

Best Practices for Transparency in Machine Generated Personalization

Laura Schelenz, Avi Segal, Kobi Gal*

ABSTRACT

Machine generated personalization is increasingly used in online systems. Personalization is intended to provide users with relevant content, products, and solutions that address their respective needs and preferences. However, users are becoming increasingly vulnerable to online manipulation due to algorithmic advancements and lack of transparency. Such manipulation decreases users' levels of trust, autonomy, and satisfaction concerning the systems with which they interact. Increasing transparency is an important goal for personalization based systems and system designers benefit from guidance in implementing transparency in their systems.

In this work we combine insights from technology ethics and computer science to generate a list of transparency best practices for machine generated personalization. We further develop a checklist to be used by designers to evaluate and increase the transparency of their algorithmic systems. Adopting a designer perspective, we apply the checklist to prominent online services and discuss its advantages and shortcomings. We encourage researchers to adopt the checklist and work towards a consensus-based tool for measuring transparency in the personalization community.

CCS CONCEPTS

• **General and reference** → Design; • **Computing methodologies** → *Artificial intelligence*; • **Information systems** → **Personalization**; **Personalization**; • **Social and professional topics** → *Codes of ethics*.

KEYWORDS

transparency, responsible personalization, personalized transparency, checklist, ethics

ACM Reference Format:

Laura Schelenz, Avi Segal, Kobi Gal. 2020. Best Practices for Transparency in Machine Generated Personalization. In *Adjunct Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20 Adjunct)*, July 14–17, 2020, Genoa, Italy. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3386392.3397593>

*Laura Schelenz, University of Tübingen, Germany: laura.schelenz@uni-tuebingen.de; Avi Segal, Ben-Gurion University of the Negev, Israel: avisegal@gmail.com; Kobi Gal, The University of Edinburgh, UK: kgal@inf.ed.ac.uk

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UMAP '20 Adjunct, July 14–17, 2020, Genoa, Italy

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7950-2/20/07...\$15.00

<https://doi.org/10.1145/3386392.3397593>

1 INTRODUCTION

Recent years saw significant increase in personalization approaches for online systems [33, 44]. Such personalization can be used to direct users' attention to relevant content [25], increase their motivation when working online [29], improve their performance [24], extend their engagement [34] and more. These approaches rely on social theories of human behavior (e.g. [10, 16]) as well as on machine learning based abilities to predict human reaction to various interventions [20, 30].

Yet, personalization technology that focuses on maximizing system designers goals runs the risk of marginalizing users¹. Personalized recommendations usually attempt to influence a person's decision-making. When such influences are hidden and subtly try to persuade users (maybe even against their expressed goals), this constitutes a form of manipulation [38]. Subverting a person's decision-making abilities reduces their autonomy. Especially with regard to personalized advertisement, personalization can exploit users' vulnerabilities [37] and may even threaten democratic processes [14].

Applying transparency to the design of personalized content can help address these challenges. First, transparency can balance power asymmetry, empowering users while curtailing the influence of companies on customer behavior [27]. Second, transparency can increase user autonomy. For example, recommender systems usually filter content according to preference models that easily create a feedback loop [42]. When users lack exposure to information diversity, their autonomy and ability to make independent decisions is impacted [28]. Third, transparency can boost privacy rights and user trust in algorithmic systems. Users can only give meaningful informed consent when they understand the risks of algorithmic decision-making [28]. Fourth, transparency can increase subjects' ability to understand the cause of decisions made by algorithms and assess whether a decision-making process is fair and non-discriminatory [4, p.2].

The computer science community has affirmed the importance of transparency in its profession. The ACM Code of Ethics reads: "The entire computing profession benefits when the ethical decision-making process is accountable to and transparent to all stakeholders" [5]. Also political bodies identify transparency as a pivotal principle in Artificial Intelligence based software [17, 22].

Especially with the advent of legal frameworks that prescribe transparency in data collection, processing, and storage [3], system designers require increased awareness and guidance about transparency in their systems [7]. Recent and emerging scholarship on explainable AI underlines the importance of transparency in computer systems [23, 31, 46]. Additionally, attempts to operationalize ethics principles for AI respond to the increased call for practical guidance [19].

