An Operational Measure of Information Leakage

Ibrahim Issa[1], Sudeep Kamath[2], and Aaron B. Wagner[1]
[1]School of Electrical and Computer Engineering, Cornell University, Ithaca, New York
[2] Department of Electrical Engineering, Princeton University, Princeton, New Jersey
Emails: ii47@cornell.edu, sukamath@princeton.edu, wagner@cornell.edu

***Abstract*—Given two discrete random variables X and Y, an operational approach is undertaken to quantify the "leakage" of information from X to Y. The resulting measure $\mathcal{L}(X \to Y)$ is called *maximal leakage*, and is defined as the multiplicative increase, upon observing Y, of the probability of correctly guessing a randomized function of X, maximized over all such randomized functions. It is shown to be equal to the Sibson mutual information of order infinity, giving the latter operational significance. Its resulting properties are consistent with an axiomatic view of a leakage measure; for example, it satisfies the data processing inequality, it is asymmetric, and it is additive over independent pairs of random variables. Moreover, it is shown that the definition is robust in several respects: allowing for several guesses or requiring the guess to be only within a certain distance of the true function value does not change the resulting measure.**

***Index Terms*—Leakage, Privacy, Sibson mutual information, Inference**

I. Introduction

Given two discrete random variables X and Y, how much information does Y *leak* about X? This basic question arises in many secrecy and privacy problems, in which X represents sensitive information and Y represents information available to an adversary. A quantitative answer to that question is necessary to assess the performance of privacy systems for which Y cannot be made independent of X, which is often the case in practice. For example, a curator might want to reveal statistical data about a given population without compromising the privacy of its individuals [1]–[3]. An adversary could also gain access to Y through a side channel [4]–[7], or through a wiretap [8,9]. Perfectly securing these channels, if even possible, could be highly detrimental to the performance of the underlying system.

Moreover, the question is interesting from a purely theoretical point of view, as it is akin to fundamental questions in information theory such as "how much information does X contain?" The fact that Shannon answered the latter question with the entropy of X, $H(X)$, might explain why many works [8]–[15] have adopted equivocation, $H(X|Y)$, as a measure of privacy. However, this choice overlooks the context in which these questions were posed. Whereas Shannon's motivating problem was finding the minimum number of bits required to *describe* X, the goal in privacy problems is different. It also fails to capture the fact that the adversary might be interested in functions of X, or other random variables dependent on X [16,17], rather than X itself.

In this paper, we give an operational definition of leakage that is motivated by the setup of a *guessing* adversary. More specifically, upon observing Y, the adversary tries to guess a (possibly randomized) function of X. Leakage for a specific function is considered to be the logarithm of the ratio of the probability of a correct guess when Y is observed, to the probability of a correct guess when it is not (i.e., a blind guess). Maximal leakage, which we denote by $\mathcal{L}(X \to Y)$, is then defined as the maximum leakage over all such randomized functions. This maximization, which is formally over discrete random variables U for which the Markov chain $U - X - Y$ holds, represents a worst-case analysis on the function of interest U, and models scenarios in which the conditional distribution $P_{U|X}$ is unknown. It is also inspired by the strong data processing constant [18].

Although the maximization is an infinite-dimensional problem, $\mathcal{L}(X \to Y)$ admits a closed-form solution. It turns out to equal the Sibson mutual information of order infinity $I_\infty(X;Y)$ [19,20], endowing it with an operational significance. Several desirable properties for a leakage measure then follow: it is zero if and only if X and Y are independent, it is not symmetric, it satisfies the data processing inequality, and it is additive over independent pairs $\{(X_i, Y_i)\}$.

