Privacy Against Brute-Force Inference Attacks

Seyed Ali Osia*, Borzoo Rassouli†, Hamed Haddadi‡, Hamid R. Rabiee*, Deniz Gündüz ‡
* Sharif University of Technology, osia@ce.sharif.edu, rabiee@sharif.edu
† University of Essex, b.rassouli@essex.ac.uk
‡ Imperial College London, {h.haddadi, d.gunduz}@imperial.ac.uk

***Abstract*—Privacy-preserving data release is about disclosing information about useful data while retaining the privacy of sensitive data. Assuming that the sensitive data is threatened by a brute-force adversary, we define *Guessing Leakage* as a measure of privacy, based on the concept of *guessing*. After investigating the properties of this measure, we derive the optimal utility-privacy trade-off via a linear program with any f-information adopted as the utility measure, and show that the optimal utility is a concave and piece-wise linear function of the privacy-leakage budget.**

I. Introduction

Large-scale data collection and analysis is a key component of many recent technological advances such as autonomous driving, online health monitoring, reliable energy grid, and intelligent IoT systems. While these data-driven applications provide better services by efficiently processing user data in massive scales, collection and sharing of personal data is increasingly creating privacy risks, especially with the advances in machine learning and data mining algorithms. Thus, developing privacy-preserving data release mechanisms that jointly consider the utility from shared data with the associated privacy leakage is key to the wide scale adoption of some of these emerging technologies. As a classic example, publishing a general purpose database, or even releasing its aggregate statistics, may threaten an individual's privacy, which has led to the introduction of widely used techniques such as k-anonymity [1], and differential privacy [2].

The information-theoretic formulation of privacy provides a general statistical framework to model both the utility and privacy, and allows investigating their trade-off as an optimization problem. Modeling the available dataset with random variable (r.v.) Y and the private/sensitive latent variable with X, the data that should be released, denoted by U is obtained as the output of a privacy-preserving statistical kernel (transformation), and can be obtained as the solution of an optimization problem. Mutual information is the most common measure of both utility and privacy, whereby the trade-off becomes the *privacy funnel* [3]. However, there is no optimal/universal definition of privacy and it can either abstract away from the adversarial threats [1], [4], or depend on the vulnerability of the sensitive data to adversarial attacks [5], [6]. In this paper, we assume that the adversary is a brute-force attacker, i.e., it performs an exhaustive trial and error attack over all the possible realizations of the sensitive data. There are many examples of such brute-force attacks in real life, where criminals steal private information by determining a cipher key via an exhaustive search [7], or checking several potential shortened URLs to discover active links [8].

We assume that a brute-force adversary performs a number of guesses to successfully determine the value of private data X, modeled as a r.v. with finite support. A privacy-preserving mechanism is built against this brute-force adversary. To this end, we define *guessing leakage*, and investigate its properties as a privacy measure. Afterwards, we formulate the utility-privacy trade-off as a non-convex optimization problem, and derive the optimal data release mechanism via a linear program (LP).

The problem of guessing is a well-established area in information theory. In [9], Massey proposed a lower bound on the minimum expected number of guesses needed to find X in terms of its Shannon entropy $H(X)$, while in [10], different moments of the guessing function are lower bounded in terms of the Renyi entropy of X, and the result is used to analyze the computational complexity of sequential decoding. Later works, such as [11], apply large deviation techniques to more general scenarios in this context. The problem of guessing also appears in Shannon-theoretic cryptography in [12], i.e., encryption against a brute-force wiretapper. The problem of computational security against a guessing attacker has been addressed in [7], [13]. The practical challenges of guessing passwords is studied in [14], [15]. However, limited attention has been devoted to the effect of data sharing on guessability of private information from a privacy-preserving point of view. In [16], a sub-optimal utility-privacy trade-off is found by using the lower bound provided in [10].

Notations. R.v.'s are denoted by capital letters, and their realizations by lower case letters. Matrices and vectors are denoted by bold capital and bold lower case letters, respectively. For a positive integer n, we define $[n] \triangleq \{1, 2, \dots, n\}$. For a finite alphabet $\mathcal{X}$, the probability simplex $\mathcal{P}(\mathcal{X})$ is the standard $(|\mathcal{X}|-1)$-simplex. Furthermore, to each probability mass function (pmf) $p_X(\cdot)$ corresponds a probability vector $\mathbf{p}_X \in \mathcal{P}(\mathcal{X})$, whose i-th element is $p_X(x_i)$ $(i \in [|\mathcal{X}|])$. Likewise, for a pair of r.v.'s (X, Y) with joint pmf $p_{X,Y}$, the probability vector $\mathbf{p}_{X|y}$ corresponds to the conditional pmf $p_{X|Y}(\cdot|y), \forall y \in \mathcal{Y}$, and $\mathbf{P}_{X|Y}$ is an $|\mathcal{X}| \times |\mathcal{Y}|$ matrix with columns $\mathbf{p}_{X|y}, \forall y \in \mathcal{Y}$. Statistical independence between X and Y is shown as $X \perp\!\!\!\perp Y$. For a convex function f such that $f(1) = 0$, and the probability mass functions p, q on $\mathcal{X}$, the

