

From the Information Bottleneck to the Privacy Funnel

Ali Makhdoumi
MIT
makhdoum@mit.edu

Salman Salamatian
EPFL
salman.salamatian@epfl.ch

Nadia Fawaz
Technicolor
nadia.fawaz@technicolor.com

Muriel Médard
MIT
medard@mit.edu

Abstract—We focus on the privacy-utility trade-off encountered by users who wish to disclose some information to an analyst, that is correlated with their private data, in the hope of receiving some utility. We rely on a general privacy statistical inference framework, under which data is transformed before it is disclosed, according to a probabilistic privacy mapping. We show that when the log-loss is introduced in this framework in both the privacy metric and the distortion metric, the privacy leakage and the utility constraint can be reduced to the mutual information between private data and disclosed data, and between non-private data and disclosed data respectively. We justify the relevance and generality of the privacy metric under the log-loss by proving that the inference threat under any bounded cost function can be upperbounded by an explicit function of the mutual information between private data and disclosed data. We then show that the privacy-utility tradeoff under the log-loss can be cast as the non-convex *Privacy Funnel* optimization, and we leverage its connection to the Information Bottleneck, to provide a greedy algorithm that is locally optimal. We evaluate its performance on the US census dataset. Finally, we characterize the optimal privacy mapping for the Gaussian Privacy Funnel.

I. INTRODUCTION

We consider a setting in which users have two kinds of data that are correlated: some data that each user would like to remain private and some non-private data that he is willing to disclose to an analyst and from which he will derive some utility. The analyst is a legitimate receiver of the disclosed data, which he will use to provide utility to the user, but he can also adversarially exploit it to infer the user's private data. This creates a tension between privacy and utility requirements. To reduce the inference threat on private data while maintaining utility, each user's non-private data is transformed before it is disclosed, according to a probabilistic privacy mapping. The design of the privacy mapping should balance the tradeoff between the utility of the disclosed data, and the privacy of the private data: it should keep the disclosed transformed data as much informative as possible about the non-private data, while leaking as little information as possible about the private data.

The framework for privacy against inference attacks in [1] proposes to design the privacy mapping as the solution to an optimization minimizing the inference threat subject to a utility constraint. Our approach relies on this framework, and makes the following three contributions. First, we show that when the log-loss is introduced in this framework in both the privacy metric and the distortion metric, the privacy leakage reduces to the mutual information between private data and disclosed data, while the utility requirement is modeled by the mutual information between non-private data and disclosed data. We justify the relevance and generality of the privacy

metric under the log-loss by proving that the inference threat, defined in [1] as the inference cost gain, under any bounded cost function can be upperbounded by an explicit function of the mutual information between private data and disclosed data. We then show that the privacy-utility tradeoff under the log-loss can be cast as the *Privacy Funnel* optimization, and study its connection to the Information Bottleneck [2]. Second, for general distributions, the privacy funnel optimization being a non-convex problem, we provide a greedy algorithm for the Privacy Funnel that is locally optimal by leveraging connections to the Information Bottleneck method [3], [2], and evaluate its performance on real-world data. Third, we study the Gaussian Privacy Funnel, where the user data has a Gaussian distribution and the mapping is also a Gaussian mapping, and we characterize the optimal privacy mapping.

Related Work: Several works, such as [4], [5], [6], [7], [8], have studied the issue of keeping some information private while disclosing some correlated information, by distorting the information disclosed. Differential privacy [6], [7] was introduced to answer queries on statistical databases in a privacy-preserving manner, by minimizing the chances of identification of the database records. One line of work in information theoretic privacy [4], [8] studies the trade-off between privacy and utility, where they consider expected distortion as a measure of utility and equivocation as a measure of privacy. Reference [8] 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. These approaches are different from our approach in three main ways. First, we do not consider a communication problem where the rate needs to be bounded. Second, we use the average amount of bits as a measure of both utility and privacy (log-loss distortion or mutual information). Third, we design the mapping from non-private data to disclosed data in non-asymptotic regime, whereas [8] designs an encoder-decoder to achieve certain asymptotic rate. The wire-tap channel, introduced in [9], focuses on designing the encoder and decoder to release information and protect private information from an eavesdropper, where utility is measured in terms of error probability of the decoded message, and secrecy is measured in terms of normalized mutual information. Our setting differs from the wire-tap channel, as it does not involve a third-party eavesdropper, but the analyst is both a legitimate receiver of the disclosed data, and a potential adversary as it can use it to try to infer private data. Moreover, we focus on

the privacy mapping design (channel design), with different measures of privacy and utility.

