Privacy-Utility Tradeoff under Statistical Uncertainty

Ali Makhdoumi
MIT, Cambridge, MA 02139
Email: makhdoum@mit.edu

Nadia Fawaz
Technicolor, Palo Alto, CA 94301
Email: nadia.fawaz@technicolor.com

***Abstract*—We focus on the privacy-accuracy tradeoff encountered by a user who wishes to release some data to an analyst, that is correlated with his private data, in the hope of receiving some utility. We rely on a general statistical inference framework, under which data is distorted before its release, according to a probabilistic privacy mechanism designed under utility constraints. Using recent results on *maximal correlation* and *hypercontractivity of Markov processes*, we first propose novel techniques to design utility-aware privacy mechanisms against inference attacks, when only partial statistical knowledge of the prior distribution linking private data and data to be released is available. We then propose optimal privacy mechanisms in the class of additive noise mechanisms, for both continuous and discrete released data, whose design requires only knowledge of second-order moments of the data to be released. We then turn our attention to multi-agent systems, where multiple data releases occur, and use tensorization results of maximal correlation to analyze how privacy guarantees compose after *collusion* or *composition*. Finally, we show the relationship between different existing privacy metrics, in particular divergence privacy, and differential privacy.**

I. Introduction

A. Motivation

In the era of Big Data, the collection and mining of user data has become a fast growing and common practice by a large number of private and public institutions. These include for instance tech companies, who exploit user data to offer personalized services to their customers, government agencies, who rely on data to address a variety of challenges, e.g. national security, national health, budget and fund allocation, or medical institutions, who analyze data to discover the origins and potential cures to diseases. In some cases, the collection, the analysis, or the sharing of a user's data with third parties is performed without the user's consent or awareness. In other cases, data is released voluntarily by a user to a specific entity, in order to get a service in return, e.g. product ratings released to get recommendations. In either case, privacy risks arise as some of the collected data may be deemed sensitive by the user, e.g. political convictions, health status, income level, or may seem harmless at first sight, e.g. product ratings, yet lead to the inference of more sensitive data with which it is correlated. The latter threat refers to an inference attack— inferring private data by exploiting its correlation with publicly released data— and is the main focus of this paper.

We consider the setting in [1], where a user has two kinds of data that are correlated: some data that he would like to remain private, and some non-private data that he is willing to release to an analyst and from which he will derive some utility. The analyst is a legitimate receiver of the released data, which he will use to provide utility to the user, but can also illegitimately exploit to infer the user's private data. This creates a tension between privacy and utility requirements. To reduce the inference threat while maintaining utility, data is distorted before its release, according to a privacy-preserving mechanism designed under utility constraint. We model the privacy-utility tradeoff according to the general framework for privacy against statistical inference introduced in [1]. The optimal privacy mapping is the solution of a convex optimization problem, which minimizes the information leakage— modeled under the log-loss by the mutual information between private data and released data— subject to a utility constraint— average distortion between original and distorted data [1], [2].

B. Challenges and Contributions

Our contributions address the following challenges in the design of utility-aware privacy mechanisms.

Design under partial statistical knowledge of the prior: Finding the optimal privacy mapping as the solution to the optimization problem in [1], [2] relies on the fundamental assumption that the prior joint distribution that links private data and data to be released is known and can be fed as an input to the optimization. In practice, the true prior distribution may not be known, but rather some prior statistics may be estimated from a set of sample data that can be observed. For example, the prior joint distribution could be estimated from a set of users who do not have privacy concerns and publicly release both their private and non-private data. Alternatively when the private data cannot be observed, the marginal distribution of the data to be released, or simply its second order statistics, may be estimated from a set of users who only release their non-private data. The statistics estimated from this set of samples can then be used to design the privacy mechanism that will be applied to new users, who are concerned about their privacy. In practice, there may also exist a mismatch between the estimated prior statistics and the true prior statistics, due for example to a small number of observable samples, or to the incompleteness of samples.

Our first contribution consists in proposing methods to design utility-aware privacy mechanisms when only partial statistical knowledge of the prior is available. More precisely, using recent information theoretic results on *Maximal correlation* and *hypercontractivity of Markov process*, we first provide an upper-bound on the information leakage, that decouples the intrinsic dependencies between the private data and the non-private data, from the designed dependencies between the non-private data and the distorted released data. A fundamental property of this method is that it allows the design of privacy mechanisms with only knowledge of the prior marginal of the non-private data— or even without any knowledge of the prior distributions— instead of requiring full knowledge of the joint prior of the private data and non-private data. As a second contribution, we provide privacy mechanisms that are optimal in the class of additive noise mechanisms, namely the Gaussian mechanism for continuous non-private data, and the discrete noise mechanism for discrete data.

Privacy of multi-agent systems— Collusion and Composition: New challenges in the design of privacy mechanisms arise when multiple data releases to one or several agents occur. We address the question of how privacy guarantees compose under multiple releases. The rules of composition of privacy guarantees help in addressing the issue of colluding agents, who share together data that was released to them individually in order to improve their inference of private data. Composition rules also help in the design of privacy mechanisms by allowing to break the joint design of a privacy mechanism for multiple pieces of data into several simpler design problems for individual pieces of data. We first focus on the case where the releases are related to the same private data, and then extend the analysis to the case where the releases are related to different but correlated pieces of private data. Using tensorization results of maximal correlation, we reason about the cumulative privacy guarantees of the union of the releases, and show that under some conditions, privacy guarantees compose according to simple rules.

Comparison of privacy metrics: We compare and show the relationship between our privacy metric and different existing privacy metrics, in particular divergence privacy, differential privacy, and information privacy. We provide examples on the differences in the privacy-accuracy tradeoffs achieved under these different notions. Finally, we show that although differential privacy is an elegant and well-studied privacy metric, it does not guarantee a small probability of error in inferring the private data given the released distorted data, on the contrary to divergence privacy.