¹<https://uxdesign.cc/user-experience-vs-business-goals-finding-the-balance-7507ac85b0a9>

We take a three step approach to developing best practices for transparency in machine generated personalization: 1) developing a new definition of transparency for algorithmic systems by drawing on prior art, 2) deriving best practices for the implementation of transparency in machine generated personalization, and 3) translating these best practices into questions for system designers to be used as a reflection and assessment tool. The outcome is an online checklist for open usage based on best practices which constitute ethical guidelines for system designers. The checklist can be used by systems designers to evaluate and operationalize transparency in their systems. It encourages (self-)reflection and an ongoing exchange with the personalization community towards ethically responsible personalization.

2 TRANSPARENCY DEFINITION

To generate a list of best practices, we began by asking: What is transparency in the context of AI systems? According to Turilli and Floridi [40], transparency is not a principle of ethics per se, but a practice that can achieve ethics goals such as autonomy and accountability. Adopting this understanding, we investigated views on transparency from technology ethics, the philosophy of technology, computer sciences, as well as ethics guidelines and legal documents. This literature review and qualitative analysis of work on "transparency" allowed us to formulate the following definition for the computer science community:

Transparency is a practice of system design that centers on the disclosure of information to users, whereas this information should be understandable to the respective user and provide insights about the system. Specifically, the information disclosed should enable the user to understand why and how the system may produce or why and how it has produced a certain outcome (e.g. why a user received a certain personalized recommendation).

The first important component of transparency is the notion that the user of a system must be able to comprehend the information disclosed to them. According to Chromnik et al. [13], transparency is an enabling condition for the user to "understand the cause of a decision." Ananny and Crawford [8] describe transparency as a form of "seeing" an actor-network. Transparency then means not merely looking inside a system but across systems, and explaining a model as it interacts with other actors in an algorithmic system [8]. Floridi et al. [17] understand transparency as explainability, whereas explainability incorporates both intelligibility (being able to grasp how a model works) and accountability (understanding who is responsible for it). Following Vakarelov and Rogerson [41], transparency means communication of information under two conditions: information must be a) sufficient and b) accessible, i.e. the user must be able to comprehend and act upon the information. According to the GDPR [3], information about data collection and processing must be provided in "clear and plain language and it should not contain unfair terms" [3, p. 8]. Here, we can see how transparency is a relational practice. Whether the information provided is transparent depends on the individual user or data subject, their cognitive abilities, their language skills, and epistemic conditions [41].

Another crucial element of transparency is information disclosure about deliberation or decision-making processes. The IEEE Guideline for "Ethically Aligned Design" states that transparency

means the possibility to ascertain why a certain decision was made [1]. For Turilli and Floridi [40], disclosing information refers to communication about the deliberation process because it reveals the values that guide organizations in their everyday practices and illustrate how they make decisions. Hence, the disclosure of the dataset may be less relevant than the actual factors (such as inferences made from the data) that inform a model and its effects on users [39]. Also Zerilli et al. [45] argue that, similar to explanations in human decision-making, a system should reveal factors in decision-making and how they are weighted. Dahl [15] even argues that it is not necessary to reveal the inner working of a model but to provide key details about how the results came about, e.g. offering expert testimony about how the system usually works. Burrell [12] suggests that reducing opacity of classifications means "exposing some of the logic of this classification."

Finally, there can be an element of participation in transparency. The user is expected to assess the system with regard to its trustworthiness based on the information that is disclosed. Furthermore, the user may become active in choosing between different models, i.e. different options of personalization [35]. The user is thus becoming involved in the process of transparency.

3 TRANSPARENCY BEST PRACTICES

From our definition of transparency, we derived nine principles of transparency for responsible personalization with particular relevance ascribed to three practices. These practices reflect the three core elements of transparency: information provided must be understandable to users, information must be disclosed about why and how a model reaches a certain outcome, and users should have a say in personalization processes. The best practices further reflect additional needs for information about the data collection processes, the composition of datasets, the functionality of a model, the responsibilities for the model or system, and how the model may interact with other models across algorithmic systems. The best practices are necessarily generic, which allows system designers in the personalization community to adapt them to their context.