The log-loss distortion has been studied in [10] as a measure of distortion in the context of multi-terminal source coding. Log-loss as measure of distortion is also studied in [11] where they show that log-loss satisfies certain properties that leads to the Information Bottleneck method [2]. Finally, for an overview of the central role of the log-loss distortion in prediction, we refer the reader to [12].

Outline: In Section II, we introduce the privacy-utility trade-off against Inference attacks. In Section III, we describe the privacy funnel method and show properties of log-loss metric, and then characterize the privacy-disclosure trade-off as the privacy funnel optimization. In Section IV, we provide a greedy algorithm to design the privacy mapping and evaluate it on real-world data. In Section V, we characterize the optimal Gaussian privacy mapping for the Gaussian privacy funnel. Proofs omitted due to lack of space can be found in [13].

Notations: Throughout the paper, X denotes a random variable over alphabet \mathcal{X} with distribution P_X . All random variables are assumed to be discrete, unless mentioned otherwise.

II. PRIVACY-UTILITY AGAINST INFERENCE ATTACKS

In this background section, we first describe the setting, and the privacy and utility metrics introduced in the framework for privacy against inference attacks in [1]. Then, we recall how the privacy-utility trade-off can be cast into an optimization.

A. Setting: We consider a setting where a user has some private data, represented by the random variable $S \in \mathcal{S}$, which is correlated with some non-private data $X \in \mathcal{X}$, that the user wishes to share with an analyst¹. The correlation between S and X is captured by the joint distribution $P_{S,X}$. Due to this correlation, releasing X to the analyst would enable him to draw some inference on the private data S . 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 by $Y \in \mathcal{Y}$. The distorted data Y is generated by passing X through a conditional distribution $P_{Y|X}$, called the privacy mapping. Throughout the paper, we assume $S \rightarrow X \rightarrow Y$ form a Markov chain. Therefore, once $P_{Y|X}$ is found, the joint distribution $P_{S,X,Y}$ is defined.

The analyst is a legitimate recipient of data Y , which it can use to provide utility to the user, e.g. some personalized service. However, the analyst can also act as an adversary by using Y to illegitimately infer private data S . The privacy mapping aims at balancing the tradeoff between utility and privacy: the privacy mapping should be designed to decrease the inference threat on private S by reducing the dependency between Y and S , while at the same time preserving the utility of Y , by maintaining the dependency between Y and X .

B. Privacy Metric: We consider the inference threat model introduced in [1], in which the analyst performs an adversarial inference attack on the private data S . More precisely, the analyst selects a distribution q , from the set \mathcal{P}_S of all probability distributions over \mathcal{S} , that minimizes an expected inference

cost function $C(S, q)$. In other words, the analyst chooses in an adversarial way a belief distribution q over the private variables S prior to observing Y , $q_0^* = \arg \min_{q \in \mathcal{P}_S} \mathbb{E}_{P_S}[C(S, q)]$, and a revised belief distribution after observing $Y = y$, $q_y^* = \arg \min_{q \in \mathcal{P}_S} \mathbb{E}_{P_{S|Y}}[C(S, q)|Y = y]$. This models a very broad class of adversaries that perform statistical inference. Using the chosen belief distribution q , the analyst can produce an estimate of the input S , e.g. using a Maximum a Posteriori (MAP) estimator. Let c_0^* and c_y^* respectively denote the minimum average cost of inferring S without observing Y , and after observing $Y = y$:

$$c_0^* = \min_{q \in \mathcal{P}_S} \mathbb{E}_{P_S}[C(S, q)], \quad c_y^* = \min_{q \in \mathcal{P}_S} \mathbb{E}_{P_{S|Y}}[C(S, q)|Y = y].$$