C. *Related work*

In the database and cryptography literatures from which differential privacy arose (e.g. [3], [4], [5]), the focus has been algorithmic; in particular, researchers have used differential privacy to design privacy preserving mechanisms for inference algorithms, transporting, and querying data. More recent works [6], [7] focused on the relation of differential privacy with statistical inference. Other frameworks similar to differential privacy exist such as the Pufferfish framework [8], which however does not focus on utility preservation. Many approaches rely on information-theoretic techniques to model and analyze the privacy-accuracy tradeoff, such as [9], [10], [11], [12], [1], [13], [2]. Information-theoretic models [9], [10], [11] focus mainly on collective privacy for all or subsets of the entries of a database, and provide fundamental and asymptotic results on the rate-distortion-equivocation region as the number of data samples grows arbitrarily large. In contrast, the framework studied in [1], models non-asymptotic privacy guarantees in terms of the inference cost gain that an adversary achieves by observing the released output. In this work, we follow the framework in [1], and use the log-loss cost to model the inference threat as the mutual information between private data and released data, as in [2], [13].

Composition of privacy guarantees under differential privacy has been studied, e.g. [5], [14]. The focus in this paper is on privacy of multi-agent systems under an information-theoretic privacy metric.

The organization of paper is as follows. In Section II, we give the problem formulation. In Section III, we propose several methods to design the privacy mapping under statistical uncertainty. In section IV, we study the design of optimal mapping in the class of additive noise. In Section V, we consider privacy of multi-agent systems, and finally compare different privacy metrics in Section VI. Due to lack of space, we omit proofs, and refer the reader to the full version of the paper [15].

II. Privacy-Accuracy tradeoff

In this section, we first describe the setting, and the privacy and accuracy metrics. Then, we characterize the privacy-accuracy tradeoff in terms of an optimization problem. We give a lower bound on the probability of error in inferring the private data from the released data, and conclude with the challenges in the design of privacy mappings with partial knowledge of the prior distribution of private and non-private data.

A. Divergence Privacy and Accuracy

We consider a setting where the user has some private data, represented by the random variable $S \in \mathcal{S}$, which is correlated with some non-private data $X \in \mathcal{X}$. The correlation between S and X is captured by the joint distribution $P_{S,X}$, while P_X denotes the marginal distribution of X. X can be either discrete, or continuous, in which case we assume it has a probability density function. To reduce the inference threat on S that would arise from the observation of X, rather than releasing X, the user releases a distorted version of X, denoted $Y \in \mathcal{Y}$. The distorted data Y is obtained by passing X through a conditional distribution $P_{Y|X}$, called the privacy mapping. Throughout the paper, we assume $S \to X \to Y$ form a Markov chain. Therefore, once we define $P_{Y|X}$, we have the joint distributions $P_{S,X,Y} = P_{Y|X}P_{S,X}$, and $P_{S,Y}$.

We first define our privacy metric.

Definition 1. *Assume $S \to X \to Y$. A conditional distribution $P_{Y|X}$ is called* ϵ-divergence private *if the distributions $P_{S,Y}$ and P_Y resulting from the joint distribution $P_{S,X,Y} = P_{Y|X}P_{S,X}$ satisfy*

$$D(P_{S,Y}||P_SP_Y) = \epsilon H(S), \tag{1}$$

where $D(P_{S,Y}||P_SP_Y) \triangleq \mathbb{E}\left[\log \frac{P(S|Y)}{P(S)}\right] \triangleq I(S;Y)$ and $\epsilon \in [0,1]$ is called the leakage factor*, and the mutual information $I(S;Y)$ represents the* information leakage. *We say a mechanism has* full privacy *if $\epsilon = 0$.*

The extreme case of full privacy $\epsilon = 0$ is equivalent to the statistical independence of S and Y. In the other extreme case $\epsilon = 1$, no uncertainty is left on S from the observation of Y.

We define accuracy as follows:

Definition 2. *Let $d : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}^+$ be a distortion measure. A conditional distribution $P_{Y|X}$ is called D-accurate if $\mathbb{E}[d(X,Y)] \leq D$.*

Definition 3. *A privacy mapping $P_{Y|X}$ is called $(\epsilon, D)-$* divergence-distortion private *if its leakage factor and expected distortion are not greater than ϵ and D, respectively.*

There exists a tradeoff between the leakage factor ϵ and the distortion level D achieved by a privacy mapping. Given the joint distribution $P_{S,X}$, the optimal privacy mapping is defined as the conditional distribution achieving the minimum objective of

$$\begin{aligned} &\min_{P_{Y|X}} I(S;Y) \\ &\text{s.t. } \mathbb{E}[d(X,Y)] \leq D \end{aligned} \tag{2}$$

The optimization problem (2) was introduced in [2], [1], and shown to be convex. Next, we give an example of the optimization given in (2) and its solution.

Example 1. Assume that S has a Bernoulli distribution Bern($\frac{1}{2}$), and that X is the result of passing S through a binary symmetric channel BSC(p) with transition probability $p \leq \frac{1}{2}$. Under the Hamming distance, the distortion constraint becomes $P[X \neq Y] \leq D$. It can be shown that the minimum objective of Optimization (2) is $I(S;Y) = 1 - h(p * D)$, where $p * D = p(1-D) + (1-p)D$, and $h(.)$ denotes the entropy of a Bernoulli random variable. The optimal privacy mapping generates Y by passing X into a channel BSC(D). Note that full privacy is only possible in two trivial cases: either $p = \frac{1}{2}$, i.e. the non-private X is independent from the private S, and there is no privacy problem to begin with; or $D = \frac{1}{2}$, i.e. the released Y is independent from the original non-private X, in which case no utility is preserved in the released data Y.

