Poster: Leveraging Prior Knowledge Asymmetries in the Design of IoT Privacy-Preserving Mechanisms

Nazanin Takbiri, Virat Shejwalkar, Amir Houmansadr, Dennis Goeckel, Hossein Pishro-Nik

***Abstract*—The Internet of Things (IoT) promises to improve user utility by tuning applications to user behavior, but the revealing of the characteristics of a user's behavior presents a significant privacy risk. Recently, the technique of remapping has been introduced in the privacy literature. Remapping exploits asymmetries in knowledge and/or sophistication between the intended application and the adversary; in particular, the user publishes a more accurate version of her data than they might have otherwise because a sophisticated adversary could obtain that accurate version anyway. Here, after introducing the system model, we first demonstrate the mechanism behind the remapping technique. Next, we characterize the loss in privacy when the user lacks knowledge of the accuracy of the adversary's statistical model; this loss in privacy occurs both because the adversary obtains a more accurate view of the user data than expected and because the adversary can exploit the remapping to improve their statistical model more than would have been possible when remapping is not employed. Finally, we introduce a random remapping approach as a countermeasure; in particular, for a given utility, the random remapping approach makes it difficult for the adversary to improve their statistical model.**

I. Introduction

Emerging technologies such as IoT promise to revolutionize users' lives by adapting to each user's specific needs and habits as gleaned from their data traces. However, this necessitates that the data of an immense number of users are interconnected, thus posing intrinsic threats to user privacy and leaving sensitive information vulnerable. There has been significant work on privacy-preserving mechanisms (PPMs). Obfuscation is one of the main PPMs which enhances privacy by using misleading, false, or ambiguous information. Note that obfuscation degrades system utility while enhancing privacy [1].

Recently, a new method termed "remapping", which is similar to posterior data processing in database systems, has emerged as an effective method to exploit asymmetries in the privacy problem to substantially improve system utility without a corresponding loss in privacy for a PPM that employs obfuscation [2]. In particular, remapping is employed in scenarios where a friend (e.g. an IoT application) exists that does not have prior statistical information about user behavior, whereas the adversary in the environment has perfect statistical information about the user's behavior. This may occur, for example, when each intended recipient is either naive or only looking at a single datum or small set of data from the user, whereas the adversary is sophisticated and has access to data across a large time period from the user.

Authors are with University of Massachusetts, Amherst, MA, 01003 USA e-mail: (ntakbiri@umass.edu, vshejwalkar@cs.umass.edu).

In such a case, the adversary can use their statistical advantage to obtain a better estimate of the user's data than the friend. Remapping recognizes this fact and reveals a more accurate version of the data that the adversary would have been able to obtain anyway using her statistical advantage, so there is no loss in privacy, but which will improve the accuracy for the user's friend. Hence, by recognizing this asymmetry, utility has been improved at no loss in privacy versus a scheme that did not do remapping; a simple example of the mechanism is demonstrated in Section II.

Contributions: Here, we explore important aspects of remapping that have not been considered. As acknowledged briefly in [2], a risk of remapping is that it relies critically on accurate knowledge of the adversary's statistical model. In particular, if the adversary does not have accurate statistical information and the user employs remapping, we discuss that privacy is compromised in two separate ways: (i) the adversary obtains a more accurate version of the data than they would have had without remapping; and (ii) the adversary is able to improve their statistical knowledge of the users' data beyond what they would have been able to do without remapping. Interestingly, we will see that the second type of leakage is increased if the obfuscation noise is increased. We provide the first analysis of the loss of privacy due to each of these factors. After analyzing the loss in privacy under standard remapping [2], we next turn to countermeasures. We introduce a random remapping algorithm, where data points are independently remapped with some probability. For a given utility for the intended recipient, this approach greatly complicates model improvement at the adversary versus deterministic remapping approaches, thus greatly improving the privacy-utility trade-off.

II. Framework

As shown in Figure 1, there exist an "intended" friend (e.g., an IoT application) who lacks prior statistical knowledge about the user behavior and a "sophisticated" adversary who has knowledge about the prior behavior of the user (π_{Adv}). The adversary observes the noisy reported data Y and uses it to find the estimate $\widetilde{X}_{Adv}$, which denotes the estimate of the adversary given her observed data (Y) and her knowledge of the prior about the user (π_{Adv}), as $\widetilde{X}_{Adv} = \mathbb{E}[X|Y, \pi_{Adv}]$. As a result, there exist asymmetries in knowledge and/or sophistication between the intended friend and the adversary. The remapping technique introduced by [2] exploits these asymmetries to

publish a more accurate version of data that the sophisticated adversary would have been able to obtain anyway.

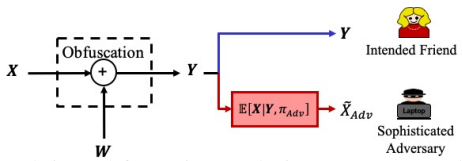

Fig. 1: Applying obfuscation technique to a user' data points.

As shown in Figure 2, each reported data is remapped into the best possible data point according to the perfect prior information of the adversary.

