Communicating Differential Privacy Models by Illustrations: A Survey and In-Depth Interview Study

Chuhao Wu[1], Tianhao Wang[2], Robert W. Proctor[2], Ninghui Li[2], Jeremiah Blocki[2], and Aiping Xiong[1]
[1]The Pennsylvania State University, [2]Purdue University

Differential privacy (DP) techniques have been applied to protect individuals' data privacy and ensure utility. However, whether users will understand and trust DP and its different models remains unclear. The current work evaluated users' feedback of proposed illustrations of three DP models: Central DP (Dwork et al., 2006), Local DP (Erlingsson et al., 2014), and Shuffler DP (Bittau et al., 2017). We conducted an online survey study with 30 participants and an in-depth interview study with an additional six participants.

The survey was designed on Qualtrics and distributed online through Amazon Mechanical Turk. After informed consent, the survey started with a description of the location data-collection scenario. In the scenario, we introduced DP to address the re-identification of anonymized location data. Then, the text description and visualization of each model were presented in a randomized order, except that the Local DP was always presented before the Shuffler DP since the former serves as the basis for the latter. Participants answered one comprehension question for each model, then ranked the order of the three models based on the perceived level of usefulness and security/privacy. We asked participants to select the DP model for data sharing in two scenarios. Participants then explained their selection decisions with answers to an open-ended question. At the end of the survey, participants filled out their demographic information. The survey took a median of 7.5 minutes to complete, and the payment was $1.50 for each participant. Results showed that participants had difficulty understanding and differentiating the security/privacy aspects of the three models. Also, it was not easy for them to distinguish the utility gap between Local DP and Shuffler DP. Participants' data-sharing decisions suggest that they grasped the implication of each model in general. However, responses to the open-ended questions revealed misunderstandings from a few participants.

To gain more insights into how the textual description and visualization helped or failed to help people understand each model, we conducted an interview study. The interview protocol followed a semi-structured design with guiding questions, including participants' general impression of each model and suggestions to improve the text descriptions and illustrations. We recruited six participants through emailing acquaintances who had limited knowledge or prior experience with any DP technique. The interviews were conducted virtually and audio recorded. Participants completed the survey before joining the interview session. Two of the authors conducted a thematic analysis (Braun & Clarke, 2006) using the audio transcripts from the interviews. Three themes were identified: 1) some key visualization about data perturbation failed to capture participant's attention; 2) it was difficult to compare the Local DP and the Shuffler DP from the security/privacy aspect; and 3) participants mentioned that they considered things more than data accuracy when evaluating the model's usefulness.

The current work can provide insights into designing illustrations to communicate DP models effectively. However, due to the qualitative nature and the small sample size, future studies are needed to address the identified issues more comprehensively.

Acknowledgments. This work was partly supported by the NSF award #1931443.

References

1. Bittau, A., Erlingsson, Ú., Maniatis, P., Mironov, I., Raghunathan, A., Lie, D., Rudominer, M., Kode, U., Tinnes, J., & Seefeld, B. (2017). Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th Symposium on Operating Systems Principles* (pp. 441–459). ACM.
2. Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, *3*, 77-101.
3. Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In S. Halevi & T. Rabin (Eds.), *Theory of Cryptography* (pp. 265–284). Springer.
4. Erlingsson, Ú., Pihur, V., & Korolova, A. (2014). RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications (CCS '14)* (pp. 1054–1067). ACM.