B. Inference Defeat through Divergence Privacy

One natural and related question is whether a privacy mapping designed to minimize information leakage by solving Optimization (2), also provides guarantees on the probability of error in inferring S from the observation of Y. In the following proposition, we provide a lower bound on the error probability in inferring S from Y, based on Fano's inequality. It should be noted that this bound holds for any inference algorithm used by the adversary.

Proposition 1. *Assume $|\mathcal{S}| > 2$ and $I(S;Y) \leq \epsilon H(S)$. Let $\hat{S}$ be an estimator of S based on the observation Y, possibly randomized. Then*

$$P[\hat{S}(Y) \neq S] \geq \frac{(1-\epsilon)H(S) - 1}{\log(|\mathcal{S}| - 1)}. \tag{3}$$

For $|\mathcal{S}| = 2$, we have $h(P_e) \geq (1-\epsilon)H(S)$.

III. Achieving Privacy

A. Partial Knowledge of the statistics of S and X

In practice, we may not have access to the exact joint probability distribution $P_{S,X}$, but rather have some partial knowledge of the joint statistics of S and X. Consequently, finding the exact solution of Optimization (2) may not be possible. This raises the question of the design of privacy mappings under partial statistical knowledge of the prior. We consider the following two cases.

Case 1– Know marginal P_X, unknown joint $P_{S,X}$: The optimal privacy mapping under this limited statistical knowledge is the conditional distribution minimizing the objective

$$\min_{P_{Y|X}} \max_{P_{S|X}} I(S;Y),$$

under the same constraints as in (2).

Case 2– No information about $P_{S,X}$: In the absence of any knowledge on the statistics of S and/or X, the optimal privacy mapping is obtained by minimizing

$$\min_{P_{Y|X}} \max_{P_{S,X}} I(S;Y),$$

under the same constraints as in (2).

In Section III-B, we introduce statistical inference techniques based on *maximal correlation*, and prove a separability result on the privacy guarantees. Namely, we provide an upper bound the information leakage in terms of $I(S;X)$ times a maximal correlation factor determined by the conditional distribution $P_{Y|X}$. This result allows us to design privacy preserving mappings without full knowledge of the joint distribution $P_{S,X}$.

B. Privacy via Maximal correlation

As a preliminary, we give the following definition.

Definition 4. *([16]) Given a joint distribution $P_{X,Y}$, we define $S^*(X;Y) = \sup_{Q_X \neq P_X} \frac{D(Q_Y||P_Y)}{D(Q_X||P_X)}$, where Q_Y is the marginal probability of Y resulting from the joint distribution $P_{Y|X}Q_X$.*

Note that $S^*(X;Y) \leq 1$ by the data processing inequality for divergence (see [16]). We have

Theorem 1. *([17]) If $S \to X \to Y$ form a Markov chain, then $I(S;Y) \leq S^*(X;Y)I(S;X)$, and the bound is tight as we vary S. In other words, assuming $I(S;X) \neq 0$, we have*

$$\sup_{S: S \to X \to Y} \frac{I(S;Y)}{I(S;X)} = S^*(X;Y). \tag{4}$$

Theorem 1 decouples the dependency of Y and S into two terms, one relating S and X, and one relating X and Y. Thus, one can upper bound the leakage even without knowing $P_{S,X}$. The application of this result in our problem is the following: Assume we are in a regime that $P_{S,X}$ is not known. If we do not use any privacy mapping, then we have $\frac{I(S;X)}{H(S)}-$ divergence privacy. If we design the privacy mapping with $S^*(X;Y) = s$ for some $s \in [0,1]$, then we obtain $s\frac{I(S;X)}{H(S)}-$ divergence privacy. Since we do not know the distribution on (S,X), irrespective to the joint distribution of (S,X), we can guarantee $s-$ divergence privacy. Therefore, the problem becomes to find $P_{Y|X}$, minimizing the following objective function.

$$\begin{aligned} &\min_{P_{Y|X}} \max_{P_X} S^*(X;Y) \\ &\mathbb{E}[d(X,Y)] \leq D. \end{aligned} \tag{5}$$

In order to study this optimization problem in more details, we need to review some results on maximal correlation. Maximal correlation is a measure of correlation between two random variables with applications both in information theory and computer science. We recall its definition, and give its relation to $S^*(X;Y)$.

Definition 5. *([18]) Given two random variables X and Y, the maximal correlation of (X,Y) is*

$$\rho_m(X;Y) = \max_{(f(X),g(Y)) \in \mathcal{T}} \mathbb{E}[f(X)g(Y)], \tag{6}$$

where $\mathcal{T}$ is the collection of pairs of real-valued random variables $f(X)$ and $g(Y)$ such that $\mathbb{E}[f(X)] = \mathbb{E}[g(Y)] = 0$ and $\mathbb{E}[f(X)^2] = \mathbb{E}[g(Y)^2] = 1$.

This measure was first introduced by Hirschfeld [18] and then studied by Rényi [19] and Ahlswede [16]. Recently, [17], [20] studied the maximal correlation and gave a geometric interpretation of this quantity. The following is a result of [16].

$$\max_{P_X} \rho_m^2(X;Y) = \max_{P_X} S^*(X;Y). \tag{7}$$

Substituting (7) in (5), the privacy preserving mapping is the solution of

$$\begin{aligned} &\min_{P_{Y|X}} \max_{P_X} \rho_m^2(X;Y) \\ &\mathbb{E}[d(X;Y)] \leq D. \end{aligned} \tag{8}$$

It is shown in [21] that maximal correlation, $\rho_m(X;Y)$ is characterized by the second largest singular value of the matrix Q with entries $Q_{x,y} = \frac{P(x,y)}{\sqrt{P(x)P(y)}}$. This optimization can be solved by *power iteration* algorithm for finding singular values of a matrix. Two quantities $S^*(X;Y)$ and $\rho_m^2(X;Y)$ have a close relation with each other. Two sufficient conditions under which $S^*(X;Y) = \rho_m^2(X;Y)$ are given in Theorem 7 of [16]. In particular, if the supremum in (4) is not achievable, then $\rho_m^2(X;Y) = S^*(X;Y)$. Next, we give an example of such case.

