Comparing Perception of Disclosure of Different Types of Information Related to Automated Tools

Vanessa Bracamonte and Takamasa Isohara
KDDI Research, Inc., Saitama, Japan

Keywords: Algorithmic Transparency, Automated Tools, Personal Data, User Study.

Abstract: Transparency has been identified as an important influence on users' perception of algorithm-based automated tools. Research in algorithmic transparency has mostly focused on two types of information: disclosure related to the personal data collected from users and explanation of how algorithms work. However, the development and use of automated tools also involve other types of information that could be the subject of disclosure. In this study, we compare perception of providing information about data provenance and human involvement, in addition to personal data processing and algorithm explanation. We conducted a user experiment and compared the disclosure of these four types of information, for two types of automated apps that process personal information. The results indicate that disclosure of information about data provenance and human involvement is perceived to be as important as personal data processing information. In addition, the relative importance of explanations about the algorithm, compared to other types of information, depended on the type of app. Finally, perception of the usefulness and accessibility of the information did not vary between types of information or app, but participants considered they would be able to understand explanations about the algorithm more than other types of information.

1 INTRODUCTION

In recent years, there has been increased interest in the transparency of automated tools. Existing research in technology that processes personal information has emphasized the importance of communicating to users about the collection, storage and processing of their data, and the methods for how to provide this information in a usable manner have also been widely investigated (Janic et al., 2013; Schaub et al., 2015; Murmann and Fischer-Hübner, 2017). For automated tools, the use of algorithms has meant that providing transparency also includes providing explanations of the logic and outcome of machine learning models (Adadi and Berrada, 2018). Research has evaluated how to provide explanations in an understandable and useful way, and how those explanations affect user perception (Wang et al., 2016; Mittelstadt et al., 2019; Kunkel et al., 2019; Nourani et al., 2019).

Although most of the focus has been on data processing and algorithm explanation, algorithmic transparency also involves other types of information. Among these is information about the provenance of the data used to train those algorithms, and information about the role of humans (for example, as reviewers or evaluators) in the algorithmic process (Diakopoulos and Koliska, 2017). These aspects of algorithmic transparency related to automated tools are often investigated separately, and how users view the disclosure of these different types of information compared to each other has not often been considered. In this paper, we conduct a preliminary study to compare the perception of importance, usefulness, accessibility and understanding of the disclosure of different types of information related to algorithm-based automated tools.

2 RELATED WORK

Research has proposed that algorithmic transparency consists of different dimensions (Diakopoulos and Koliska, 2017). These dimensions include information about personal data processing, algorithm explanation, the data used to train the algorithm and the extent of human involvement in the process. How users view the disclosure of each of these types of information, and how to provide transparency about it in a usable and understandable manner has been investigated separately.

Bracamonte, V. and Isohara, T.
Comparing Perception of Disclosure of Different Types of Information Related to Automated Tools.
DOI: 10.5220/0010901200003120
In *Proceedings of the 8th International Conference on Information Systems Security and Privacy (ICISSP 2022)*, pages 544-550
ISBN: 978-989-758-553-1; ISSN: 2184-4356

Transparency of personal data processing, including collection and storage practices, has been researched in the context of the information provided in privacy policies. Extensive research has been conducted on how to disclose information about data processing practices in the privacy notices that providers give to users (Janic et al., 2013; Schaub et al., 2015; Murmann and Fischer-Hübner, 2017; Tesfay et al., 2018). When using AI-based technology such as smart assistants, some user demographics such as older users indicate that they are not aware of the policies related to the privacy and security of their data (Bonilla and Martin-Hammond, 2020). These users report being concerned about this type of information, but they also indicated that they do not know where to find it (Bonilla and Martin-Hammond, 2020).

Transparency of the algorithm itself is also considered important to improve perception of automated services. When the algorithms are opaque, users can become frustrated (Eslami et al., 2019). Therefore, research is being conducted on how to provide explanations of algorithms (Adadi and Berrada, 2018) and how to improve the usability of these explanations (Wang et al., 2016; Mittelstadt et al., 2019; Kunkel et al., 2019; Nourani et al., 2019). On the other hand, research also indicates that the level of transparency related to how an algorithm works can have a negative effect on trust (Kizilcec, 2016).

