inputs i_1, i_2, i_3 , the activity a is obtained as follows: for every possible activity, we compute the feedback by performing fuzzy inference. In 80% of cases, a is randomly selected from those resulting in the highest feedback. In the remaining 20% of cases, a is randomly selected from all activities with positive feedback (i.e., with feedback greater than 5, the middle feedback value).

For every participant we obtain a data set made of 5000 data points, for a total of 15 data sets (see Dell'Anna and Jamshidnejad, 2022 for details). Each data set provides a representation of the preferences (i.e., the preferred activities) of one of the participants given different possible crisp values of the environmental input variables. Table 6

Table 7

The average accuracy, number of rules, feedback and repetitions obtained with the tested EFSs on the 15 data sets. FWAda and SKMOEFS are abbreviations of FWAAdaBoostSOFIES and SKMOEFS-MPAES_RCS, respectively.

	accuracy	# rules	$feedback_s$	$feedback_m$	repetitions
<i>random</i>	0.04 ± 0.02	–	6.49 ± 0.16	7.24 ± 0.32	0.57 ± 0.20
EFS4SAR	0.14 ± 0.05	20.00 ± 0	7.32 ± 0.38	7.92 ± 0.35	1.49 ± 0.37
SAFLS	0.44 ± 0.16	20.53 ± 2.14	8.72 ± 0.43	6.43 ± 1.82	4.70 ± 1.62
FWAada	0.34 ± 0.16	28.40 ± 13.08	8.37 ± 0.39	7.84 ± 1.88	3.16 ± 1.72
MEEFIS	0.40 ± 0.14	27.57 ± 6.00	8.60 ± 0.43	6.23 ± 2.54	5.53 ± 2.99
PSO-ALMMo	0.08 ± 0.05	28.40 ± 13.08	6.93 ± 0.42	7.35 ± 1.26	3.10 ± 0.90
SKMOEFS	0.42 ± 0.15	11.49 ± 4.93	8.64 ± 0.42	6.85 ± 2.24	4.71 ± 1.82
ANFIS-GA	0.07 ± 0.05	28.40 ± 13.08	6.78 ± 0.52	6.51 ± 1.85	3.89 ± 1.35
ANFIS-PSO	0.07 ± 0.04	28.40 ± 13.08	6.82 ± 0.44	6.60 ± 1.78	3.74 ± 1.36
ALLMo-0	0.22 ± 0.10	28.40 ± 13.08	7.83 ± 0.31	8.05 ± 0.67	1.53 ± 0.80
eTS-LS-SVM	0.10 ± 0.12	1.00 ± 0	6.94 ± 1.63	2.36 ± 1.95	10.48 ± 3.19

reports a summary of the number of resulting activities (interpreted as classes) per data set.

Experimental results. Table 7 reports the results in terms of the *feedback* and *repetitions* metrics described in Section 4.1 and in terms of accuracy on the 15 test sets (obtained via standard 80/20 splitting of the data sets) and number of rules identified by the EFSs. As per Section 4.1, we repeat every simulation 5 times to obtain multiple data points for every simulation. $feedback_s$ and $feedback_m$ are obtained by means of the fuzzy patients according to the preferences of the participants of the survey. In the case of $feedback_s$, the fuzzy patients never change their original feedback about activities (analogously to the dashed red curve in Fig. 4(b)). In the case of $feedback_m$, instead, the patients also have a memory and each patient changes its feedback about an activity over time according to its randomly sampled *type* as described in Section 4.1. Row *random* in Table 7 reports, for comparison, the results obtained with a classifier that randomly assigns one of the possible classes to each data point.