Example 2. Let $X \sim \text{Bern}(\frac{1}{2})$ and $Y = X + N$ (mode 2), where $N \sim \text{Bern}(D)$ and $X \perp\!\!\!\perp N$. It is shown in [20] that, $S^*(X;Y) = \rho_m^2(X;Y) = (1-2D)^2$. Using this bound where $S \sim \text{Bern}(\frac{1}{2})$, $X = S + \text{Bern}(p)$, and $Y = X + \text{Bern}(D)$, we obtain $I(S;Y) \leq (1-2D)^2(1-h(p))$. Compare this to what we showed in Example 1: $I(S;Y) = 1 - h(p * D)$. Here, $(1-2D)^2$ is the injected privacy term obtained by the privacy mapping $P_{Y|X}$ and $1-h(p)$ is the intrinsic information/privacy term, quantifying the relation between X and S.

Next, we consider the case where only the marginal distribution P_X is given and we do not have access to $P_{S,X}$. We wish to design $P_{Y|X}$. The optimization problem is given by

$$\min_{P_{Y|X}: \mathbb{E}[d(X;Y)] \leq D} S^*(X;Y) \tag{9}$$

Now, consider the following optimization problem by replacing $S^*(X;Y)$ by $\rho_m^2(X;Y)$.

$$\min_{P_{Y|X}:\mathbb{E}[d(X;Y)]\leq D} \rho_m^2(X;Y) \tag{10}$$

We solve this optimization, and if the final solution satisfies $S^*(X;Y) = \rho_m^2(X;Y)$, then we have the solution to (9). In particular, if one of the conditions given in [16] holds, then we have the solution to (9). Next, we reformulate the constraint set in (10).

Theorem 2. *Given a distribution P_X, let $\sqrt{P_X}$ denote a vector with entries equal to square root of entires of P_X. If Q is a $|\mathcal{X}|\times|\mathcal{Y}|$ matrix satisfying the following constraints: 1) $Q \geq 0$ (entry-wise), 2) $QQ^t\sqrt{P_X} = \sqrt{P_X}$, then $P_{Y|X}$ (and $P_{X,Y}$) can be found uniquely such that $Q_{x,y} = \frac{P(x,y)}{\sqrt{P(x)}\sqrt{P(y)}}$.*

By Theorem 2, Optimization (10) can be cast as

$$\begin{aligned} &\min \lambda_2(Q) \\ &Q:\ QQ^t\sqrt{P_X} = \sqrt{P_X}, \\ &\qquad \mathbb{E}[d(X;Y)] \leq D,\ Q \geq 0 (\text{entry-wise}), \end{aligned} \tag{11}$$

where $\lambda_2(Q)$ denotes the second largest singular value of Q and expectation is over the joint probability induced by matrix Q. Note that the constraints are quadratic in the entries of Q. As an example of distortion constraint, $\mathbb{P}[X = Y] = \text{tr}\left(\mathcal{D}(\sqrt{P_X})Q\mathcal{D}(Q^t\sqrt{P_X})\right) \geq 1-D$ is quadratic in Q, where $\mathcal{D}(v)$ is a diagonal matrix with entries of v on the diagonal. Once we find Q, then $P_{Y|X}$ follows from that. The following shows under some conditions the optimization given in (11) is a convex optimization.

Corollary 1. *For a given P_X and P_Y, the optimization given in* (11) *is a convex optimization.*

IV. Additive Noise Privacy Mappings

Designing a privacy mapping requires characterizing $p_{Y|X}$ for all possible pairs $(x,y) \in \mathcal{X}\times\mathcal{Y}$, i.e. solving the optimization over $|\mathcal{X}|\,|\mathcal{Y}|$ variables. When $\mathcal{Y} = \mathcal{X}$, and $|\mathcal{X}|$ is large, solving the optimization over $|\mathcal{X}|^2$ variables may become intractable.

In this Section, we restrict our attention to the class of additive noise mappings, i.e. $Y = X + N$. We propose class-optimal privacy mappings for continuous and discrete variables X. The first advantage of additive noise mappings is their simplicity, and the fact that their design requires solving an optimization with a smaller number of variables, namely the noise distribution parameter. Moreover, the design of the optimal noise does not require full knowledge of the priors, but only knowledge of second order moments of X, namely the variance or covariance matrix of X.

A. Gaussian Mechanism

Consider a continuous X. We show that the optimal additive-noise mapping, under l_2-distortion, requires only knowledge of VAR(X) (or covariance matrix for multi-dimensional X), but not of P_X.

Since $S \to X \to Y$, we have $I(S;Y) \leq I(X;Y)$. For a given distortion, D, assume the minimum objective of the following rate distortion problem

$$\min_{P_{Y|X}:\mathbb{E}[d(X,Y)]\leq D} I(X;Y),$$

is I^*. Therefore, we have $I(D) \leq I^*$, where $I(D)$ is the minimum objective of (2). Using Theorem 1 for $Y \to X \to S$ (this follows from $S \to X \to Y$), we can further bound $I(D)$. In particular, since $I(S;Y) \leq I(X;Y)S^*(X;S)$, we obtain $I(D) \leq S^*(X;S)I^*$. Note that if $X = f(S)$ is a deterministic function of S, then $I(S;Y) = I(X;Y)$ and the bound is tight (this happens for instance in linear regression when $X = AS$ for some matrix A).

First, we show that among all privacy mechanisms in this class, adding Gaussian noise is optimal. Let $X \in \mathbb{R}^n$. Without loss of generality, we assume $\mathbb{E}[X] = 0$. Denote the covariance matrix of X by C_X. Let $Y = X+N$, where N is a multi-dimensional noise independent from X, with mean 0 and covariance matrix C_N.