We provide a conditional probability law $P_{U|X}$ that achieves the maximum and depends on the joint probability P_{XY} only through its $X-$marginal, P_X. In particular, $P_{U|X}$ is such that: for distinct x's, the supports of $P_{U|X=x}$'s are disjoint, and each $P_{U|X=x}$ effectively "shatters" the atom x into (almost) uniformly distributed u's to get an (almost) uniform marginal P_U. Moreover, we show that, in general, there is no deterministic law $P_{U|X}$ that achieves the maximum. Indeed, we could have X and Y such that $\mathcal{L}(X \to Y) > 0$, whereas observing Y does not affect the probability of guessing any deterministic function of X.

Furthermore, we show that the definition of maximal leakage is robust in several respects. In the definition of $\mathcal{L}(X \to Y)$, we allow the adversary *one* guess only. A natural extension would be to allow for, say, k guesses for some integer k. This is particularly relevant for privacy problems. For example, if U is a password to some system, then an adversary is typically allowed several wrong guesses before he/she is possibly locked out. We call the modified measure k-maximal leakage, and denote it by $\mathcal{L}^{(k)}(X \to Y)$. We show that, in fact, the two definitions are equivalent for all k.

Finally, we consider the case in which the adversary only needs the guess to be within a certain distance of the true function value, according to an arbitrary distance metric. As such, the random variable U, over which we are optimizing, now lives in a given metric space $\mathcal{U}$ and is no longer restricted to be discrete. We call this modified measure maximal locational leakage, and we denote it by $\mathcal{L}_{\mathcal{U}}(X \to Y)$. We show that $\mathcal{L}_{\mathcal{U}}(X \to Y) \leq \mathcal{L}(X \to Y)$, and equality holds under an unboundedness condition on the metric space $\mathcal{U}$.

II. Related Leakage Metrics

The literature on leakage and privacy measures is vast, spanning the fields of information theory, computer science, and computer security. The closest to our work comes from computer security in [21]–[24]. In particular, [21] defines leakage from X to Y as the logarithm of the multiplicative increase, upon observing Y, of the probability of guessing X *itself* correctly, neglecting that the adversary might be interested in certain functions of X. [22] considers a *worst case* approach, and maximizes the previous quantity over all distributions on the alphabet of X (while $P_{Y|X}$ is fixed). The resulting quantity turns out to equal $\mathcal{L}(X \to Y)$ and is called in the computer security literature "maximal leakage" as well, or min-capacity (the term min-capacity is slightly misleading as $\mathcal{L}(X \to Y)$ is in fact greater than or equal to the Shannon capacity of the channel defined by $P_{Y|X}$ [23]. However, the term "min" was used because the min-entropy appears when computing probabilities of correct guesses). It is denoted by $ML(P_{Y|X})$, and its properties were further studied in [23,24]. [25,26] investigate relationships between maximal leakage and differential privacy [27], which is the most widely adopted metric in database security. Roughly speaking, differential privacy requires that, for any two *neighboring* databases, the probabilities of any given output do not differ significantly.

Another connected line of work stems from cryptography, and in particular from the notion of *semantic security* [28] which considers the security of encryption schemes. First, [28] introduces the notion of "advantage" for a given function of the messages. It is the *additive* increase, upon observing the encrypted message (i.e., the ciphertext), of the probability of correctly guessing the value of the function. In our framework, "advantage" is defined as the multiplicative increase. Since one is typically interested in securing hard-to-guess functions for which the probability of a correct guess is small (since, otherwise, we are already "doomed"), the multiplicative increase is arguably more descriptive of the change. It also makes more intuitive sense when viewing leakage in terms of leaked *bits*. Semantic security then requires that, for an adversary that can work only for a polynomial (in the length of the message) amount of time, the advantage is negligible for all *deterministic* functions that are computable in polynomial time, and for all input distributions.