f-divergence is defined as[1] $D_f(p||q) \triangleq \sum_{x\in\mathcal{X}} q(x) f(\frac{p(x)}{q(x)})$. Finally, for a pair $(X,Y) \sim p_{X,Y}$, the f-information is defined as $I_f(X;Y) \triangleq D_f(p_{X,Y}||p_X \cdot p_Y)$.

II. System model

A. Preliminaries

Consider a triplet of discrete r.v.'s $(X,Y,W) \in \mathcal{X}\times\mathcal{Y}\times\mathcal{W}$, with finite alphabets, and distributed according to $p_{X,Y,W}$. Let Y and X denote the useful data to be revealed, and the sensitive data to be concealed, respectively. As in [17], W denotes what the user/curator directly observes, which may be a noisy representation of the pair (X,Y). Assume that the *privacy mapping/data release mechanism* takes W as input and maps it to the *released data* denoted by U. In this scenario, $(X,Y)-W-U$ form a Markov chain, and the privacy mapping is captured by the conditional distribution $p_{U|W}$.

The aim of the privacy mapping is to simultaneously preserve the fidelity of Y and the privacy of X by the release of U. In this paper, we measure the utility of the release by the f-information between Y and U, i.e., $I_f(Y;U)$, which incorporates mutual information as a special case, and define the privacy measure in the sequel.

B. Guessing

Consider the problem of guessing the realization of a discrete r.v. $X \in \mathcal{X}$ by asking questions of the form "Is X equal to x?", until the answer is "Yes." As in [10], a *guessing function/strategy* of X is denoted by a bijection $G(X): \mathcal{X} \to [|\mathcal{X}|]$. Hence, for a given guessing strategy $G(\cdot)$, $G(x)$ represents the number of required guesses when $X=x$. Likewise, for a pair of r.v.'s (X,Y), $G(X|Y): \mathcal{X}\times\mathcal{Y} \to [|\mathcal{X}|]$ is the guessing function of X given Y, i.e., $G(x|y)$ denotes the number of guesses required to determine $X=x$, when $Y=y$. For a given pmf p_X, let G^* denote the minimizer(s)[2] of $\mathbb{E}[G(X)]$, which guesses the value of X in decreasing order of probabilities as shown in [9]. The guessing entropy is defined as $H_G(X) \triangleq \mathbb{E}[G^*(X)]$, which, from the guesser's point of view, measures the uncertainty in X, as it denotes the minimum average number of guesses required to determine its realization. We have $H_G(X) \in [1, \frac{|\mathcal{X}|+1}{2}]$, where the minimum or maximum are attained when X is, respectively, deterministic or uniformly distributed over $\mathcal{X}$. Throughout the paper, the guessing entropy $H_G(X)$ and $H_G(\mathbf{p}_X)$ are written interchangeably[3].

[1]We assume that p is absolutely continuous with respect to q, i.e., $q(x)=0$ implies $p(x)=0$.

[2]It can be readily verified that G^* is unique if and only if $p_X(x_i) \neq p_X(x_j),\ \forall i \neq j$.

[3]It is noteworthy to mention that $H_G(X)$ as the minimum number of guesses is not consistent with intuition when $|\mathcal{X}|=1$ or 2, as in the former no guessing is needed, while $H_G(X)=1$, and in the latter, one guess is sufficient to determine the value of X, while $H_G(X)>1$. Although a more exact notion would be $H_G(X) \triangleq \min\{\mathbb{E}[G^*(X)], |\mathcal{X}|-1\}$, one can always justify $H_G(X)$ as the minimum number of guesses by emphasizing on the need of receiving an affirmative answer, i.e., until the answer is "Yes". Having said that, the analysis of this paper remains valid in either cases.