Table 1 shows the list of the best practices as well as the sources on which these practices build. It also identifies the relevant system architecture components for each best practice based on the Input-Processing-Output architecture model [9]. We extend this architecture with a "Control" component to represent the control given to the user over the system's personalization behavior. We define user control as the possibility of users to interact with the system to adjust elements thereof to their respective needs and preferences. It is important that users not only "feel" that they have control because this can put them at risk of exploitation. If users think that they have control, they might feel encouraged to share more data [39]. Possible misuse of transparency measures and the potential abuse of our approach are addressed in the discussion.

4 CHECKLIST

Based on the definition and best practices, we have defined a checklist for system designers to assess the transparency of machine generated personalization. The checklist codifies transparency practices as developed in the previous section. For instance, practice number 3 requires "disclosing relevant and detailed information

No.	Component	Description of transparency standard	Sources
1	Input, Processing, Output, Control	Disclosing accessible and actionable information, meaning that the user can comprehend and act upon the information	[3, 41]
2	Input, Processing	Disclosing relevant and detailed information about data collection and processing; notification about the data collected for personalization, information about pre-processing and possible biases in the dataset	[3, 5, 11]
3	Processing	Disclosing relevant and detailed information about the goals of the designer/system, the reasoning of a system, the factors and criteria used (potentially also how they are weighted), as well as the inferences made to reach an algorithmic decision	[1, 4, 12, 13, 17, 36, 39, 40, 45]
4	Processing	If possible, providing expert testimony (e.g. by a member of the design team) about how a system works and reaches a certain outcome	[15]
5	Processing	If possible, disclosing information about how a model may affect the user and how it may interact with other models across systems	[8]
6	Output	Disclosing that a machine is communicating with the user and not a real person	[2]
7	Output	Disclosing information about those responsible for the model (e.g. name of the company or designer)	[17]
8	Control	Proposing alternative choices for user interaction with the system, e.g. different options for personalization	[35]
9	Control	Providing the user with opportunities to adjust personalization or specify their goals as these goals are expected to drive personalization	[21]

Table 1: Transparency Best Practices for Machine Generated Personalization

about the goals of the designer/system." We reframed this best practice into a question and asked at the top of the checklist: "Does the system inform the user about the purpose of personalization?" In this fashion, we formulated checklist questions for all the Input, Processing, Output and Control best practices identified in Table 1. In this process, we prioritize some best practices that were overwhelmingly affirmed by the literature. We note the qualitative nature of our methodology and expect future work to extend our approach to a consensus-based tool for measuring transparency in the personalization community.

The resulting checklist is presented at Table 2. The checklist’s web version is given at: <http://tiny.cc/evxckz>.

The checklist includes a total of 23 questions. After filling it online, the system designer can download a PDF file with their responses. They can also print an empty copy of the checklist to be filled offline if needed. We note that the checklist is supplied as an assessment tool for system designers, enabling them to identify areas in their systems which suffer from lack of transparency. Ideally, a system designer has implemented transparency so that they can check yes for every question. However, the goal should not be to score high on the checklist but rather to honestly reflect and decide on priorities and next steps.

5 CASE STUDY: APPLYING THE CHECKLIST

We performed an initial application of the proposed checklist as a reflective and assessment tool for the following online services that use personalization: Facebook, Netflix, YouTube, Spotify, and Amazon. For each of these destinations, we took a system designer’s point of view, and asked "(how) are the transparency elements from the checklist supported on this particular site?". For this assessment

we adopted the checklist and examined the above web services using one of the authors account on these sites. Specifically, we checked the information available to registered users on the sites including the privacy policy, the legal terms and conditions and other information that is shared with the user and covers any of the checklist elements. We answered each checklist question for each site with a "yes", "no" or "partial" reply.