Although there has been extensive research conducted on the disclosure of personal data processing and explanations of algorithms, there has been less focus on other aspects of algorithmic transparency. With regards to human involvement in algorithmic processes, this involvement occurs in different stages of the development and use of automated tools (Bechmann and Bowker, 2019; Tubaro et al., 2020), and it can have an impact on the privacy of users. Therefore, this is important information for the purposes of transparency. Users report that they expect human oversight in the review of decisions of certain automated tools (Kaushik et al., 2021). However, depending on the context of use, involvement of human reviewers carries a privacy risk, and therefore transparency in the processes used for this oversight are required (Kaushik et al., 2021).

With regards to transparency of data provenance, that is, disclosure of information about the data used to train algorithms, research has been conducted on how to provide this information (Gebru et al., 2021). Although this type of information disclosure is often aimed at developer users, it can also influence end-user perception of an automated system (Anik and Bunt, 2021).

All of these types of information are important to inform users about algorithm-based tools, but they are not often considered together. In this work, we evaluate how the disclosure of these types of information is perceived relative to each other.

3 METHOD

3.1 Experiment Design

We designed an experiment to compare how disclosure of different types of information about automated tools is perceived, when those tools process personal information that may be sensitive or private. The experiment was designed with a between-subjects factor of type of app, and a within-subjects factor of type of information disclosed.

For the between-subjects factor of type of app, we used two hypothetical apps that had an objective to help the user. We chose to use hypothetical AI-based apps for privacy and eHealth, as these types of apps would need to process potentially sensitive personal information in order to fulfill their objective. One was an app for determining whether private or sensitive information was contained in social media content (PrivApp) and the other was an app for inferring and reducing the level of stress through a game (eHealthApp). These hypothetical apps were based on research proposals for detection of privacy sensitive information (Tesfay et al., 2018) and stress (Garcia-Ceja et al., 2016). The choice of using two different hypothetical apps for the experiment was done in order to evaluate whether a different level of sensitivity in the data processed and results of the different apps had an effect on the perception of information disclosure.

The within-subjects factor of type of information disclosed had the following levels: algorithm explanation, data processing, data provenance, and human involvement. As mentioned, these type of tools process personal data, infer personal data as output through the use of an algorithm and are trained by using data from other users. The first three aspects of algorithm transparency are present in most AI-based tools. And for the types of hypothetical apps in the experiment, the involvement of humans in the process of improving the algorithm or to debug problems is a reasonable assumption.

3.2 Questionnaire

The questionnaire was structured as follows. We first presented participants with a description of the hypothetical apps. For the social media privacy app (Pri-

vApp), we asked participants to imagine a mobile app that could automatically analyze the text and images that they wanted to post on social media. The purpose of the app was described as "to protect (their) privacy". We indicated that the app would ask them for their desired privacy protection level when they first started, and that they could write text or upload photos as usual. If the app detected that the content of the text or images included information that could be considered private or sensitive, it would show a message. Finally, the app would remember the user's choice and use it for future decisions.

For the stress reduction app (eHealthApp), we described the app as able to automatically analyze the user's playing style in games, with the purpose of reducing their stress level. We indicated that the app would ask them for their current stress level and information such as age, and then they would play some games as usual. The app would detect whether their stress level appeared to have been reduced and would show a message. The app would then remember if the user stopped or continued playing and use that choice in the future. Finally, we also described each app as free, that it did not have ads, and that the use of the app was voluntary, not obligatory.

After presenting this description, we asked participants open-ended questions about the app: *"What is the purpose of the app?"*, *"What will the app ask when you first start?"*, *"What will the app detect?"*. These questions served to check participants' attention and to verify that the they had understood the characteristics of the app that was described. The attention questions were shown in random order. We also asked participants about whether they would use the app (*"I would use this app in my daily life."*) or thought other people would use it (*"I can think of people I know who would use this app."*).