First, we note that the most recent benchmark EFSs (e.g., SAFLS (Gu and Shen, 2021), SKMOEFS (Gallo et al., 2020), MEEFIS (Gu, 2021)) provide, in general, higher accuracy than older ones (e.g., ALMMo-0 (Angelov et al., 2018), eTS-LS-SVM (Komijani et al., 2012)) and of EFS4SAR. This indicates that the EFSs successfully identify rules that classify data points in the same cluster with the same class. In our experiments, this also implies, in the case of $feedback_s$, high feedback received from the patients, since the EFSs effectively suggest the activities preferred by the patients. These results, however, also illustrate that for similar data points (i.e., in similar circumstances, e.g., when it's afternoon of a busy day), benchmark EFSs tend to repeat the same suggestion, i.e., they tend to classify two similar data points with the same class. This behavior, reasonable for EFSs that are intended for accurate classification, motivates our work: existing EFSs are effective but they lack originality, since they do not account for repetitions and diversity of suggestions. However, in a real setting these aspects are essential for creative and effective SARs, because continued repetitions of similar activities is expected to lead to a decrease of the effectiveness and feedback received from patients (e.g., as a result of increased boredom due to the repetitions). This is evident when looking at $feedback_m$, where patients take into account the previously suggested activities in providing their feedback. Here, we note that the feedback of the benchmark EFSs tends to drop due to the high number of repetitions.

The results obtained with EFS4SAR, differently, are in line with the results reported in Section 4.2. While EFS4SAR shows a lower feedback compared to the benchmark EFSs in the case of $feedback_s$, it also provides a lower number of repetitions, i.e., for similar data points, EFS4SAR proposes a variety of (sometimes less preferred) activities. The proposed activities are personalized to patients (comparing EFS4SAR against *random*, we see that the feedback with EFS4SAR is consistently higher), while also accounting for the need of regular stimulation of the cognition of the patients. Proposing sometimes less

preferred activities is important because, for example, a patient may hate a certain physical exercise, whereas her/his therapist recommends it for improvement of the patient's mental/physical health. Furthermore, when considering $feedback_m$, we see that except for ALMMo-0, the feedback received by EFS4SAR is higher than the other EFSs.

ALMMo-0 (Angelov et al., 2018) had the most similar performance to EFS4SAR, i.e., few repetitions, which led to lower $feedback_s$ and higher $feedback_m$. ALMMo-0, however, does not explicitly account for creativity in determining the most opportune class, and the results cannot be attributed to an explicit intention to achieve high effectiveness with low repetitions. Differently, the results of EFS4SAR can be clearly explained by means of its parameters related to the importance of low repetitions and high diversity of the suggestions. By modifying at run time these parameters, furthermore, EFS4SAR permits to adjust the degree of creativity based on the evolving needs of the patients.

In terms of interpretability, we note that EFSs that determine the number of rules based on the number of classes (e.g., FWAda (Gu and Angelov, 2021), PSO-ALMMo (Gu et al., 2021), ALMMo-0 (Angelov et al., 2018)) have higher number of rules than other approaches. If the number of possible activities to suggest to the patient grows, the approach of these systems may lead to lower interpretability, and to possibly lower performances if little training data is available. SKMOEFS (Gallo et al., 2020), instead, confirms its effectiveness, as a MOEFS, in optimizing the two conflicting goals of accuracy and number of rules, providing in the case of $feedback_s$ one of the highest accuracy and lowest number of rules. In the context of decision-making of SARs, however, the number of rules should not only be minimized for the sake of interpretability, but also related to the personality and needs of a patient. For example, a SAR should have more rules when interacting with a very stubborn patient who does not easily accept suggestions (so that the robots can try more alternatives), while in other cases it may be sufficient to have only few rules. For this reason, in EFS4SAR, similarly to ANFIS-GA and ANFIS-PSO (Moayedi et al., 2020), the number of c-rules is a *parameter* that can be fixed or varied over time. In our experiments, EFS4SAR used 20 rules for all data sets, illustrating that satisfactory results can be obtained with a low number of rules also when the number of classes/activities is higher.

5. Conclusions and future work

Socially Assistive Robots (SARs) are increasingly used in dementia care to help implementing nonpharmaceutical therapeutic interventions. In this paper, we introduced EFS4SAR, an Evolving Fuzzy logic System for Socially Assistive Robots that combines fuzzy logic systems with evolutionary algorithms to realize a rule-based decision-making approach for SARs that autonomously evolves over time in a creative and therapeutically effective way. EFS4SAR evolves the rules that steer the decision-making of the SAR taking into account the indications of the therapists and caregivers and the experience obtained while interacting with the patients. This leads to a SAR that can provide

long-term personalization and decision-making that accommodate the evolving needs of the patients as unique individuals. Additionally, the use of genetic algorithms and of a diversity factor in the evolution of the rules allows the SAR to determine rules that are not only aligned with the preferences of the patients but are also creative, thus suitable for stimulating and challenging the divergent thinking and cognition of the patients.