Proposition 2. *Assume P_X is unknown in the design of privacy mapping and we only know VAR$(X) = \sigma_X^2$ for some σ_X. Also consider the class of privacy preserving schemes obtained by adding independent noise , N, to the signal, X. The noise has zero mean and variance (ℓ_2 norm distortion) equal to σ_N^2 for some σ_N. We show that, Gaussian noise is the best, in the following sense:*
$\max_{P_X:X\perp\!\!\!\perp N_G,\ VAR(X)=\sigma_X^2} I(X;X+N_G) \leq \max_{P_X:X\perp\!\!\!\perp N,\ VAR(X)=\sigma_X^2} I(X;X+N)$,
where N_G is Gaussian with zero mean and same variance as N. This implies that, the worst-case information leakage using N_G is not greater than worst-case information leakage using N.

For a given C_X and distortion level, D, Gaussian mechanism is as following:

1) take the eigen-value decomposition of C_X.
2) the covariance matrix of N_G, C_N, has eigen-vectors aligned with the eigen-vectors of C_X. Moreover, the corresponding eigen-values of C_N are given by solving
$$\begin{aligned} \min_{\sigma_i:\ 1\leq i\leq n} &\prod_{i=1}^{n} \frac{\sigma_i + \lambda_i}{\sigma_i} \\ \text{s.t.} &\sum_{i=1}^{n} \sigma_i \leq D, \end{aligned} \tag{12}$$
where λ_is and σ_is ($1 \leq \lambda_i \leq n$) denote the eigen-values of C_X and C_N, respectively.
3) let $Y = X + N_G$ where $N_G \sim \mathcal{N}(0, C_N)$. The distortion is given by $\sum_{i=1}^n \mathbb{E}[(Y_i - X_i)^2] = \text{tr}(C_N) = \sum_{i=1}^n \sigma_i \leq D$.

In the following theorem we prove that, the proposed mechanism is optimal.

Theorem 3. *Assuming ℓ_2-norm distortion and a distortion level, D, the covariance matrix of the optimum noise in the Gaussian mechanism has eigen-vectors aligned with the eigen-vectors of C_X. Also, the eigen-values are obtained by solving* (12)*.*

Optimization (12) can be solved with an approach similar to reverse water-filling ([22, Chapter 10.3]).

Example 3. Assume X is a deterministic real-valued function of S, $X = f(S)$ and that, $\text{VAR}(X) = \sigma_X^2$. Because of $S \rightarrow X \rightarrow Y$, we have $I(X;Y) = I(S;Y)$. Let $N \sim \mathcal{N}(0, \sigma^2)$ and $Y = X + N$. For any ϵ, we can achieve $(\epsilon, D)-$ divergence-distortion privacy, where $D = \frac{\sigma_X^2}{e^{2\epsilon H(S)} - 1}$. Note that this analysis works only for $\epsilon > 0$. Once we want to have perfect privacy, i.e. $\epsilon = 0$, then this scheme chooses $\sigma_N^2 = \infty$. In practice, this means that, Y is independent from X. If we assume $Y = \mathbb{E}[X]$ (a deterministic value), then $I(Y;S) = 0$ and $\mathbb{E}[d(X,Y)] = \text{VAR}(X)$. Therefore, For distortion level greater than or equal to $\text{VAR}(X)$, the deterministic mechanism that sets $Y = \mathbb{E}[X]$ achieves $\epsilon = 0$.

B. Discrete Noise Mechanism

In this section we consider discrete random variable, X, where $\mathcal{X} = \mathbb{Z}$. The case were $\mathcal{X}$ is a finite set is very similar. Again, we bound $I(X;Y)$ in order to bound $I(S;Y)$. Let the the distortion measure to be ℓ_p norm, i.e., the distortion between X and Y to be $\mathbb{E}[|X - Y|^p]^{\frac{1}{p}}$ for some $1 \leq p \leq \infty$.

Definition 6. *For a given $1 \leq p \leq \infty$, among all random variables with ℓ_p norm less than or equal to D, denote the distribution with the maximum entropy by $P^*_{p,D}$. Also denote the maximum entropy by $H^*(p, D)$.*

Next, we characterize $P^*_{p,D}$ and its entropy.

Lemma 1. *For any $1 \leq p \leq \infty$, $P^*_{p,D}$ is given by $P^*_{p,D}[Z = i] = AB^{-|i|^p}$, where A and B are chosen such that $\sum_{i=-\infty}^{\infty} AB^{-|i|^p} = 1$ and $\mathbb{E}[|Z|^p]^{\frac{1}{p}} = D$. Moreover, we have $H^*(p, D) = -\log A + (\log B) D^p$.*

Proposition 3. *Assume P_X is unknown in the design of privacy mapping and we only know $\mathbb{E}[|X|^p]^{\frac{1}{p}} = \mu_X$ for some μ_X. Also consider the class of privacy preserving schemes obtained by adding independent noise , N, to the signal, X. The noise has zero mean and ℓ_p norm distortion equal to μ_N, i.e. $\mathbb{E}[|N|^p]^{\frac{1}{p}} = \mu_N$ for some μ_N. We show that, a noise with $P^*_{p,D}$ distribution is the best, in the following sense:*

$\max_{P_X : X \perp\!\!\!\perp N_p,\ \mathbb{E}[|X|^p] = \mu_X^p} I(X; X + N_p) \leq \max_{P_X : X \perp\!\!\!\perp N,\ \mathbb{E}[|X|^p] = \mu_X^p} I(X; X + N)$,

where N_p is a noise with $P^*_{p,D}$ distribution and N is a random variable such that $\mathbb{E}[N_p] = \mathbb{E}[N] = 0$ and $\mathbb{E}[|N_p|^p] = \mathbb{E}[|N|^p] = \mu_N^p$. This implies that, the worst-case information leakage using N_p is not greater than worst-case information leakage using N.