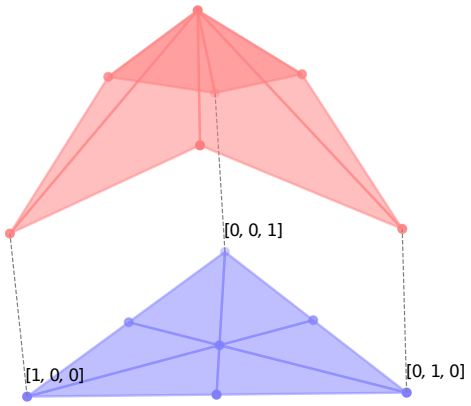

Fig. 1. Guessing entropy for a ternary r.v. X. The probability simplex $\mathcal{P}(\mathcal{X})$ is divided into $|\mathcal{X}|!=6$ rank partitions.

Definition 1. For a probability vector $\mathbf{p} \in \mathcal{P}(\mathcal{X})$, Let the *rank vector* of $\mathbf{p}$, denoted by $\mathbf{r}_\mathbf{p}$, be a vector that labels the elements of $[|\mathcal{X}|]$ according to their order induced by sorting $\mathbf{p}$ in descending order. For example, for $\mathbf{p}_1 = \begin{bmatrix}0.6 & 0.1 & 0.3\end{bmatrix}^T$, we have $\mathbf{r}_{\mathbf{p}_1} = \begin{bmatrix}1 & 3 & 2\end{bmatrix}^T$. For $\mathbf{p}_2 = \begin{bmatrix}0.5 & 0.25 & 0.25\end{bmatrix}^T$, either of $\begin{bmatrix}1 & 2 & 3\end{bmatrix}^T$ or $\begin{bmatrix}1 & 3 & 2\end{bmatrix}^T$ could be a candidate[4] for $\mathbf{r}_{\mathbf{p}_2}$.

Definition 2. The *rank partition* $\mathcal{P}_\mathbf{r}(\mathcal{X})$ is the set of all probability vectors $\mathbf{p} \in \mathcal{P}(\mathcal{X})$ with rank vector $\mathbf{r}$, where $\mathbf{r}$ is an arbitrary permutation of the elements in $[|\mathcal{X}|]$. Hence, the probability simplex $\mathcal{P}(\mathcal{X})$ can be divided into $|\mathcal{X}|!$ equal-rank partitions.

Proposition 1. *The guessing entropy $H_G(X)$ is a piece-wise linear and concave functional (see Figure 1) of $p_X(\cdot)$.*

Proof. Let $X \sim \mathbf{p}$, and it is immediate that $H_G(X) = \mathbf{r}_\mathbf{p}^T \cdot \mathbf{p}$, which proves the piece-wise linearity of $H_G(X)$ in $\mathbf{p}$. Furthermore, we have $\mathbf{r}_\mathbf{p}^T \cdot \mathbf{p} \leq \mathbf{r}^T \cdot \mathbf{p}$ for any rank vector $\mathbf{r} \neq \mathbf{r}_\mathbf{p}$, which follows from [9]. Hence, the concavity is proved as follows. For $\lambda \in [0,1]$ and two arbitrary $\mathbf{p},\mathbf{q} \in \mathcal{P}(\mathcal{X})$, let $\tilde{\mathbf{p}} \triangleq \lambda\mathbf{p} + (1-\lambda)\mathbf{q}$. We have,

$$\begin{aligned} H_G(\tilde{\mathbf{p}}) &= \mathbf{r}_{\tilde{\mathbf{p}}}^T \cdot \tilde{\mathbf{p}} \\ &= \lambda \mathbf{r}_{\tilde{\mathbf{p}}}^T \cdot \mathbf{p} + (1-\lambda)\mathbf{r}_{\tilde{\mathbf{p}}}^T \cdot \mathbf{q} \\ &\geq \lambda \mathbf{r}_{\mathbf{p}}^T \cdot \mathbf{p} + (1-\lambda)\mathbf{r}_{\mathbf{q}}^T \cdot \mathbf{q} \\ &= \lambda H_G(\mathbf{p}) + (1-\lambda) H_G(\mathbf{q}), \end{aligned}$$

where the inequality is strict if and only if $\mathbf{p},\mathbf{q}$ belong to different rank partitions. □

C. Brute-force inference attack

A brute-force adversary aims at inferring the private data, i.e., X, via trial and error, or equivalently, guessing. Here, we

[4]It can be verified that by imposing the constraint $r_i < r_j$ $(i \neq j)$ if and only if $p_i - p_j + \mathbb{1}\{p_i = p_j\}(j-i) > 0$, i.e., equal probabilities being ranked based on their index order, the rank vector is well defined. Nonetheless, the (possible) ambiguity in the definition of $\mathbf{r}$ is not problematic here.

adopt a model similar to the cost function-based inference in [18] by defining the cost of inference attack as the number of trials which lead to a correct guess. Prior to observing any realization of the released data U, the adversary uses the optimal guessing strategy $G^*(\cdot)$ to minimize its cost of inference, i.e., $C_0^* = H_G(X)$. After observing $U = u$, the adversary can update its belief about the private data as the posterior $p_{X|U}(\cdot|u)$, which in turn leads to an updated[5] guessing strategy $G^*(\cdot|u)$. Therefore, the minimum cost of inference is now $C_u^* = H_G(X|U = u)$. Considering the average additive gain (where the average is over U) that the adversary sees in its inference cost, we are ready to define the privacy leakage measure and investigate some of its properties as follows.