Thanks to the observation of Y , the analyst obtains an average gain in inference cost of $\Delta C = c_0^* - \mathbb{E}_{P_Y}[c_y^*]$. The average inference cost gain ΔC was proposed as a general privacy metric in [1], as it measures the improvement in the quality of the inference of private data S due to the observation of Y . The privacy mapping $P_{Y|X}$ should aim at reducing ΔC , or in other words it should aim at bringing the inference cost given Y closer to the initial inference cost.

C. Accuracy Metric: The privacy mapping should maintain the utility of the distorted data Y . In the framework proposed in [1], the utility requirement is modeled by a constraint on the average distortion $\mathbb{E}_{P_{X,Y}}[d(X, Y)] \leq D$, for some distortion measure $d : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^+$, and some distortion level $D \geq 0$. Assuming that the distortion measure d is a function of X and Y , but not of their statistical properties, the average distortion $\mathbb{E}_{P_{X,Y}}[d(X, Y)]$ is linear in $P_{Y|X}$. Consequently, the distortion constraint is a linear constraint in $P_{Y|X}$.

D. Privacy-Accuracy Tradeoff: The optimal privacy mapping for a given distortion level D is obtained as the solution of the following optimization

$$\min_{P_{Y|X} : \mathbb{E}_{P_{X,Y}}[d(X, Y)] \leq D} \Delta C \quad (1)$$

If ΔC is convex in $P_{Y|X}$, then (1) is a convex optimization, since the constraint $\mathbb{E}_{P_{X,Y}}[d(X, Y)]$ is linear in $P_{Y|X}$.

III. THE PRIVACY FUNNEL METHOD

In this section, we focus on the privacy-utility framework when the log-loss is used in both the privacy metric and in the distortion metric. We justify the relevance of the log-loss in such a framework, and characterize the resulting privacy-disclosure tradeoff as the Privacy Funnel optimization. Finally, we show how the Privacy Funnel is related to the Information Bottleneck [2], and how algorithms developed for the latter can inform the design of algorithms for the former.

A. Privacy metric under log-loss: In this section, we focus on the threat model under the log-loss cost function. We first recall that, under this cost-function, the privacy leakage can be measured by the mutual information $I(S; Y)$ between the private variable S and the variable Y . We then justify the relevance and the generality of the use of the log-loss in the threat model, by showing that the inference cost gain for any bounded cost function can be upperbounded by a function of the mutual information between S and Y .

¹Both S and X can be either a single variable or a vector of variables.

Under the log-loss cost function $C(s, q) = -\log q(s)$, $\forall s \in \mathcal{S}$, the privacy leakage can be measured by the mutual information $I(S; Y)$, as stated in the following lemma.

Lemma 1 ([1]). *The average inference cost gain under the log-loss cost function $C(s, q) = -\log q(s)$, is the mutual information between S and Y : $\Delta C = I(S; Y)$.*

We now justify the relevance and the generality of the use of the log-loss in the threat model. More precisely, in Theorem 1 below, we prove that for any bounded cost function $C(S, q)$, the associated inference cost gain ΔC can be upperbounded by an explicit constant factor of $\sqrt{I(S; Y)}$. Thus, controlling the cost gain under the log-loss, so that it does not exceed a target privacy level, is sufficient to ensure that the privacy threat under a different bounded cost function is also controlled. Therefore, the design of the privacy mapping can be focused on minimizing the privacy leakage as measured by $I(S; Y)$.

Theorem 1. *Let $L = \sup_{s \in \mathcal{S}, q \in \mathcal{P}_S} |C(s, q)| < \infty$. We have $\Delta C = c_0^* - \mathbb{E}_{P_Y}[c_Y^*] \leq 2\sqrt{2}L\sqrt{I(S; Y)}$.*

The proof of Theorem 1 requires the following lemma.

