



Delft University of Technology

Tensions in transparent urban AI designing a smart electric vehicle charge point

Alfrink, Kars; Keller, Ianus; Doorn, Neelke; Kortuem, Gerd

DOI

[10.1007/s00146-022-01436-9](https://doi.org/10.1007/s00146-022-01436-9)

Publication date

2022

Document Version

Final published version

Published in

AI and Society

Citation (APA)

Alfrink, K., Keller, I., Doorn, N., & Kortuem, G. (2022). Tensions in transparent urban AI: designing a smart electric vehicle charge point. *AI and Society*, 38(3), 1049-1065. <https://doi.org/10.1007/s00146-022-01436-9>

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.



Tensions in transparent urban AI: designing a smart electric vehicle charge point

Kars Alfrink¹ · Ianus Keller² · Neelke Doorn³ · Gerd Kortuem¹

Received: 29 April 2021 / Accepted: 10 March 2022
© The Author(s) 2022

Abstract

The increasing use of artificial intelligence (AI) by public actors has led to a push for more transparency. Previous research has conceptualized AI transparency as knowledge that empowers citizens and experts to make informed choices about the use and governance of AI. Conversely, in this paper, we critically examine if transparency-as-knowledge is an appropriate concept for a public realm where private interests intersect with democratic concerns. We conduct a practice-based design research study in which we prototype and evaluate a transparent smart electric vehicle charge point, and investigate experts' and citizens' understanding of AI transparency. We find that citizens experience transparency as burdensome; experts hope transparency ensures acceptance, while citizens are mostly indifferent to AI; and with absent means of control, citizens question transparency's relevance. The tensions we identify suggest transparency cannot be reduced to a product feature, but should be seen as a mediator of debate between experts and citizens.

Keywords Urban AI · Artificial intelligence · Transparency · Electric vehicles

1 Introduction

Digital technologies such as big data, sensor networks and artificial intelligence (AI) are becoming increasingly important in the control of urban infrastructure, and public administration more broadly (Chiusi et al. 2020; Crawford et al. 2019). However, it is now widely recognized such AI systems may lead to unfair outcomes, even if they have been designed with the best intentions (Eubanks 2018; Ranchorás 2020). These concerns have prompted researchers,

governments and civil society groups to formulate ethical principles for deployment and use of AI, emphasizing values such as transparency, fairness and accountability (Jobin et al. 2019; Mittelstadt et al. 2016; Tsamados et al. 2021). Likewise, some cities have started to embrace a digital rights agenda and are formulating principles and policies to govern public AI systems [e.g., (The Cities Coalition for Digital Rights 2021)].

Many ethical and policy frameworks see *transparency* as an important prerequisite for ensuring fairness and public acceptance (Brauneis and Goodman 2018; Stoyanovich and Howe 2018). Empirical research in human–computer interaction (HCI) has focused on identifying which forms of user interface-level transparency are most effective for increasing user understanding and trust (Abdul et al. 2018). In this HCI-research, transparency is typically framed as a form of objective knowledge that empowers people to make informed choices about how best to use and govern AI systems. However, researchers have started to point out theoretical and practical limitations of the transparency ideal (Ananny and Crawford 2018), and the importance of considering the human experience of AI transparency (Alvarado and Waern 2018; Vakarelov and Rogerson 2020). What is more, in case of *public* AI systems, such as those controlling urban infrastructure, i.e., “urban AI”, the relationship

✉ Kars Alfrink
c.p.alfrink@tudelft.nl

Ianus Keller
a.i.keller@tudelft.nl

Neelke Doorn
n.doorn@tudelft.nl

Gerd Kortuem
g.w.kortuem@tudelft.nl

¹ Sustainable Design Engineering, TU Delft, Delft, The Netherlands

² Human Centered Design, TU Delft, Delft, The Netherlands

³ Values, Technology and Innovation, TU Delft, Delft, The Netherlands

between users and those who design, develop and govern systems is different from commercial settings: These systems effectively *enact policy* (Shaw and Graham 2017; Kitchin et al. 2017), and users are not simply consumers, but also *citizens* who are entitled to democratic control over policy, AI-enacted or otherwise.

Therefore, our aim is to examine the degree to which transparency-as-knowledge is a suitable concept for urban AI systems in both an empirical and critical way. We contribute to the ongoing discussion of transparent AI by investigating diverging conceptions of transparency between those who design, develop and govern urban AI systems (hereafter “experts”), and users of those same systems (“citizens”).

We focus on smart electric vehicle (EV) charging as an empirical ground for studying transparency in urban AI systems. Smart EV charging serves as a useful example of how urban AI shapes the lived experience of cities, and of city-making itself. Smart EV charging facilities augment and mediate both public spaces and travel spaces. In this context, many stakeholders consider transparency an essential ingredient for ensuring public acceptance (Döbelt et al. 2015; Madhu et al. 2019; Fabianek et al. 2020). Using a practice-based design research approach (Koskinen 2011), we collaborated with commercial companies and the municipality of Amsterdam to prototype and evaluate a *transparent* smart EV charge point which provides EV drivers with explanations of smart charging decisions.

Our findings shed light on several tensions between motivations experts have for providing transparency, such as social acceptance, and attitudes and expectations citizens have towards urban AI systems, such as indifference or a desire for control.

In what follows, we first briefly provide context on smart EV charging, and the Amsterdam design project that formed the basis for our empirical work. We then summarize work on transparency in HCI design research, philosophy of technology, and the social sciences of big data and AI. Subsequently, we describe the field study we undertook with the design project prototype. Following this, we offer six narrative themes to capture our findings with regards to expert understanding and citizen experience of transparency. In the concluding discussion, we contextualize these findings in light of the literature, and examine the main points of tension between expert understanding and citizen experience.

2 Background and motivation

2.1 Smart electric vehicle charging

Electric mobility is seen by many cities as a key way to improve efficiency and equity of the flow of goods and people, and to reduce negative externalities including air

pollution and climate change (Geels 2012). However, in OECD countries, there are indicators electric grid capacity is not sufficient to support the growing number of EVs (Muratori 2019). In general, this is not an issue of overall energy availability, but of limited grid capacity (Huang and Kockelman 2020). This concern is especially relevant for local distribution grids in cities and neighborhoods where EVs are particularly prevalent and where charging sessions are clustered around peak times. If demand for charging exceeds supply, not every vehicle can be charged, and choices need to be made: who will be charged first, and who last? For this reason, energy network providers have started to deploy “smart charging” solutions, which make timing and capacity of charging dependent on factors such as grid capacity, electricity demand and availability of renewable energy. Smart charging allows for dynamic management of demand by curtailing the rate and amount of electricity EVs can charge when connected to a charge point (Wang et al. 2016; Mehta et al. 2018; Moghaddam et al. 2018; Frendo et al. 2019). The use of AI in governing the grid and charge points makes it possible to increase the number of EVs by more than 60% without having to upgrade physical grid infrastructure (Ofgem 2018).

Of course, EVs are not an unambiguously positive development, nor should the transition to EVs be considered inevitable. In fact, EVs are a contested subject involving many social, political and ethical debates. To name but a few concerns: EVs may perpetuate existing car culture, increased electricity needs may not be met by renewable sources, battery production depends on exploitation of limited mineral resources with adverse social and ecological consequences, and EV battery recycling itself can cause pollution (Ortar and Ryghaug 2019).

Smart grid solutions may reinforce and accelerate practices producing energy demand peaks, rather than contributing to more sustainable ways of living (Strengers 2012). A focus on solving the problem of demand also distracts from rethinking everyday practices requiring energy in the first place (Strengers 2014, 2013). In other words, smart EV charging can be seen as a form of “technological solutionism” (Morozov 2013), where social ills are framed as problems to be fixed by means of technology, while avoiding structural change (Foth et al. 2021).

In any case, smart charging solutions significantly alter the EV charging experience: EVs may charge slower than expected; drivers may be disadvantaged by receiving less electricity or slower charging rates than other drivers, even if both cars are plugged in at the same time and charge point. It may also have unexpected side effects such as some neighborhoods receiving less electricity than others. In short, use of AI makes EV charging less predictable. From the perspective of experts, this threatens social acceptance. Transparency promises to contribute

to people's understanding of and trust in smart EV charging systems.

2.2 'The transparent charging station'

As of July 2021, the city of Amsterdam operates 2503 charging stations, or 4974 charging points (Gemeente Amsterdam 2021). Of these, at the time of our study, 452 stations (904 points) were part of a smart charging system called Flexpower (Gemeente Amsterdam 2020). This system increases charge speed when solar energy is available, and decreases speed around peak times when the grid is used more intensively.