Next, we again indicated to the participants that the app worked by analyzing their data and by applying an algorithm to get results. We explained that we would ask their opinion on 4 types of information that could be disclosed, on what happens to their data and how the app worked. After that, we presented the participants with a description of each type of information, and gave examples of what the disclosure included. The description and examples of each type of information was followed by questions on their opinion of the disclosure of that type of information. Participants viewed and answered questions about all 4 types of information, in random order.

For each type of information disclosed we asked participants about their opinion of its importance (*"In general, it is very important that the app gives this type of information to users."*), its usefulness for deciding whether to use the app (*"This type of information would be useful for me to decide whether to use the app."*) and their perceived information gathering capacity in terms of accessibility (*"If I wanted to find this type of information, I would know exactly where to look."*) and understanding (*"This type of information would be too technical for me to understand."* (Reverse coded)). These last two items were adapted from (Griffin et al., 2008). The questions were presented in random order. The responses were measured on a 7-point Likert scale, ranging from *Strongly disagree* to *Strongly agree*. We also asked participants about where they would expect to find each type of information (*"If you wanted to know this type of information, where would you first look for it?"*.

We asked participants about their perception of the app in terms of sensitivity of the data collection (*"The data that this app would collect is very sensitive and/or private."*) and sensitivity of its results (*"The results of this app would be very sensitive and/or private."*). Finally, we asked questions about participants' use of automated apps ("AI-based apps") (*"Do you currently use any AI-based apps that have access to your data or personal information?"*) and asked them to name the apps, if possible.

3.3 Data Collection

We used Amazon Mechanical Turk to collect responses. We limited participation to workers from the USA, with a 99% approval rate, and that had worked on at least 1000 tasks. We estimated the survey response time at 10 minutes, and compensated participants with US$1.7. The actual survey time obtained after collecting the responses was 10.7 minutes, which resulted in a compensation rate of $9.5/h. The survey ran from October 29 to November 1, 2021.

3.4 Limitations

The methodology used for this study has a number of limitations. First, the study was vignette-based and the participants were only provided with description of the hypothetical app and the type of information that would be disclosed. Second, we survey participants recruited on Amazon Mechanical Turk, and these participants can have a higher level concern about privacy than the general population (Kang et al., 2014). Finally, we used single items questions which may not entirely capture the participants' perception of transparency aspects.

4 RESULTS

We initially collected 220 responses. After rejecting responses with answers that were unrelated to the questions, the sample for analysis was 183 responses. The gender distribution of the sample was 63 female (34%) and 120 male (66%) participants. The age distribution was: 20-29 years-old, 31 participants (17%); 30-39 years-old, 83 participants (45%); 40-49 years-old, 33 participants (18%); 50-59 years-old, 27 participants (15%); and 60+ years-old, 9 participants.

The results of two-sample Mann-Whitney U tests showed a significant difference for the perceived sensitivity of data collection (W = 5424, p <0.001) and the sensitivity of results (W = 5241, p <0.01) between the two types of apps. For data collection sensitivity, the PrivApp had a mean = 5.6, median = 5; and the eHealthApp had a mean = 4.8, median = 4. For result sensitivity, the PrivApp had a mean = 5.5, median = 5; and the eHealthApp had a mean = 4.8, median = 4. The results show that the apps were perceived as having a different level of sensitivity. However, we had expected that the eHealthApp would be perceived as more sensitive, as it related to the users' health (stress level). Instead, the results show that both data collection and results of the eHealthApp were considered less sensitive and/or private compared to the PrivApp (Figure 1).

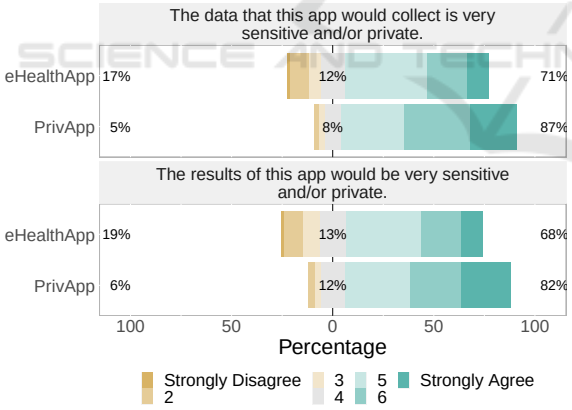

Figure 1: Distribution of the answers to the question of sensitivity of data collection (top) and results (bottom) for each type of app.