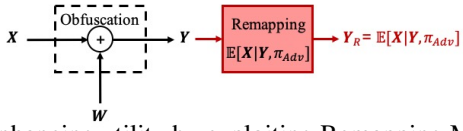

Fig. 2: Enhancing utility by exploiting Remapping Mechanism

Case* 1*: Perfect Knowledge of the Adversary: In this section, we assume the adversary knows the exact value of the prior distribution of the user (π_{Adv}). Without remapping, the user's intended friend, which is oblivious to the prior knowledge of the user, observes only the noisy data (Y). In comparison to the user's friend, the sophisticated adversary obtains $\widetilde{X}_{Adv} = \mathbb{E}[X|Y,\pi_{Adv}]$. However, when the remapping technique is employed, both the adversary and the user's friend observe the same reported data, $\widetilde{X}_{Adv} = \widetilde{X}_{Friend} = Y_R = \mathbb{E}[X|Y,\pi_{Adv}]$. Since the intended friend is oblivious to the prior statistical knowledge about the user behavior, the MSE of the adversary is always smaller than or equal to the MSE of the friend . We can conclude the remapping technique provides the best utility among techniques satisfying the same level of privacy in the case of perfect knowledge of the adversary.

Case* 2*: Imperfect Knowledge of the Adversary: Here, we assume the adversary has a noisy version of the prior information, as might be obtained from a learning set of limited length. Now, if remapping is not employed, the user's intended friend observes the reported data (Y). In contrast, the sophisticated adversary uses both Y and $\widehat{\pi}_{Adv}$ to improve her knowledge not only about the true data (X) but also about the distribution of the true data ($\widehat{\pi}_{Adv}$). In addition, if the remapping technique is employed, our friend observes the remapped data. However, the adversary observes not only Y_R, but also $\widehat{\pi}_{Adv}$, and uses both of them to estimate $\widetilde{\pi}_{Adv}$ and $\widetilde{X}_{Adv}$. We can conclude that increasing the obfuscation noise somewhat surprisingly increases the leakage about the distribution of the true data (π_{Adv}) when remapping is employed. Note that $Y_R = \mathbb{E}[X|Y,\widehat{\pi}_{Adv}]$ depends on two parameters: 1) $\widehat{\pi}_{Adv}$ and 2) $Y = X+W$; thus, if we increase the obfuscation noise , Y_R relies less on Y and more on $\widehat{\pi}_{Adv}$. Now in the extreme case, where the amount of noise goes to infinity, the observed data (Y) is useless and, as a result, $Y_R = \mathbb{E}[X|Y,\widehat{\pi}_{Adv}] = \pi_{Adv}$. Hence, remapping leaks complete information about the statistical model.

III. Randomized Remapping

As derived in Section II, the remapping technique can leak a significant amount of information about the distribution of the true data (π_{Adv}) if the adversary does not have the perfect prior for the behavior of the user. Here, we introduce a new technique called randomized remapping to improve privacy. This technique provides a trade-off between the leakage of the distribution of the true data (π_{Adv}) and the leakage of true data (X). In the randomized remapping, we have an unfair coin where the probability of a head is equal to p_H. For each data point, we toss the coin and if a head is observed, the remapped data (Y_R) is released, and if a tail is observed, the noisy version of the data (Y) is released. Thus, the randomized remapping provides a much better trade-off compared to standard remapping. The value of p_H is a design parameter, so, based on the application requirements and privacy requirements, the appropriate amount of value of p_H should be chosen.

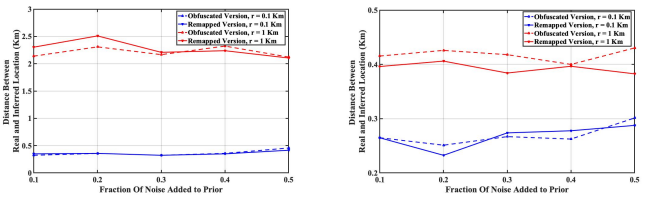

a) High entropy prior. b) Low entropy prior.

Fig. 3: Privacy leakage demonstration due to remapping: With significant utility improvement due to remapping, privacy leakage is also significant. In both (a), (b), for smaller obfuscation radius, remapping does not improve utility significantly (<30%), so we do not see clear leakage. For larger obfuscation radius, 1Km, both utility improvement (>80%) and privacy leakage are significant for the low entropy case (b), but not for the high entropy case (a).

IV. Sketch of the Results

To demonstrate the extent of leakage in real world scenarios, experiments on the Gowalla data set have been performed. The data contains users' check-ins with details such as geo-coordinates of locations, time, and location-ids. Geo-indistinguishability is the realization of differential privacy in the location privacy domain. As shown in Figure 3, we experiment with highly distributed versus highly concentrated priors. With highly concentrated priors, we observe significant leakage due to remapping i.e., adversarial inference using a noisy prior on an obfuscated location is very different from that using a remapped location. However, for a highly distributed prior, there is not significant leakage due to remapping, since, as expected, remapping does not improve utility very much.

References

[1] A. E. Miguel, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, "Geo-indistinguishability: Differential privacy for location-based systems," in *Proceedings of the 2013 ACM Conference on Computer and Communications Security*. ACM, 2013.

[2] K. Chatzikokolakis, E. ElSalamouny, and C. Palamidessi, "Efficient utility improvement for location privacy," *Proceedings on Privacy Enhancing Technologie*, vol. 2017, no. 4, pp. 210–231, 2017.