Prompted by rising public concern about the risks of the Internet of Things and AI, in 2016 electric grid operator Alliander¹ and EV charging knowledge institute ElaadNL,² commissioned a design study from design agency The Incredible Machine³ to develop ways of making smart charging transparent for EV drivers. The outcome was the *Transparent Charging Station*, a speculative design prototype of a smart charge point using a video game metaphor for visualizing automated charging decisions (Turel et al. 2017). A key aspect of the *Transparent Charging Station* is the use of priority schemes: for example, shared EVs would get priority access to charge faster, sooner and more than non-shared private vehicles. The design study received significant public interest but also raised questions about the meaning, viability and utility of transparency in the context of a street-level public service.

A follow-up project, *UI for Smart EV Charging*, was initiated in 2019 by the same knowledge institute and design agency, who were joined by the municipality of Amsterdam and the Amsterdam Institute for Advanced Metropolitan Solutions.⁴ The aim was to develop a prototype transparency interface inspired by but distinct from the *Transparent Charging Station* speculative prototype. The project built on the newly formulated digital agenda of the city of Amsterdam entitled *A Digital City for and by Everyone*, which lays out values and ambitions for a "free and inclusive digital city" in which the digital rights of all residents are protected (Gemeente Amsterdam 2019). This design was aimed at a solution compatible with existing Flexpower charge points in an effort to further study the technical feasibility, usefulness, usability and desirability of transparency provided through a screen-based user interface. The first and last author agreed to become part of this project group to consult during the design phase, and to lead the evaluation of

the design solution. Simultaneously, we pursued our independent research agenda into the varying conceptions of transparency by major direct stakeholders involved in urban AI projects.

3 Related work

Transparency is a widely held and discussed moral and political value, especially in settings where informed consent, accountability and deliberation are emphasized. In the context of AI, in particular when developed using machine learning (ML), transparency commonly refers to visibility and accessibility of information related to a system's functioning. Opacity of AI systems can stem from a variety of sources: deliberate secrecy by system developers and operators; lack of technical literacy of the observer; or technical properties of systems themselves (Burrell 2016). In particular, transparent AI aims to provide *explanations* of model behavior. Such explanations can be arrived at by developing models that are *interpretable* by humans, for example because they are rule-based. When models are developed with techniques producing opaque or "black box" models resisting human interpretation, explanations can still be produced in a post-hoc fashion by means of a supplemental explanation model (Zhang and Chen 2020; Kim et al. 2020).

In debates around social and ethical ramifications of AI, transparency has quickly become a central if contested notion. Many view transparency as a desirable value, either for moral reasons, or because it aids understanding and increases trust. Others point out AI systems resist straightforward explanation due to their socio-technical nature, and warn against transparency shifting responsibility from system developers to users.

Surveying the literature in HCI design research, philosophy of technology, and interdisciplinary work on big data and AI, we can identify this same emphasis on the relationship between transparency, understanding, and trust. There is also a growing body of critical work exploring transparency's limits.

3.1 Transparency, understanding and trust

The main vehicle for creating transparency of AI systems on the level of user interfaces is through so-called "explanations", informational and/or interactive elements communicating some aspect of an AI system's workings. Various kinds of explanations can contribute to people's understanding of an AI system (Rader et al. 2018).

Explanation completeness and soundness impact the fidelity of end users' mental models. Explanations with a high level of completeness have the lowest perceived cost and highest benefit. However, this favorable cost-benefit

¹ <https://www.alliander.com>.

² <https://www.elaad.nl>.

³ <http://www.the-incredible-machine.com>.

⁴ <https://www.ams-institute.org>.

perception does hinge on users being able to adjust system behavior (Kulesza et al. 2013). Furthermore, when users feel they are able to form an adequate mental model from simply interacting with systems, explanations are less likely to be considered beneficial, because they take attention away from primary tasks (Bunt et al. 2012).

Increasing transparency by providing explanations can improve people's trust in AI systems (Kizilcec 2016; Eslami et al. 2018; Binns et al. 2018). User literacy of AI systems may mediate the degree to which explanations increase trust (Shin 2021). There does not appear to be a single best way of explaining a system to increase trust (Binns et al. 2018). In some cases, trust only increases as a result from explanations when user expectations have been violated by system behavior (Kizilcec 2016). There may also be a bell-curved relationship between information amount, and user trust. Providing too much information can actually *erode* trust (Kizilcec 2016). Trust does not appear to be impacted by people's objective understanding of systems, nor by the form of explanation used (Cheng et al. 2019). There is some evidence explanations need not even be *truthful* to increase user trust (Eiband et al. 2019).

User trust may also be impacted significantly by their attitudes to the larger systems that form the context of automated decision-making. For example, looking at the application of AI in child welfare services, Brown et al. (2019) find people's distrust of non-automated systems increases their discomfort with AI.

A tension exists between making people aware of AI's functioning and preventing them from developing behaviors at odds with system developer goals. A level of obfuscation is necessary to prevent bad actors from gaming the system, whereas a lack of transparency reduces people's sense of control and makes them unsure about how their behavior might impact outcomes (Alvarado and Waern 2018; Jhaver et al. 2018; Eslami et al. 2019).

3.2 Critiques of transparency

In philosophy of technology and interdisciplinary work on the social implications of big data and AI, critical efforts have explored the limitations of the transparency ideal.

The language of transparency suggests we are removing things obscuring our view, while in fact transparency requires active production of information (Menéndez-Viso 2009). However, more data do not necessarily lead to better understanding. In our current age, it is not a lack of information but a sheer abundance of data obscuring our understanding (Caduff 2017). Furthermore, when we strive to make automated decisions explainable, we should be wary of the distinction between appearing transparent and actually being transparent. The latter requires *actionable* information, that is to say, information humans can use as a resource for their

own decision-making (Vakarelov and Rogerson 2020). Publishing (non-actionable) information in an effort to merely appear transparent can be a form of “tokenism” or “engagement theatre” in that it does not actually increase democratic control over urban AI systems (Kamols et al. 2021; Monno and Khakee 2012). Indeed, large tech companies use transparency initiatives at least in part to stave off government regulation (Grandinetti 2021). Another risk of focusing on transparency is that it makes us less likely to consider if we want an AI to determine a particular aspect of our lives at all (De Laat 2019).

Transparent AI is often treated as an issue best dealt with behind closed doors by experts. AI presents challenges for traditional HCI design in general (Holmquist 2017), and participatory design approaches in particular (Bratteteig and Verne 2018). However, a small but growing body of work seeks to bridge the gap between advocating for abstract principles and supporting design choices situated in context (Aizenberg and Hoven 2020) and opening up AI development processes to a broader range of stakeholders (Krafft et al. 2021). In this way, users and citizens gain control over ways in which transparency is implemented in particular AI systems so they support their needs.

The transparency ideal can reinforce a neoliberal model of human agency, in which perfectly informed and fully consenting individuals make rational decisions that in aggregate produce improved social outcomes (Ananny and Crawford 2018). This model reduces citizenship to one of consumer choice. If on the basis of information provided a person disagrees with a system's functioning, they are expected to defect to a competing but sufficiently equivalent service. In the language of Hirschman (1970), this is the “exit” option. The alternative is “voice”: expressing disagreement to effect change. Hirschman suggests we focus on the latter because making space for and responding to feedback increases “loyalty”. Since we are dealing with public AI systems, relying on exit alone is problematic because citizens should have a say in the operation of these systems, and limiting participation to “voting with your feet” infringes on people's right to the city (Shaw and Graham 2017; Sadowski and Pasquale 2015; Cardullo and Kitchin 2017; Foth et al. 2015).

Making AI system data and models visible is not the same as holding whole socio-technical assemblages accountable. For this, it is necessary to see who has power to change systems, and to be able to experiment with changes ourselves (Ananny and Crawford 2018; Hollanek 2020; Dourish 2016). Socio-technical complexity of urban AI systems may exceed individual capacity, in which case it can only be understood collectively (Innerarity 2021). Transparency can also not account for cases in which a system's behavior deviates from design intent due to adversarial attacks (Descampe et al. 2021). Furthermore, there are ways of increasing accountability that do not depend on transparency at all.

