Contestable City Algorithms

Kars Alfrink
Delft University of Technology
Delft, NL
c.p.alfrink@tudelft.nl

Thijs Turel
Amsterdam Institute of Advanced Metropolitan Solutions
Amsterdam, NL
thijs.turel@ams-institute.nl

Gerd Kortuem
Delft University of Technology
Delft, NL
g.w.kortuem@tudelft.nl

CHI 2020 Extended Abstracts, April 25–30, 2020, Honolulu, HI, USA.

ACM ISBN 978-1-4503-6819-3/20/04.
DOI: https://doi.org/10.1145/3334480.XXXXXXX
*update the above block & DOI per your rightsreview confirmation (provided after acceptance)

Abstract

This paper explores if and how ordinary citizens should be able to influence the operation of public algorithmic decision-making systems. We present 'smart electric vehicle charging' as a novel case study for automated decision making and present a design study of a 'Transparent Smart Charge Point' as a way to explore the possibilities, limitations and ethical issues of algorithmic decision-making systems in the public realm. We conclude with a reflection on contestability as a critical aspect of future public decision-making systems.

Author Keywords

Smart city; algorithms; ethics; fairness; transparency; contestability; electric vehicles; smart charging.

CSS Concepts

• Human-centered computing~Human computer interaction (HCI); • Human-centered computing~Participatory design • Human-centered computing~Interface design prototyping • Social and professional topics~Socio-technical systems • Social and professional topics~Corporate surveillance • Computing methodologies~Artificial intelligence • Computing methodologies~Machine learning
https://dl.acm.org/ccs/ccs_flat.cfm

Introduction

Cities across the globe are increasingly using intelligent systems and automated decision-making to provide city services, inform policies and tackle urban challenges. This includes intelligent traffic management systems, crowd management systems, predictive policing, among others. At the heart of such systems are algorithms that subject citizens to automated decision-making processes to achieve outcomes that are considered beneficial by municipalities or public organizations (such as limiting car traffic to neighborhoods or increasing safety around public events) [1].

While most decisions made by such systems are minor in their impact, such as whether to issue a traffic fine, others have potentially grave consequences. It is now widely recognized that automated systems used by the government to deliver public services often fall short for the people they are intended to serve, even if such systems have been designed with the best intentions [2].

A considerable amount of work has already been done to encourage or require good practice in the development and use of algorithmic systems. For example, participatory design approaches [3] and value-sensitive design (VSD) methods [4] are increasingly used to include citizens directly or indirectly impacted by such systems [5]. Another approach is to improve the transparency of algorithmic systems. Brauneis and Goodman [6] state, "In the public sector, the opacity of algorithmic decision making is particularly problematic both because governmental decisions may be especially weighty, and because democratically-elected governments bear special duties of accountability." In this context, making systems transparent to ordinary people is seen as a way to increase accountability [7]–[10]. Meanwhile, data protection laws like the GDPR mandate certain practices around the use of personal data and effectively create a "right to explanation," whereby a user has a legal right to ask for an explanation of an algorithmic decision that was made about them [11]–[13].

However, current approaches have serious shortcomings: participatory approaches require benevolent organizations willing to listen to citizen concerns and are ill-suited in addressing emerging issues that arise during ongoing operations of system after deployment. Despite a recent focus on algorithmic transparency [14]–[23], research thus far has indicated few ways for citizens and non-experts to meaningfully understand the complex and unpredictable nature of algorithmic systems.

In the following, we explore if and how ordinary citizens should be able to understand and influence the operation of public algorithmic decision-making systems. We first discuss 'smart electric vehicle charging' as a novel case study for automated decision making and present a design study of a 'Transparent Smart Charge Point' that aims to help electric vehicles drivers to understand smart charging prioritization decisions. We use this design study to explore the possibilities, limitations and ethical issues of algorithmic decision-making systems in the public realm. Based on this, we will then reflect on the need of contestability as a critical aspect of future public algorithmic decision-making systems.

Case Study: Smart Electric Vehicle Charging

As the number of electric vehicles on the road is steadily rising there is a growing concern that the capacity of the electrical grid is not sufficient to supply all electric vehicles with enough energy. This concern is especially relevant for local distribution grids in cities and neighborhoods where electric cars are particularly prevalent and where charging sessions are clustered around peak times. If demand for charging exceeds supply, not every car can be charged, and choices need to be made: who will be charged first, and who last?