To conduct a preliminary comparison between the different sites and between the different sections of the checklist for each site, we then computed the percentage of "Yes" and "Partial" replies for each checklist section. For each checklist question, we gave a "Yes" reply a value of 1 and a "Partial" reply a value of 0.5. We then summed these values for each section and divided it by the total number of questions in the corresponding section. Figure 1 presents the result of this comparison. We discuss these results in the next section.

6 DISCUSSION

The major advantage of the transparency checklist is that it helps system designers understand where they are strong on transparency and where improvements are needed. Looking at Figure 1, we notice that existing online systems primarily focus on realizing transparency in the "Input" category, i.e. with regard to data collection and the handling of user data. They are particularly weak in providing information about why and how models bring about certain personalization ("Processing"). They also lack participatory elements such as offering the user different options of personalization or allowing the user to supply feedback ("Control").

This trend to follow best practices of data or "Input" transparency may be attributable to the rise of data protection laws such as the GDPR. System designers so far pay less attention to transparency

General:
Does the system inform the user about the purpose of personalization?
Does the system inform the user who developed the technology and is liable in cases of wrongdoing?
Does the system inform the user about their rights under data protection law?
Does the system inform the user about possible risks of engaging with the system?
Input:
Have users given informed consent about the collection, processing, and storage of their data?
Does the system inform the user about the fact that data is collected for personalization?
Does the system inform the user about which data is collected to produce personalized content for them?
Does the system inform the user about pre-processing done with the data collected for personalization purposes?
Does the system inform the user if their data is used and shared beyond the goals of personalization?
Processing:
Does the system inform the user about the kind of data that is processed to create a certain personalized item?
Does the system explain to the user why they are receiving a certain personalization?
Does the system inform the user about the behavioral models underlying the personalization system?
Does the system inform the user about possible constraints of the model such that may result from pre-processing or biases in the dataset?
Output:
Does the system present information to the user in a location where they can notice it and access it easily?
Does the system provide information to the user in a comprehensible way and can they act upon this information?
Does the system provide the user with information in a clear and simple language that avoids technical terms?
Does the system make it clear to the user that they interact with a machine?
Control:
Does the system provide the user with the opportunity to specify their goals which are then used for personalization?
Does the system provide the user with different options as to the personalized content they receive?
Does the system provide the user with opt-in and opt-out options (e.g. for data collection)?
If applicable, can the user adjust frequency and timing of personalized content?
Does the user have a say in which data or models are used for personalization?
Does the system encourage the user to give feedback and express their opinion about the personalization mechanisms used (type, frequency, duration, etc.)?

Table 2: Transparency Checklist

about the reasoning and underlying logic of personalization. This is a severe shortcoming as ethics literature clearly identifies the need to disclose information about how a certain outcome (personalization) emerged. We suspect that transparency about the reasoning of a system will gain relevance in the future. In fact, there is an ongoing debate whether the GDPR even provides a legal right to receive an explanation for algorithmic decision-making [43].

Literature also points to the need for user control to fulfill transparency [35]. Our application of the checklist points to significant shortcomings in existing systems in the realm of "user control." As a system designer, having applied the checklist and seen some blind spots, one would now be able to make a deliberate decision about whether to increase user control in one's own system.

For instance, being responsible for the personalization systems in Figure 1, designers might want to invest in transparency with regard to the how and why of personalized recommendations. The reasoning behind personalized recommendations may not be clear to the user. Did the services factor age, gender, frequency of engagement, previous likes and shares in other social media platforms? What assumptions were made about the user that motivated the designers to present them a certain option? How does the social network of the user affect the recommendations they receive?

Information about these questions and others can help users make autonomous decisions about their usage and consumption. It is possible that some users want to explore content beyond their age group to broaden their perspectives. Similarly, some users may not agree with gender-based recommendations and thus, information about how gender factors into personalized content may encourage them to explore the service for higher diversity. Paired with the opportunity to provide feedback and adjust the system's personalization behavior (user control), transparency practices enable a deeper and more meaningful engagement with the service at hand. Users can thus grasp how a system behavior affects their personal choices and can better adjust this behavior for their needs.