Separate two-sample Mann-Whitney U tests were also used to test the differences in whether the participant thought they or people they know would use the app. The results show that there were no significant differences for either of these variables between the PrivApp and eHealthApp (p >0.05). For both variables, the mean was above the neutral point of the scale, indicating a positive opinion about these type of apps: mean of use for themselves = 4.8, median = 5; mean of use for others = 5.6, median = 6.

For evaluating the main variables of the survey, we used a 2x4 repeated measures ANOVA. Specifically, we used the Aligned Rank Transform (ART) ANOVA (Wobbrock et al., 2011) (R implementation (Kay et al., 2021)) as a non-parametrical analysis for the single item variables. For the variable of importance of providing the information, the results showed a significant main effect of the type of information (p = 0.03), as well as a significant interaction effect (p = 0.04). We conducted a contrast analysis for the interaction effect, with Tukey correction. However, none of the individual contrasts were significant, which may be due to lack of power for the analysis. As can be visualized in the interaction plot for the variable (Figure 2), providing information was considered highly important for all types. The mean of importance of disclosure for the eHealthApp was consistently higher than for PrivApp, except for information about algorithm explanation where the relationship is inverted.

For the variable of usefulness of information disclosure, the results show that there were no significant differences for the main effects or the interaction. We observe that similarly to importance, participants consider that disclosure of all types of information would be highly useful for making a decision.

With regards to whether participants expected to be able to find the information (accessibility), the results showed no significant differences for the main effects or the interaction. For the expectation of being able to understand the information, the results showed a significant difference in the main effect of type of information disclosed. The contrast analysis, with Tukey correction, showed that participants considered that they would be able to understand the information about how the algorithm works more than information about how their personal data was processed (p = 0.04), data provenance (p= 0.03) and human involvement (p = 0.04). In general, participants were optimistic that they would find and understand these types of information.

The results in general show that, as expected, there was a positive perception towards information disclosure. Participants considered that providing information of every type was important and that all of these types of information would be useful for them to decide about the app. Although the perceived sensitivity level was significantly different, with the data and results of PrivApp being perceived as more sensitive, there were no significant results for the main effect of type of app for any of the variables. However, we observe that perception of disclosure of information about the algorithm had an inverted relationship com-

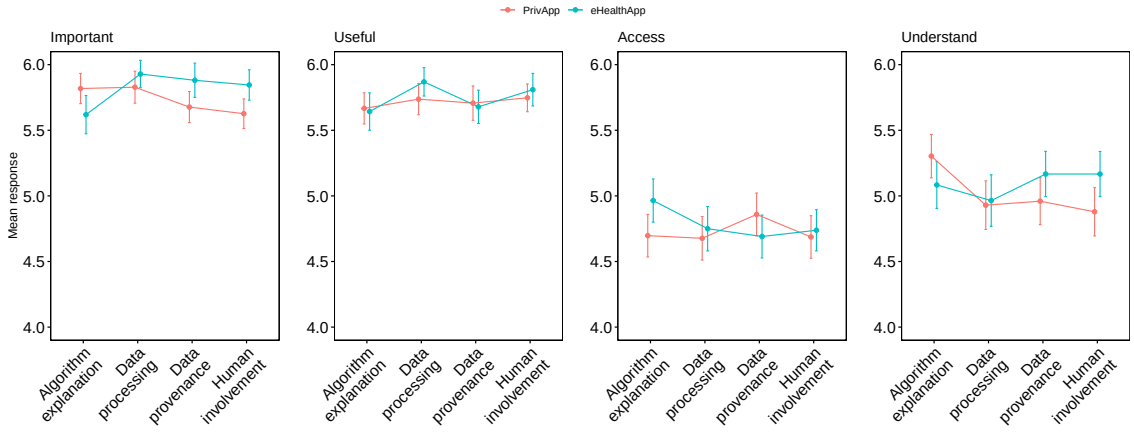

Figure 2: Interaction plots for the mean and standard error for the variables of interest. The response scale is partially represented in the plots, and ranges from 4: Neither agree nor disagree (neutral point) to 6: Agree.

pared to other types of information, between the two types of apps. In future research, we will investigate what aspects of the app affect the perception of algorithm explanation disclosure.