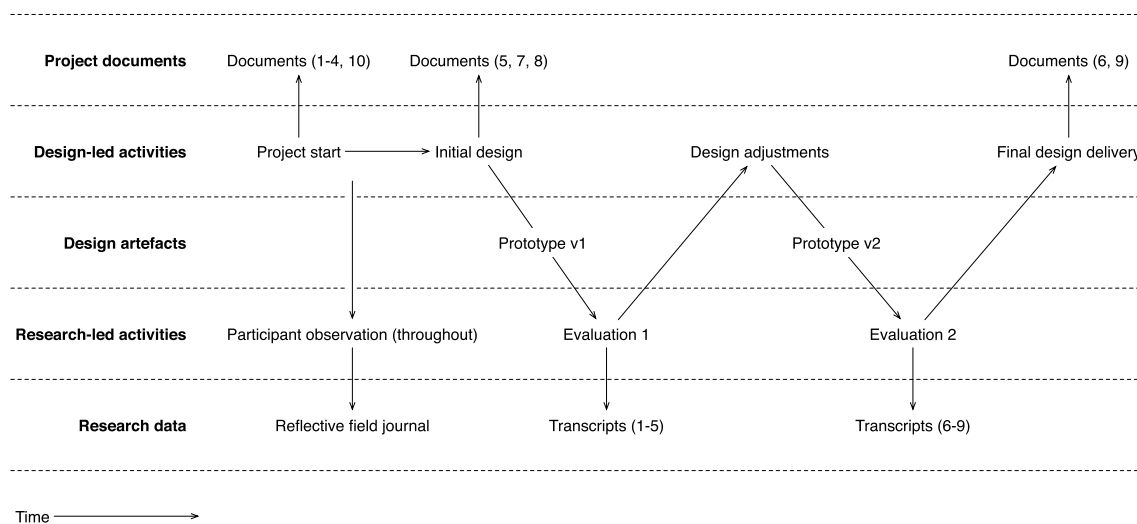


Fig. 1 Overview of ‘UI for Smart EV Charging’ project structure

One example is to introduce ways for AI systems to exercise “discretion”, to diverge from baked-in policy in cases where user dissatisfaction with system behavior is detected (Alkhatib and Bernstein 2019). Another is to include means for decision subjects to *contest* AI system decisions, to make them responsive to requests for human intervention (Walmsley 2021; Vaccaro et al. 2019; Sarra 2020; Almada 2019; Hirsch et al. 2017).

Empirical work in HCI indicates transparency through user interface-level explanations can contribute to understanding and trust. However, both understanding and trust achieved in this way are highly contingent and may not even be justified in the objective sense. At the same time, critical work points to the limitations of the transparency ideal, often questioning the motivations of system designers and developers, pointing out how their understanding and valuing of transparency may not be the same as users. When conflicting mental models and values are glossed over, design processes and outcomes are likely to suffer. Since public AI systems enact policy and in such settings, users are also citizens, the ways in which transparency interfaces mediate the relationship between experts and citizens should be considered together. We therefore argue it is necessary to improve our understanding of the varying conceptions of transparency by the major direct stakeholders involved.

4 Methods

Our overall research approach is qualitative-interpretive. To investigate how experts understand transparency, and how citizens experience a transparent AI system, we conducted what Koskinen (2011) describes as a practice-based design research study in the “field” mode. In this approach, design

methods such as interventions with prototypes in real-world settings are used to generate research data. In addition, we draw on participatory action research (Kemmis and McTaggart 2006) for our active involvement in the industry project that produced the prototype. Our research consists of two main activities: (1) participation as design experts in an industry project to observe how experts conceptualize and implement transparent urban AI; and (2) evaluation in the field of the resulting design to understand how citizens experience urban AI transparency, as implemented in a transparent smart EV charge point. These activities we undertook as part of the *UI for Smart EV Charging* project (Sect. 2.2). Figure 1 provides an overview of the project structure.

Data collected consisted of project documents, field observations, and interviews. The first author was present at all meetings of the design team to observe and participate in the discussions. A reflective field journal was kept and documents produced during this phase, such as the design agency’s project proposal and slide decks used during presentations (D1 through D10), were stored for future analysis. Analysis was performed using reflexive thematic analysis (Braun and Clarke 2006).

4.1 Design process and evaluation

The design and evaluation were done in sequence. The design phase was led by the design agency. Their starting points were their previous experience on the preceding speculative project, the project proposal drafted by the project partners to acquire funding (D10), and a requirements document developed by the project partners (D4). In the initial design phase, five EV drivers were interviewed to acquire insight into user needs (D7). A typical user journey for EV charging was mapped (D5). The first author spent



Fig. 2 The design prototype was evaluated with EV drivers recruited on the spot at a fast charging facility

a day at the design agency to ideate various approaches for the design. Following an exploration of various design options and feedback from the consortium, a final design was chosen. This design was developed into a high-fidelity non-functional prototype.

For evaluation, a fast charging facility centrally located in the Netherlands was selected as the field site. We set up the prototype next to fast charge points (Fig. 2) and invited people who came to charge their car to participate in the study. If they agreed, we went through an information sheet and consent form. To improve ecological validity, we did not provide participants with details on how the system operates, beyond telling them we were testing the design of a “smart” EV charge point, that adjusts speed based on

a number of (unspecified) factors. Subsequently, we asked them to perform the task of charging their car using our prototype. While they did so, we invited them to think out loud and occasionally prompted them with open-ended questions. After completing the task, we followed up with a semi-structured interview to dig deeper into their experience with the prototype. All sessions were recorded using video and audio. Still photographs were also taken. Furthermore, researchers took hand-written notes while observing. Overall, we conducted two rounds of one-day long design evaluations: round one included five participants (P1 through P5; 1 female, 4 male), whereas round two included four participants (P6 through P9; 1 female, 3 male). Audio recordings of the evaluations were subsequently transcribed for further analysis. All participant and some document quotes in this paper’s results section were translated from Dutch by the first author.

4.2 Prototype

The prototype consists of a 1:1 scale cardboard replica of the charge points in use in Amsterdam. The signage on the stations is reproduced and ports have been added for actual charge connectors to fit into. A 12.9-inch tablet is attached to the top of the charge point for the transparency interface to run on.

Figure 3 shows a selection of screens from the prototype (translated from Dutch by the first author). The basic structure consists of: (1) an idle screen, (2) a screen shown once charging has started, and (3) a screen shown after charging




Fig. 3 Key screens of prototype v2

has concluded. We distinguish two types of screen elements: those supporting the task of charging (e.g., a prompt to swipe a card to begin) and elements aiming to make the smart charging system *transparent*, i.e., explanations.

The prototype screens were created in a graphics package and data reflecting the imagined scenario of use was added. The screens were collected in presentation software so it was possible to advance them using a concealed wireless remote control in response to participant actions. We created two versions of the user interface design, v1 and v2. V2 addressed some basic usability issues detected during the initial round of evaluations. These usability fixes aside, both versions of the prototype are identical.

4.2.1 'Rules' as explanations

The main means of providing transparency is a set of elements that together list the "rules" governing system behavior. For each rule, its currently active state is displayed along with a short descriptive name. In v1, the other possible states are also immediately shown and a few lines of additional explanatory text are included. In v2, each rule can be tapped to reveal a modal box including the additional text, the other possible states and a graph or diagram offering a visual explanation. Once charging starts, each rule also includes an indication of how it impacts the charge speed. V1 uses amperes (A) as the indicator of charge speed (actually the unit of current). V2 instead uses kilowatt (kW) (the unit of power).

The screen displayed when a charging session is finished uses the conceit of a cash register receipt to show how much the user had charged in total, expressed in kWh. The receipt also shows any changes to each rule that may have occurred during charging while the user was away. A QR-code and a unique URL for the charge session are also displayed and some text next to it explains the code can be scanned or the URL accessed to receive a digital copy of this receipt.

4.3 Analysis

Two datasets were analyzed: dataset 1 compiles documents produced during the design project; dataset 2 compiles data from prototype evaluation sessions with EV drivers. Analysis was done using the qualitative data analysis software *Atlas.ti*. The data were first coded inductively by the first author. Codes were repeatedly refined and grouped into an initial set of themes. The second author independently coded a subset of the data, and a refined set of themes was jointly developed. The first author also checked the codes and themes with the commercial collaborators. Finally, the themes were once more condensed into the final, smaller, richer and more narrative set presented in this paper.

5 Results

We generated six themes related to accounts of transparency in the data. Two themes for experts were derived from design project documents: (X1) truthful information produces transparency; and (X2) transparency enables fairness assessment. Four themes for citizens were derived from prototype evaluation session transcripts: (C1) transparency mediates concern; (C2) transparency is burdensome; (C3) transparency invites strategic behavior and (C4) transparency evokes desire for control. Almost all the data were included in the themes.

5.1 Expert understanding of transparency

5.1.1 Theme X1: truthful information produces transparency

Experts talk about transparency as something created by providing *truthful* information about "automated decisions" (D4). The issue with these decisions is that they are *opaque*, hidden inside "black boxes" (D10).

However, algorithms that currently control smart city objects are "black boxes": the public is affected by their decisions, but does not know what factors are taken into consideration and how they are weighed against each other to reach a decision. (D10)

Here, we get a glimpse of what decisions an AI system makes: it weighs various factors against each other. However, throughout the documents, *decision-making* and *prioritization* are used interchangeably. We can also see that not only decisions, but motivations for them must be made transparent.

Prioritization appears to produce dilemmas. Some people will lose and others will win out in resource distribution.