There are several variants of semantic security. In particular, entropic security [29,30] drops the computational bounds (on the adversary and the considered functions), but restricts its attention to input distributions with high min-entropy. [31] introduces semantic security to the wiretap channel, and does not restrict it to computationally bounded adversaries, nor deterministic polynomial-time computable functions. For a given encryption scheme, [31] then upper and lower-bounds the advantage of semantic security in terms of – what the authors call – mutual information security advantage, which is defined as the maximum, over all input distributions, of the mutual information between the message and the output of the channel whose input is the encryption of the message. Moreover, for discrete random variables X and Y, [32] upper-bounds the advantage over all deterministic functions in terms of their maximal correlation, which inspired [33] to use the latter quantity as a secrecy metric. [32], inspired by the correspondence analysis literature [34], also generalized maximal correlation to k-correlation, which is defined as the sum of the k largest principal inertial components of the joint distribution P_{XY}.

Finally, rate-distortion-based approaches to privacy metrics can be found in [35]–[39]. Although the particular metrics differ among those works (e.g., expected distortion, probability of a guess satisfying the distortion constraint, etc.), they all assume that there is a known distortion function up to which the adversary is interested in the sensitive information X. For further discussion of privacy metrics, we refer the reader to [40], which categorizes over eighty such metrics.

III. Maximal Leakage

Let X and Y be two discrete random variables, with alphabets $\mathcal{X}$ and $\mathcal{Y}$ respectively. We denote by P_{XY} the joint distribution of (X,Y).

Definition 1 (Maximal Leakage): Given a joint distribution P_{XY} on finite alphabets $\mathcal{X}$ and $\mathcal{Y}$, the maximal leakage from X to Y is defined as

$$\mathcal{L}(X\to Y) = \sup_{U-X-Y-\hat{U}} \log \frac{\mathbf{Pr}\Big(U=\hat{U}\Big)}{\max_{u\in\mathcal{U}} P_U(u)},$$

where U and $\hat{U}$ take values in the same finite alphabet.

We can rewrite $\mathcal{L}(X\to Y)$ as

$$\mathcal{L}(X\to Y) = \sup_{U-X-Y} \log \frac{\sum_{y\in\mathcal{Y}} \max_{u\in\mathcal{U}} P_{UY}(u,y)}{\max_{u\in\mathcal{U}} P_U(u)}. \tag{1}$$

Our main theorem is the characterization of maximal leakage as follows.

Theorem 1: For any joint distribution P_{XY} on finite alphabets $\mathcal{X}$ and $\mathcal{Y}$, the maximal leakage from X to Y is given by

$$\mathcal{L}(X\to Y) = \log \sum_{y\in\mathcal{Y}} \max_{\substack{x\in\mathcal{X}:\\ P_X(x)>0}} P_{Y|X}(y|x).$$

Proof: Assume, without loss of generality, that $P_X(x) > 0$ for all $x \in \mathcal{X}$. Note that the right-hand side is equal to $I_\infty(X;Y)$ [19,20]. To show that $\mathcal{L}(X \to Y) \le I_\infty(X;Y)$, consider any U satisfying $U - X - Y$. Let

$$\mathcal{L}(X\to Y)[U] = \log \frac{\sum_{y\in\mathcal{Y}} \max_{u\in\mathcal{U}} P_{UY}(u,y)}{\max_{u\in\mathcal{U}} P_U(u)}, \tag{2}$$

so that $\mathcal{L}(X\to Y) = \sup_{U:U-X-Y} \mathcal{L}(X\to Y)[U]$. Then,

$$\begin{aligned}
&\sum_{y\in\mathcal{Y}} \max_{u\in\mathcal{U}} P_{UY}(u,y)\\
&= \sum_{y\in\mathcal{Y}} \max_{u\in\mathcal{U}} \sum_{x\in\mathcal{X}} P_X(x)P_{U|X}(u|x)P_{Y|X}(y|x)\\
&\le \sum_{y\in\mathcal{Y}} \max_{u\in\mathcal{U}} \sum_{x\in\mathcal{X}} P_X(x)P_{U|X}(u|x) \max_{x'\in\mathcal{X}} P_{Y|X}(y|x')\\
&= \sum_{y\in\mathcal{Y}} \left(\max_{x'\in\mathcal{X}} P_{Y|X}(y|x')\right) \max_{u\in\mathcal{U}} \sum_{x\in\mathcal{X}} P_X(x)P_{U|X}(u|x)\\
&= \sum_{y\in\mathcal{Y}} \max_{x\in\mathcal{X}} P_{Y|X}(y|x) \max_{u\in\mathcal{U}} P_U(u).
\end{aligned}$$