Definition 3. The *guessing leakage* is defined as:

$$GL(X \longrightarrow U) \triangleq H_G(X) - H_G(X|U). \tag{1}$$

Proposition 2. *We have*

$$GL(X \longrightarrow U) \geq 0, \tag{2}$$

with equality if and only if the rank vector associated with $p_{X|U}(\cdot|u)$, i.e., $\mathbf{r}_{\mathbf{p}_{X|u}}$, does not change with u, $\forall u \in \mathcal{U}$. Note that this is a weaker condition than statistical independence.

Proof. From Proposition 1, we have

$$\begin{aligned} H_G(X|U) &= \sum p(u) H_G(\mathbf{p}_{X|u}) \\ &\leq H_G\left(\sum p(u)\mathbf{p}_{X|u}\right) \\ &= H_G(X). \end{aligned}$$

where the inequality is tight if and only if all the $\mathbf{p}_{X|u}$'s ($\forall u$) belong to the same rank partition. Alternatively, (2) can be proved by [10, Corollary 1]. □

Remark 1. *The difference in (1) is reminiscent of the mutual information expanded in terms of entropies, which follows from its logarithmic nature. Nonetheless, one can observe that $GL(X \longrightarrow U)$ has no relation with mutual information in terms of one being a lower/upper bound to the other. Let X, U be binary random variables with $X|\{U = u_i\} \sim Bern(\frac{i}{i+3})$, $i = 1, 2$. It is immediate that $X \not\perp U$, and hence, $I(X;U) > 0$. Also, since we have $\mathbf{r}_{\mathbf{p}_{X|u_1}} = \mathbf{r}_{\mathbf{p}_{X|u_2}}$, we get $GL(X \longrightarrow U) = 0$. Therefore, $I(X;U) > GL(X \longrightarrow U)$. In another example, let $X = U$, which results in $I(X;U) = H(X)$, and $GL(X \longrightarrow U) = H_G(X) - 1$. According to [9, Section III], we can have distributions for which $H(X)$ is vanishingly small, while $H_G(X)$ is large. As a result, $I(X;U) < GL(X \longrightarrow U)$.*

A privacy measure can be investigated in terms of data processing inequalities.

Definition 4. A privacy measure $J(X;U)$ is said to satisfy the post-processing inequality [17] if and only if for any Markov chain $X - U - U'$, we have $J(X;U') \leq J(X;U)$, and the linkage inequality, if and only if $J(X;U') \leq J(U;U')$.

Theorem 1. *Guessing leakage satisfies the post-processing inequality, but not the linkage inequality.*

Proof. Consider the Markov chain $X - U - U'$. We have

$$\begin{aligned} H_G(X|U') &= \sum_{u'} p(u') H_G(X|U' = u') \\ &= \sum_{u'} p(u') H_G\left(\sum_u p(u|u')\mathbf{p}_{X|u}\right) \\ &\geq \sum_{u'} p(u') \sum_u p(u|u') H_G(\mathbf{p}_{X|u}) \qquad (3) \\ &= \sum_u p(u) H_G(\mathbf{p}_{X|u}) \\ &= H_G(X|U), \end{aligned}$$

where (3) follows from the concavity of H_G. Hence, we have

$$GL(X \longrightarrow U') \leq GL(X \longrightarrow U).$$

In other words, no independent processing of the released data, U, can increase the privacy leakage.

To show that GL does not satisfy the linkage inequality, let $U' \sim \text{Bern}(\theta)$ for some $\theta \in (0, 1)$. Also, let $\mathcal{U} = [3]$ with $\mathbf{p}_{U|u'_0} = [.35\ .4\ .25]^T$ and $\mathbf{p}_{U|u'_1} = [.3\ .6\ .1]^T$. By setting $X = U \bmod 2$, we have $X - U - U'$ form a Markov chain, and $GL(X \longrightarrow U') > GL(U \longrightarrow U') = 0$. □