Lemma 2. *Let $C(s, q)$ be a bounded cost function such that $L = \sup_{s \in \mathcal{S}, q \in \mathcal{P}_S} |C(s, q)| < \infty$. For any given $y \in \mathcal{Y}$,*

$$\mathbb{E}_{P_{S|Y}}[C(S, q_0^*) - C(S, q_y^*)|Y = y] \leq 2\sqrt{2}L\sqrt{D(P_{S|Y=y}|P_S)}.$$

The proof of Lemma 2 follows from Pinsker's inequality and is given in the full version [13]. We now prove Theorem 1.

proof of Theorem 1: We have

$$\begin{aligned} \Delta C &= \mathbb{E}_{P_S}[C(S, q_0^*)] - \mathbb{E}_{P_Y}[\mathbb{E}_{P_{S|Y}}[C(S, q_y^*)|Y = y]] \\ &= \mathbb{E}_{P_Y}[\mathbb{E}_{P_{S|Y}}[C(S, q_0^*) - C(S, q_y^*)|Y = y]] \\ &\leq 2\sqrt{2}L\mathbb{E}_{P_Y}[D(P_{S|Y=y}|P_S)] \leq 2\sqrt{2}L\sqrt{I(S; Y)}, \end{aligned}$$

where the last step follows from concavity of square root function and the one before that follows from Lemma 2. ■

Using Fano's lemma, the average cost gain under log-loss can also be interpreted as a bound on the probability of error an adversary would make when inferring S from Y [15]. It is to be noted that ΔC is an averaged quantity. A formulation for the worst-case privacy leakage under log-loss is introduced in [1], but we mention that at perfect privacy, *i.e.* when $\Delta C = 0$, both average and worst case privacy amount to statistical independence between S and Y when using the log-loss. In practice, it has been shown in [16] and [17] that utility can be maintained while enforcing perfect privacy for demographic attributes while releasing data for recommendation. Further empirical evidence in [14] suggests that the average cost gain is a good approximation to the worst case scenario.

B. Accuracy metric under log-loss: Consider the log-loss distortion defined as $d(x, y) = -\log P(X = x|Y = y)$, which is a function of x and y as well as $P_{Y|X}$. Using log-loss, the average distortion is $\mathbb{E}[d(X, Y)] = \mathbb{E}_{P_{X, Y}}[-\log P_{X|Y}] = H(X|Y)$ that can be minimized by designing the mapping $P_{Y|X}$. Thus, the constraint $\mathbb{E}[d(X, Y)] \leq D$ would be $H(X|Y) \leq D$ for a given distortion level, D . Given P_X , and therefore $H(X)$, and assuming that $R = H(X) - D$, the

distortion constraint can be rewritten as $I(X; Y) \geq R$, that is the same as the constraint of (2). It should be noted that the average distortion under the log-loss is not linear in $P_{Y|X}$.

For a given P_{SX} and $P_{Y|X}$, where $S \rightarrow X \rightarrow Y$, we define the *disclosure* to be the mutual information between X and Y .

C. Privacy-Disclosure Trade-off: There is a trade-off between the information that the user shares about X and the information that the user keeps private about S . We pass X through a randomized mapping $P_{Y|X}$ and reveal Y to the analyst. The purpose of this mapping is to make Y informative about X and to make Y uninformative about S . Given P_{SX} , we design the privacy mapping $P_{Y|X}$ to maximize the amount of information $I(X; Y)$ that user disclose about the public information, X , while minimizing the collateral information about the private variable S measured by $I(S; Y)$.

The trade-off between disclosure and privacy in the design of the privacy mapping is represented by the following optimization, that we refer to as the *Privacy Funnel*:

$$\min_{P_{Y|X}: I(X; Y) \geq R} I(S; Y). \quad (2)$$

For a given disclosure level R , among all feasible privacy mappings $P_{Y|X}$ satisfying $I(X; Y) \geq R$, the privacy funnel selects the one that minimizes $I(S; Y)$. Note that $I(X; Y)$ is convex in $P_{Y|X}$ and since $P_{Y|S}$ is linear in $P_{Y|X}$ and $I(S; Y)$ is convex in $P_{Y|S}$, the objective function $I(S; Y)$ is convex in $P_{Y|X}$. However, because of the constraint $I(X; Y) \geq R$, the Privacy Funnel (2) is not a convex optimization [18, Chap. 4].

