A Differential Privacy Mechanism with Network Effects for Crowdsourcing Systems

Extended Abstract

Yuan Luo
Imperial College London
London, United Kingdom
y.luo@imperial.ac.uk

Nicholas R. Jennings
Imperial College London
London, United Kingdom
n.jennings@imperial.ac.uk

ABSTRACT

In crowdsourcing systems, it is important for the crowdsource campaign initiator to incentivize users to share their data to produce results of the desired computational accuracy. This problem becomes especially challenging when users are concerned about the privacy of their data. To overcome this challenge, existing work often aims to provide users with differential privacy guarantees to incentivize privacy-sensitive users to share their data. However, this work neglects the network effect that a user enjoys greater privacy protection when he aligns his participation behaviour with that of other users. To explore the network effect and provide a suitable differential privacy guarantee, we design PINE (Privacy Incentivization with Network Effects). PINE is a mechanism that maximizes the initiator's payoff while providing participating users with privacy protections.

KEYWORDS

Crowdsourcing systems; Incentive mechanism; Differential Privacy; Network Effects

ACM Reference Format:
Yuan Luo and Nicholas R. Jennings. 2018. A Differential Privacy Mechanism with Network Effects for Crowdsourcing Systems. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

A growing number of crowdsourcing applications aggregate users' personal data for decision-making purposes. Examples include movie rating systems such as Netflix and OK.com and traffic monitoring systems such as INRIX and WAZE. Users who share their personal data can certainly derive benefits from such information aggregation. Furthermore, the benefit a user derives typically increases with the number of participants in the system. Hence, crowdsourcing applications need to recruit and maintain large numbers of users to guarantee the accuracy of their results.

However, the increasing use of individuals' data has been accompanied by growing concerns about their privacy. For example, Narayanan and Shmatikov [16] were able to recover private information such as movie viewing history and political preferences from Netflix's published movie ratings. Ma et al. [15] re-identified individuals' journeys from anonymized location data published by the ShangHai Grid system. This risk of exposing private information deters privacy-sensitive users [11]; therefore appropriate means for encouraging participation are central to the design of crowdsourcing systems [1, 12]. In particular, it is important to develop methods for incentivizing such users to share their data with the crowdsource campaign initiator ("initiator" hereafter) who will process and publish the computational result[1]. To this end, recent work has started to explore the question by relating differential privacy guarantees to questions of incentives in mechanism design (*e.g.*, [2, 9, 18]).

In more detail, a guarantee of *differential privacy* provides a degree of privacy protection to participating users when an initiator computes a population's data and publishes the computational result [4]. Specifically, a differentially private mechanism involves publishing a noisy version of the computational result, where the noise level corresponds to a differential privacy guarantee (more noise means greater privacy). At this time, several real-world applications have or will adopt differential privacy mechanisms to collect users' data. For example, the Chrome web browser uses RAPPOR [6], a differentially private mechanism, to track the distribution of users' browser configuration behaviour. Apple deploys local differential privacy in its iOS system for collecting their users' data [19].

To incentivize users to share their personal data, one of the simplest ways is to pay them for using it. However, determining the correct price is challenging: low payments may not attract sufficient participants, causing insufficient data for an accurate computation, while high payments may make the system uneconomic. Existing incentive mechanisms with differential privacy guarantees aim to elicit each user's valuation for the various levels of guarantee and then tailor payments based on the elicited values to incentivize sufficient participants to produce accurate results and pay the minimum possible cost [9, 11, 17].

In contrast, in crowdsourcing applications, an initiator usually recruits thousands of users [10] so it would be prohibitively expensive to negotiate compensation individually with each of them. Furthermore, a user may be unable to state a quantitative privacy loss but is usually able to say whether or not he is willing to share his data [7]. Finally, it is more difficult to identify a user's identity in a large population than in a smaller one. Thus, the more users contributing their personal data, the more difficult it is to identify a specific user's identity and therefore the more users should be willing to contribute their data. This *network effect* [5, 14] means a

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

[1] We refer to the initiator who collects and processes the data as "she" and individual users as "he".

user's participation decision may be correlated with that of other users. In practice, openPDS/safeAnswers [3], a Personal Data Store, uses this network effect to hide individuals' personal data in a population so better preserving users' privacy. Unfortunately, most existing work does not consider this, which means such mechanisms are likely to be overpaying to achieve a given level of privacy. Ghosh and Ligett [8] do study the impact of the network effect on users' participation decisions. However, they only focus on interactions among different users without considering the impact of the initiator's payment on the users' strategies. Thus, they cannot help the initiator to achieve the target computational accuracy with minimum cost.