Energy networks providers have thus started to deploy 'smart charging' solutions to dynamically manage demand that curtail the rate and amount of electricity each vehicle can charge. The term 'smart charging' refers to the process of making the timing and capacity of the charging of an electric vehicle dependent on factors such as grid capacity, electricity demand and the availability of sustainable energy. According to the UK energy regulator Ofgem, smart charging could allow at least 60 per cent more electric vehicles to be charged without having to upgrade the physical grid infrastructure [24], while reducing the need for new power stations and accommodating more renewable power generation such as wind and solar.

While present-day smart charging pilots make use of relatively simple algorithms with only one or two variables as input (such as current electricity demand and supply), it is expected that decisions made by the grid and charging infrastructure will quickly become more complex. Industry is envisioning algorithms that will integrate parameters such as battery capacity of each vehicle, participation in a shared vehicle scheme, number of drivers, predicted daily electricity demand

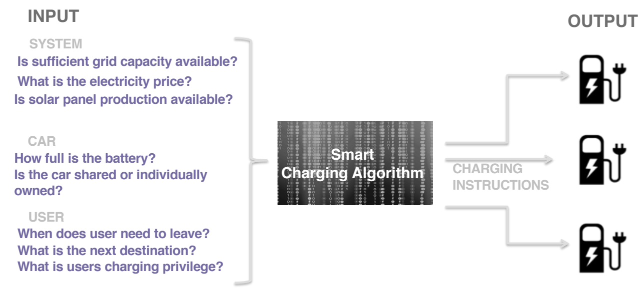

Figure 1. Future Smart Charging Parameters

etc., using (personal) data from various sources to make an automated decision on whom to charge and when to charge (Figure 1). More advanced solutions use the car battery as intermediate electricity store and use machine learning algorithms to optimize electricity distribution of EV charging.

Drivers will experience the outcome of these automatic decisions in various ways: electric vehicles may charge much slower than expected or drivers may receive less electricity or slower charging rates than other drivers (even if both cars are plugged in at the same time at the same charging station). Algorithms may also have unexpected side-effects such as some neighborhoods having less electricity available for charging than others.

Responsible Smart Charging

The use of smart charging algorithms raises a range of ethical issues:

- *Fairness*: are the benefits and drawbacks of smart charging systems experienced equally by all citizens?
- *Bias*: are certain groups of people systematically disadvantaged (for example drivers of smaller cars or drivers living in certain neighborhoods)?
- *Transparency*: are stakeholders (individual citizens, regulators, municipalities) able to understand how algorithmic smart charging decisions are made?
- *Accountability*: are charge point operator, grid operators and all other stakeholders involved in designing and operating smart charging systems (Figure 2) able to account for how and why the system operates in a way it does? Are regulators and municipalities able to effectively govern such system and prevent public harm?
- *Power*: who gets to decide how smart charging systems should operate?

At the moment, all these issues are largely unexplored.

Design Study: The Transparent Charge Point

As a starting point to explore these issues, a group of companies, researchers and designers[1] teamed up in a design exercise to explore the meaning of transparency prototype of a 'Transparent Smart Charge Point' that makes charging decisions visible to drivers (Figure 3).

The transparent charge point is a 2-meter-tall charging station with a large display for visualizing the charging priorities of connected vehicles. The display uses a 'Tetris'-like visualization metaphor where each

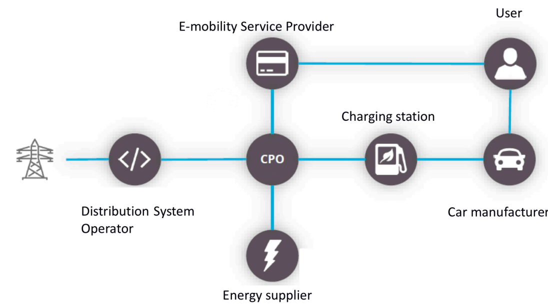

Figure 2. Commercial Stakeholders Involved in Operating Smart Charging Systems (source [25])

horizontal line represents the charging rate of each connected vehicle at a particular 5-minute time interval (Figure 4).