D. Connection to the Information Bottleneck Method: The information bottleneck method, introduced in [2], considers the setting where a variable X is to be compressed, while maintaining the information it bears about another correlated variable S . The information bottleneck method is a technique generalizing rate-distortion, as it seeks to optimize the tradeoff between the compression length of X and the accuracy of the information preserved about S in the compressed output Y . The information bottleneck optimization [2] is: $\min_{P_{Y|X}: I(S; Y) \geq C} I(X; Y)$, for some constant C . The information bottleneck optimization bears some resemblance to the privacy funnel (2), but is indeed a different optimization. In the information bottleneck, the mapping $P_{Y|X}$ is designed to make $I(X; Y)$ small and $I(S; Y)$ large, whereas in the privacy funnel we have the reverse situation, *i.e.*, the privacy mapping aims to make $I(X; Y)$ large and $I(S; Y)$ small.

Several techniques were developed to solve the information bottleneck problem such as alternating iteration [2] and agglomerative information bottleneck [3]. A question we examined is whether algorithms developed to solve the information bottleneck optimization could be adapted to solve the privacy funnel optimization. The alternating iteration algorithm [2] finds a stationary point of the Lagrangian of information bottleneck optimization defined as $\mathcal{L} = I(X; Y) - \beta I(S; Y)$ for some β . The stationary point can be a local minimum, which addresses the information bottleneck, or a local maximum in which case it addresses the privacy funnel. However, there is no guarantee on the convergence of this alternating algorithm to either a local minimum or a local maximum.

Algorithm 1 Greedy algorithm-privacy funnel

Input: $R, P_{S,X}$
Initialization: $\mathcal{Y} = \mathcal{X}, P_{Y|X}(y|x) = \mathbf{1}\{y = x\}$.
while there exists i', j' such that $I(X; Y^{i'-j'}) \geq R$ **do**
among those i', j' , let
 $\{y_i, y_j\} = \arg \max_{y_{i'}, y_{j'} \in \mathcal{Y}} I(S; Y) - I(S; Y^{i'-j'})$
merge: $\{y_i, y_j\} \rightarrow y_{ij}$
update: $\mathcal{Y} = \{\mathcal{Y} \setminus \{y_i, y_j\}\} \cup \{y_{ij}\}$ and $P_{Y|X}$
Output: $P_{Y|X}$

Thus, in Section IV we develop a new greedy algorithm that is guaranteed to converge to a solution of the privacy funnel.

IV. ALGORITHM FOR THE PRIVACY FUNNEL

We showed that the privacy funnel (2) optimization is not a convex optimization. In this section, we provide a greedy algorithm to solve (2) and we evaluate it on real-world data.

Greedy Algorithm: Suppose the constraint $I(X; Y) \geq R$ is given for some $R \leq H(X)$. We wish to find $P_{Y|X}$ that minimizes $I(S; Y)$. Note that for $\mathcal{Y} = \mathcal{X}$ and $P_{Y|X}(y|x) = \mathbf{1}\{x = y\}$ (where $\mathbf{1}\{x = y\} = 1$ if and only if $x = y$), the condition $I(X; Y) \geq R$ is satisfied because $I(X; Y) = H(X) \geq R$. However, $I(S; Y)$ might be too large. The idea is to merge two elements of \mathcal{Y} to make $I(S; Y)$ smaller, while satisfying $I(X; Y) \geq R$. This method is motivated by agglomerative information method introduced in [3]. We merge y_i and y_j and denote the merged element by y_{ij} . We then update $P_{Y|X}$ as $p(y_{ij}|x) = p(y_i|x) + p(y_j|x)$, for all $x \in \mathcal{X}$. After merging, we also have $p(y_{ij}) = p(y_i) + p(y_j)$. Consider the row stochastic matrix P as $P_{x,y} = P_{Y|X}(y|x)$ for all $x \in \mathcal{X}$ and all $y \in \mathcal{Y}$. In Algorithm (1) we start with P as an identity matrix and then at each iteration we delete two columns of P (corresponding to y_i and y_j) and add their summation as a new column (corresponding to y_{ij}) to P . Thus, the resulting matrix at the end contains only zeros and ones, determining all $x \in \mathcal{X}$ and all $y \in \mathcal{Y}$ such that $P_{Y|X}(y|x) = 1$. Let Y^{i-j} be the resulting Y from merging y_i and y_j . Algorithm (1) is a greedy algorithm that uses this idea in order to solve optimization (2). One need to calculate $I(S; Y) - I(S; Y^{i-j})$ and $I(X; Y) - I(X; Y^{i-j})$ at each iteration of Algorithm (1). Proposition 1 shows an efficient way to calculate them.