The discrete noise mechanism is as following:

1) For a distortion measure ℓ_p ($1 \leq p \leq \infty$) and a distortion level, D, find $P^*_{p,D}$ as in Lemma 1.
2) Let $Y = X + Z$, where $Z \sim P^*_{p,D}$. We have $d(X,Y) = \mathbb{E}[|Y - X|^p]^{\frac{1}{p}} = \mathbb{E}[|Z|^p]^{\frac{1}{p}} \leq D$.

Next, we analyze the mutual information, $I(X;Y)$. Using Minkowski's inequality, we have $\mathbb{E}[|Y|^p]^{\frac{1}{p}} = \mathbb{E}[|X + Z|^p]^{\frac{1}{p}} \leq \mathbb{E}[|X|^p]^{\frac{1}{p}} + \mathbb{E}[|Z|^p]^{\frac{1}{p}} = \mu_X + D$. Therefore, we obtain $I(X;Y) = H(X+Z) - H(Z) \leq H^*(p, \mu_X + D) - H^*(p, D)$.

V. Privacy of Multi-agent Systems

New challenges in the design of privacy preserving mappings arise when multiple data releases to one or several agents occur. In this section, we address these challenges by bounding the information leakage under collusion/composition based on the information leakage of individual schemes. We first define the challenges under collusion and composition.

Collusion: The private data S is correlated with both non-private data X_1 and X_2. Two privacy mappings are applied to these non-private data to produce two distorted data, Y_1 and Y_2 that are then released to two agents. We wish to analyze the cumulative privacy guarantees on S when the agents share Y_1 and Y_2. We focus on the case where the two privacy mappings are designed in a decentralized fashion: Each privacy mapping is designed to protect against the inference of S from each of the released data separately.

Composition: The private data S is correlated with the non-private data X_1 and X_2 through the joint probability distribution P_{S,X_1,X_2}. Assume that we are able to design separately two privacy mappings that transform X_1 (resp. X_2) into Y_1 (resp. Y_2). An agent asks for the pair (X_1, X_2). We wish to combine the two separate privacy mappings mentioned previously to generate a privacy mapping for the pair (X_1, X_2), that still guarantees a certain level of privacy, without having to design jointly from scratch the new privacy mapping for the pair (X_1, X_2). Composition allows to make the design simpler, by breaking one large optimization with many variables into several smaller optimizations with fewer variables.

Both collusion and composition problems can be captured by the following setting: a private variable S is correlated with X_1 and X_2. We perform two separate privacy mappings on X_1 and X_2 to obtain Y_1 and Y_2, respectively. $P_{Y_1|X_1}$ and $P_{Y_2|X_2}$ are designed with given distortion levels, and the individual information leakages are $I(S;Y_1)$ and $I(S;Y_2)$. Assume that Y_1 and Y_2 are combined together into a pair (Y_1, Y_2), either by colluding agents, or by the privacy agent through composition. We want to analyze the privacy level under this combination of information.

Lemma 2. *Assume Y_1, Y_2, and S form a Markov chain in any order. If the privacy preserving mappings leak $I(Y_1;S)$ and $I(Y_2;S)$ bits, then under collusion/composition at most $I(Y_1;S) + I(Y_2;S)$ amount of information is leaked. In other words, $I(Y_1,Y_2;S) \leq I(Y_1;S) + I(Y_2;S)$. Moreover, if $S \rightarrow Y_1 \rightarrow Y_2$, then $I(S;Y_1,Y_2) \leq I(Y_1;S)$. If $S \rightarrow Y_2 \rightarrow Y_1$, then $I(S;Y_1,Y_2) \leq I(Y_2;S)$.*

Note that in general, the collusion/composition might lead to full recovery of S. For instance, let S, Y_1, and Y_2 be three $\text{Bern}(\frac{1}{2})$ random variables such that $S = Y_1 \oplus Y_2$ and $Y_1 \perp\!\!\!\perp Y_2$. Then, we have $I(Y_1;S) = I(Y_2;S) = 0$, whereas $I(Y_1,Y_2;S) = 1$ bit and S is fully recoverable from (Y_1,Y_2). Another example is when $Y_1 = S + N$ where N is some noise and $Y_2 = S - N$. We can fully recover $S = \frac{Y_1+Y_2}{2}$.

Next, we use the results on maximal correlation to upper bound the amount of information leakage in the presence of collusion/composition.

Theorem 4. *Let $P_{Y_1|X_1}$ and $P_{Y_2|X_2}$ be designed separately, i.e., $P_{Y_1,Y_2|X_1,X_2} = P_{Y_1|X_1}P_{Y_2|X_2}$. Let $\lambda = \max\{S^*(X_1;Y_1), S^*(X_2;Y_2)\}$. If $I(Y_1;Y_2) \geq \lambda I(X_1;X_2)$, then we have $I(S;Y_1,Y_2) \leq I(S;X_1,X_2)\max\{S^*(X_1;Y_1), S^*(X_2;Y_2)\}$.*

Therefore, if both mappings are designed separately with small maximal correlation, then after collusion/composition we can bound the amount of information leakage.