The bottom line represents that current distribution of electricity between connected vehicles (in the shown example three vehicles represented by red, green and blue colors are connected). The lines above the bottom line represent the planned future distribution of electricity between connected vehicles, with the top-most line representing the distribution two hours in the future. The display on the right side of Figure 4 indicates that three vehicles are currently connected, but only blue and green are currently charging. The vehicle indicated by red will start charging in about 50 minutes, after the blue car is charged. Because of expected variations in electricity availability at the

[1] The Transparent Charge Point was commissioned by Dutch companies Alliander and Elaad and designed by The Incredible Machine with involvement of one the authors (see also [25]).

Figure 3 Transparent Smart Charge Point (source [25])

charge point, the charging rate varies, as indicated by the varying width of the colored columns.

The example also indicates that at present the green car and the blue car receive roughly the same amount of energy (bottom line), while in about 10 minutes the blue car will receive a far greater share of the available electricity than the green car. The decision of how much electricity each car is able to charge at each time is determined by an (envisioned) algorithm.

This design study was informed by participatory design sessions where industry representatives, researchers and citizens were asked to envision future prioritization algorithms. The final design study uses a prioritizing scheme that classifies drivers based on their occupation or social standing. For instance, participants of the workshop expected medical professionals to receive higher charging priorities than those with non-medical professions. Participants even considered 'non-ethical' aspects, like being a welfare recipient, as potential future parameters for a smart charging system.

Discussion

The Transparent Smart Charge Point is a speculative critical design exercise and is not intended to represent a viable or even desirable future solution, but instead was created to facilitate a public debate. It allowed us to explore the design space and potential implications of future smart charging solutions which are currently developed by the energy industry.

We are currently in the process of setting up a functional trial of a smart charging system and are conducting sessions with citizens to gauge their opinions about smart charging in general and automated algorithmic decision making in particular. Early impressions make clear that making drivers aware of the functioning of charging algorithms during charging sessions has limitations related to information overload and awareness of the presence of algorithms. In addition, we discovered that some drivers who have an interest in understanding how their charge point operates have doubts that they would be able to influence its behavior. In the words of one of the participants: "I doubt if knowing the algorithm makes sense, there's probably nothing that can be done about it." This leads us to the question of contestability: why and in which way should ordinary citizens be able to influence the operation of algorithmic decision systems?

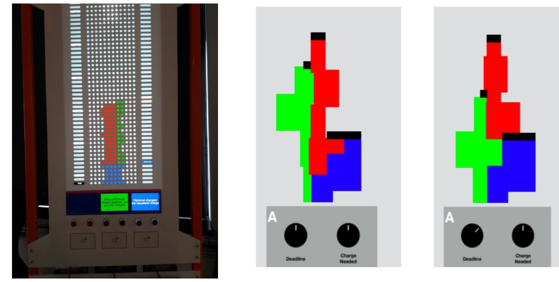

Figure 4. Transparent Charge Point Display with Future Charging Schedule (source [25])

From Transparency to Contestability

Defining Contestability

In everyday language, contestability is understood in two ways. First, "to contest" is to oppose an action, possibly because the action is perceived as mistaken or simply wrong. Second, to contest is to compete for power over something. These two meanings do not exclude each other. Furthermore, it is important to distinguish between contested and contestable. People can express disagreement with the workings of a smart EV charging system, or seek to gain influence over those workings, making it contested. But it is only when the system itself has the capability of responding to such opposition or is opened up to make room for the influence of multiple stakeholders, that it becomes contestable.

The Need for Contestability of Smart Public Infrastructure

The current crop of smart EV charging systems optimizes across a network as a whole. The system in our case study operates on data about power grid usage and weather patterns. In the near term, real-time control of individual charging sessions is expected to be made possible, which could for example enable the privileging of vehicles in a car-sharing scheme over those that are privately owned.

As the sophistication of smart charging algorithms increases, their predictability for individual users declines and the likelihood of mistaken or unintended actions increases. Likewise, as stakeholders seek to incentivize certain behaviors over others by encoding policy into algorithms, the need for public accountability increases.

TYPES OF IMPACTS

The impacts of a smart EV charging system on drivers can range from the immediately obvious to things that are hard or impossible to predict by an individual user.