Proposition 1. For a given joint distribution $P_{S,X,Y} = P_{S,X}P_{Y|X}$, we have $I(S; Y) - I(S; Y^{i-j}) =$

$$p(y_{ij})H\left(\frac{p(y_i)P_{S|Y=y_j} + p(y_j)P_{S|Y=y_i}}{p(y_{ij})}\right) - (p(y_i)H(P_{S|Y=y_i}) + p(y_j)H(P_{S|Y=y_j})).$$

We also have $I(X; Y) - I(X; Y^{i-j}) =$

$$p(y_{ij})H\left(\frac{p(y_i)P_{X|Y=y_j} + p(y_j)P_{X|Y=y_i}}{p(y_{ij})}\right) - (p(y_i)H(P_{X|Y=y_i}) + p(y_j)H(P_{X|Y=y_j})).$$

The proof is given in the full version of paper [13]. Proposition 1 shows that the difference in the mutual information after merging changes only if the new variable, y_{ij} , is involved.

Algorithm 2 Greedy algorithm-information bottleneck

Input: $\Delta, P_{S,X}$
Initialization: $\mathcal{Y} = \mathcal{X}, P_{Y|X}(y|x) = \mathbf{1}\{y = x\}$
while there exists i', j' such that $I(S; Y^{i'-j'}) \geq \Delta$ **do**
among those i', j' , let
 $\{y_i, y_j\} = \arg \max_{y_{i'}, y_{j'} \in \mathcal{Y}} I(X; Y) - I(X; Y^{i'-j'})$
merge: $\{y_i, y_j\} \rightarrow y_{ij}$
update: $\mathcal{Y} = \{\mathcal{Y} \setminus \{y_i, y_j\}\} \cup \{y_{ij}\}$ and $P_{Y|X}$
Output: $P_{Y|X}$

The greedy algorithm is locally optimal at every step since we minimize $I(S; Y)$. However, there is no guarantee that such a greedy algorithm induces a global optimal privacy mapping.

Note 1. The minimum of $I(S; Y)$ in (2) is a decreasing function of $I(X; Y)$ and is achieved for a mapping $P_{Y|X}$ that satisfies $I(X; Y) = R$ (if possible due to discrete alphabets). For a given mutual information, R , there are many conditional probability distributions, $P_{Y|X}$, achieving $I(X; Y) = R$. Among which there is one that gives the minimum $I(S; Y)$ and one that gives the maximum $I(S; Y)$. We can modify the greedy algorithm so that it converges to a local maximum of $I(S; Y)$ for a given $I(X; Y) = R$. The algorithm which we call *greedy algorithm-information bottleneck* is given in Algorithm (2). Algorithm (1) and Algorithm (2) allow us to approximately characterize the range of values $I(S; Y)$ can take for a given value of $I(X; Y)$ as being those between the local minimum and the local maximum. Interestingly, by observing the gap between the local maximum and the local minimum, we have a relative idea on the effectiveness of the greedy algorithm, i.e., if the difference is significant it means a negligent mapping may lie anywhere between those values, possibly leading to a much higher privacy threat.