Therefore, $\mathcal{L}(X\to Y)[U] \le I_\infty(X;Y)$ for all $P_{U|X}$, hence $\mathcal{L}(X\to Y) \le I_\infty(X;Y)$.

For the reverse inequality, we construct a $P_{U|X}$ for which $\mathcal{L}(X\to Y)[U] = I_\infty(X;Y)$. To that end, let $p^\star = \min_{x\in\mathcal{X}} P_X(x)$. For each $x\in\mathcal{X}$, let $k(x) = P_X(x)/p^\star$, and let $\mathcal{U} = \bigcup_{x\in\mathcal{X}} \{(x,1),(x,2),\ldots,(x,\lceil k(x)\rceil)\}$. For each $u=(i_u,j_u)\in\mathcal{U}$ and $x\in\mathcal{X}$, let $P_{U|X}(u|x)$ be:

$$\begin{aligned}
&P_{U|X}((i_u,j_u)|x)\\
&= \begin{cases} \frac{p^\star}{P_X(x)}, & i_u = x, \quad 1\le j_u \le \lfloor k(x)\rfloor,\\ 1-\frac{(\lceil k(x)\rceil -1)p^\star}{P_X(x)}, & i_u = x, \quad j_u = \lceil k(x)\rceil,\\ 0, & i_u \ne x, \quad 1\le j_u\le \lceil k(i_u)\rceil. \end{cases}
\end{aligned}$$

Remark 1: It is easy to check that if $\lfloor k(x)\rfloor = \lceil k(x)\rceil$, then the corresponding formulas are equal.

Then, for each $((i_u,j_u),x)\in\mathcal{U}\times\mathcal{X}$,

$$\begin{aligned}
&P_{UX}((i_u,j_u),x)\\
&=\begin{cases} p^\star, & i_u = x, 1\le j_u\le\lfloor k(x)\rfloor,\\ P_X(x)-(\lceil k(x)\rceil-1)p^\star, & i_u=x, j_u=\lceil k(x)\rceil,\\ 0, & i_u\ne x, 1\le j_u\le\lceil k(i_u)\rceil. \end{cases}
\end{aligned}$$

As mentioned in the introduction, the supports of $P_{U|X=x}$ are disjoint for distinct x's, and each x is effectively shattered into shards of probability $p^\star$. Now, note that

$$\max_{u\in\mathcal{U}} P_U(u) = \max_{(i_u,j_u)\in\mathcal{U}} P_{UX}((i_u,j_u),i_u) = p^\star. \tag{3}$$

Now, consider any $(u,y)\in\mathcal{U}\times\mathcal{Y}$. We have

$$\begin{aligned}
&P_{UY}((i_u,j_u),y)\\
&= \sum_{x\in\mathcal{X}} P_X(x)P_{U|X}((i_u,j_u)|x)P_{Y|X}(y|x)\\
&= P_X(i_u)P_{U|X}((i_u,j_u)|i_u)P_{Y|X}(y|i_u)\\
&=\begin{cases} p^\star P_{Y|X}(y|i_u), & 1\le j_u\le\lfloor k(i_u)\rfloor,\\ (P_X(x)-(\lceil k(x)\rceil-1)p^\star)P_{Y|X}(y|i_u), & j_u=\lceil k(i_u)\rceil. \end{cases}
\end{aligned}$$