We evaluated our proposed approaches by simulating interactions between the SAR and hundreds of artificial agents characterized by different preferences and exhibiting different patterns of behaviors. We quantified the results via two measures that reflect the standard definition of creativity: the originality and therapeutic effectiveness of the activities suggested by the SAR to the patients. Results show that the proposed approaches can effectively learn the preferences of the patients while guaranteeing high diversity in the suggestions, and that the proposal is robust to noise in the feedback received from the patients. Moreover, we compared EFS4SAR against 9 state-of-the-art evolving fuzzy systems using 15 data sets based on real-world human preferences about daily activities. The experiments illustrate that the activities suggested by the SAR employing EFS4SAR result in similar effectiveness as other systems, while possessing higher originality.

In the future, we intend to conduct a more exhaustive experimentation based on more refined and realistic models of patients. We plan to study the scalability of the proposed solution, including a comparison of our system with different encodings of the chromosomes, and with different fuzzy operators for fuzzy inference. Additionally, future work includes the extension of the evolutionary approach to accommodate the evolution of the membership functions that characterize the variables of the fuzzy logic systems; the integration of our approach with the extensive complementary state-of-the-art approaches for short-term personalization and adaptation (e.g., adapting the level of difficulty of a game while playing); and allowing for evolution of the a-rules, in addition to the c-rules. Finally, a clearly important future direction consists in the evaluation of the proposed approaches with human subjects (including patients) interacting with a real robot which will suggest and perform real therapeutic approaches. In this direction, the hyper-parameters of the algorithms will need to be optimally tuned (potentially iteratively after acquiring evidence from the experimental settings) to provide optimal performance, and the robot will need to identify and consider in its decision-making social cues and practices to include personal and institutional norms for dementia patients, and collect actual feedback measurements related to the therapeutic outcome, e.g., in terms of decreased measured stress, or increased cognitive capabilities over time.

Ethics statement

This work involved human subjects in its research. Approval of all ethical and experimental procedures and protocols was granted by the Human Research Ethics Committee of TU Delft under Approval No. 1841.

CRedit authorship contribution statement

Davide Dell'Anna: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing. **Anahita Jamshidnejad:** Conceptualization, Methodology, Validation, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research has been supported by the NWO - Open Competition Domain Science XS project “Creative Social Robots with Evolutionary AI for Dementia Therapy” (OCENW.XS2.055), which has been financed by the Netherlands Organisation for Scientific Research (NWO).