As with many practices, transparency has limits. Providing information does not guarantee that we understand a model, e.g. due to lack of resources, human capital [28], and basic digital or technical literacy [12]. Disclosing information can also confuse users rather than adding to clarity [8]. A particular concern here is that systems may provide scores of information that cannot be processed by users and that may encourage trust without increasing user autonomy and control. Transparency may further clash with important ethics principles such as privacy. Full disclosure of input or output data may put users at risk of being re-identified. Other concerns

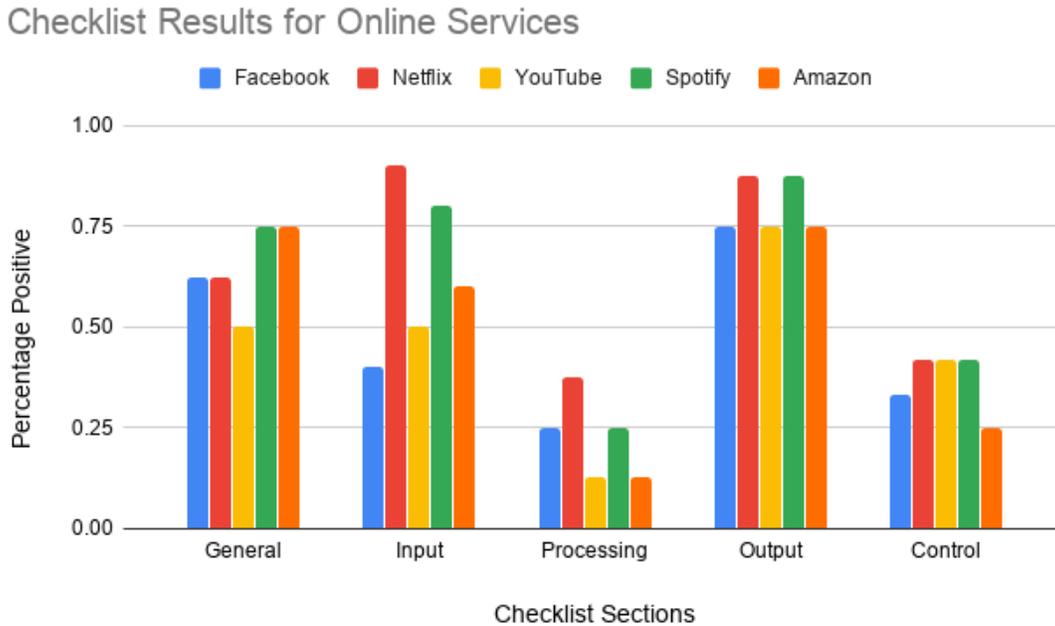


Figure 1: Preliminary checklist, online sites: Y-axis is the percentage of positive and partial replies in each checklist section

are protecting business interests (e.g. proprietary information) as well as using transparency to “game the system,” i.e. users which manipulate their input data to receive the desired outcome [27]. These limitations also put a checklist in perspective as the level of transparency depends on the unique use case. A checklist further may be misused to prove compliance with transparency best practices without meaningfully changing practices.

Another significant issue concerns the relationship of information to the user. Transparency is a relational practice: the same information may make something transparent to one group or individual but not to others [41]. It follows that transparency must be configured to the individual user. In fact, we may need a personalization technology to fulfill the transparency best practices for machine generated personalization [26, 32]. A relational approach to transparency best practices also requires considering the implications of networks on personalization outcomes. How is a user affected by other users’ behavioral patterns? How can we disclose information beyond individual personalization without violating users’ privacy?

Finally, while an ethics perspective promotes user control and meaningful transparency, it is not certain that users desire transparency and control. From privacy research, we know that users claim privacy to be an important issue for them but rarely take steps to protect their data (“privacy paradox”) [18]. Similar dynamics may apply to transparency. Nevertheless, users should have the opportunity to take advantage of transparency. System designers then have an ethical responsibility to implement transparency best practices.