Data Set: The US 1994 Census dataset [19] is a well-known dataset in the machine learning community, which is a sample of the US population from 1994. For each of the entries, it contains features such as age, work-class, education, gender, and native country, as well as an income category. The income level is a binary variable which determines whether the income is above or below USD 50000, gender is a binary variable, education level is a variable with four categories, age is a variable divided into seven categories. For our purposes, we consider the private attributes $S = (\text{age, income level})$ and the attributes to be released as $X = (\text{age, gender, education level})$. The goal of the privacy mapping is to release a modified version of attributes Y which is informative about X but that renders the inference of S based on Y hard.

Numerical Results: In Fig. 1, we plot the minimum and maximum of $I(S; Y)$ for a given $I(X; Y)$. This figure is based on US 1994 census data set described before. The top curve shows the maximum of $I(S; Y)$ versus $I(X; Y)$, using Algorithm (2). The bottom curve shows the minimum of $I(S; Y)$ versus $I(X; Y)$, using Algorithm (1). The area between the two curves shows the possible pairs of $(I(X; Y), I(S; Y))$ as $P_{Y|X}$ varies (a subset of possible pairs, since the algorithms are sub-optimal). Indeed, we will design the mapping to lie




Fig. 1. Maximum and minimum of $I(S;Y)$ for a given $I(X;Y)$.

on the bottom curve. For a given R , if we design the mapping negligently, we may have $I(S;Y)$ on the top curve instead of the bottom curve.

V. THE GAUSSIAN CASE

In this section, assuming Gaussian $P_{S|X}$, we find the optimal Gaussian privacy mapping, $P_{Y|X}$. Let Σ_x and $\Sigma_{x,s}$ denote the covariance matrices of P_X and $P_{X,S}$. Let P_X be an n -dimensional Gaussian distribution and $P_{Y|X}$ be a Gaussian conditional distribution such that (X, Y) is jointly Gaussian. We can write Y in the innovation form, i.e., $Y = AX + Z$, where A is a full-rank $t \times n$ matrix, Z is a zero-mean Gaussian random variable independent from X (Theorem 2.3 of [20]), and t is the dimension of Y . Therefore, in the design of $P_{Y|X}$, we only require to find the matrix A , and the co-variance of Z . We use the approach of [21] to solve the information bottleneck problem for Gaussian case.

Remark 1. Let (S, X) have a jointly Gaussian distribution. For any $S = s$, the conditional distribution $P_{X|S=s}$ is a Gaussian distribution with co-variance $\Sigma_x - \Sigma_{x,s}\Sigma_s^{-1}\Sigma_{x,s}^t$, which we denote by $\Sigma_{x|s}$ (see [20], chapter 2).

Consider the Lagrangian of the optimization given in (2) as $\mathcal{L} = I(S;Y) - \beta I(X;Y)$, for some $\beta \in [0, 1]$. Consider the optimization $\min_{P_{Y|X}} I(S;Y) - \beta I(X;Y)$. We will find the optimal A and Z that achieves the optimal value for any β . By varying β , this provides the curve of $I(S;Y)$ versus $I(X;Y)$.

Theorem 2. Consider the optimization $\min_{P_{Y|X}} I(S;Y) - \beta I(X;Y)$. The optimal solution is characterized as $Y = AX + Z$, where $A = \text{diag}(M_1, \dots, M_t)V = MV$, $Z \sim \mathcal{N}(0, I)$, and V is a matrix containing the t left eigenvectors of $\Sigma_{x|s}\Sigma_x^{-1}$ corresponding to the t largest eigen-values for some t . M is the solution of the following optimization problem

$$\min_M \frac{1}{2}(1 - \beta) \log |M\Gamma M^t + I| - \frac{1}{2} \log |M\Lambda\Gamma M^t + I|,$$

where $\Gamma = \text{diag}(\Gamma_1, \dots, \Gamma_t) = V\Sigma_x V^t$ (we will show that $V\Sigma_x V^t$ is diagonal) and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_t)$ are the corresponding eigen-values of $\Sigma_{x|s}\Sigma_x^{-1}$.

