Exposing and Mitigating Privacy Loss in Crowdsourced Survey Platforms

Thivya Kandappu, Vijay Sivaraman
University of New South Wales, Australia
t.kandappu@student.unsw.edu.au
vijay@unsw.edu.au

Arik Friedman, Roksana Boreli
NICTA, Australia
arik.friedman@nicta.com.au
roksana.boreli@nicta.com.au

ABSTRACT

Crowdsourcing platforms such as Amazon Mechanical Turk and Google Consumer Surveys can profile users based on their inputs to online surveys. In this work we first demonstrate how easily user privacy can be compromised by collating information from multiple surveys. We then propose, develop, and evaluate a crowdsourcing survey platform called Loki that allows users to control their privacy loss via at-source obfuscation.

Categories and Subject Descriptors

K.6.5 [**Management of Computing and Information Systems**]: Security and Protection; C.2.4 [**Computer-Communication Networks**]: Distributed Systems

Keywords

Privacy; Surveys; Obfuscation

1. INTRODUCTION

Academic and market researchers are increasingly using online platforms for crowdsourcing sruvey information from online users. For example, the Amazon Mechanical Turk (AMT) platform [1] pays volunteers to take surveys, while the Google Consumer Surveys platform [7] restricts user access to premium content via a "surveywall". Users participating in either system can lose privacy with each personal fact or opinion that they reveal in the course of the surveys they participate in. This privacy loss can accumulate over time, with potential social and legal consequences for the users.

Our first contribution is to show that an entity can easily use an existing crowdsourcing platform (AMT in our case) to de-anonymize users and obtain sensitive private information, in a short time at very low cost, by correlating responses across multiple surveys. Our second contribution is to design, prototype, and evaluate an alternative crowdsourcing platform that allows users to obfuscate their answers at-source to control their privacy loss without having to trust any external entity. All user experiments undertaken in this work were conducted with ethics approval (number 08/2013/19) obtained from the UNSW Human Research Ethics Advisory Panel 'H'.

CoNEXT Student Workshop'13, December 9, 2013, Santa Barbara, CA, USA.

http://dx.doi.org/10.1145/2537148.2537150.

2. EXPOSING PRIVACY LOSS

To illustrate how easily privacy is compromised in a crowdsourcing platform, we launched the following series of survey tasks in AMT (launched through a third-party aggregator called CrowdFlower [3]):

1. The first survey queried users for their opinions on astrology services, and in the process obtained their star-sign and day/month of birth.
2. In the second survey we conducted a market research of online match-making services, and obtained the users' gender and year of birth.
3. In the third survey we asked about mobile phone coverage and obtained users' zip code information.

We designed our surveys with sufficient redundancy to help us identify and filter out users who gave random responses. Further, these surveys were posted independently over several days, and users were unlikely to have known that they were conducted by the same entity. Note that although AMT does not reveal the name or personal details of any user, it reports back to the surveyor a unique ID that is constant across the surveys taken by a user. We therefore had the potential ability to identify users who took all the three surveys above, and de-anonymize them using their date of birth, gender, and zip code (previous studies [11, 6] have shown the effectiveness of using these attributes in re-identification of individuals).

We then launched a fourth survey asking users to anonymously tell us about their smoking habits and coughing frequency. Of the 400 unique users who took our surveys, 72 could be de-anonymized from the first three surveys, and we could infer the respiratory health (and likelihood of tuberculosis) for 18 of these de-anonymized individuals from the fourth survey using their unique ID, resulting in a serious breach of privacy. Our experiment took only a few days and cost less than $30; we can only imagine what the scale of privacy loss could be, were this experiment to be conducted by entities with larger resources.

We followed up the above experiment with another survey in which we asked users if they would participate if they knew they could be de-anonymized and profiled. Out of

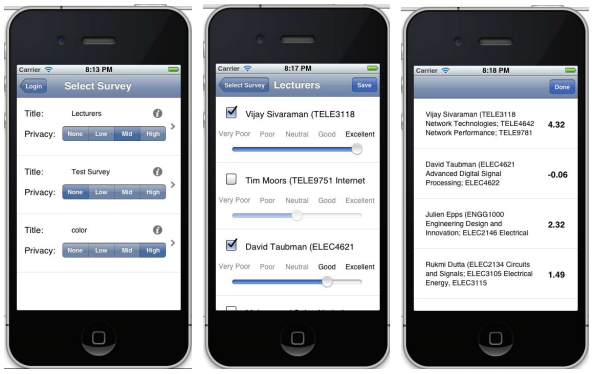

(a) Surveys (b) Questions (c) Responses

Figure 1: Screenshots of iPhone App showing (a) list of surveys and choice of privacy levels, (b) ratings-based questions , and (c) obfuscated user responses.

100 users who took this survey, 73 (including 15 of the 18 users above whose respiratory health could potentially be de-anonymized) responded that they did not know they could be profiled, and indicated that they would not participate if they knew they were being profiled. Our experiments illustrate that users can be profiled easily and at low cost, despite their disapproval of such practices.