Corollary 2. *In collusion or composition, assume that we wanted to guarantee $\epsilon-$divergence privacy after combination of information. We design both mappings such that both of the corresponding maximal correlations to the mappings are bounded by ϵ. Thus, using Theorem 4, we have $I(S;Y_1,Y_2) \leq \epsilon I(S;X_1,X_1) \leq \epsilon H(S)$. Therefore, both of the individual mappings are $\epsilon-$divergence private and after collusion/composition we still have a $\epsilon-$divergence private setting.*

Note 1. *Assume we have two private random variables S_1 and S_2 each of them correlated with X_1 and X_2, respectively. We distort X_1 and X_2 to obtain Y_1 and Y_2, respectively. An agent has access to Y_1 and Y_2 and wishes to discover (S_1,S_2). Similarly, we can show that:*

if $P_{Y_1|X_1}$ and $P_{Y_2|X_2}$ be designed separately, i.e., $P_{Y_1,Y_2|X_1,X_2} = P_{Y_1|X_1}P_{Y_2|X_2}$. Let $\lambda = \max\{S^*(X_1;Y_1), S^*(X_2;Y_2)\}$. If $I(Y_1;Y_2) \geq \lambda I(X_1;X_2)$, then we obtain

$$I(S_1,S_2;Y_1,Y_2) \leq I(S_1,S_2;X_1,X_2)\max\{S^*(X_1;Y_1), S^*(X_2;Y_2)\}.$$

VI. Comparison of Privacy Metrics

In this section, we compare the existing privacy measures in the literature. In particular, we compare divergence privacy and differential privacy and show that, while divergence privacy guarantees a small probability of inferring private random variable based on the released data (Proposition 1), differential privacy fails to guarantee. Let $S = (S_1, \ldots, S_n)$ and $S \rightarrow X \rightarrow Y$. Next, we give the notions of privacy in the literature and compare them.

Definition 7.

- Differential privacy*([4]): For a given ϵ, $P_{Y|S}$ is $\epsilon-$ differentially private if $\sup_{y,s,s':\ s\sim s'} \frac{P(y\in A|s)}{P(y\in A|s')} \leq e^{\epsilon}$, for any measurable set A, where $s \sim s'$ denote neighboring. The notion of neighboring can have multiple definitions, e.g. Hamming distance 1 (differ in a single coordinate), or ℓ_p distance below a threshold. In this paper, we use the former definition.*
- Strong differential privacy*([23]): For a given ϵ, $P_{Y|S}$ is $\epsilon-$strongly differential private if $\sup_{y,s,s'} \frac{P(y\in A|s)}{P(y\in A|s')} \leq e^{\epsilon}$, for any measurable set A and s and s'. This definition is related to local differential privacy ([23]). This is stronger than differential privacy, because we relaxed the neighboring assumption.*
- Information privacy *([1]): For a given ϵ, $P_{Y|S}$ is $\epsilon-$information private if $e^{-\epsilon} \leq \frac{P(s\in B|y\in A)}{P(s\in B)} \leq e^{\epsilon}$, for any measurable sets A and B.*
- Worst-case divergence privacy*: For a given ϵ, $P_{Y|S}$ is worst-case $\epsilon-$divergence private if $H(S) - \min_y H(S|Y=y) = \epsilon H(S)$*
- (ϵ,δ)-differential privacy*([24]): For any given ϵ and δ, $P_{Y|S}$ is $(\epsilon,\delta)-$ differentially private if $P(y\in A|s) \leq P(y\in A|s')e^{\epsilon} + \delta$, for any measurable set A and neighboring s and s'.*

Next, we compare the definitions given above.

Proposition 4. *We have the following relation between the privacy metrics.*

1) *$\epsilon-$ strong differential privacy $\Rightarrow$ $\epsilon-$ information privacy*
2) *$\epsilon-$ information privacy $\Rightarrow$ $2\epsilon-$ strong differential privacy*
3) *$\epsilon-$ information privacy $\Rightarrow$ $\frac{\epsilon}{H(S)}-$ worst-case divergence privacy*
4) *$\frac{\epsilon}{H(S)}-$ worst-case divergence privacy $\Rightarrow$ $\frac{\epsilon}{H(S)}-$ divergence privacy*
5) *$\epsilon-$ differential privacy $\Rightarrow$ $(\epsilon,\delta)-$ differential privacy for any $\delta \geq 0$.*

In the sequel, we give two examples comparing differential privacy with divergence privacy. In the first example, we focus on the probability of recovering the private data given that we satisfy these notions of privacy. In the second example we show the difference between the Gaussian mechanisms tailored to $(\epsilon,\delta)-$ differential privacy and divergence privacy.

In the next example we consider the particular case of counting query. This example is partially studied in [1]. We show that, using differential privacy, full detection of the private data is possible. On the other hand, using divergence privacy, the probability of detecting the private data is small.

Example 4. Let $S_1, \dots, S_n$ be binary correlated random variables and let $X = \sum_{i=1}^{n} S_i$. Assume $S_1, \dots, S_n$ are correlated in a way that, $S_1 \geq \cdots \geq S_n$. Therefore, knowing X, we can exactly recover $S = (S_1, \dots, S_n)$. Also, assume S_is $(1 \leq i \leq n)$ are correlated in a way that $P(X = ki) = \frac{1}{1+n/k}$, for $i \in \{0, 1, \dots, n/k\}$ (assume, $n = 0$ mode k). $P(Y|S)$ is $\epsilon-$ differentially private if we add Laplacian noise to X, i.e., $Y = X + \text{Lap}(\frac{1}{\epsilon})$ ([4]). Fix ϵ and let $n = k^k$, where k goes to infinity. It is shown that error probability in detecting X (and S) is approximately $P_e = e^{\frac{-k\epsilon}{2}}$ ([1]), which is very small for large enough k. Thus, differential privacy does not guarantee a small probability of detecting S. Note that, the divergence privacy factor is approximately $\frac{I(S;Y)}{H(S)} = 1 - e^{\frac{-k\epsilon}{2}}$, which is very close to one and this is the reason for large detection probability. Now, consider Gaussian mechanism, where we add Gaussian noise instead of Laplacian noise. In this scheme, the variance of the Gaussian noise depends on the correlation in the data S via the variance of X, σ_X^2. We have $\sigma_X^2 \approx \frac{1}{12} k^{2k}$, where $\approx$ denotes that, the ratio goes to 1 as k goes to infinity. Let N be a Gaussian distribution with a variance satisfying: $\frac{\sigma_X^2}{\sigma_N^2} \approx k^{2\epsilon(k-1)}$. Adding this noise to X, the leakage factor is less than or equal to ϵ. Moreover, $P_e \geq \frac{(1-\epsilon)\log(1+n/k)}{\log n} \overset{k\to\infty}{\to} 1 - \epsilon$. The reason for this large error probability is mainly: using Gaussian mechanism, we partially take into consideration the prior correlation of S, by using σ_X^2 in the design of noise.