Then, for a given $y\in\mathcal{Y}$,

$$\begin{aligned}
\max_{(i_u,j_u)\in\mathcal{U}} P_{UY}((i_u,j_u),y) &= \max_{(i_u,1)\in\mathcal{U}} p^\star P_{Y|X}(y|i_u)\\
&= \max_{x\in\mathcal{X}} p^\star P_{Y|X}(y|x).
\end{aligned}\tag{4}$$

Finally, we get

$$\mathcal{L}(X\to Y) \ge \mathcal{L}(X\to Y)[U] = \log\sum_{y\in\mathcal{Y}} \max_{x\in\mathcal{X}} P_{Y|X}(y|x),$$

where the inequality follows from the definition, and the equality follows from equations (2), (3), and (4). ∎

The above result and analysis sheds light on the reason behind the equivalence between $\mathcal{L}(X\to Y)$ and $ML(P_{Y|X})$. First, $\mathcal{L}(X\to Y)$ depends on P_X only through its support. Moreover, on the one hand, the maximizer for $ML(P_{Y|X})$ is always the uniform distribution [22]; and on the other hand, for a uniform P_X, the above optimizing $P_{U|X}$ is simply the identity map, which is the function of interest in [21,22].

In light of this, one might wonder if there is always a deterministic map $P_{U|X}$ that achieves $\mathcal{L}(X\to Y)$. This is, however, not true in general. Suppose P_{XY} satisfies the following condition: there exists $x^\star \in \mathcal{X}$ such that for all $y \in \mathcal{Y}$, $P_{X|Y}(x^\star|y) \geq 1/2$. Then, for any deterministic function f, $\mathcal{L}(X\to Y)[f(X)]$ (cf. (2)) is zero since $f(x^\star)$ is always the optimal choice for the adversary, with and without the observation of Y. The above condition, however, is not sufficient for X and Y to be independent. The equivalence between $\mathcal{L}(X \to Y)$ and $I_\infty(X;Y)$ implies, on the other hand, that the independence of X and Y is necessary for maximal leakage to be zero, (see Corollary 2). Due to its usefulness, we state this equivalence as a separate corollary.

Corollary 1: For any joint distribution P_{XY} on finite alphabets $\mathcal{X}$ and $\mathcal{Y}$,

$$\mathcal{L}(X\to Y) = I_\infty(X;Y),$$

where $I_\infty(X;Y)$ is the Sibson mutual information of order infinity.

Sibson's $I_\alpha(X;Y)$ ($\alpha \geq 0$) is an extension of the concept of Renyi entropy $H_\alpha(X)$ and Renyi divergence $D_\alpha(P||Q)$. Although there are other possible extensions, [19] argues for the adoptions of Sibson's definition. Our result could be seen as also supporting that claim (more recently, $I_\infty(X;Y)$ has been used as a complexity measure in the study of communication complexity [41]).

For binary-valued X, say $\mathcal{X} = \{0,1\}$, Sibson [20] showed that

$$I_\infty(X;Y) = \log\left(1 + \frac{1}{2}\|P_{Y|X}(\cdot|1) - P_{Y|X}(\cdot|0)\|\right),$$

where $\|.\|_1$ is the L-1 distance. The term inside the log is twice the probability of success in binary hypothesis testing, which sheds light on why $I_\infty(X;Y)$ arises as maximal leakage. We evaluate $\mathcal{L}(X \to Y)$ for some special cases.

Example 1: If $X \sim \text{Ber}(q)$, $0 < q < 1$, and Y is the output of a BSC with parameter p, $0 \leq p \leq 1/2$, then $\mathcal{L}(X\to Y) = \log(2(1-p))$.

Example 2: If $X \sim \text{Ber}(q)$, $0 < q < 1$, and Y is the output of a BEC with parameter ϵ, $0 \leq \epsilon \leq 1$, then $\mathcal{L}(X\to Y) = \log(2-\epsilon)$, and $\mathcal{L}(Y\to X) = \log 2$.

Example 3: For any deterministic law $P_{Y|X}$, $\mathcal{L}(X\to Y) = \log|\{y : P_Y(y) > 0\}|$.