Another observation from these results is that participants considered information about data provenance and human involvement as important as information about the processing of their personal data. We hypothesized that there would be an order to the perceived importance and usefulness of these aspects. These results indicate that participants consider them equally, at least in this hypothetical scenario, although this may be a general effect of a high perceived value of transparency.

With regards to the question of where participants would first look for each type of information, the distribution of responses are shown in Figure 3. The majority of participants indicated that they would first look for any type of information in the app's privacy policy. Privacy policies are the most well-known notice for users (Reidenberg et al., 2015), and therefore it follows that it would be their first choice when looking for further information. However, privacy policies do not often contain disclosure of aspects of algorithm transparency other than data processing. Research on more usable versions or alternatives to privacy policies includes proposals for how to categorize the information that these privacy policies contain (Wilson et al., 2016; Zaeem et al., 2018; Tesfay et al., 2018). The proposed categorizations are based on privacy research and on regulation such as the EU's GDPR, and do not include categories related to other aspects of algorithm transparency. In addition, we observe that the second most frequent choice was the FAQ of the app. This form of information disclosure is usually structured as question-and-answer and, as its name suggest, addresses questions related to topics that generate users' interest. Its structure may be perceived as more accessible to users, in comparison to the terms and conditions document or even a user manual. These results give an indication of where and how participants expect the disclosure of this information, and also hint at logistical challenges of providing it.

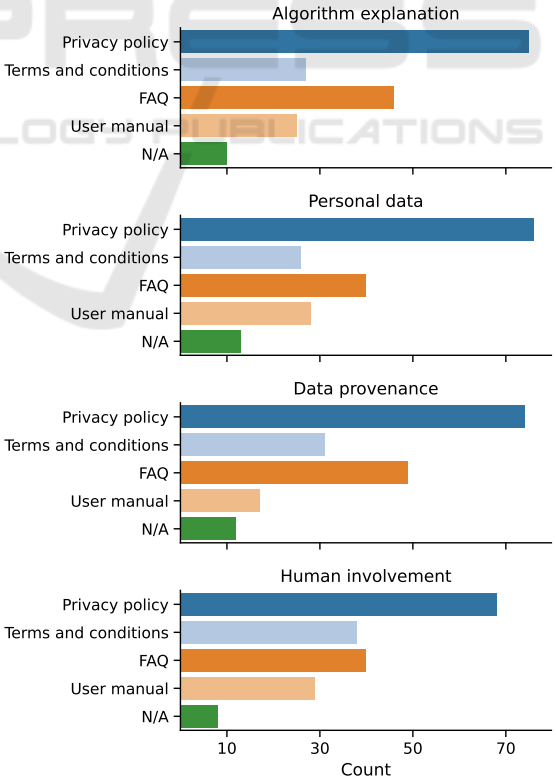

Figure 3: Distribution of the answers to the question of where the participants would first look for each type of information.

Finally, with regards to the participants use of AI-based apps, 111 participants (61%) indicated that they did not use these kind of apps, 38 participants (21%) indicated that they did and 34 participants (18%) were unsure. The apps most frequently mentioned by participants that indicated they used AI-based apps were smart assistants such as Google Assistant, Amazon Alexa and Apple Siri. Participants also mentioned the names of the devices that include these assistants (e.g Amazon Echo which includes Amazon Alexa). Other type of app frequently mentioned were fitness apps (e.g. Fitbit). Social media websites were also mentioned: "facebook, instagram, pretty much anything nowadays". Similar answers about AI-apps being perceived as widespread were also reported by other participants.