$(\epsilon, \delta)-$differential privacy metric can be achieved by adding Gaussian noise to the signal, X ([24]). In the next example we compare the mechanism given in [24] with our Gaussian mechanism.

Example 5. It is shown that, by adding Gaussian noise with variance $\sigma^2 \geq \frac{1}{\epsilon^2} 2\log(2/\delta)$ we can achieve $(\epsilon, \delta)-$differential privacy ([24]). This scheme results in a distortion $D \geq \frac{1}{\epsilon^2} 2\log(2/\delta)$ and the leakage of information $L \leq \frac{1}{2}\log\left(1 + \frac{\sigma_X^2}{\frac{1}{\epsilon^2} 2\log(2/\delta)}\right)$. A qualitative way for comparison is to state that: Using (ϵ, δ) differential privacy Gaussian mechanism, we would require a large distortion to achieve a small leakage. On the other hand, using a divergence privacy Gaussian mechanism given in IV-A, a scheme that leaks L bits with minimum distortion, D, achieves any $(\epsilon, \delta)-$ differential privacy, where $\frac{1}{\epsilon^2} 2\log(\frac{2}{\delta}) = \frac{\sigma_X^2}{e^{2L}-1}$.

References

[1] F. du Pin Calmon and N. Fawaz, "Privacy against statistical inference," in *Allerton Conference on Communication, Control, and Computing, Allerton*, 2012.

[2] D. Rebollo-Monedero, J. Forné, and J. Domingo-Ferrer, "From t-closeness-like privacy to postrandomization via information theory," *IEEE Trans. on Knowledge and Data Engineering*, vol. 22, no. 11, 2010.

[3] C. Dwork, "Differential privacy," in *Automata, Languages and Programming*. Springer, 2006.

[4] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography*. Springer, 2006.

[5] C. Dwork, G. N. Rothblum, and S. Vadhan, "Boosting and differential privacy," in *IEEE FOCS*, 2010.

[6] L. Wasserman and S. Zhou, "A statistical framework for differential privacy," *Journal of the American Statistical Association*, vol. 105, no. 489, 2010.

[7] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Local privacy and statistical minimax rates," *arXiv preprint arXiv:1302.3203*, 2013.

[8] D. Kifer and A. Machanavajjhala, "A rigorous and customizable framework for privacy," in *ACM PODS*, 2012.

[9] I. S. Reed, "Information theory and privacy in data banks," in *Proceedings of the June 4-8, 1973, national computer conference and exposition*. ACM, 1973.

[10] H. Yamamoto, "A source coding problem for sources with additional outputs to keep secret from the receiver or wiretappers (corresp.)," *Information Theory, IEEE Transactions on*, vol. 29, no. 6, 1983.

[11] L. Sankar, S. R. Rajagopalan, and H. V. Poor, "Utility-privacy tradeoff in databases: An information-theoretic approach," *IEEE Trans. Inf. Forensics Security*, vol. 8, no. 6, June 2013. [Online]. Available: http://arxiv.org/abs/1102.3751

[12] A. Evfimievski, J. Gehrke, and R. Srikant, "Limiting privacy breaches in privacy preserving data mining," in *ACM PODS*, 2003.

[13] S. Salamatian, A. Zhang, F. du Pin Calmon, S. Bhamidipati, N. Fawaz, B. Kveton, P. Oliveira, and N. Taft, "How to hide the elephant- or the donkey- in the room: Practical privacy against statistical inference for large data," in *IEEE GlobalSIP*, 2013.

[14] A. Blum, K. Ligett, and A. Roth, "A learning theory approach to non-interactive database privacy," in *ACM STOC*, 2008.

[15] A. Makhdoumi and N. Fawaz, "Privacy-utility tradeoff under statistical uncertainty," *ArXiv e-prints*, 2013. [Online]. Available: http://arxiv.org/

[16] R. Ahlswede and P. Gács, "Spreading of sets in product spaces and hypercontraction of the markov operator," *The Annals of Probability*, 1976.

[17] V. Anantharam, A. Gohari, S. Kamath, and C. Nair, "On maximal correlation, hypercontractivity, and the data processing inequality studied by erkip and cover," *arXiv preprint arXiv:1304.6133*, 2013.

[18] H. O. Hirschfeld, "A connection between correlation and contingency," in *Proceedings of the Cambridge Philosophical Society*, vol. 31. Cambridge Univ Press, 1935.

[19] A. Rényi, "On measures of dependence," *Acta Mathematica Hungarica*, vol. 10, no. 3, 1959.

[20] S. Kamath and V. Anantharam, "Non-interactive simulation of joint distributions: The hirschfeld-gebelein-rényi maximal correlation and the hypercontractivity ribbon," in *Allerton Conference on Communication, Control, and Computing, Allerton*, 2012.

[21] H. S. Witsenhausen, "On sequences of pairs of dependent random variables," *SIAM Journal on Applied Mathematics*, vol. 28, no. 1, 1975.

[22] T. M. Cover and J. A. Thomas, *Elements of information theory*. Wiley-interscience, 2012.

[23] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith, "What can we learn privately?" *SIAM Journal on Computing*, vol. 40, no. 3, 2011.

[24] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: Privacy via distributed noise generation," in *Advances in Cryptology-EUROCRYPT 2006*. Springer, 2006.