On one end of the spectrum, a person may notice that their car is charged slower than the one connected to the other point of the station. This may be due to the fact that they are driving a privately-owned car and the other car is shared. In this case policy is enacted correctly but it is something a person did not expect, does not understand or does not agree with. correctly but it is something a person did not expect, does not understand or does not agree with.

On the other end of the spectrum, a person may be systematically disadvantaged because of life patterns that are out of step with the norms presumed by the system’s design. For example, a schedule that optimizes charging around a typical 5-day workweek and 8-hour workdays may interfere with those people that work odd hours. Similarly, dense urban areas with high grid usage may be disadvantaged by stations that are always charging slower, whereas less dense areas suffer no such impacts.

In all of these cases, to be responsive to on-the-ground reality of the system in action, a public smart EV charging system needs to be made contestable. Elsewhere, in the context of labor relations, it has been argued that contestability needs transparency, a channel for voice and a forum for evaluation [26].

CHANNELS FOR VOICE, FORUMS FOR EVALUATION

When mistaken happen or when the system makes decisions that drivers do not agree with, the perceived legitimacy of the smart charging system declines. In such cases, people can respond in roughly one of two ways. One is to opt out or “exit” from the system altogether. The other type of response is to express disagreement in the hopes of influencing the policies that shape the system's behavior. In the context of public smart EV charging stations, for both system operators and users, opting out is undesirable or even impossible. This leaves us with expressing disagreement, or “voice”. Adding channels for voice has the benefit of decreasing the likelihood of people opting out, it can also increase the perceived legitimacy of the system [27].

A second element needed for contestability is a forum for the evaluation of contestations. This can be understood in at least two ways. On the one hand, it can be thought of as the capacity of a system to respond to user feedback in the moment, exercising something akin to the discretion exercised by street-level bureaucrats when enacting official policy [28]. In the context of smart EV charging, maybe it is possible for an individual charge station to deviate from the default charging speed for a single session if requested by a user, provided global network conditions allow for it. Such types of “slack” would go a long way towards alleviating some of the frustrations that arise from being subjected to automated decision-making.

On the other hand, for other types of contestation, those that demand more permanent change on the scale of a part or all of the network, it will likely be necessary to involve humans in the loop. With smart EV charging, which is characterized by a complex network of stakeholders each controlling part of the system’s behavior (Figure 2), it is not immediately clear where such a forum should be located. Here, a level of impartiality or separation of power on the part of the forum may be desirable. It may also be necessary to create a platform for a system’s users to identify shared grievances and collectively argue for changes; something roughly analogous to a labor union.

Transparency as a ground for contestation
As the design study illustrates, transparency may offer some insight into the workings of such a system, but in the event of a mishap, it does nothing for people who want to address the issue by changing the system. Contestability, on the other hand, does not necessitate full comprehension of an algorithm by an end-user. Instead, it addresses the problem of legitimacy by adding "channels for voice" as well as forums for evaluation, so that the loop between the policy that has shaped the algorithm and its on-the-ground actual unfolding can be closed.

Furthermore, contestability recasts transparency as a resource for contestation. This helps us address the issue of people’s already scarce attention in a saturated informational environment. From an information ethics perspective, quality instead of quantity should be the goal [29]. Contestability provides us with a specific qualitative measure to use as a way of determining which information to include in the transparency “layer” of smart public infrastructure.

Contestability Challanges
Smart EV charging highlights a limitation of algorithmic transparency: Systems such as these are complex socio-technical systems that resist straightforwardly making visible that which is hidden [30]. It is tempting to think of the thing that needs to be made transparent as a monolithic "algorithm" but in reality, the behavior

of a smart EV charging system is determined as much by hardware and people as it is by software.

For example, in our case study, the actual speed at which a public charge station can charge is related to the fuses it contains, which are in turn governed by safety norms. Public charge stations are retrofitted with a different fusing scheme to enable the wider range of charging speeds. We have witnessed that design efforts aimed at making smart EV charging stations more transparent start out with a somewhat naive perspective on the task at hand -- understand "the algorithm" and then visualize it in a graspable manner for end-users. But soon enough, the actual complexities of the system as a whole come into view, and designers find themselves having to make judgment calls about how far to go with the transparency ideal. Should people know about these fuses, for example? What guidelines do we use to answer such questions?