When a city service is scarce, prioritization is required. By using smart applications, cities need to make the prioritization beforehand and program it explicitly. It's possible to prioritize on: Target groups, like citizens, disabled, professionals, etc; Shared vehicles; Price; Time slots. (D10)

This notion of dilemmas connects to one of the driving motivations for pursuing transparent smart charging. It is not so much a moral imperative, but a pragmatic one. The concern is that opacity threatens *acceptance* of EV driving and charging by citizens.

Visibility [of] the automatic decision making in the smart charging process can help the adoption of this new technology. (D10)

The conceptual metaphor used by stakeholders to describe how transparency is achieved is sometimes explicitly vision-based. Apparently, automated choices can be made *visible*.

The Transparent Charging Station will provide insight into this by making the underlying choices of the algorithm visible on the display. (D1)

5.1.2 Theme X2: transparency enables fairness assessment

Using transparency, users should be able to determine if they have been fairly treated. Fairness assessment is impacted by design choices. For example, at one point during the design process, the design agency emphasized they had moved away from determining fairness by comparison.

It is not about understanding fairness by comparing your treatment to that of fellow chargers, it is about whether you think the (choice for) parameters and weights is fair. (D6)

Fairness is also invoked on the level of messaging. One of the aims of the design project is to convey a “positive message” (D4) about the municipality’s role in the transition to fully electric driving in the city.

The core of the message is that the interests of different parties are fairly represented in order to arrive at solutions that work for citizens, government and private parties as smoothly as possible. (D4)

Project members agreed this message should be conveyed using a “positive tone of voice” (D4). When discussing tone of voice, fairness is once again invoked, although it may also be understood as truthfulness, because the Dutch word for both truthful and fair is the same (“eerlijk”) and it is not entirely clear from context which meaning is intended here.

The design and all communication around it are based on a positive tone of voice (truthful, predictable, not too difficult, positive connotation, municipality listens, no algorithmic doom scenarios). (D4)

We see truthfulness and fairness recur on different levels throughout the project. Truthfulness is seen as a quality of information provided by the system, producing transparency about how a person is treated. This treatment can be more or less fair, the assessment of which is enabled by truthful information. The system is imagined to convey a message that people are indeed being treated fairly. That every party with an interest smart EV charging, citizens included, is given fair consideration. Finally, fairness and truthfulness are (ambiguously) invoked as a desired *tone* of messages.

5.2 Citizen experience of a transparent AI system

5.2.1 Theme C1: transparency mediates concern

In general, participants were welcoming of automated decision-making in the EV charging process. Many responded positively to the notion of using automation to optimize EV charging towards what could be described as common interests: a stable electric grid, a fair distribution of power, and sustainability in general terms.

In addition, many people seemed more or less indifferent to the presence of automated decision-making. For example, when asked if any automated decisions had been made, P6 responded “Yes, but based on what was already there.” By which they meant, the system was simply responding to the inputs it sensed in the environment. P5 commented they were sure there were “technicians who have thought about it ...” In other words, they put their faith in the expertise of the people who built the system. P6 simply stated “I take it the way it is.”

One of the most striking statements for us was when P7 said “I don’t think I should be able to make a choice about that,” referring to trade-offs between collective interest at the expense of individual efficiency.

People’s indifference to AI may be in part due to the fact that when charging in the city, less is at stake compared to say a fast charging session. Charge speed is slow, session duration is short, and out of all charging options (at home, at work, at a fast charging facility, in the city) public charge points are the least depended on. P7: “on the one hand when I’m going to run an errand and I’m done within the hour yeah then I don’t care how fast.” Any charge received while parking is considered a bonus.

A few participants *did* express concerns about situations in which they would be disadvantaged by the system, and the impossibility of making a one-time exception. For example, when they were in a hurry, or when they were forced to charge during peak hours.

The strongest reactions *against* automated decision-making related to the shared car priority feature. Many participants latched on to this, while ignoring most of the other rules made transparent. P2: “Shared car has priority. I don’t like that but okay. Sustainable of course.” Some recognized it would be beneficial for sustainability reasons, so they did accept its rationale. But none of the participants were shared car drivers themselves. Some participants wondered about what was considered a shared car, who determined this, and how the system would deal with, for example, shared cars from outside of Amsterdam.

P2, a resident of Amsterdam, made a connection between shared car priority and local politics, which recently had taken a more left-leaning, progressive turn than in years before. They expressed fear of politicians pushing for more

extreme forms of shared car priority at the expense of private car owners.

Well ... See if it will be that way later ... You will get a political decision. Politicians are going to say yes but ... If there is a shared car, the other cannot get in and so on, you have to let three shared cars go first. Especially with Femke⁵ in Amsterdam, I am a little afraid of that. (P2)

The comment was made somewhat in jest, but it does stand in remarkable contrast to the general indifference to automated decision-making which we have tried to capture thus far.

This discomfort with treating some EVs differently from others may be due in part to the scarcity of charge points. It can be a challenge to find a free spot. If one ends up next to a shared car and charges slower as a result, it feels unfair.

Well, that shared car [priority] makes me go gosh darn it ... I find it very annoying. At a busy time I am ... Racing through the city and all the stations were occupied in the neighborhood. Then I arrive here and then I actually get punished a bit more. Then they would have had to put a few more stations in Zuid [affluent city area] ... We all have a Tesla. So that's a bit complicated. (P2)

There appears to be a relationship between people's attitude towards automated decision-making, and the purposes to which it is put. This is different from the narrative about people being suspicious of all automation, regardless of where it is applied. It also sheds a different light on in which cases transparency is necessary or desired.

5.2.2 Theme C2: transparency is burdensome

With this theme, we capture how EV charging is often a sub-optimal experience, made worse by the additional demands transparency puts on people.

First, EV charging in general is an error-prone activity. Poor design and engineering of charge points and wider infrastructure frequently lead to failed charging attempts. In our prototype evaluations, most participants started a charge session in the “wrong” way, even though instructions were listed on the opening screen. They typically made their way through it in a trial and error fashion: swiping a card and plugging in a connector in succession until the system progressed to the next state, not taking time to read any instructions beforehand. It is a usability truism users do not read, and people charging an EV are clearly no exception. P7

acknowledged as much when she responded to the explanation by saying “So anyway, that only makes sense if you read it very carefully.” It is likely this situation is *even worse* outside of a prototype evaluation because when EV charging at a public charge point, people are likely to be in a rush. They might have someone else waiting to use the same charge point, and in any case, they will probably have somewhere else to be. Therefore, as P7 pointed out, they are not inclined to study a user interface at length when they are setting up their EV for charging.

In the city I am not going to do that ... I think. Certainly not when I go shopping ... usually it is like let's get it over with and then you want to go on again ... I would be very interested in how it works, but I would rather see that afterwards. (P7)

A final source of unease is uncertainty over the amount of charge delivered. While charging, prototype v1 displayed the real-time charge speed in amperes (amp), a measure of current. This was a largely meaningless indicator for participants.

This – 12 ampere doesn't tell me very much. I really benefit from seeing where I am at now and how much time it will take me to get to 100%. So what percentage am I at and how much time does it take me to get to 100%. I think that's important. (P5)

V2 switched to kilowatt (kW), a measure of power. Participants could at least extrapolate from this real-time measure to an expected amount of energy received at session end. Most participants were also able to translate a session's worth of energy charged measured in kWh to range, because they had memorized the capacity of their EV battery. Or, they compared the listed amount to what they knew a fast charge point delivers. Needless to say, all this mental arithmetic meant more work for participants, and although some did take pride in their ability to perform it, most were perfectly happy to offload all of it onto a system.

For all of these reasons, it should come as no surprise many participants reported feeling overwhelmed by explanations. P9: “There is already a lot of information on it, I must say.” Participants do not welcome additional demands put upon them by this information when all they want to do is charge their EV. P3: “I think it's a lot of information. ... I just want to charge.” Additional information lead to confusion. P7: “I think it's too much info. Honestly, it's confusing. From the start I actually think I see way too much.” This confusion is caused at least in part because participants think they are expected to act on it somehow.

Apparently, participants are focused on completing the task of EV charging with confidence and minimal hassle. The information added to the interface in the interest of transparency does not directly support task completion.

⁵ Femke Halsema, at the time mayor of Amsterdam and former leader of the national Green Left party.

Since, as we captured with the previous theme, participants are largely indifferent to AI, this information is experienced mostly as a burden.

This is kind of competing for my attention. ... What I want to know, just step 1, 2, 3, 4 of my actions but this here is a lot of information that makes me go "what should I do with it?" (P9)

5.2.3 Theme C3: transparency invites strategic behavior

Participants expressed intent to adapt their behavior to the system, something that could be considered an unintended side effect of providing more transparency. Participants suggested they might pick a different time to charge, so that they would benefit from increased speed during off-peak hours, or would enjoy extra speed when the sun was out. Another reason for changing behavior particularly in relation to solar power availability was sustainability.