The proportion of participants that indicated that they were not sure whether they were using an AI-based app was similar to the proportion of those participants that had answered in the affirmative. These participants mentioned specific apps in their answers, along with comments that indicated that uncertainty: "I am unsure if Google maps (and similar apps) are considered AI-based, but I am relatively certain that this app has access to some of my private data, such as location information." Other participants reported a general assumption, without mentioning specific examples: "I'm not sure, but I'm assuming that some of the apps I use do use AI and collect my data. I just don't know for sure to name them.", and "I chose I'm not sure for a reason; I don't know if any of the apps that I have employ some degree of ai that can access my personal information." The question addressed AI-based apps that accessed personal information in particular, but as the comments indicate, some participants are uncertain of whether the apps were AI-based and whether the apps made use of their personal information. The answers to this question suggest that this information is not clearly communicated to some participants, even for apps from well-known companies.

5 CONCLUSION

In this paper, we conducted an experiment to investigate how the disclosure of four types of information related to automated tools is perceived. In the experiment, we described two hypothetical apps that have to process personal information in order to help the user. The results showed that participants considered information about data provenance and human involvement as important as information about personal data processing. The importance of disclosure of information about how the algorithm works, compared to other types of information, appeared to depend on the type of app. On the other hand, participants opinion of their own ability to understand information was higher for explanations about the algorithm than for other types of information.

In this preliminary work, we have focused on perception related to the type of information that would be disclosed to the user. Forms of transparency, such as transparency of final decisions, transparency in rationale and transparency in process (de Fine Licht and de Fine Licht, 2020), should also be considered in future research. Research indicates that providing information about the reason for a decision can improve the perception of fairness in a process (de Fine Licht et al., 2011). However, some forms of transparency can also result in riskier behavior (Acquisti et al., 2013; Adjerid et al., 2013). Future work will evaluate how type of information, forms of transparency and level of detail interact to affect perception of automated apps, and investigate how to provide transparency in a way that results in more protection for users.

REFERENCES

Acquisti, A., Adjerid, I., and Brandimarte, L. (2013). Gone in 15 Seconds: The Limits of Privacy Transparency and Control. *IEEE Security Privacy*, 11(4):72–74.

Adadi, A. and Berrada, M. (2018). Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, 6:52138–52160.

Adjerid, I., Acquisti, A., Brandimarte, L., and Loewenstein, G. (2013). Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, SOUPS '13, pages 1–11, New York, NY, USA. Association for Computing Machinery.

Anik, A. I. and Bunt, A. (2021). Data-Centric Explanations: Explaining Training Data of Machine Learning Systems to Promote Transparency. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, number 75, pages 1–13. Association for Computing Machinery, New York, NY, USA.

Bechmann, A. and Bowker, G. C. (2019). Unsupervised by any other name: Hidden layers of knowledge production in artificial intelligence on social media. *Big Data & Society*, 6(1):2053951718819569.

Bonilla, K. and Martin-Hammond, A. (2020). Older adults' perceptions of intelligent voice assistant privacy, transparency, and online privacy guidelines. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*.

de Fine Licht, J., Naurin, D., Esaiasson, P., and Gilljam, M. (2011). Does transparency generate legitimacy? an experimental study of procedure acceptance of open-

and closed-door decision-making. *QoG Working Paper Series*, 8:1–32.

de Fine Licht, K. and de Fine Licht, J. (2020). Artificial intelligence, transparency, and public decision-making. *AI & SOCIETY*, 35(4):917–926.

Diakopoulos, N. and Koliska, M. (2017). Algorithmic Transparency in the News Media. *Digital Journalism*, 5(7):809–828.

Eslami, M., Vaccaro, K., Lee, M. K., Elazari Bar On, A., Gilbert, E., and Karahalios, K. (2019). User Attitudes towards Algorithmic Opacity and Transparency in Online Reviewing Platforms. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–14. Association for Computing Machinery, New York, NY, USA.

Garcia-Ceja, E., Osmani, V., and Mayora, O. (2016). Automatic Stress Detection in Working Environments From Smartphones' Accelerometer Data: A First Step. *IEEE Journal of Biomedical and Health Informatics*, 20(4):1053–1060.

Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., III, H. D., and Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12):86–92.

Griffin, R. J., Yang, Z., Ter Huurne, E., Boerner, F., Ortiz, S., and Dunwoody, S. (2008). After the flood: Anger, attribution, and the seeking of information. *Science Communication*, 29(3):285–315.