III. Utility-privacy trade-off

Having defined the utility and privacy measures, the utility-privacy trade-off can be written as[6]:

$$t(\epsilon) \triangleq \sup_{\substack{p_{U|W}: \\ (X,Y)-W-U \\ GL(X \longrightarrow U) \leq \epsilon}} I_f(Y;U), \tag{4}$$

where $\epsilon \in [0, GL(X \longrightarrow W)]$. By the routine application of cardinality bounding techniques [19], it becomes sufficient to have $|\mathcal{U}| \leq |\mathcal{W}| + 1$. Furthermore, the supremum can be replaced by maximum, since a continuous objective function attains its supremum over a compact set. Also, searching over $p_{U|W}$ can be equivalent[7] to searching over the pair $(p_U, \mathbf{p}_{W|u})$, such that $\mathbf{p}_W = \sum_u p_U(u)\mathbf{p}_{W|u}$. Therefore, the problem reduces to

$$t(\epsilon) = \max_{\substack{p(\cdot), \mathbf{p}_{W|u} \in \mathcal{P}(\mathcal{W}): \\ \sum p(u)\mathbf{p}_{W|u} = \mathbf{p}_W, \\ \sum p(u) H_G\left(\mathbf{P}_{X|W}\mathbf{p}_{W|u}\right) \geq H_G(X) - \epsilon}} \sum_u p(u) D_f\left(\mathbf{P}_{Y|W}\mathbf{p}_{W|u} || \mathbf{p}_Y\right). \tag{5}$$

For an arbitrary rank vector $\mathbf{r}$, let $\mathcal{Q}_{\mathbf{r}}(\mathcal{W})$ denote the inverse image of $\mathcal{P}_{\mathbf{r}}(\mathcal{X})$ under $\mathbf{P}_{X|W}$, which is a linear transformation from $\mathcal{P}(\mathcal{W})$ to $\mathcal{P}(\mathcal{X})$. For a given $\mathbf{r}$, $\mathcal{Q}_{\mathbf{r}}(\mathcal{W})$ is a convex polytope with a finite number of extreme points, since it can

[5]Note that knowledge of the posterior is not necessary in the sense that the adversary could also get the same updated strategy by only knowing the rank vector associated to it, i.e., $\mathbf{r}_{\mathbf{p}_{X|u}}$.

[6]We assume that $Y \not\perp W$, since otherwise, $t(\epsilon) = 0$, $\forall \epsilon$.

[7]Trivially, for any pair $(p_U, \mathbf{p}_{W|u})$, $(X, Y) - W - U$ can be constructed.

be written as the intersection of a finite number of closed half-spaces in $\mathcal{P}(\mathcal{W})$[8]. For example, with $\mathbf{r} = [1\ 3\ 2]^T$, we have $\mathcal{Q}_{\mathbf{r}}(\mathcal{W}) = \{\mathbf{p} \in \mathcal{P}(\mathcal{W}) | \mathbf{v}_1 \cdot \mathbf{p} \geq \mathbf{v}_3 \cdot \mathbf{p},\ \mathbf{v}_3 \cdot \mathbf{p} \geq \mathbf{v}_2 \cdot \mathbf{p}\}$, which forms the intersection of $\mathcal{P}(\mathcal{W})$ with two closed half-spaces, where $\mathbf{v}_i$ denotes the ith row of $\mathbf{P}_{X|W}$. Let $\mathbb{Q}_r$ denote the set of extreme points of $\mathcal{Q}_{\mathbf{r}}(\mathcal{W})$. As a result, any element of $\mathcal{Q}_{\mathbf{r}}(\mathcal{W})$ can be written as a convex combination of the elements of $\mathbb{Q}_r$.

In what follows, we show that there is no loss of optimality in replacing $\mathbf{p}_{W|u} \in \mathcal{P}(\mathcal{W})$ in (5) with $\mathbf{p}_{W|u} \in \mathbb{Q}$, where $\mathbb{Q} \triangleq \cup_r \mathbb{Q}_r$. It is already known that the f-divergence, $D_f(p||q)$, is convex in (p, q), and in p for a fixed q, which implies the convexity of the objective function of (5) in $\mathbf{p}_{W|u}$. For an arbitrary $\mathbf{p}_{W|u} \in \mathcal{P}(\mathcal{W})$, we have $\mathbf{p}_{W|u} \in \mathcal{Q}_{\mathbf{r}}(\mathcal{W})$ for some $\mathbf{r}$. Writing this $\mathbf{p}_{W|u}$ as a convex combination of the elements of $\mathbb{Q}_r$ does not alter $H_G(\cdot)$, which follows from the piece-wise linearity of $H_G(\cdot)$, i.e., $H_G\left(\mathbf{P}_{X|W}\mathbf{z}\right)$ is linear in $\mathbf{z}$ for $\mathbf{z} \in \mathcal{Q}_{\mathbf{r}}(\mathcal{W})$. Furthermore, this does not decrease the objective function, which is a direct consequence of the convexity of $D_f\left(\mathbf{P}_{Y|W}\mathbf{z}||\mathbf{p}_Y\right)$ in $\mathbf{z}$. Therefore, in (5), $\mathbf{p}_{W|u} \in \mathcal{P}(\mathcal{W})$ can be replaced with $\mathbf{p}_{W|u} \in \mathbb{Q}$, which leads to the following theorem.