Some participants were driven less by a desire to be more sustainable and were more interested in reaping economic benefits. For example, P5, while discussing the "receipt" displayed at the end of a charging session, talked about how they were most interested in changing their behavior so they would *pay less*.

What I would like to see here is ... What have I paid? ... What I can do better to charge better next time. Now I'm looking at this, but I don't immediately see what I can do about it. Do I have to go to the green bar? How can I influence it? (P5)

Regardless of whether a participant was looking to improve charging speed, sustainability, or cost, availability of more information about smart charging system operation appears to inspire an intent to optimize behavior.

It should be noted not all participants were as keen on changing their behavior. Most participants happily speculate about what *other* people might do with the information provided. P8: "Well then people are aware and they charge at a certain time when they have the choice." But when we put them on the spot, they frequently admitted they did not expect to change anything about their behavior *themselves*. This could be because benefits do not outweigh additional effort. It could also be because, particularly in the context of city charging, time and place of a charging session is strongly dictated by circumstances, such as availability of free parking spots. As both P1 and P8 stated: "If I have to charge, I have to charge."

5.2.4 Theme C4: transparency evokes desire for control

Many participants interpreted explanations *themselves* as things to be interacted with. When confronted with the

opening screen, P1 asked "But do I have to choose something or not?" and while looking at the status indicators said: "What those are? ... A choice, I can tap on." (P1) And P3 commented "I think that's very interesting for you. That people look at such a screen in this way. They are staring at it going 'what the hell should I choose.'"

Most participants expected to be able to change things to their advantage, not only in terms of charge speed but also in terms of price: "Here I can say 'shared car yes or no'" (P4) "What I see here I could determine myself ... Determine the best price-quality ratio." (P5) "That you can interact. So if I do that, it will be cheaper or more expensive. Or it will go faster or slower. You kind of think that." (P2).

In some cases participants wanted to have the choice to be altruistic. P1 expressed a desire to decide for themselves if they would indeed give priority to the shared car connected to the charge point they were using: "I gave something to the shared car. So I've done my good deed for the day." It should be noted being nice to shared cars would really depend on circumstances. "Shared car priority. So here you can choose how nice you want to be. But maybe also if you are in a hurry then you are in a hurry." (P1).

Those who were more or less indifferent to AI typically did not respond too strongly to the revelation they in fact could not exercise any control. Some participants, however, did respond with some chagrin, such as P3 here:

Well, it is clear to me that I actually decide almost nothing and the system decides everything for me. I find that strange... Because I... Well, the system itself makes decisions while I don't know what exactly is happening. (P3)

P3 clearly did not experience the sense of agency expected to result from transparency. Other participants warned against not including user choice because that would lead to rejection from users: "I think that people will respond negatively more quickly maybe ... Because then you can only grumble about it." (P1).

Furthermore, lack of control in some cases leads to participants questioning the value of including any explanations at all.

Just now I had the idea that I could opt for green energy. Or I can choose to deliver back to the net. So I thought I could make choices. But basically it's just plug in and all kinds of stuff happens. But I have no influence on it. So then I think why should I have all that information? If I have no influence on it, what should I do with it? (P3)

When asked how they would deal with automated decisions they disagree with, the majority of participants seemed somewhat resigned to accepting the system's functioning as it is presented to them. Although one might expect some




Fig. 4 Tensions between expert motivations and citizen experiences. Experts implement AI to achieve energy infrastructure reliability and sustainability. A need for legal compliance and societal acceptance drives efforts to make AI transparent. This transparency increases citizen awareness, which triggers several responses: concern over

optimization targets, increased cognitive burden, various forms of adaptive behavior, and an increased desire for control. Each response impacts citizen trust in positive or negative ways, which in turn affects social acceptance desired by experts

participants would want to make their disagreement known to system operators, most stated they would simply try to find a different (non-smart) charge point.

If I know I have no influence on things right now, because I can't change any settings. When I know that a decision is made for me and I do not like them at that time, or faster loading or yeah, that car is charging really fast, it is a shared car, well, in that case I can choose to to find another charge point. (P4)

Making attempts to influence system developers hardly came up. P1, when asked what they would do if they disagreed with how things worked, said: “Well I tweet. I’ll put it on twitter.” This same participant felt if you cannot make any choices, if everything is automated, this is bad, because “then you can only grumble about it.”

6 Discussion

Using a reflective thematic analysis of design process documents and prototype evaluation transcripts, we have captured ways in which a group of experts understand transparency, and how the transparent urban AI system resulting from their efforts, a transparent smart EV charge point, is experienced by citizens. In this next section, we reflect on three tensions constructed from a comparison between the two groups of narrative themes (Fig. 4). These are: (1) information quality over quantity; (2) level of concern; (3) sense of control.

6.1 Tension #1: information quality over quantity

We have found that according to experts, transparency is created by providing truthful information about automated decisions (X1). However, the belief that being truthful leads to increased trust is not born out by previous

research. For example, placebo explanations can still improve trust (Eiband et al. 2019). In addition, explanations can be satisfying to users without necessarily being truthful (Eslami et al. 2018). What is more, users typically cannot ascertain *correctness* of system output from an explanation alone (Rader et al. 2018).

Experts believe that, because automated decisions might benefit some more than others, and because AI is by its nature hidden, they need to be made visible (X1). This can be seen as an example of setting too high standards of explanation for machine decisions (Zerilli et al. 2019). A vast number of decisions are made by city governments benefiting some more than others, and a lot of city governance recedes from the view of ordinary citizens. Not all such decisions are made visible to the extent pursued in this project. On the other hand, it can be argued AI systems demand a higher degree of transparency precisely *because* of their technical nature. When ML techniques produce proper black-box models, such systems are fundamentally less predictable and more opaque than a human equivalent (Günther and Kasirzadeh 2021).

The project can be considered an example of how pursuing transparency sidelines the question “should we be using AI at all?” (De Laat 2019). Experts did not question the decision to use an AI system in EV charging. The pursuit of transparency serves to support the ongoing use of AI.

Experts in our study talk about transparency in terms of *making the hidden visible*, as if an AI is something that can be seen and there is merely something obstructing our view in need of removal. However, what we see them *do* in our project is something quite different, and supports the notion that transparency requires active production of information (Menéndez-Viso 2009). The language our experts use is also at odds with the view that understanding a system requires more than seeing inside it, that it requires being able to

change systems, and seeing how they behave in relation to their environment (Ananny and Crawford 2018).

Information provided in the interest of transparency is frequently experienced by citizens as burdensome (C2). It is *not* perceived to be supportive of charging an EV, a task rather error-prone and stress-inducing to boot.

Citizens apparently perceive the benefits of engaging with explanations to not outweigh the costs. This can be due to explanations lacking completeness (Kulesza et al. 2013), although the interface in question did provide explanations of a range of system aspects. A more likely reason is that citizens believe they can form an adequate mental model of the system by simply *using* it, where “adequate” may be next to no model at all (Bunt et al. 2012).

In other words, information *quality* over quantity is the actual moral problem of transparency (Menéndez-Viso 2009). The information produced in the interest of transparency can occlude as much as it reveals, and in the process add to anxiety already felt in response to being subjected to automated decision-making (Ananny and Crawford 2018; Caduff 2017). Others have argued for a distinction between explanations in direct support of tasks and explanations of AI decisions with only an indirect connection to user actions (Rader et al. 2018). Our project more closely matches the latter category. Perhaps, when AI is indirectly connected to user actions, explanations should be made subordinate to information which is in support of tasks, and made available upon explicit user request, or when a high likelihood of a need for explanations has been detected through implicit signals.

There exists a tension here between experts’ desire to make citizens aware of the presence of an AI and the aim of old-fashioned user interface design in support of user tasks. In a world where AI is part of an increasing number of systems, we should ask ourselves if user interfaces are the proper location for raising awareness.

This point of tension is of particular importance when dealing with street-level touch-points of urban AI systems, because people’s attention tends to be even more limited due to pressures of everyday urban activities.

6.2 Tension #2: level of concern

Experts pursue transparency because citizens may reject use of automated decision-making in public infrastructure (X1). The position of experts here mirrors the idea that transparency has no moral content (Menéndez-Viso 2009), it is a means towards an end. For our experts, the goal is acceptance by the public, EV drivers in particular. Various factors may lead to this acceptance, but the reasoning in the project seems to be that once citizens understand the workings of a system, they can ascertain its fairness, and once they know it is fair, they will accept the system. However, the

literature we have reviewed provides a mixed view of the relationships between transparency, understanding, fairness and trust (Binns 2018; Cheng et al. 2019; Kizilcec 2016). What this means is that if acceptance is the aim, relying on transparency is likely insufficient or may even backfire, for example when an excess of information decreases people’s trust (Kizilcec 2016). It may also be the case that a limited form of transparency offering justification of decisions is more likely to increase perceived legitimacy of a system, rather than the more extensive transparency of the automated decision-making process *itself* attempted in this project (Fine Licht and Fine Licht 2020).