Designing a mechanism that exploits network effects and provides a differential privacy guarantee is non-trivial and has not been done before. It is complex as the privacy guarantee level couples with an initiator's payment strategy and all potential users' strategies. Moreover, users' participation decisions also affect the initiator's computational accuracy. The initiator's computational accuracy is here defined as the difference between the initiator's published result and the true computation on the full set of data from the entire population [13]. Hence, one cannot design a mechanism by simply combining existing techniques.

Against this background, this paper outlines a novel mechanism, which we call "PINE" (Privacy Incentivization with Network Effects). PINE maximizes the initiator's payoff, defined as the difference between the payments made to the privacy-sensitive users and the computational accuracy of the population's data[2].

2 MODEL AND OBSERVATIONS

We consider a setting where an initiator wants to learn and publish some statistic about the population. Learning the average rating for a movie and knowing the average traffic speed for a road are two examples. The interactions between the initiator and the users are as follows. The initiator first announces the noise level, the number of potential users, the target error, and a payment that is calculated based on PINE. Then, each user makes his participation decision based on his privacy sensitivity, his prior belief about the privacy sensitivities of the population, and the payment announced by the initiator. If a user decides to join the survey, he report his private data to the initiator. After collecting all participants' reporting data, the initiator infers the statistic of the population by computing on the collected dataset and gets the result. The initiator publishes the result by adding the noise that is determined by PINE.

Our results give us the following observations:

(a) When the noise added to the published result is low, only the users with no privacy sensitivity will join and the optimal payment is zero. Intuitively, it is easier to infer the true computational result from the published computational result in the low noise setting than in the high noise one. With a decrease in the added noise level, a user's privacy revealing risk increases. Thus, the initiator needs to pay a large amount of money to incentivize users to participate. However, when the noise level is too low, the benefit regarding the accuracy brought by incentivizing users' participation is smaller than the payment made to the users. In such cases, the initiator is willing to give up the accuracy in return for low cost.

(b) With a high noise level, the optimal payment decreases with the number of potential users. When the noise level is high, a user's privacy revealing risk decreases and the initiator can incentivize users to participate without paying too much. When the number of potential users increases, the number of participants also increases given the same noise level and payment. Due to the network effect, the more users participating in this survey, the more difficult it is to identify whether a given user's private data is included in the computation dataset. This means more users are willing to participate and the initiator can decrease the payment.

(c) With a median noise level, the optimal payment first increases and then decreases with the number of potential users. The intuition is as follows. When the number of potential users is not so large, the network effect's benefit is less obvious as the number of participants is small. Under a median noise level, the initiator has to increase the payment to incentivize as many participants as possible from the pool of potential users to guarantee the accuracy. When the potential users' number is large enough, the network effect's impact allows the initiator to decrease the payment without jeopardizing the computational accuracy.

3 CONCLUSION AND FUTURE WORK

We outline a new model that can be used to underpin the design of incentives for crowdsourcing applications. Our mechanism provides a theoretically sound mechanism that incentivizes sufficient numbers of users to share their data so that a given level of computational accuracy can be guaranteed with minimum payment.

There are several possible directions to extend this work. First, we assume that there is no correlation between a user's evaluation for his privacy loss and his private data. While this is justifiable in several online settings such as collecting data regarding watching movie behaviour, mobile applications for traffic monitoring, or polls in online communities, there are also several other settings such as computations on medical data where this assumption is unlikely to hold. Exploring these settings where a user's evaluation for his privacy loss depends on his private data will lead to a different formulation and result. Considering an untrustworthy initiator scenario where the initiator will collude with any third party/government body for her own personal benefit and announcing false information to users is another possible direction. Although letting individual users outweigh their benefits by contributing their data and add corresponding amount of noise prior to the release of the data to the initiator could prevent the collusion between the initiator and any third party, our numerical study show that the performance is worse than PINE. How to design a mechanism to prevent this collusion without jeopardizing the benefits of the initiator and users deserve careful study in our future work.