Theorem 2. *The utility-privacy trade-off in (5) can be solved by a linear program (LP).*

Proof. Let $\{\mathbf{q}_i\}_{i=1}^k$ denote the elements of $\mathbb{Q}$. The problem reduces to

$$\begin{aligned} t(\epsilon) = \max_{p(\cdot)\geq 0} & \sum_{i=1}^{k} p(u_i)\ D_f\left(\mathbf{P}_{Y|W}\mathbf{q}_i||\mathbf{p}_Y\right) \\ \text{s.t.} & \sum_{i=1}^{k} p(u_i)\ \mathbf{r}^T_{(\mathbf{P}_{X|W}\mathbf{q}_i)}\mathbf{P}_{X|W}\mathbf{q}_i \geq H_G(X) - \epsilon, \\ & \sum_{i=1}^{k} p(u_i)\mathbf{q}_i = \mathbf{p}_W, \end{aligned} \tag{6}$$

which is an LP[9]. □

Corollary 1. *The utility-privacy trade-off in (4) is a concave and piece-wise linear function of ϵ (see Fig. 2).*

Proof. This follows from the LP sensitivity analysis [20, Lemma 2]. □

Remark 2. *If the elements of $\mathbf{p}_X$ are distinct, it is always possible to reveal some information with no privacy leakage, i.e., $t(0) > 0$. Moreover, if all the columns of $\mathbf{P}_{X|W}$ belong to the same rank partition, no data release can cause any leakage of privacy, i.e., $t(\epsilon) = I_f(W;W), \forall \epsilon$.*

Proof. If the elements of $\mathbf{p}_X$ are distinct, we have $\mathbf{p}_X \in \text{int}(\mathcal{P}_{\mathbf{r}}(\mathcal{X}))$ for some $\mathbf{r}$. Since Y and W are not independent, there exists a vector $\mathbf{v}$ such that i) $\mathbf{p}_Y \neq \mathbf{P}_{Y|W}(\mathbf{p}_W + \eta\mathbf{v})$ for sufficiently small $\eta > 0$ and ii) $\mathbf{1}^T \cdot \mathbf{v} = 0$. Let $U' \sim \text{Bern}(\frac{1}{2})$, and set $\mathbf{p}_{W|u'_i} = \mathbf{p}_W + (-1)^i\eta\mathbf{v}$ for $i = 1, 2$. As a result, $p_{Y,U'} \neq p_Y \cdot p_{U'}$, which results in $I_f(Y;U') > 0$. Also, for sufficiently small η, $\mathbf{p}_{X|u'_i}$ $(= \mathbf{P}_{X|W}\mathbf{p}_{W|u'_i}, i = 1, 2)$ lie in the same rank partition as $\mathbf{p}_X$, which results in $GL(X \longrightarrow U') = 0$. Therefore, we have $t(0) \geq I_f(Y;U') > 0$ The second claim is proved by setting $U = W$, using the data-processing inequality for f-information, and noting that $GL(X \longrightarrow W) = 0$. □

Remark 3. *In order to specify the LP for a particular setting, it is sufficient to identify the elements of $\mathbb{Q}$. The procedure of finding the extreme points of a convex polytope is a classical problem, which is omitted here due to lack of space. In the special case of $W = X$, i.e., when the curator has direct access only to the sensitive data, these points are already known as $(1, 0, \ldots, 0)^T$, $(\frac{1}{2}, \frac{1}{2}, 0, \ldots, 0)^T$, $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, \ldots, 0)^T$ and so on, and their permutations, which are nothing but the extreme points of all the rank partitions of $\mathcal{P}(\mathcal{W})$.*

Remark 4. *Theorem 2 is a direct consequence of the convexity of the objective function, and the piece-wise linearity of the privacy measure. Hence, a similar analysis applies when the utility is measured by the minimum mean square error (MMSE), minimum probability of error, or $GL(Y \longrightarrow U)$.*