If we consider responses of some citizens to the explanations provided, it appears trust may have *decreased* instead of improved. Transparency enabled them to see some of the decisions affecting charge speed, and in some cases had them wondering about system developer motivations to include those factors, most notably in the case of the shared car priority feature. Being critical is not necessarily the same as being distrustful, but uncritical acceptance is unlikely to be the result of transparency. This suggests transparency efforts should be prepared for public debate with users around issues uncovered by transparency. Furthermore, such opportunities for “voice” may in fact *increase* trust (Hirschman 1970).

The experience of citizens is characterized by an overall acceptance of, or even indifference to, the presence of automated decision-making (C1). AI is seen as a convenient way of optimizing towards broadly shared collective interests such as electric grid stability and sustainability.

This echoes others’ findings that people find AI useful, but crucially also that the AI systems people find *most* useful is not necessarily the one they find most fair (Dolin et al. 2018). Potentially, also in our project, people are making trade-offs between fairness and usefulness. This could go some way towards explaining our participants’ general indifference to the use of AI in EV charging.

It is only when something is at stake (illustrated in our study by the shared car priority feature) that citizens start to question AI (C1). Others have pointed out different people respond to transparency in different ways. As a result, people also *trust* a given system in different ways (Ananny and Crawford 2018). The various motivations our participants have for EV driving (e.g., saving money, protecting the environment) may influence how they respond to AI features exposed to them in the interface. Prior experiences with organizations deploying AI influence the extent to which people trust AI (Brown et al. 2019). Similarly, the extent to which participants consider local government, power companies, etc. to be “on their side” also figures into how they perceive use of AI for automated decisions.

AI opacity *as such* is hardly ever an issue for citizens. This suggests experts should focus more on contested issues

(such as air quality, parking space, congestion) and how automated decision-making *interacts* with those. A typical line of reasoning is that technology can improve on these issues, and should, therefore, be welcomed. This is certainly a general driver behind the push for electric mobility in Amsterdam. At the same time it is felt AI lacks transparency and that this should be fixed. What is *not* considered is how a person's view on an issue like spatial justice may affect the extent to which they welcome EV charging, and by extension a smart charging system, *regardless* of how transparent it is. This suggests transparency efforts should be focused more towards those matters that are actually contested, and how AI *mediates* those issues.

6.3 Tension #3: sense of control

Experts believe explanations are actionable by citizens (X2). Experts presume explanations make it possible for citizens to assess the *fairness* of decisions by evaluating inputs, processes and outcomes of “the AI”, by having access to a justification for the AI's design, and by knowing who “owns” the AI.

However, actionability is influenced not only by content but also by the format of explanations. For an explanation to be actionable, it must be usable as “currency” in a person's decision-making process (Vakarelov and Rogerson 2020). Having the ability to assess fairness by itself does not equip a person to *act* on that assessment. The current transparent EV charge point design addresses this issue through the previously described “receipt” feature (Sect. 4.2.1). However, none of our participants spontaneously suggested they would use those resources if they disagreed with the system's functioning.

Citizens intend to use explanations as a resource for adapting behavior towards altruistic or egoistic ends (C3). The tension between transparency and the possibility of “gaming” behavior, as well as concerns over exposure of intellectual property has been noted previously (Alvarado and Waern 2018; Jhaver et al. 2018). In our project, we did not see experts express such worries.

The fact that citizens indicate desire to change behavior in response to explanations suggests information *is* actionable to some degree. However, actionable information by itself does not appear to provide sufficient means of influence over system behavior.

Explanations created expectations of user control, an ability to override automated decisions (C4). The absence of control leads some participants to question the relevance of explanations.

Others have shown transparency paired with *control* can alleviate anxiety caused by automated decisions (Jhaver et al. 2018). Users can have favorable cost–benefit perceptions of explanations if they are able to act on provided information

by adjusting system behavior (Kulesza et al. 2013). Possibly, our participants' responses to the transparency interface would have been quite different had it also offered means of directly or indirectly influencing the operation of the AI. This underscores the fact that explanations by themselves are not always perceived as sufficiently actionable. Furthermore, it is also possible that, rather than alleviating frustrations around lack of control through something akin to seamful design (Alvarado and Waern 2018; Chalmers and Galani 2004), further opening up of systems actually produces even greater anxiety in users who are already overwhelmed by explanations.

This desire for control appears to be at odds with the fact that experts are not willing or able to offer direct control over system behavior and anticipate explanations alone to be a sufficient form of accountability. Citizens do not want to, or believe they are not able to, petition experts for changes to system behavior, despite the presence of explanations that could be leveraged for purposes of recourse. This suggests more explicit channels for voice should be made available in or around touchpoints of AI systems.

In case of disagreements with automated decisions, most citizens opt to defect to an alternative means of charging, rather than try and influence policies shaping system behavior (C4).

Some have argued transparency invokes a neoliberal model of agency (Ananny and Crawford 2018; Cardullo and Kitchin 2017; Kitchin 2019). We think our project is a clear illustration of this logic in action. Our participants did not appear to feel they had a substantial say in the operation of the system. This could be because there were no clear “channels for voice”. It could also be because people's lack of “loyalty” to the organization deploying the system made them disinclined to go through the trouble of acting on their disagreements (Hirschman 1970; Centivany and Glushko 2016). In any case, we feel that in our project, transparency alone was insufficient for creating any form of agency beyond EV drivers simply not “buying” the services of the charge point they disagree with, a reduction of citizen participation to consumer choice. The implications of this for experts is that, if shared control over system behavior is desired, “voice” needs to be put on the table. Citizens need to be able to discuss and debate the significance of information they are provided through transparency, in a dialog with system designers, developers and operators. This suggests efforts to make urban AI more transparent should be paired with more participatory and collaborative approaches to city governance and policy-making (Dourish 2010; Foth 2018; Frauenberger et al. 2018).

In closing, we argue these tensions show transparency should not be seen as a property of technology, but must be understood as a communicative process between experts and citizens, who are more than mere users. AI systems mediate this process,

inviting some actions, and inhibiting others (Verbeek 2006). Understanding of a system is not the product of simply receiving and processing information. Understanding emerges from debate between stakeholders, and is always provisional. All three tensions we identify (information quantity, level of concern, desire for control) in various ways point to the need for additional channels for voice through which this debate can be facilitated.

Some of our findings relate quite specifically to the vagaries of smart EV charging. An example would be the challenge of finding an intuitive measure of charge speed. Furthermore, the class of AI we worked on is a deterministic one. The design solutions pursued in our project may not be viable when dealing with stochastic systems. At the moment, EV driving is accessible mostly to significantly affluent or professionally employed people. Certainly, our sample of citizens skews highly educated. It is likely this informs their attitudes to issues such as sustainability and automation. Our expert-citizen distinction leaves out a lot of other relevant stakeholders to investigate. Indeed, our citizens were all direct stakeholders and users of the system. There are plenty of citizens who are not EV drivers, but who are likely to be impacted by the roll-out of EV charging infrastructure in various ways. For example, by reduced neighborhood parking space, or by continued prioritization of road space for cars. There are also a good deal of stakeholders who have relevant expertise, but who are not considered “experts” in the sense we have been using here, that is to say those stakeholders who have some level of formal influence over the shape the system takes.

7 Conclusion

We have presented findings from a practice-based design research study investigating diverging conceptions of transparency by expert and citizen stakeholders of an urban AI system. Our expert participants believe transparency is achieved by providing truthful information about automated decisions. They expect citizens are able to assess system fairness using this information, and that they can act on this information. Meanwhile, our citizen participants are largely indifferent to AI, and they primarily experience explanations as burdensome, and question their relevance if they are not accompanied with the ability to override system decisions.

Transparency is a growing topic of interest in HCI design research, and in public discourse, it is commonly invoked as a solution to negative effects of AI opacity. As a result, transparency has also been taken up as a desirable system property in urban AI systems development. Our findings illustrate it is necessary to remain critical of assumptions driving the pursuit of transparency in user interfaces of AI systems. Transparency puts additional cognitive demands on people, and shifts responsibility of ensuring fairness onto them, reinforcing a neoliberal model of agency.

For these reasons, we believe transparency should be reframed. It should not be seen as a property of a system through which information flows from experts to individual users. Rather, transparency must be seen as a communicative process between experts *and* citizens, mediated by AI systems.