7 CONCLUSION AND FUTURE WORK

In this work, we developed a transparency definition, best practices, and a checklist for system designers to better implement transparency in machine generated personalization. We applied the checklist to prominent online services that use personalization and found that systems lack transparency with regard to “processing” and “user control.” System designers may want to spend more time explaining why and how their systems reach a certain outcome (personalization) and provide more options for users to adjust personalization. In this context, we note that transparency is a relational practice and information should be personalized to ensure that diverse users understand the disclosed information.

While we propose a first transparency and user control checklist, we recognize that it may be amended in the future. Ideally, the items in the checklist should be discussed by experts in the field and present a consensus of the personalization community [6]. We encourage system designers to provide feedback on the checklist, and suggest the organization of workshops to develop tangible design solutions that implement transparency in personalization. Finally, more conceptual work on transparency and related concepts such as understandability, explainability, controllability, and accountability can help advance the discourse on responsible personalization.

8 ACKNOWLEDGEMENTS

This project has received funding from the European Union’s Horizon 2020 WeNet project (<https://www.internetofus.eu/>) under grant agreement No 823783. The authors would like to kindly thank PD Dr. Jessica Heesen for her valuable comments on draft versions of the paper.

REFERENCES

- [1] 2019. Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems. <https://standards.ieee.org/content/ieee-standards/en/industry-connections/ec/autonomous-systems.html>
- [2] 2020. <https://www.brookings.edu/research/the-case-for-ai-transparency-requirements/>.
- [3] 27 April 2016. General Data Protection Regulation (REGULATION (EU) 2016/679).
- [4] Behnouth Abdollahi and Olfa Nasraoui. C 2018. Transparency in Fair Machine Learning: the Case of Explainable Recommender Systems. In *Human and Machine Learning*, Jianlong Zhou and Fang Chen (Eds.). Human-computer interaction series, Vol. 44. Springer, Cham, Switzerland, 21–35. https://doi.org/10.1007/978-3-319-90403-0_2
- [5] ACM Council. 2018. ACM Code of Ethics and Professional Conduct. <https://ethics.acm.org/code-of-ethics/>
- [6] Balazs Aczel, Barnabas Szasz, Alexandra Sarafoglou, Zoltan Kekecs, Šimon Kucharský, Daniel Benjamin, Christopher D Chambers, Agneta Fisher, Andrew Gelman, Morton A Gernsbacher, et al. 2019. A consensus-based transparency checklist. *Nature human behaviour* (2019), 1–3.
- [7] Salema Amershi, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, Eric Horvitz, Dan Weld, Mihaela Vorvoreanu, Adam Fournery, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, and Paul N. Bennett. 2019. Guidelines for Human-AI Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, Stephen Brewster, Geraldine Fitzpatrick, Anna Cox, and Vassilis Kostakos (Eds.). ACM Press, New York, New York, USA, 1–13. <https://doi.org/10.1145/3290605.3300233>
- [8] Mike Ananny and Kate Crawford. 2018. Seeing Without Knowing: Limitations of the Transparency Ideal and its Application to Algorithmic Accountability. *New Media & Society* 20, 3 (2018), 973–989. <https://doi.org/10.1177/1461444816676645>
- [9] Sebastian Boell and Dubravka Cercez-Kecmanov. 2012. Conceptualizing Information Systems: From input-Processing-Output devices to Sociomaterial Apparatuses. (2012).
- [10] Ali Borji and Laurent Itti. 2012. State-of-the-art in visual attention modeling. *IEEE transactions on pattern analysis and machine intelligence* 35, 1 (2012), 185–207.