Remark 5. *Guessing leakage is defined in (1) as the additive gain in the inference cost of a brute-force attacker. Alternatively, one could define the following multiplicative gain as the privacy-leakage measure:*

$$GL_m(X \longrightarrow U) \triangleq \frac{H_G(X)}{H_G(X|U)}, \tag{7}$$

which belongs to $[1, \frac{|\mathcal{X}|+1}{2}]$. It can be readily verified that the two theorems of this paper remain valid when (7) replaces (1). When the adversary is a memoryless guesser, i.e., each new guess is independent of the previous guesses, we have

$$\log_2 GL_m(X \longrightarrow U) = I_{\frac{1}{2}}(X;U),$$

where $I_\alpha(X;U)$ denotes the Arimoto mutual information of order $\alpha(\geq 0)$, which follows from [21, Theorem 1]. However, this does not hold in general, and is only pertinent to the special case of a memoryless guesser.

IV. Numerical Results

Suppose an attacker has access to a list of user IDs that potentially belong to individuals with highly sensitive job positions. The goal of the attacker is to discover whether any of these IDs are related to a person working for a list of presumed government agencies. Each organization has a registered domain name, hence the attacker can build all of the possible email addresses by combining user IDs with domain names, and exhaustively checking the existence of emails, which can potentially be used for consequent social engineering. Assuming that checking a large number of emails is a risky task for the attacker, they will try to minimize the number of attempts. Again, assuming the attacker has access

[8]Depending on the transformation $\mathbf{P}_{X|W}$, $\mathcal{Q}_{\mathbf{r}}(\mathcal{W})$ could be an empty set for some values of $\mathbf{r}$. Note that the empty set can be viewed as a null polytope.

[9]Note that the constraint $\sum_i p(u_i) = 1$ is redundant, as it is implicitly implied by $\sum_i p(u_i)\mathbf{q}_i = \mathbf{p}_W$.

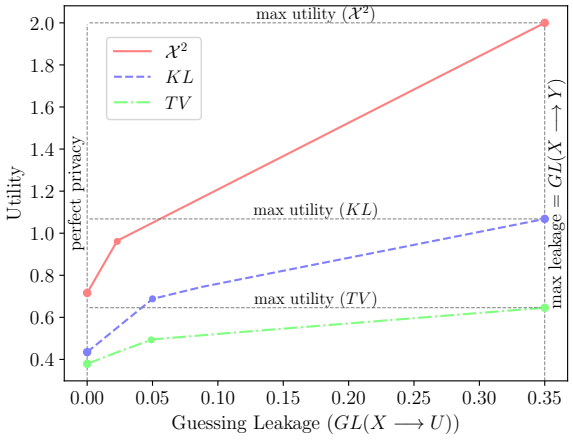

Fig. 2. Optimal utility-privacy trade-off for three different utility measures.

to the location check-ins of users in a social network and also knows the exact location of all branches of the organizations, they can wisely reorder the emails and check them in a manner that minimizes the number of trials.

Based on this scenario, suppose we are interested in three organizations, x_1, x_2 and x_3, with 500, 300 and 200 employees, respectively. Let X denote the organization variable with probability vector $\mathbf{p}_X = [0.5, 0.3, 0.2]$. Considering 100 user IDs and without any side information, the average number of trials is 170 ($= 100 * H_G(X)$). Now assume we have 3 cities, y_1, y_2 and y_3, and each organization has offices in 2 of them (for example, x_1 has 300 employees in y_1 and 200 employees in y_2, x_2 has 150 employees in y_1 and 150 employees in y_3, and x_3 has 60 employees in y_2 and 140 employees in y_3.). In this way, we have defined $P_{Y|X}$ where Y denotes the city variable. If the attacker knows the city variable for all users, the minimum number of trials is $100 * H_G(X|Y) = 135$, which is less than 170 ($GL(X \longrightarrow Y) = 0.35$). Hence, a privacy preserving mapping is needed for check-ins of the people with sensitive positions, which can be done by intentionally creating wrong check-ins. This is equivalent to sampling U from $P_{U|Y}$, which obviously increases the level of privacy while it decreases the correctness of location check-ins (utility). Therefore, we should build a trade-off between utility and leakage. Figure 2 shows the optimal utility-privacy trade-off for this example, when the utility is measured by three variants of f-information corresponding to the $\mathcal{X}^2$-divergence, KL-divergence, and total variation (TV) distance.

V. Conclusions and Future Work

We considered the problem of privacy against brute-force adversaries. By investigating the properties of guessing entropy, we introduced *guessing leakage* as a privacy measure. We studied the optimal utility-privacy trade-off with f-information as the utility measure, and showed it to be the solution of an LP. Unless the curator has direct access only to the private data, we need to identify the extreme points of a convex polytope, whose complexity grows exponentially. Hence, sub-optimal algorithms are to be sought, such as restricting the search space to the set of all deterministic mappings, which is the subject of our ongoing work. Another practical direction is to address this problem when the curator is uncertain (or even unaware) of the underlying distribution.