As a second example, the higher rate of charging as a result of the availability of solar energy isn't actually the direct result of the sun being out or not. In fact, it turns out at the moment it is governed by the weather prediction for the city as a whole, chopped up into 15-minute segments. So here we have an "algorithm" that is the result of a programmer translating policy requirements into something that can actually operate on the current hardware infrastructure. Again, the socio-technical nature of the system comes into view, and now, we need to once again ask ourselves whether we have a responsibility to people to uphold the illusion of intelligence -- our system tracking the actual availability of solar energy -- or if people would be better served knowing the algorithm they are subjected to in reality, and the malleability of said algorithm.

All of these questions are important because in our view transparency on its own is insufficient, but it is necessary for contestability. It is not a cure-all, but transparency does provide the basis for contestation. What the smart EV charging case study shows is that what to include or leave out of the picture is a political choice as much as anything else. There is no straightforward design heuristic that will tell us in every case which elements to include. Instead, the design process in part becomes a process by which it is collectively determined which aspects of a system can and cannot be contested.

The distributed nature of these systems surfaces another challenge for contestability, namely that there is more than one stakeholder determining the policies that shape the behavior of the system. Even if we limit ourselves to the software that makes up the system, we are not dealing with one single codebase, but several software systems talking to each other that together result in the smart EV charging system that people interact with. If we add contestability to this mix, questions arise about where contestations should flow, and who is responsible for evaluating their merits. Furthermore, if changes to the system are required that transcend the part directly owned by the evaluating agent, how are these prioritized across the ongoing development and operation efforts of the various stakeholders?

Emerging Research Questions

Having established the reasons for contestability and some of the related issues within the context of smart

EV charging (and more broadly in the context of public socio-technical systems), a number of research questions emerge. These are all grounded in our ongoing design-led research.

First of all, we need a better understanding of how people view and understand smart public infrastructure. Specifically, for contestability, a question is what concrete harm people experience now but also in the future.

Second, as has already been identified, there is the matter of what to include and exclude in the transparency layer of a contestable system. What criteria should we use to determine those elements that are crucial for as a basis for contestation? This goes hand in hand with a consideration of which parts of a system should be contestable in the first place. A related question is if there are any barriers to making particular elements contestable. Or, to put it in another way, what is the view of each stakeholder in smart public infrastructure on the relative merits of contestability?

Third, we feel the distribution of algorithmic decision-making capabilities across the various parts that make up a system needs to be pulled into the design process. Because those decisions impact the capacity for the system as a whole to respond to contestations. The question then becomes how to determine the optimal site of automated decision-making in a larger system. Here it may also be relevant to focus on the significance of the fact that we are focusing on public infrastructure. Is a different model of contestability necessary when we are dealing with systems that are (partially) publicly owned or for public use?

Fourth, turning to people's contestations itself, the conceptual notion of "channels for voice" needs to be made concrete in actual system features. What are the most suitable modalities for expressing disagreement? When and where should these be offered? How can we move beyond a model of contestation that is based on individuals voicing their concerns? How can people who share a similar fate find and connect to each other, and how can they collectively make themselves heard?

Fifth, the matter of evaluating and responding to contestations should also made part of the design effort. What categories of dissent should a system respond to automatically, and what categories need to be handled through human intervention before flowing into a system's ongoing design, development and operation? What resources do non-experts need to be a part of this evaluation? How can a system be given discretionary capabilities so that it can deviate from standard policy in individual cases when local circumstances warrant it? How can we measure the relative contestability of a system?

Together, these questions constitute the beginnings of a research agenda for the design of contestable smart public infrastructure.

References

[1] S. Ranchordás, “Nudging citizens through technology in smart cities,” *Int. Rev. Law, Comput. Technol.*, 2019.

[2] H. Lebovits, “Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor,” *Public Integr.*, 2019.

[3] E. B.N. Sanders, “From user-centered to participatory design approaches,” 2002.

[4] B. Friedman, D. G. Hendry, and A. Borning, “A Survey of Value Sensitive Design Methods,” *Found. Trends® Human–Computer Interact.*, 2017.

[5] D. Gooch *et al.*, “AmplifyingQuiet voices: Challenges and opportunities for participatory design at an urban scale,” *ACM Trans. Comput. Interact.*, 2018.

[6] R. Brauneis and E. P. Goodman, “Algorithmic Transparency for the Smart City,” *Ssrn*, vol. 103, pp. 103–176, 2017.