References

- Abdi, J., Al-Hindawi, A., Ng, T., Vizcaychipi, M.P., 2018. Scoping review on the use of socially assistive robot technology in elderly care. *BMJ Open* 8 (2), e018815.
- Angelov, P.P., Filev, D.P., 2004. An approach to online identification of Takagi-Sugeno fuzzy models. *IEEE Trans. Syst. Man Cybern. B* 34 (1), 484–498.
- Angelov, P.P., Filev, D.P., 2005. Simpl_eTS: a simplified method for learning evolving Takagi-Sugeno fuzzy models. In: *Proceedings of the IEEE International Conference on Fuzzy Systems, FUZZ-IEEE 2005*, Reno, Nevada, USA, May 22–25, 2005. IEEE, pp. 1068–1073.
- Angelov, P.P., Gu, X., Príncipe, J.C., 2018. Autonomous learning multimodel systems from data streams. *IEEE Trans. Fuzzy Syst.* 26 (4), 2213–2224.
- Anon, 0000. WHO | Integrated Care for Older People, URL <https://www.who.int/publications/i/item/WHO-FWC-ALC-19.1>.
- Antonelli, M., Ducange, P., Marcelloni, F., 2014. A fast and efficient multi-objective evolutionary learning scheme for fuzzy rule-based classifiers. *Inform. Sci.* 283, 36–54.
- Ascensão, T., Jamshidnejad, A., 2022. Autonomous socially assistive drones performing personalized dance movement therapy: An adaptive fuzzy-logic-based control approach for interaction with humans. *IEEE Access* 10, 15746–15770.
- Bai, Y., Wang, D., 2006. Fundamentals of fuzzy logic control — Fuzzy sets, fuzzy rules and defuzzifications. In: Bai, Y., Zhuang, H., Wang, D. (Eds.), *Advanced Fuzzy Logic Technologies in Industrial Applications*. Advances in Industrial Control. Springer London, London, pp. 17–36.
- Bao, R., Rong, H., Angelov, P.P., Chen, B., Wong, P., 2018. Correntropy-based evolving fuzzy neural system. *IEEE Trans. Fuzzy Syst.* 26 (3), 1324–1338.
- Broekens, J., Heerink, M., Rosendal, H., 2009. Assistive social robots in elderly care: a review. *Gerontechnology* 8 (2), 94–103.
- Brüggemann, S., Bitmead, R.R., 2021. Exponential convergence of recursive least squares with forgetting factor for multiple-output systems. *Automatica* 124, 109389.
- Bruno, B., Mastrogianni, F., Pecora, F., Sgorbissa, A., Saffiotti, A., 2017. A framework for culture-aware robots based on fuzzy logic. In: *Proceedings of the 2017 IEEE International Conference on Fuzzy Systems, FUZZ-IEEE 2017*, Naples, Italy, July 9–12, 2017. IEEE, pp. 1–6.
- Chua, T.W., Tan, W.W., 2011. Non-singleton genetic fuzzy logic system for arrhythmias classification. *Eng. Appl. Artif. Intell.* 24 (2), 251–259.
- Clabaugh, C., Mahajan, K., Jain, S., Pakkar, R., Becerra, D., Shi, Z., Deng, E., Lee, R., Ragusa, G., Mataric, M.J., 2019. Long-term personalization of an in-home socially assistive robot for children with autism spectrum disorders. *Front. Robot. AI* 6, 110.
- Cluzel, F., Yannou, B., Dihlmann, M., 2012. Using evolutionary design to interactively sketch car silhouettes and stimulate designer's creativity. *Eng. Appl. Artif. Intell.* 25 (7), 1413–1424.
- de Jesús Rubio, J., Cruz, D.R., Elias, I., Ochoa, G., Balcázar, R., Aguilar, A., 2019. ANFIS system for classification of brain signals. *J. Intell. Fuzzy Systems* 37 (3), 4033–4041.
- Dell'Anna, D., Jamshidnejad, A., 2022. Supplementary Material for “Evolving Fuzzy Logic Systems for Creative Personalized Socially Assistive Robots”. Zenodo, <http://dx.doi.org/10.5281/zenodo.6673720>.
- Epp, T.D., 2003. Person-centred dementia care: A vision to be refined. *Can. Alzheimer Dis. Rev.* 5 (3), 14–19.
- Fazzolari, M., Alcalá, R., Herrera, F., 2014. A multi-objective evolutionary method for learning granularities based on fuzzy discretization to improve the accuracy-complexity trade-off of fuzzy rule-based classification systems: D-MOFARC algorithm. *Appl. Soft Comput.* 24, 470–481.
- Fazzolari, M., Alcalá, R., Nojima, Y., Ishibuchi, H., Herrera, F., 2013. A review of the application of multiobjective evolutionary fuzzy systems: Current status and further directions. *IEEE Trans. Fuzzy Syst.* 21 (1), 45–65.
- Gad, A.F., 2021. PyGAD: An Intuitive Genetic Algorithm Python Library. *arXiv:2106.06158*.
- Gallo, G., Ferrari, V., Marcelloni, F., Ducange, P., 2020. SK-MOEFs: a library in python for designing accurate and explainable fuzzy models. In: Lesot, M., Vieira, S.M., Reformat, M.Z., ao Paulo Carvalho, J., Wilbik, A., Bouchon-Meunier, B., Yager, R.R. (Eds.), *Information Processing and Management of Uncertainty in Knowledge-Based Systems - 18th International Conference, IPMU 2020*, Lisbon, Portugal, June 15–19, 2020, Proceedings, Part III. In: *Communications in Computer and Information Science*, vol. 1239, Springer, pp. 68–81.
- Garibaldi, J.M., 2019. The need for fuzzy AI. *IEEE CAA J. Autom. Sin.* 6 (3), 610–622.