- [11] Joy Buolamwini and Timnit Gebru. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. *Proceedings of Machine Learning Research* 81 (2018), 1–15. <http://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf>
- [12] Jenna Burrell. 2016. How the Machine “Thinks”: Understanding Opacity in Machine Learning Algorithms. *Big Data & Society* 3, 1 (2016), 205395171562251. <https://doi.org/10.1177/2053951715622512>
- [13] Michael Chromnik, Malin Eiband, Sarah Theres Völk, and Daniel Buschek. 20 March 2019. Dark Patterns of Explainability, Transparency, and User Control for Intelligent Systems.
- [14] Crain and Nadler. 2019. Political Manipulation and Internet Advertising Infrastructure. *Journal of Information Policy* 9 (2019), 370. <https://doi.org/10.5325/jinfopoli.9.2019.0370>
- [15] E. S. Dahl. 2018. Appraising Black-Boxed Technology: the Positive Prospects. *Philosophy & Technology* 31, 4 (2018), 571–591. <https://doi.org/10.1007/s13347-017-0275-1>
- [16] Edward L Deci and Richard M Ryan. 2002. Overview of self-determination theory: An organismic dialectical perspective. *Handbook of self-determination research* (2002), 3–33.
- [17] Luciano Floridi, Josh Cowsls, Monica Beltrametti, Raja Chatila, Patrice Chazerand, Virginia Dignum, Christoph Luetge, Robert Madelin, Ugo Pagallo, Francesca Rossi, Burkhard Schafer, Peggy Valcke, and Effy Vayena. 2018. AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. *Minds and Machines* 28, 4 (2018), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- [18] Nina Gerber, Paul Gerber, and Melanie Volkamer. 2018. Explaining the Privacy Paradox: A Systematic Review of Literature Investigating Privacy Attitude and Behavior. *Computers & Security* 77 (2018), 226–261. <https://doi.org/10.1016/j.cose.2018.04.002>
- [19] AI Ethics Impact Group. [n.d.].
- [20] Jason S Hartford, James R Wright, and Kevin Leyton-Brown. 2016. Deep learning for predicting human strategic behavior. In *Advances in Neural Information Processing Systems*. 2424–2432.
- [21] Jeffrey Heer. 2019. Agency plus automation: Designing artificial intelligence into interactive systems. *Proceedings of the National Academy of Sciences* 116, 6 (2019), 1844–1850.
- [22] High-Level Expert Group on Artificial Intelligence. [n.d.]. Ethics Guidelines for Trustworthy AI. <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>
- [23] Joana Hois, Dimitra Theofanou-Fuelbier, and Alischa Janine Junk. 2019. How to Achieve Explainability and Transparency in Human AI Interaction. In *HCI INTERNATIONAL 2019 - POSTERS*, Constantine Stephanidis (Ed.). Communications in Computer and Information Science, Vol. 1033. Springer, [Place of publication not identified], 177–183. https://doi.org/10.1007/978-3-030-23528-4_25
- [24] Corey Jackson, Gabriel Mugar, Kevin Crowston, and Carsten Østerlund. 2016. Encouraging Work in Citizen Science: Experiments in Goal Setting and Anchoring. In *Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion*. 297–300.
- [25] Sucheta V Kolekar, Radhika M Pai, and Manohara Pai MM. 2019. Rule based adaptive user interface for adaptive E-learning system. *Education and Information Technologies* 24, 1 (2019), 613–641.
- [26] Pigi Kouki, James Schaffer, Jay Pujara, John O’Donovan, and Lise Getoor. 2019. Personalized Explanations for Hybrid Recommender Systems. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*, Wai-Tat Fu, Shimei Pan, Oliver Brdiczka, Polo Chau, and Gaelle Calvary (Eds.). ACM, New York, NY, USA, 379–390. <https://doi.org/10.1145/3301275.3302306>
- [27] Bruno Lepri, Nuria Oliver, Emmanuel Letouzé, Alex Pentland, and Patrick Vinck. 2018. Fair, Transparent, and Accountable Algorithmic Decision-making Processes. *Philosophy & Technology* 31, 4 (2018), 611–627. <https://doi.org/10.1007/s13347-017-0279-x>