Acknowledgements We would like to thank our field study participants for their rich responses, as well as the *UI for Smart EV Charging* project consortium members for being supportive of our participation: the Amsterdam Institute for Advanced Metropolitan Solutions, ElaadNL, The Incredible Machine, and the municipality of Amsterdam.

Author contributions All the authors contributed to the study conception and design. Material preparation, data collection and analysis were performed by KA. Supplemental analysis was performed by IK. The first draft of the manuscript was written by KA and all the authors commented on previous versions of the manuscript. All the authors read and approved the final manuscript.

Funding This research was supported by a grant from the Dutch National Research Council NWO (grant no. CISC.CC.018).

Availability of data and material Materials and data will be made available at the 4TU.ResearchData repository under <https://doi.org/10.4121/14501130>.

Code availability Not applicable.

Declarations

Conflict of interest Not applicable.

Ethics approval Approval was obtained from the university’s Human Research Ethics Committee.

Consent to participate All participants in this study gave their informed consent.

Consent for publication All participants in this study gave their informed consent.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abdul A, Vermeulen J, Wang D, Brian BY, Lim Y, Kankanhalli M (2018) Trends and trajectories for explainable, accountable and intelligible systems: an HCI research agenda. In: Conference on human factors in computing systems—proceedings, 1–18 April 2018
- Aizenberg E, van den Hoven J (2020) Designing for human rights in AI. *Big Data Soc* 7(2):205395172094956
- Alkhatib A, Bernstein M (2019) Street-level algorithms: a theory at the gaps between policy and decisions. In: Conference on human factors in computing systems—proceedings
- Almada M (2019) Human Intervention in automated decision-making: toward the construction of contestable systems. In: Proceedings of the 17th international conference on artificial intelligence and law, ICAIL 2019, pp 2–11
- Alvarado O, Waern A (2018) Towards algorithmic experience: initial efforts for social media contexts. In: Conference on human factors in computing systems—proceedings, 1–9 Apr 2018
- Ananny M, Crawford K (2018) Seeing without knowing: limitations of the transparency ideal and its application to algorithmic accountability. *New Media Soc* 20(3):973–989
- Binns R (2018) Algorithmic accountability and public reason. *Philos Technol* 31(4):543–556
- Binns R, Van Kleek M, Veale M, Lyngs U, Zhao J, Shadbolt N (2018) ‘It’s reducing a human being to a percentage’: perceptions of justice in algorithmic decisions. In: Conference on human factors in computing systems—proceedings, 1–14 Apr 2018
- Bratteteig T, Verne G (2018) Does AI make PD obsolete?: Exploring challenges from artificial intelligence to participatory design. In: Proceedings of the 15th participatory design conference: short papers, situated actions, workshops and tutorial, vol 2, pp 1–5. ACM, Hasselt
- Braun V, Clarke V (2006) Using thematic analysis in psychology. *Qual Res Psychol* 3(2):77–101
- Brauneis R, Goodman EP (2018) Algorithmic transparency for the smart city. *Yale JL Tech* 20:103
- Brown A, Chouldechova A, Putnam-Hornstein E, Tobin A, Vaithianathan R (2019) Toward algorithmic accountability in public services, pp 1–12
- Bunt A, Lount M, Lauzon C (2012) Are explanations always important? p 169
- Burrell J (2016) How the machine ‘thinks’: understanding opacity in machine learning algorithms. *Big Data Soc* 3(1):1–12
- Caduff C (2017) Targets in the cloud: on transparency and other shadows. *Sci Technol Human Values* 42(2):315–319
- Cardullo P, Kitchin R (2017) Being a ‘citizen’ in the smart city: up and down the scaffold of smart citizen participation, pp 1–24
- Centivany A, Glushko B (2016) “Popcorn tastes good”: participatory policymaking and reddit’s “AMAgeddon”. In: Conference on human factors in computing systems—proceedings, pp 1126–37
- Chalmers M, Galani A (2004) Seamless interweaving: heterogeneity in the theory and design of interactive systems. In Proceedings of the 2004 conference on designing interactive systems processes, practices, methods, and techniques—DIS ’04, 243. ACM Press, Cambridge
- Cheng H-F, Wang R, Zhang Z, O’Connell F, Gray T, Maxwell Harper F, Zhu H (2019) Explaining decision-making algorithms through UI: strategies to help non-expert stakeholders. In: Proceedings of the 2019 CHI conference on human factors in computing systems, Paper 559. Association for Computing Machinery, Glasgow. <https://doi.org/10.1145/3290605.3300789>
- Chiusi F, Fischer S, Kayser-Bril N, Spielkamp N (2020) Automating society report 2020. Algorithm Watch. <https://automatingsociety.algorithmwatch.org>
- Crawford K, Dobbe R, Dryer T, Fried G, Green B, Kazianus E, Kak A et al (2019) AI now 2019 report. AI Now Institute, New York. https://ainowinstitute.org/AI_Now_2019_Report.html
- De Laat PB (2019) The disciplinary power of predictive algorithms: a foucauldian perspective. *Ethics Inf Technol* 21(4):319–329
- de Fine Licht K, de Fine Licht J (2020) Artificial intelligence, transparency, and public decision-making: why explanations are key when trying to produce perceived legitimacy. *AI Soc* 35(4):917–926
- Descampe A, Massart C, Poelman S, Standaert F-X, Standaert O (2021) Automated news recommendation in front of adversarial examples and the technical limits of transparency in algorithmic accountability. *AI Soc*. <https://doi.org/10.1007/s00146-021-01159-3>
- Döbelt S, Jung M, Busch M, Tscheligi M (2015) Consumers’ privacy concerns and implications for a privacy preserving smart grid architecture—results of an Austrian study. *Energy Res Soc Sci* 9(September):137–145
- Dolin C, Weinshel B, Shan S, Hahn CM, Choi E, Mazurek ML, Ur B (2018) Unpacking perceptions of data-driven inferences underlying online targeting and personalization. In: Conference on human factors in computing systems—proceedings, 1–12 Apr 2018
- Dourish P (2010) HCI and environmental sustainability: the politics of design and the design of politics. In: Proceedings of the 8th ACM conference on designing interactive systems—DIS ’10, 1. ACM Press, Aarhus
- Dourish P (2016) Algorithms and their others: algorithmic culture in context. *Big Data Soc* 3(2):1–11
- Eiband M, Buschek D, Kremer A, Hussmann H (2019) The impact of placebic explanations on trust in intelligent systems. In: Conference on human factors in computing systems—proceedings
- Eslami M, Krishna Kumaran SR, Sandvig C, Karahalios K, Vaccaro K, Lee MK, On AEB, Gilbert E, Karahalios K (2018) Communicating algorithmic process in online behavioral advertising. In: Conference on human factors in computing systems—proceedings, 1–13 Apr 2018
- Eslami M, Vaccaro K, Lee MK, On AEB, Gilbert E, Karahalios K (2019) User attitudes towards algorithmic opacity and transparency in online reviewing platforms. In: Conference on human factors in computing systems—proceedings, pp 1–14
- Eubanks V (2018) Automating inequality: how high-tech tools profile, police, and punish the poor. St. Martin’s Press
- Fabianek P, Will C, Wolff S, Madlener R (2020) Green and regional? A multi-criteria assessment framework for the provision of green electricity for electric vehicles in Germany. *Transp Res Part D: Transp Environ* 87(October):102504
- Foth M (2018) Participatory urban informatics: towards citizen—ability. *Smart Sustain Built Environ* 7(1):4–19
- Foth M, Brynskov M, Ojala T, Sassen S (eds) (2015) Citizen’s right to the digital city: urban interfaces, activism, and placemaking. Springer, Singapore Heidelberg New York Dordrecht London
- Foth M, Mann M, Bedford L, Fieuw W, Walters R (2021) A capitalocentric review of technology for sustainable development: the case for more-than-human design. In: Global information society watch 2020: technology, the environment and a sustainable world: responses from the global south, pp 78–82
- Frauenberger C, Foth M, Fitzpatrick G (2018) On scale, dialectics, and affect: pathways for proliferating participatory design. In: Proceedings of the 15th participatory design conference: full papers, vol 1, pp 1–13. ACM, Hasselt
- Fredero O, Gaertner N, Stuckenschmidt H (2019) Real-time smart charging based on precomputed schedules. *IEEE Trans Smart Grid* 10(6):6921–6932
- Geels FW (2012) A socio-technical analysis of low-carbon transitions: introducing the multi-level perspective into transport studies. *J Transp Geogr* 24:471–482