- Geyer-Schulz, A., 1995. Holland classifier systems. In: Griffiths, M., Whitehouse, D. (Eds.), *Proceedings of the 1995 International Conference on Applied Programming Languages, APL 1995*, San Antonio, Texas, USA, June 4-8, 1995. ACM, pp. 43–55.
- Giannoglou, V.G., Stavrakoudis, D.G., Theocharis, J.B., Petridis, V., 2015. Genetic fuzzy rule based classification systems for coronary plaque characterization based on intravascular ultrasound images. *Eng. Appl. Artif. Intell.* 38, 203–220.
- Gu, X., 2021. Multilayer ensemble evolving fuzzy inference system. *IEEE Trans. Fuzzy Syst.* 29 (8), 2425–2431.
- Gu, X., Angelov, P.P., 2021. Multi-class fuzzily weighted adaptive boosting-based self-organising fuzzy inference ensemble systems for classification. *IEEE Trans. Fuzzy Syst.* 1, (Early Access).
- Gu, X., Shen, Q., 2021. A self-adaptive fuzzy learning system for streaming data prediction. *Inform. Sci.* 579, 623–647.
- Gu, X., Shen, Q., Angelov, P.P., 2021. Particle swarm optimized autonomous learning fuzzy system. *IEEE Trans. Cybern.* 51 (11), 5352–5363.
- Hannemann, B.T., 2006. Creativity with dementia patients. *Can creativity and art stimulate dementia patients positively?* *Gerontology* 52 (1), 59–65.
- Hemminghaus, J., Kopp, S., 2017. Towards adaptive social behavior generation for assistive robots using reinforcement learning. In: Mutlu, B., Tscheligi, M., Weiss, A., Young, J.E. (Eds.), *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction, HRI 2017*, Vienna, Austria, March 6-9, 2017. ACM pp. 332–340.
- Hidalgo, D., Cervantes, L., Castillo, O., Melin, P., Martínez-Soto, R., 2020. Fuzzy parameter adaptation in genetic algorithms for the optimization of fuzzy integrators in modular neural networks for multimodal biometry. *Comput. Syst.* 24, 1093–1105.
- Hornby, G., Globus, A., Linden, D., Lohn, J., 2006. Automated antenna design with evolutionary algorithms. In: *Space 2006*, San Jose, California. p. 7242.
- Jang, J.S., Lee, J.S., Yoo, D.H., 2015. Effects of spaced retrieval training with errorless learning in the rehabilitation of patients with dementia. *J. Phys. Ther. Sci.* 27 (9), 2735–2738.
- Jost, L., 2006. Entropy and diversity. *Oikos* 113 (2), 363–375.
- Kasabov, N.K., Song, Q., 2002. DENFIS: Dynamic evolving neural-fuzzy inference system and its application for time-series prediction. *IEEE Trans. Fuzzy Syst.* 10 (2), 144–154.
- Kazil, J., Masad, D., Crooks, A., 2020. Utilizing python for agent-based modeling: The mesa framework. In: Thomson, R., Bisgin, H., Dancy, C., Hyder, A., Hussain, M. (Eds.), *Social, Cultural, and Behavioral Modeling*. Springer International Publishing, Cham, pp. 308–317.
- Kennedy, J., Eberhart, R., 1995. Particle swarm optimization. In: *Proceedings of International Conference on Neural Networks (ICNN'95)*, Perth, WA, Australia, November 27 - December 1, 1995. IEEE, pp. 1942–1948.
- Komijani, M., Lucas, C., Araabi, B.N., Kalhor, A., 2012. Introducing evolving takagi-sugeno method based on local least squares support vector machine models. *Evol. Syst.* 3 (2), 81–93.
- Koshiyama, A.S., Tanscheit, R., Vellasco, M.M., 2019. Automatic synthesis of fuzzy systems: An evolutionary overview with a genetic programming perspective. *Wiley Interdiscipl. Rev.: Data Min. Knowl. Discov.* 9 (2), e1251.
- Kumar, A., 2013. Encoding schemes in genetic algorithm. *Int. J. Adv. Res. IT Eng.* 2 (3), 1–7.
- Lehman, J., Clune, J., Misevic, D., Adami, C., Altenberg, L., Beaulieu, J., Bentley, P.J., Bernard, S., Beslon, G., Bryson, D.M., et al., 2020. The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *Artif. Life* 26 (2), 274–306.