[2] Although we demonstrate PINE in crowdsourcing systems, it can also be used in other applications that require users to send information to data aggregators performing monitoring or control tasks (*e.g.*, cloud-computing systems, smart grid systems, and database-assisted TV white space networks).

REFERENCES

[1] Praphul Chandra, Sujit Gujar, and Yadati Narahari. 2017. Referral-Embedded Provision Point Mechanisms for Crowdfunding of Public Projects. In *Proceedings*

of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS). 642–650.

[2] Yiling Chen, Stephen Chong, Ian A Kash, Tal Moran, and Salil Vadhan. 2016. Truthful mechanisms for agents that value privacy. *ACM Transactions on Economics and Computation* 4, 3 (2016), 13.

[3] Yves-Alexandre de Montjoye, Erez Shmueli, Samuel S Wang, and Alex Sandy Pentland. 2014. openpds: Protecting the privacy of metadata through safeanswers. *PloS one* 9, 7 (2014), e98790.

[4] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*. Springer, 265–284.

[5] David Easley and Jon Kleinberg. 2010. *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge University Press.

[6] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *ACM SIGSAC Conference on Computer and Communications Security*. 1054–1067.

[7] Arpita Ghosh and Robert Kleinberg. 2014. Optimal contest design for simple agents. In *ACM conference on Economics and computation (EC)*. 913–930.

[8] Arpita Ghosh and Katrina Ligett. 2013. Privacy and coordination: Computing on databases with endogenous participation. In *ACM conference on Electronic Commerce (EC)*. 543–560.

[9] Arpita Ghosh and Aaron Roth. 2015. Selling privacy at auction. *Games and Economic Behavior* 91 (2015), 334–346.

[10] Shweta Jain, Ganesh Ghalme, Satyanath Bhat, Sujit Gujar, and Y Narahari. 2016. A deterministic MAB mechanism for crowdsourcing with logarithmic regret and immediate payments. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 86–94.

[11] Haiming Jin, Lu Su, Houping Xiao, and Klara Nahrstedt. 2016. Inception: Incentivizing privacy-preserving data aggregation for mobile crowd sensing systems. In *ACM Mobile Ad Hoc Networking and Computing (MobiHoc)*, Vol. 16. 341–350.

[12] Omer Lev, Maria Polukarov, Yoram Bachrach, and Jeffrey S Rosenschein. 2013. Mergers and collusion in all-pay auctions and crowdsourcing contests. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 675–682.

[13] Yang Liu and Yiling Chen. 2017. Sequential Peer Prediction: Learning to Elicit Effort using Posted Prices.. In *AAAI Conference on Artificial Intelligence (AAAI)*. 607–613.

[14] Yuan Luo, Lin Gao, and Jianwei Huang. 2016. An integrated spectrum and information market for green cognitive communications. *IEEE Journal on Selected Areas in Communications* 34, 12 (2016), 3326–3338.

[15] Chris YT Ma, David KY Yau, Nung Kwan Yip, and Nageswara SV Rao. 2013. Privacy vulnerability of published anonymous mobility traces. *IEEE/ACM Transactions on Networking* 21, 3 (2013), 720–733.

[16] Arvind Narayanan and Vitaly Shmatikov. 2008. Robust de-anonymization of large sparse datasets. In *IEEE Symposium on Security and Privacy*. 111–125.

[17] Aaron Roth and Grant Schoenebeck. 2012. Conducting truthful surveys, cheaply. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*. 826–843.

[18] Weina Wang, Lei Ying, and Junshan Zhang. 2016. The Value of Privacy: Strategic Data Subjects, Incentive Mechanisms and Fundamental Limits. In *ACM International Conference on Measurement and Modeling of Computer Science (SIGMETRICS)*. 249–260.

[19] Apple Worldwide Developers Conference (WWDC). 2016. Engineering Privacy for Your Users. Video. (2016). Retrieved June 2016 from https://developer.apple.com/videos/play/wwdc2016/709