References

[1] L. Sweeney, "k-anonymity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, 2002.

[2] C. Dwork, "Differential privacy," in *International Colloquium on Automata, Languages and Programming*, 2006, pp. 1–12.

[3] A. Makhdoumi, S. Salamatian, N. Fawaz, and M. Médard, "From the information bottleneck to the privacy funnel," in *Information Theory Workshop (ITW), 2014 IEEE*. IEEE, 2014, pp. 501–505.

[4] B. Rassouli and D. Gündüz, "Optimal utility-privacy trade-off with total variation distance as a privacy measure," in *2018 IEEE Information Theory Workshop (ITW)*, Nov 2018, pp. 1–5.

[5] J. Liao, O. Kosut, L. Sankar, and F. P. Calmon, "A tunable measure for information leakage," in *2018 IEEE International Symposium on Information Theory (ISIT)*, June 2018, pp. 701–705.

[6] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, "Estimation efficiency under privacy constraints," *IEEE Transactions on Information Theory*, 2018.

[7] M. M. Christiansen, K. R. Duffy, F. du Pin Calmon, and M. Médard, "Multi-user guesswork and brute force security," *IEEE Transactions on Information Theory*, vol. 61, no. 12, pp. 6876–6886, 2015.

[8] C.-K. Chu, W.-T. Zhu, J. Han, J. K. Liu, J. Xu, and J. Zhou, "Security concerns in popular cloud storage services," *IEEE Pervasive Computing*, vol. 12, no. 4, pp. 50–57, 2013.

[9] J. L. Massey, "Guessing and entropy," in *Proceedings of 1994 IEEE International Symposium on Information Theory*, June 1994, pp. 204–.

[10] E. Arikan, "An inequality on guessing and its application to sequential decoding," *IEEE Transactions on Information Theory*, vol. 42, no. 1, pp. 99–105, Jan 1996.

[11] M. K. Hanawal and R. Sundaresan, "Guessing revisited: A large deviations approach," *IEEE Transactions on Information Theory*, vol. 57, no. 1, pp. 70–78, 2011.

[12] N. Merhav and E. Arikan, "The shannon cipher system with a guessing wiretapper," *IEEE Transactions on Information Theory*, vol. 45, no. 6, pp. 1860–1866, 1999.

[13] A. Beirami, R. Calderbank, K. Duffy, and M. Médard, "Quantifying computational security subject to source constraints, guesswork and inscrutability," in *Information Theory (ISIT), 2015 IEEE International Symposium on*. IEEE, 2015, pp. 2757–2761.

[14] P. G. Kelley, S. Komanduri, M. L. Mazurek, R. Shay, T. Vidas, L. Bauer, N. Christin, L. F. Cranor, and J. Lopez, "Guess again (and again and again): Measuring password strength by simulating password-cracking algorithms," in *Security and Privacy (SP), 2012 IEEE Symposium on*. IEEE, 2012, pp. 523–537.

[15] B. Ur, S. M. Segreti, L. Bauer, N. Christin, L. F. Cranor, S. Komanduri, D. Kurilova, M. L. Mazurek, W. Melicher, and R. Shay, "Measuring real-world accuracies and biases in modeling password guessability." in *USENIX Security Symposium*, 2015, pp. 463–481.

[16] Y. Rachlin, K. Probst, and R. Ghani, "Maximizing privacy under data distortion constraints in noise perturbation methods," in *Privacy, security, and trust in KDD*. Springer, 2009, pp. 92–110.

[17] Y. Wang, Y. O. Basciftci, and P. Ishwar, "Privacy-utility tradeoffs under constrained data release mechanisms," *CoRR*, vol. abs/1710.09295, 2017. [Online]. Available: http://arxiv.org/abs/1710.09295

[18] F. du Pin Calmon and N. Fawaz, "Privacy against statistical inference," in *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*. IEEE, 2012, pp. 1401–1408.

[19] A. E. Gamal and Y.-H. Kim, *Network Information Theory*. New York, NY, USA: Cambridge University Press, 2012.

[20] B. Jansen, J. De Jong, C. Roos, and T. Terlaky, "Sensitivity analysis in linear programming: just be careful!" *European Journal of Operational Research*, vol. 101, no. 1, pp. 15–28, 1997.

[21] W. Huleihel, S. Salamatian, and M. Médard, "Guessing with limited memory," in *2017 IEEE International Symposium on Information Theory (ISIT)*, June 2017, pp. 2253–2257.