- Liberati, G., Raffone, A., Belardinelli, M.O., 2012. Cognitive reserve and its implications for rehabilitation and alzheimer's disease. *Cogn. Process.* 13 (1), 1–12.
- Liu, P., Glas, D.F., Kanda, T., Ishiguro, H., Hagita, N., 2014. How to train your robot - teaching service robots to reproduce human social behavior. In: *The 23rd IEEE International Symposium on Robot and Human Interactive Communication, IEEE RO-MAN 2014*, Edinburgh, UK, August 25-29, 2014. IEEE, pp. 961–968.
- Lu, T., 2015. Genetic-algorithm-based type reduction algorithm for interval type-2 fuzzy logic controllers. *Eng. Appl. Artif. Intell.* 42, 36–44.
- Lughofer, E., 2008. FLEXFIS: A robust incremental learning approach for evolving Takagi-Sugeno fuzzy models. *IEEE Trans. Fuzzy Syst.* 16 (6), 1393–1410.
- Lughofer, E., 2015. Evolving fuzzy systems - fundamentals, reliability, interpretability, useability and applications. In: Rosa, A.C., Guervós, J.J.M., Dourado, A., Cadenas, J.M., Madani, K., Ruano, A.E.B., Filipe, J. (Eds.), *Proceedings of the 7th International Joint Conference on Computational Intelligence (IJCCI 2015) - Volume 1: ECTA, Lisbon, Portugal, November 12-14, 2015*. SciTePress, p. 11.
- Lughofer, E., Angelov, P.P., 2011. Handling drifts and shifts in on-line data streams with evolving fuzzy systems. *Appl. Soft Comput.* 11 (2), 2057–2068.
- Maciel, L., Ballini, R., Gomide, F.A.C., 2017. Evolving possibilistic fuzzy modeling for realized volatility forecasting with jumps. *IEEE Trans. Fuzzy Syst.* 25 (2), 302–314.
- Marti, P., Bacigalupo, M., Giusti, L., Mennecozzi, C., Shibata, T., 2006. Socially assistive robotics in the treatment of behavioural and psychological symptoms of dementia. In: *Proceedings of the First IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechanics, BioRob 2006*, Pisa, Italy, 20-22 Feb. 2006. IEEE pp. 483–488.
- Mitchell, M., 1998. *An Introduction to Genetic Algorithms*. MIT press, Cambridge, MA.
- Moayed, H., Raftari, M., Sharifi, A., Jusoh, W.A.W., Rashid, A.S.A., 2020. Optimization of ANFIS with GA and PSO estimating α ratio in driven piles. *Eng. Comput.* 36 (1), 227–238.
- Mobahi, H., Ansari, S., 2003. Fuzzy perception, emotion and expression for interactive robots. In: *Proceedings of the IEEE International Conference on Systems, Man & Cybernetics: Washington, D.C., USA, 5-8 October 2003*. IEEE, pp. 3918–3923.
- Moro, C., Nejat, G., Mihailidis, A., 2018. Learning and personalizing socially assistive robot behaviors to aid with activities of daily living. *ACM Trans. Hum.-Robot Interact. (THRI)* 7 (2), 1–25.
- Palmiero, M., Di Giacomo, D., Passafiume, D., 2012. Creativity and dementia: a review. *Cogn. Process.* 13 (3), 193–209.
- Pratama, M., Anavatti, S.G., Angelov, P.P., Lughofer, E., 2014a. PANFIS: A novel incremental learning machine. *IEEE Trans. Neural Netw. Learn. Syst.* 25 (1), 55–68.
- Pratama, M., Anavatti, S.G., Lughofer, E., 2014b. GENEFIS: Toward an effective localist network. *IEEE Trans. Fuzzy Syst.* 22 (3), 547–562.
- Pratama, M., Pedrycz, W., Lughofer, E., 2018. Evolving ensemble fuzzy classifier. *IEEE Trans. Fuzzy Syst.* 26 (5), 2552–2567.
- Ramachandran, A., Huang, C., Gartland, E., Scassellati, B., 2018. Thinking aloud with a tutoring robot to enhance learning. In: Kanda, T., Sabanovic, S., Hoffman, G., Tapus, A. (Eds.), *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction, HRI 2018*, Chicago, IL, USA, March 05-08, 2018. ACM, pp. 59–68.