3. SOLUTION DESIGN AND RESULTS

3.1 Solution Approach

Unlike prior approaches that rely on anonymization (that can be circumvented [9, 10]) or trusted third-parties [2, 8], we believe that the best entity to control privacy loss is the user. We therefore designed an (iOS and Android) app that allows users to choose their level of privacy for each survey (Fig. 1(a) shows our app interface with an option of none, low, medium, or high privacy), respond to the survey questions as usual (as a first step we focus only on ratings-based questions as shown in Fig 1(b)), and have the app obfuscate their responses prior to uploading them to the server (Fig. 1(c) shows the obfuscated ratings that are reported). Our obfuscation method adds Gaussian noise to the user's true response, with standard deviation successively larger for higher privacy level chosen by the user. We note that the underlying method of adding noise is general and can be applied to other question types (e.g., multiple-choice questions) in which the response set is countable (this excludes free-text responses). We have also developed a mathematical framework (not discussed in this paper), relying on differential-privacy [4, 5] to quantify the privacy loss, so that the cumulative privacy loss can be tracked and balanced across the user base, while ensuring sufficient accuracy of the aggregated response.

3.2 Preliminary Results

We built a prototype of our crowdsourcing survey platform, called Loki, including the front-end iPhone/iPad and Android apps, and a back-end database/ server built in Django framework (see project web-site: `http://loki.eng.unsw.edu.au/`). We trialled the system with 131 volunteers from our University, by launching a survey asking them to rate various lecturers, as shown in screen-shots in Fig. 1.

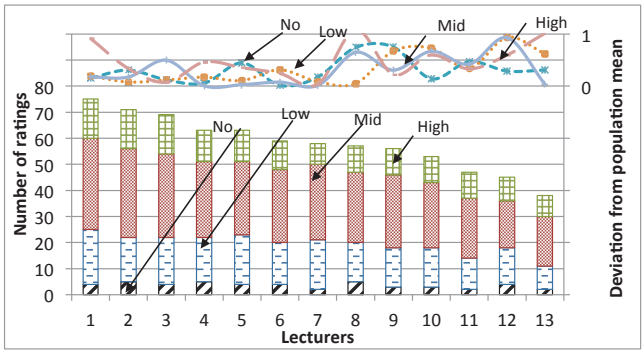

Figure 2: Variation in mean across the bins for various lecturers

We evaluated the success of our solution in two ways: (a) whether it meets user expectations of privacy, and (b) whether the results obtained are sufficiently accurate in spite of the obfuscation.

User perception: We talked to each participant immediately after the survey; most participants said that they liked the way we presented privacy options (four easy-to-understand levels), could easily see how the mechanism operated (i.e., the privacy level corresponds to the magnitude of Gaussian noise), and felt comfortable that their privacy was protected when they saw their noisy responses. Of the 131 students who took the survey, we found that 18 chose no privacy, 32 chose low privacy, 51 chose medium privacy, and 30 chose high privacy setting. We conjecture that the medium level was chosen by most users since they perceived it as a "safer" option than the extreme values.

Accuracy: We validated the accuracy of the responses using two methods – by comparing to the university-run rating mechanism (which uses a trusted third party), and by comparing the ratings across the various privacy bins in our system. Though the university does not publicly reveal the ratings of lecturers, we were able to verify that the small handful that were privately revealed to the authors of this paper corroborated well with those obtained from our system. For example, one author obtained an average score of 4.72 based on the noisy responses, only slightly higher than the average rating of 4.61 (out of 5) he has received from the (trusted third-party) university system over the past 3 years. To validate accuracy across the privacy bins, we plot in Fig. 2 the difference between the mean rating obtained from a given privacy bin and the overall mean rating. The figure also shows a histogram of the number of students rating each lecturer per privacy bin. Unsurprisingly, the accuracy of the estimated lecturer mean rating is lower when fewer users are assigned to the bin, particularly for higher privacy bins. This trade-off between accuracy and privacy is inevitable, but our study shows that even with a relatively small sample size the error is sufficiently small to make inferences.

4. CONCLUSION

We have demonstrated that current online platforms for crowdsourcing survey information allow user privacy to be easily compromised. We developed a new system that uses at-source obfuscation to empower users to control their privacy loss, while still yielding sufficiently accurate outcomes.

5. REFERENCES

[1] Amazon Mechanical Turk. `https://www.mturk.com/mturk/`.

[2] R. Chen, A. Reznichenko, P. Francis, and J. Gehrke. Towards Statistical Queries over Distributed Private User Data. In *NSDI*, 2012.

[3] CrowdFlower. `https://crowdflower.com/`.

[4] C. Dwork. Differential Privacy: A Survey of Results. In *TAMC*, 2008.

[5] C. Dwork, F. Mcsherry, K. Nissim, and A. Smith. Calibrating Noise to Sensitivity in Private Data Analysis. In *TCC*, 2006.

[6] P. Golle. Revisiting the Uniqueness of Simple Demographics in the US Population. In *ACM Workshop on Privacy in the Electronic Society*, 2006.

[7] Google Consumer Surveys. `http://www.google.com/insights/consumersurveys/home`.

[8] K. Ligett and A. Roth. Take it or Leave it: Running a Survey when Privacy Comes at a Cost. In *WINE*, 2012.

[9] A. Narayanan and V. Shmatikov. Robust De-anonymization of Large Sparse Datasets. In *IEEE Symosium on Security and privacy*, 2008.

[10] B. F. Ribeiro, W. Chen, G. Miklau, and D. F. Towsley. Analyzing Privacy in Enterprise Packet Trace Anonymization. In *NDSS*, 2008.

[11] L. Sweeney. Simple Demographics Often Identify People Uniquely. *Carnegie Mellon University, Laboratory for Internation Data Privacy, Pittsburgh, PA*, 2000.