The proof and the details of this section is given in the full version of paper [13].

VI. CONCLUSIONS

We study the privacy-utility trade-off against inference attacks when the log-loss is used both in the privacy and

utility metrics. We justify the generality of the privacy threat under the log-loss by proving that the threat under any bounded cost inference function can be upperbounded by an explicit function of the mutual information between private and disclosed data. We cast the tradeoff under the log-loss as the Privacy Funnel optimization, which is non-convex. We leverage its connection to the Information Bottleneck to design a locally-optimal greedy algorithm, that we evaluate on the US census dataset. Finally, we characterize the optimal privacy mapping for the Gaussian Privacy Funnel.

ACKNOWLEDGEMENT

The authors are grateful to Prof. Thomas Courtade, Prof. Kave Salamatian, and Prof. Tsachy Weissman for encouraging them to study the connections between privacy and the information bottleneck.

REFERENCES

- [1] F. du Pin Calmon and N. Fawaz, "Privacy against statistical inference," in *Allerton Conf. on Communication, Control, and Computing* 2012.
- [2] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," *arXiv preprint physics/0004057*, 2000.
- [3] N. Slonim and N. Tishby, "Agglomerative information bottleneck," *Proc. of Neural Information Processing Systems (NIPS-99)*, 1999.
- [4] 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.
- [5] R. Agrawal and R. Srikant, "Privacy-preserving data mining," *ACM Sigmod Record*, vol. 29, no. 2, pp. 439–450, 2000.
- [6] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography*. Springer, 2006.
- [7] C. Dwork, "Differential privacy," in *Automata, Languages and Programming*. Springer, 2006.
- [8] L. Sankar, S. R. Rajagopalan, and H. V. Poor, "Utility-privacy tradeoff in databases: An information-theoretic approach," *IEEE Trans. Inf. Forensics Security*, 2013.
- [9] A. D. Wyner, "The wire-tap channel," *Bell Syst. Tech. J.*, vol. 54, 1975.
- [10] T. A. Courtade and R. D. Wesel, "Multiterminal source coding with an entropy-based distortion measure," in *ISIT 2011*.
- [11] P. Harremoës and N. Tishby, "The information bottleneck revisited or how to choose a good distortion measure," in *ISIT 2007*.
- [12] N. Merhav and M. Feder, "Universal prediction," *Information Theory, IEEE Transactions on*, vol. 44, no. 6, pp. 2124–2147, 1998.
- [13] A. Makhdoumi, S. Salamatian, N. Fawaz, and M. Médard, "From the information bottleneck to the privacy funnel," *ArXiv e-prints*, 2014. [Online]. Available: <http://arxiv.org/>
- [14] S. Salamatian, N. Fawaz, B. Kveton, and N. Taft, "SPPM: Sparse Privacy Preserving Mappings," *Conference on Uncertainty in Artificial Intelligence*, 2014.
- [15] A. Makhdoumi and N. Fawaz, "Privacy utility trade-off under statistical uncertainty," in *Communication, Control, and Computing (Allerton), 2013 51th Annual Allerton Conference on*. IEEE, 2013.
- [16] S. Salamatian, A. Zhang, F. 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 *Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE*, 2013.
- [17] A. Zhang, S. Bhamidipati, N. Fawaz, and B. Kveton, "Priveview: Media consumption and recommendation meet privacy against inference attacks," in *IEEE Web 2.0 Security and Privacy Workshop (W2SP)*, 2014.
- [18] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [19] A. Asuncion and D. Newman, "UCI machine learning repository," 2007. [Online]. Available: <http://www.ics.uci.edu/~mllearn/MLRepository.html>
- [20] R. G. Gallager, *Discrete stochastic processes*. Kluwer Academic Publishers Boston, 1996, vol. 101.
- [21] G. Chechik, A. Globerson, N. Tishby, and Y. Weiss, "Information bottleneck for gaussian variables," in *Journal of Machine Learning Research*, 2005, pp. 165–188.