- Rebok, G.W., Ball, K., Guey, L.T., Jones, R.N., Kim, H.-Y., King, J.W., Marsiske, M., Morris, J.N., Tennstedt, S.L., Unverzagt, F.W., Willis, S.L., ACTIVE Study Group, 2014. Ten-year effects of the advanced cognitive training for independent and vital elderly cognitive training trial on cognition and everyday functioning in older adults. *J. Am. Geriatr. Soc.* 62 (1), 16–24.
- Rey, D., Neuhaus, M., 2011. Wilcoxon-signed-rank test. In: Lovric, M. (Ed.), *International Encyclopedia of Statistical Science*. Springer, Berlin, Heidelberg pp. 1658–1659.
- Rong, H., Sundararajan, N., Huang, G., Saratchandran, P., 2006. Sequential Adaptive Fuzzy Inference System (SAFIS) for nonlinear system identification and prediction. *Fuzzy Sets and Systems* 157 (9), 1260–1275.
- Rossi, S., Ferland, F., Tapus, A., 2017. User profiling and behavioral adaptation for HRI: A survey. *Pattern Recognit. Lett.* 99, 3–12.
- Runco, M.A., Jaeger, G.J., 2012. The standard definition of creativity. *Creativity Res. J.* 24 (1), 92–96.
- Sabanovic, S., Bennett, C.C., Chang, W.L., Huber, L., 2013. PARO Robot affects diverse interaction modalities in group sensory therapy for older adults with dementia. In: *Proceedings of the 13th IEEE International Conference on Rehabilitation Robotics, ICORR*, Seattle, WA, USA, June 24-26, 2013. IEEE, pp. 1–6.
- Sharif, F., Jahanbin, I., Amirad, A., Moghadam, M.H., 2018. Effectiveness of life review therapy on quality of life in the late life at day care centers of Shiraz, Iran: A randomized controlled trial. *Int. J. Commun. Nurs. Midwifery* 6 (2), 136–145.
- Sharkawy, A.B., 2010. Genetic fuzzy self-tuning PID controllers for antilock braking systems. *Eng. Appl. Artif. Intell.* 23 (7), 1041–1052.
- Skrjanc, I., Andonovski, G., Ledezma, A., Sipele, O., Iglesias, J.A., Sanchis, A., 2018. Evolving cloud-based system for the recognition of drivers' actions. *Expert Syst. Appl.* 99, 231–238.
- Skrjanc, I., Iglesias, J.A., Sanchis, A., Leite, D.F., Lughofer, E., Gomide, F.A.C., 2019. Evolving fuzzy and neuro-fuzzy approaches in clustering, regression, identification, and classification: A survey. *Inform. Sci.* 490, 344–368.
- Subramanian, K., Das, A.K., Sundaram, S., Ramasamy, S., 2014. A meta-cognitive interval type-2 fuzzy inference system and its projection based learning algorithm. *Evol. Syst.* 5 (4), 219–230.
- Syriopoulou-Delli, C.K., Gkiolnta, E., 2022. Review of assistive technology in the training of children with autism spectrum disorders. *Int. J. Dev. Disabil.* 68 (2), 73–85.
- Tapus, A., Tapus, C., Mataric, M.J., 2008. User-robot personality matching and assistive robot behavior adaptation for post-stroke rehabilitation therapy. *Intell. Serv. Robot.* 1 (2), 169–183.
- Tapus, A., Tapus, C., Mataric, M.J., 2009. The use of socially assistive robots in the design of intelligent cognitive therapies for people with dementia. In: *Proceedings of the 2009 IEEE International Conference on Rehabilitation Robotics, Kyoto, Japan, 23-26 June, 2009*. IEEE, pp. 924–929.
- Thrift, P.R., 1991. Fuzzy logic synthesis with genetic algorithms. In: Belew, R.K., Booker, L.B. (Eds.), *Proceedings of the 4th International Conference on Genetic Algorithms*, San Diego, CA, USA, July 1991. Morgan Kaufmann, pp. 509–513.
- Trautwein, S., Scharpf, A., Barisch-Fritz, B., Niermann, C., Woll, A., 2017. Effectiveness of a 16-week multimodal exercise program on individuals with dementia: study protocol for a multicenter randomized controlled trial. *JMIR Res. Prot.* 6 (3), e6792.
- Tsiakas, K., Abujelala, M., Makedon, F., 2018. Task engagement as personalization feedback for socially-assistive robots and cognitive training. *Technologies* 6 (2), 49.

