How Usable Is the Spinner-based Randomized Response Technique?

Seo Young Ko, Sriram Viswanathan, Alan Esquenazi, Swadhin Routray, Jatan Loya, Tianshi Li, and Lorrie Cranor
Carnegie Mellon University

Background & Research Question

- Randomized Response Technique (RRT) is one of the Local Differential Privacy (LDP) mechanisms for privacy-preserving sensitive data sharing without trusted administrators
  - LDP differs from DP in terms of the random noise being added at an individual level before it is sent to the server or administrator.
- **Challenges:** Difficult for non-technical users to understand, trust the RRT, and add valid noise
  - Bullek et al proposed implementing RRT using a spinner to improve comprehension and trustworthiness [1]
- **Research Question**: Can people successfully add valid noise with the spinner based randomized response technique?

Methods

- Between-group and exploratory 10-min online survey using Qualtrics
  - Recruited 60 lay persons via Prolific, 20 in each condition
- Conditional Groups:
  - Group 1: No Differential Privacy
  - Group 2: Textual RRT with automated noise
  - Group 3: Spinner RRT with user-led noise
- Utilized hypothetical online survey scenarios that asked for sensitive information
- Both qualitative and quantitative measures for comprehension, trust and comfort levels

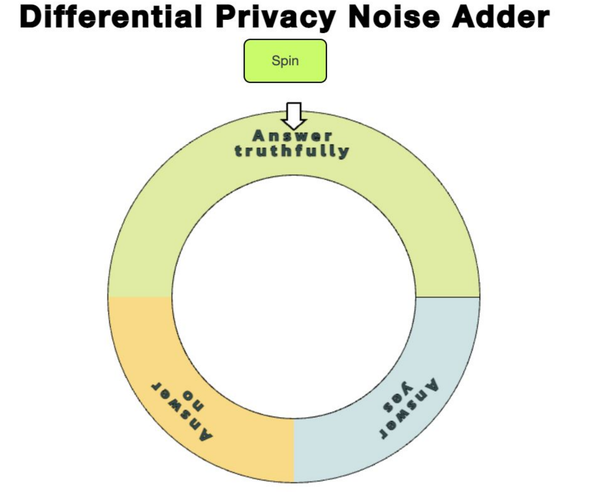

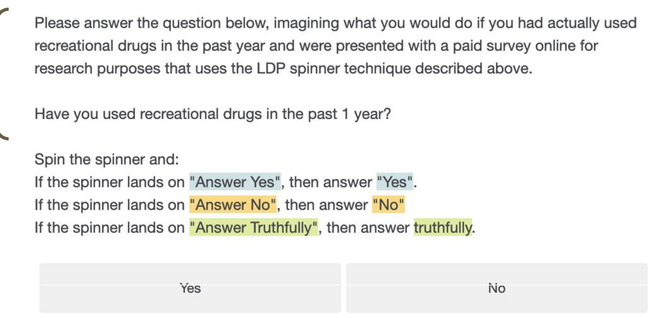

Results

- Our quantitative data analysis did not lead to any significant differences in comprehension, honesty, discomfort and trust levels between groups.
- Only 13 out of 20 of Group 3 were able to add noise correctly using the spinner
  - Even with explicit "Answer Yes" and "Answer No" responses, some failed
- Qualitative data analysis shows:
  - Some people using spinner couldn't understand the noise add mechanism and how it provides privacy (Group 3)
  - Machine-added noise (Group 2) had some complaints about the lack of transparency

Spinner Response	The number of people who correctly added noise by following the spinner prompt
"Answer Truthfully"	6 out 11 (54%)
"Answer No"	4 out of 5 (80%)
"Answer Yes"	3 out of 4 (75%)
Total	13 out of 20 (65%)

Example of qualitative responses:

"I don't fully understand how a spinner creates noise over my data"

"Just being told that it exists doesn't mean it will work"

Discussion & Future Work

- The ineffectiveness of spinner-based interface may be due to a lack of understanding or trust in it. Conducting interviews to uncover specific features for improvement will complement this work.
- How to improve the spinner interface?
  1. Provide a holistic view of DP mechanism including aggregation and data analysis instead of a narrow view of adding noise to achieve better comprehension (inherent limitation of spinner)
  2. Use offline coin-flip mechanism [2] instead of virtual spinner to improve trust

References
[1] Brooke Bullek, Stephanie Garboski, Darakhshan J. Mir, and Evan M. Peck. Towards Understanding Differential Privacy: When Do People Trust Randomized Response Technique? CHI 2017
[2] Cynthia Dwork. Differential privacy. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, Automata, Languages and Programming, pages 1–12, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg

CyLab Carnegie Mellon University Security and Privacy Institute
Carnegie Mellon University