Janic, M., Wijbenga, J. P., and Veugen, T. (2013). Transparency Enhancing Tools (TETs): An Overview. In *2013 Third Workshop on Socio-Technical Aspects in Security and Trust*, pages 18–25.

Kang, R., Brown, S., Dabbish, L., and Kiesler, S. (2014). Privacy Attitudes of Mechanical Turk Workers and the U.S. Public. In *10th Symposium On Usable Privacy and Security ({SOUPS} 2014)*, pages 37–49.

Kaushik, S., Yao, Y., Dewitte, P., and Wang, Y. (2021). "How I Know For Sure": People's Perspectives on Solely Automated Decision-Making ({SADM}). In *Seventeenth Symposium on Usable Privacy and Security ({SOUPS} 2021)*, pages 159–180.

Kay, M., Elkin, L. A., Higgins, J. J., and Wobbrock, J. O. (2021). *ARTool: Aligned Rank Transform for Nonparametric Factorial ANOVAs*.

Kizilcec, R. F. (2016). How Much Information? effects of Transparency on Trust in an Algorithmic Interface. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 2390–2395, New York, NY, USA. Association for Computing Machinery.

Kunkel, J., Donkers, T., Michael, L., Barbu, C.-M., and Ziegler, J. (2019). Let Me Explain: Impact of Personal and Impersonal Explanations on Trust in Recommender Systems. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–12. Association for Computing Machinery, New York, NY, USA.

Mittelstadt, B., Russell, C., and Wachter, S. (2019). Explaining Explanations in AI. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT* '19, pages 279–288, New York, NY, USA. Association for Computing Machinery.

Murmann, P. and Fischer-Hübner, S. (2017). Tools for Achieving Usable Ex Post Transparency: A Survey. *IEEE access : practical innovations, open solutions*, 5:22965–22991.

Nourani, M., Kabir, S., Mohseni, S., and Ragan, E. D. (2019). The Effects of Meaningful and Meaningless Explanations on Trust and Perceived System Accuracy in Intelligent Systems. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 7(1):97–105.

Reidenberg, J. R., Breaux, T., Cranor, L. F., French, B., Grannis, A., Graves, J. T., Liu, F., McDonald, A., Norton, T. B., and Ramanath, R. (2015). Disagreeable privacy policies: Mismatches between meaning and users' understanding. *Berkeley Tech. LJ*, 30:39.

Schaub, F., Balebako, R., Durity, A. L., and Cranor, L. F. (2015). A design space for effective privacy notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 1–17.

Tesfay, W. B., Hofmann, P., Nakamura, T., Kiyomoto, S., and Serna, J. (2018). PrivacyGuide: Towards an Implementation of the EU GDPR on Internet Privacy Policy Evaluation. In *Proceedings of the Fourth ACM International Workshop on Security and Privacy Analytics*, IWSPA '18, pages 15–21, New York, NY, USA. ACM.

Tubaro, P., Casilli, A. A., and Coville, M. (2020). The trainer, the verifier, the imitator: Three ways in which human platform workers support artificial intelligence. *Big Data & Society*, 7(1):2053951720919776.

Wang, N., Pynadath, D. V., and Hill, S. G. (2016). Trust Calibration Within a Human-Robot Team: Comparing Automatically Generated Explanations. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, HRI '16, pages 109–116, Piscataway, NJ, USA. IEEE Press.

Wilson, S., Schaub, F., Dara, A. A., Liu, F., Cherivirala, S., Giovanni Leon, P., Schaarup Andersen, M., Zimmeck, S., Sathyendra, K. M., Russell, N. C., B. Norton, T., Hovy, E., Reidenberg, J., and Sadeh, N. (2016). The Creation and Analysis of a Website Privacy Policy Corpus. Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 1330–1340. Association for Computational Linguistics.

Wobbrock, J. O., Findlater, L., Gergle, D., and Higgins, J. J. (2011). The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 143–146, New York, NY, USA. Association for Computing Machinery.

Zaeem, R. N., German, R. L., and Barber, K. S. (2018). PrivacyCheck: Automatic Summarization of Privacy Policies Using Data Mining. *ACM Trans. Internet Technol.*, 18(4):53:1–53:18.