[7] N. Diakopoulos, “Algorithmic accountability reporting: On the investigation of black boxes,” *Columbia Journal. Sch.*, pp. 1–33, 2014.

[8] N. Diakopoulos, “Accountability in Algorithmic Decision-making,” *Queue*, vol. 59, no. 2, pp. 56–62, Jan. 2015.

[9] AI Now, “Algorithmic Accountability Policy Toolkit,” no. October, 2018.

[10] E. Rader, K. Cotter, and J. Cho, “Explanations as Mechanisms for Supporting Algorithmic Transparency,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI ’18*, 2018, pp. 1–13.

[11] E. Bayamlıoğlu, “Transparency of Automated Decisions in the GDPR: An Attempt for Systemisation,” *Ssrn*, no. January, pp. 0–50, 2018.

[12] S. Wachter *et al.*, “Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation,” *Int. Data Priv. Law*, vol. 7, no. 4, pp. 243–265, 2017.

[13] G. Malgieri and G. Comandé, “Why a Right to Legibility of Automated Decision-Making Exists in the General Data Protection Regulation,” *Int. Data Priv. Law*, 2017.

[14] S. Wachter, B. Mittelstadt, and C. Russell, “Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR,” *SSRN Electron. J.*, pp. 1–52, 2017.

[15] M. Veale, “Logics and practices of transparency and opacity in real-world applications of public sector machine learning,” Jun. 2017.

[16] B. Mittelstadt, C. Russell, and S. Wachter, “Explaining Explanations in AI,” pp. 279–288, 2018.

[17] M. Gervasio, K. Myers, E. Yeh, and B. Adkins, “Explanation to avert surprise,” *CEUR Workshop Proc.*, vol. 2068, 2018.

[18] A. Woodruff, S. E. Fox, S. Rousso-Schindler, and J. Warshaw, “A Qualitative Exploration of Perceptions of Algorithmic Fairness,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI ’18*, 2018, pp. 1–14.

[19] M. Eslami, S. R. K. Kumaran, C. Sandvig, and K. Karahalios, “Communicating algorithmic process in online behavioral advertising,” *Conf. Hum. Factors Comput. Syst. - Proc.*, vol. 2018-April, 2018.

[20] S. Jhaver, Y. Karpfen, and J. Antin, “Algorithmic anxiety and coping strategies of airbnb hosts,” *Conf. Hum. Factors Comput. Syst. - Proc.*, vol. 2018-April, pp. 1–12, 2018.

[21] M. Eslami, K. Vaccaro, M. K. Lee, A. E. Bar On, E. Gilbert, and K. Karahalios, “User attitudes towards algorithmic opacity and transparency in online reviewing platforms,” *Conf. Hum. Factors Comput. Syst. - Proc.*, pp. 1–14, 2019.

[22] H. Rutjes, M. C. Willemsen, and W. A. IJsselsteijn, “Considerations on Explainable AI and Users’ Mental Models,” in *CHI’19 Workshop: Where is the Human? Bridging the Gap Between AI and HCI*, 2019.

[23] H. Kaur and A. C. Williams, “Building Shared Mental Models between Humans and AI for Effective Collaboration,” in *CHI*, 2019.

[24] “Getting more out of our electricity networks through reforming access and forward-looking charging arrangements,” 2018.

[25] T. Turel, D. Joskin, F. Geerts, E. Van Kaathoven, and M. Schouwenaar, “Designing a Transparent

Smart Charge Point," *Evs30*, 2017.

[26] L. Lindblom, "Consent, Contestability, and Unions," *Bus. Ethics Q.*, 2019.

[27] A. Centivany and B. Glushko, "'Popcorn tastes good': Participatory policymaking and Reddit's 'AMAgeddon,'" in *Conference on Human Factors in Computing Systems - Proceedings*, 2016.

[28] A. Alkhatib and M. Bernstein, "Street-Level Algorithms: A Theory at the Gaps Between Policy and Decisions," *Proc. 2019 CHI Conf. Hum. Factors Comput. Syst.*, pp. 1–13, 2019.

[29] L. Floridi, "Big Data and Information Quality," in *The Philosophy of Information Quality*, 2014.

[30] M. Ananny and K. Crawford, "Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability," *New Media Soc.*, 2018.