- Gemeente Amsterdam (2019) Een Digitale Stad Voor Één van Iedereen. Gemeente Amsterdam. https://assets.amsterdam.nl/publish/pages/964754/agenda_digitale_stad_v3.pdf
- Gemeente Amsterdam (2020) Nota van Beantwoording Milieuzone Amsterdam 2020. Gemeente Amsterdam. https://assets.amsterdam.nl/publish/pages/850790/nota_van_beantwoording_milieuzone_amsterdam_2020.pdf
- Gemeente Amsterdam (2021) Laad-Update Amsterdam Q2 2021. Gemeente Amsterdam. https://assets.amsterdam.nl/publish/pages/957633/02340_laadupdate_cijfers_q2-2021_v2.pdf
- Grandinetti J (2021) Examining embedded apparatuses of AI in Facebook and TikTok. *AI Soc*. <https://doi.org/10.1007/s00146-021-01270-5>
- Günther M, Kasirzadeh A (2021) Algorithmic and human decision making: for a double standard of transparency. *AI Soc*. <https://doi.org/10.1007/s00146-021-01200-5>
- Hirsch T, Merced K, Narayanan S, Imel ZE, David DC (2017) Designing contestability: interaction design, machine learning, and mental health. In: DIS 2017—Proceedings of the 2017 ACM conference on designing interactive systems, ACM Press, pp 95–99
- Hirschman AO (1970) Exit, voice, and loyalty: responses to decline in firms, organizations, and states, vol 25. Harvard University Press, Cambridge
- Hollanek T (2020) AI transparency: a matter of reconciling design with critique. *AI Soc*. <https://doi.org/10.1007/s00146-020-01110-y>
- Holmquist LE (2017) Intelligence on tap: artificial intelligence as a new design material. *Interactions* 24(4):28–33
- Huang Y, Kockelman KM (2020) Electric vehicle charging station locations: elastic demand, station congestion, and network equilibrium. *Transp Res Part D: Transp Environ* 78(January):102179
- Innerarity D (2021) Making the black box society transparent. *AI Soc* 36(3):975–981
- Jhaver S, Karpfen Y, Antin J (2018) Algorithmic anxiety and coping strategies of airbnb hosts. In: Conference on human factors in computing systems—proceedings, 1–12 Apr 2018
- Jobin A, Ienca M, Vayena E (2019) The global landscape of AI ethics guidelines. *Nat Mach Intell* 1(9):389–399
- Kamols N, Foth M, Guaralda M (2021) Beyond engagement theatre: challenging institutional constraints of participatory planning practice. *Australian Planner* 57(1):23–35
- Kemmis S, McTaggart R (2006) Participatory action research: communicative action and the public sphere. *Educ Act Res* 14(4):459–476
- Kim B, Park J, Suh J (2020) Transparency and accountability in ai decision support: explaining and visualizing convolutional neural networks for text information. *Decis Support Syst* 134(July):113302
- Kitchin R (2019) Toward a genuinely humanizing smart urbanism. In: *The right to the smart city*, pp 193–204
- Kitchin R, Coletta C, Mcardle G (2017) Urban informatics, governmentality and the logics of urban control. *SocArXiv* pp 1–21
- Kizilcec RF (2016) How much information? Effects of transparency on trust in an algorithmic interface. In: Conference on human factors in computing systems—proceedings, pp 2390–2395
- Koskinen IK (ed) (2011) Design research through practice: from the lab, field, and showroom. Morgan Kaufmann/Elsevier, Waltham
- Krafft PM, Young M, Katell M, Lee JE, Narayan S, Epstein M, Dailey D et al (2021) An action-oriented AI policy toolkit for technology audits by community advocates and activists. In: Proceedings of the 2021 ACM conference on fairness, accountability, and transparency, pp 772–81. ACM, Virtual Event Canada
- Kulesza T, Stumpf S, Burnett M, Yang S, Kwan I, Wong WK (2013) Too much, too little, or just right? Ways explanations impact end users' mental models. In: Proceedings of IEEE symposium on visual languages and human-centric computing, VL/HCC, pp 3–10
- Madhu GM, Vyjayanthi C, Modi CN (2019) A novel framework for monitoring solar PV based electric vehicle community charging station and grid frequency regulation using blockchain. In: 2019 10th international conference on computing, communication and networking technologies (ICCCNT), pp 1–7. IEEE, Kanpur
- Mehta R, Srinivasan D, Khambadkone AM, Yang J, Trivedi A (2018) Smart charging strategies for optimal integration of plug-in electric vehicles within existing distribution system infrastructure. *IEEE Trans Smart Grid* 9(1):299–312
- Menéndez-Viso A (2009) Black and white transparency: contradictions of a moral metaphor. *Ethics Inf Technol* 11(2):155–162
- Mittelstadt BD, Daniel B, Allo P, Taddeo M, Wachter S, Floridi L (2016) The ethics of algorithms: mapping the debate. *Big Data Soc* 3(2):205395171667967
- Moghaddam Z, Ahmad I, Habibi D, Phung QV (2018) Smart charging strategy for electric vehicle charging stations. *IEEE Trans Transp Electrif* 4(1):76–88
- Monno V, Khakee A (2012) Tokenism or political activism? Some reflections on participatory planning. *Int Plan Stud* 17(1):85–101
- Morozov E (2013) To save everything, click here: the folly of technological solutionism, 1st edn. PublicAffairs, New York
- Muratori M (2019) Integrating EVs in the electricity system. <https://www.osti.gov/biblio/1576477>
- Ofgem (2018) Ofgem's future insights paper 5—implications of the transition to electric vehicles. Ofgem. <https://www.ofgem.gov.uk/publications/ofgems-future-insights-paper-5-implications-transition-electric-vehicles>
- Ortar N, Ryghaug M (2019) Should all cars be electric by 2025? The electric car debate in Europe. *Sustainability (switzerland)* 11(7):1–16
- Rader E, Cotter K, Cho J (2018) Explanations as mechanisms for supporting algorithmic transparency. In: Conference on human factors in computing systems—proceedings, 1–13 Apr 2018
- Ranchordás S (2020) Nudging citizens through technology in smart cities. *Int Rev Law Comput Technol* 34(3):254–276
- Sadowski J, Pasquale F (2015) The spectrum of control: a social theory of the smart city. *First Monday*. <https://doi.org/10.5210/fm.v20i7.5903>
- Sarra C (2020) Put dialectics into the machine: protection against automatic-decision-making through a deeper understanding of *contestability by design*. *Global Jurist* 20(3):20200003
- Shaw J, Graham M (2017) An informational right to the city? Code, content, control, and the urbanization of information. *Antipode* 49(4):907–927
- Shin D (2021) How do people judge the credibility of algorithmic sources? *AI Soc*. <https://doi.org/10.1007/s00146-021-01158-4>
- Stoyanovich J, Howe B (2018) Follow the data! algorithmic transparency starts with data transparency. *New America*. December 14, 2018. <https://www.newamerica.org/pit/blog/follow-data-algorithmic-transparency-starts-data-transparency/>
- Strengers Y (2012) Peak electricity demand and social practice theories: reframing the role of change agents in the energy sector. *Energy Policy* 44(May):226–234
- Strengers Y (2013) Smart energy technologies in everyday life: smart Utopia? Palgrave Macmillan, New York
- Strengers Y (2014) Smart energy in everyday life: are you designing for resource man? *Interactions* 21(4):24–31
- The Cities Coalition for Digital Rights (2021) Cities for digital rights. Cities for digital rights. August 6, 2021. <https://citiesfordigitalrights.org>
- Tsamados A, Aggarwal N, Cows J, Morley J, Roberts H, Taddeo M, Floridi L (2021) The ethics of algorithms: key problems and solutions. *AI Soc*. <https://doi.org/10.2139/ssrn.3662302>
- Turel T, Joskin D, Geerts F, van Kaathoven E, Schouwenaar M (2017) Designing a transparent smart charge point. In: 30th international electric vehicle symposium (EVS30), vol 1. Stuttgart, Germany
- Vaccaro K, Karahalios K, Mulligan DK, Kluttz D, Hirsch T (2019) Contestability in algorithmic systems. In: Conference companion publication of the 2019 on computer supported cooperative work and social computing, pp 523–27. ACM, Austin
- Vakarelov O, Rogerson K (2020) The transparency game: government information, access, and actionability. *Philos Technol* 33(1):71–92

- Verbeek P-P (2006) Materializing morality. *Sci Technol Human Values* 31(3):361–380
- Walmsley J (2021) Artificial intelligence and the value of transparency. *AI Soc* 36(2):585–595
- Wang Q, Liu X, Du J, Kong F (2016) Smart charging for electric vehicles: a survey from the algorithmic perspective. *IEEE Commun Surv Tutor* 18(2):1500–1517
- Zerilli J, Knott A, Maclaurin J, Gavaghan C (2019) Transparency in algorithmic and human decision-making: is there a double standard? *Philos Technol* 32(4):661–683
- Zhang Y, Chen X (2020). Explainable recommendation: a survey and new perspectives. *Found Trends® Inf Retr* 14(1):1–101

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.