- [28] Brent Daniel Mittelstadt, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter, and Luciano Floridi. 2016. The Ethics of Algorithms: Mapping the Debate. *Big Data & Society* 3, 2 (2016), 205395171667967. <https://doi.org/10.1177/2053951716679679>
- [29] Conor Muldoon, Michael J O’Grady, and Gregory MP O’Hare. 2018. A survey of incentive engineering for crowdsourcing. *The Knowledge Engineering Review* 33 (2018).
- [30] Ori Plonsky, Reut Apel, Eyal Ert, Moshe Tennenholtz, David Bourgin, Joshua C Peterson, Daniel Reichman, Thomas L Griffiths, Stuart J Russell, Evan C Carter, et al. 2019. Predicting human decisions with behavioral theories and machine learning. *arXiv preprint arXiv:1904.06866* (2019).
- [31] Avi Rosenfeld and Ariella Richardson. 2019. Explainability in Human-agent Systems. *Autonomous Agents and Multi-Agent Systems* 33, 6 (2019), 673–705. <https://doi.org/10.1007/s10458-019-09408-y>
- [32] Johannes Schneider and Joshua Handali. 2019. Personalized explanation in machine learning: A conceptualization. *arXiv preprint arXiv:1901.00770* (2019).
- [33] Avi Segal, Kobi Gal, Ece Kamar, Eric Horvitz, and Grant Miller. 2018. Optimizing interventions via offline policy evaluation: studies in citizen science. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [34] Avi Segal, Ece Kamar, Eric Horvitz, Alex Bowyer, Grant Miller, et al. 2016. Intervention Strategies for Increasing Engagement in Crowdsourcing: Platform, Predictions, and Experiments. (2016).
- [35] Judith Simon. 2010. The Entanglement of Trust and Knowledge on the Web. *Ethics and Information Technology* 12, 4 (2010), 343–355. <https://doi.org/10.1007/s10676-010-9243-5>
- [36] Frode Sørmo, Jörg Cassens, and Agnar Aamodt. 2005. Explanation in Case-Based Reasoning—Perspectives and Goals. *Artificial Intelligence Review* 24, 2 (2005), 109–143. <https://doi.org/10.1007/s10462-005-4607-7>
- [37] Shaun B. Spencer. 2019. The Problem of Online Manipulation. *SSRN Electronic Journal* (2019). <https://doi.org/10.2139/ssrn.3341653>
- [38] Daniel Susser, Beate Roessler, and Helen F. Nissenbaum. 2018. Online Manipulation: Hidden Influences in a Digital World. *SSRN Electronic Journal* (2018). <https://doi.org/10.2139/ssrn.3306006>
- [39] Omer Tene and Jules Polonetsky. 2013. Big Data for All: Privacy and User Control in the Age of Analytics. *Northwestern Journal of Technology and Intellectual Property* 239 (2013).
- [40] Matteo Turilli and Luciano Floridi. 2009. The Ethics of Information Transparency. *Ethics and Information Technology* 11, 2 (2009), 105–112. <https://doi.org/10.1007/s10676-009-9187-9>
- [41] Orlin Vakarelov and Kenneth Rogerson. 2019. The Transparency Game: Government Information, Access, and Actionability. *Philosophy & Technology* 23, 1–2 (2019), 193. <https://doi.org/10.1007/s13347-019-0340-z>
- [42] Giuseppe A. Veltri. 2019. The Political Bubble. *Longitude* 93 (2019), 54–59.
- [43] Sandra Wachter, Brent Mittelstadt, and Luciano Floridi. 2016. Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation. *SSRN Electronic Journal* (2016). <https://doi.org/10.2139/ssrn.2903469>
- [44] Stav Yanovsky, Nicholas Hoernle, Omer Lev, and Kobi Gal. 2019. One Size Does Not Fit All: Badge Behavior in Q&A Sites. In *Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization*. 113–120.
- [45] John Zerilli, Alistair Knott, James Maclaurin, and Colin Gavaghan. 2019. Transparency in Algorithmic and Human Decision-Making: Is There a Double Standard? *Philosophy & Technology* 32, 4 (2019), 661–683. <https://doi.org/10.1007/s13347-018-0330-6>
- [46] Yongfeng Zhang and Xu Chen. 2020. Explainable Recommendation: A Survey and New Perspectives. *Foundations and Trends® in Information Retrieval* 14, 1 (2020), 1–101. <https://doi.org/10.1561/1500000066>