- Umbrico, A., Cesta, A., Cortellessa, G., Orlandini, A., 2020. A holistic approach to behavior adaptation for socially assistive robots. *Int. J. Soc. Robot.* 12 (3), 617–637.
- Valenzuela-Rendón, M., 1991. The fuzzy classifier system: A classifier system for continuously varying variables. In: Belew, R.K., Booker, L.B. (Eds.), *Proceedings of the 4th International Conference on Genetic Algorithms*, San Diego, CA, USA, July 1991. Morgan Kaufmann, pp. 346–353.
- Van Mierlo, L., Van der Roest, H., Meiland, F., Dröes, R., 2010. Personalized dementia care: proven effectiveness of psychosocial interventions in subgroups. *Ageing Res. Rev.* 9 (2), 163–183.
- Van Wylsberghe, A., 2016. *Healthcare Robots: Ethics, Design and Implementation*. London, Routledge.
- Vaughan, S., Wallis, M., Polit, D., Steele, M., Shum, D., Morris, N., 2014. The effects of multimodal exercise on cognitive and physical functioning and brain-derived neurotrophic factor in older women: a randomised controlled trial. *Age Ageing* 43 (5), 623–629.
- Vitiello, A., Acampora, G., Staffa, M., Siciliano, B., Rossi, S., 2017. A neuro-fuzzy-Bayesian approach for the adaptive control of robot proxemics behavior. In: *Proceedings of the 2017 IEEE International Conference on Fuzzy Systems, FUZZ-IEEE 2017*, Naples, Italy, July 9–12, 2017. IEEE, pp. 1–6.
- Warner, J., Sexauer, J., scikit-fuzzy, twmeggs, alexsavio, Unnikrishnan, A., Castelhão, G., Pontes, F.A., Uelwer, T., pd2f, laurazh, Batista, F., alexbuy, den Broeck, W.V., Song, W., Badger, T.G., Pérez, R.A.M., Power, J.F., Mishra, H., Trullols, G.O., Hörteborn, A., 99991, 2019. JDWarner/scikit-fuzzy: Scikit-Fuzzy version 0.4.2. Zenodo, <http://dx.doi.org/10.5281/zenodo.3541386>.
- Webb, N., Ruiz-Garcia, A., Elshaw, M., Palade, V., 2020. Emotion recognition from face images in an unconstrained environment for usage on social robots. In: *Proceedings of the 2020 International Joint Conference on Neural Networks, IJCNN 2020*, Glasgow, United Kingdom, July 19–24, 2020. IEEE, pp. 1–8.
- Woods, B., O'Philbin, L., Farrell, E.M., Spector, A.E., Orrell, M., 2018. Reminiscence therapy for dementia. In: *Cochrane Database of Systematic Reviews*, vol. 3. John Wiley & Sons, Ltd.