The following corollary summarizes some useful properties of $\mathcal{L}(X\to Y)$.

Corollary 2: For any joint distribution P_{XY} on finite alphabets $\mathcal{X}$ and $\mathcal{Y}$,

1) (*Data Processing Inequality*) If the Markov chain $X - Y - Z$ holds for a discrete random variable Z, then $\mathcal{L}(X\to Z) \leq \min\{\mathcal{L}(X\to Y), \mathcal{L}(Y\to Z)\}$.
2) $\mathcal{L}(X\to X) = H_0(X) = \log|\{x : P_X(x) > 0\}|$.
3) $\mathcal{L}(X\to Y) \leq \min\{\log|\mathcal{X}|, \log|\mathcal{Y}|\}$.
4) $\mathcal{L}(X\to Y) \geq I(X;Y)$.
5) $\mathcal{L}(X\to Y) = 0$ iff X and Y are independent.
6) $\mathcal{L}(X\to Y)$ is not symmetric in X and Y.
7) (*Additivity*) If $\{(X_i,Y_i)\}_{i=1}^{\ell}$ are mutually independent, then
$$\mathcal{L}(X_1^\ell\to Y_1^\ell) = \sum_{i=1}^{\ell} \mathcal{L}(X_i\to Y_i).$$
8) $\exp\{\mathcal{L}(X\to Y)\}$ is convex in $P_{Y|X}$ for fixed P_X.

Proof: Properties 1) through 4), and 7) are shown for $I_\infty(X;Y)$ [19]. 5) follows from the definition and 4). 6) is clear and is illustrated in Example 2. 8) follows from the fact that, for each $y \in \mathcal{Y}$, $\max_x P_{Y|X}(y|x)$ is convex in $P_{Y|X}$. ∎

Note that properties 1), 5), and 7) can be regarded as axiomatic for a leakage measure. Property 4) shows that a small maximal leakage is a more stringent requirement than a small mutual information. Property 8) shows that minimizing maximal leakage, for a fixed P_X, amounts to minimizing a convex function.

IV. Maximal Leakage Variations

We show the robustness of maximal leakage, by proving that variations on its definition yield the same quantity.

A. k-maximal leakage

We allow the adversary several guesses, as arises in some practical situations discussed in the introduction.

Definition 2 (k-Maximal Leakage): Given a joint distribution P_{XY} on finite alphabets $\mathcal{X}$ and $\mathcal{Y}$, and a positive integer k, the k-maximal leakage from X to Y is defined as

$$\mathcal{L}^{(k)}(X\to Y) = \sup_{U-X-Y-(\hat{U}_i)_{i=1}^k} \log \frac{\mathbf{Pr}\left(\bigvee_{i=1}^k U = \hat{U}_i\right)}{\max_{\substack{S\subseteq\mathcal{U}\\|S|\leq k}} P_U(S)}.$$

The following lemma establishes the equivalence between maximal leakage and k-maximal leakage.

Lemma 1: For any joint distribution P_{XY} on finite alphabets $\mathcal{X}$ and $\mathcal{Y}$, and any $k \in \mathbb{N}$,

$$\mathcal{L}^{(k)}(X\to Y) = \mathcal{L}(X\to Y).$$

Proof:

To show $\mathcal{L}^{(k)}(X \to Y) \geq \mathcal{L}(X \to Y)$, for any $P_{U|X}$, we construct $P_{V|X}$ such that $\mathcal{L}^{(k)}(X\to Y)[V] =$

$\mathcal{L}(X\to Y)[U]$. In particular, for a given $P_{U|X}$ and associated alphabet $\mathcal{U}$, let

$$\mathcal{V} = \bigcup_{u\in\mathcal{U}} \{(u,1),(u,2),\dots,(u,k)\},$$

and $P_{V|X}(v|x) = P_{V|X}((a_v,b_v)|x) = P_{U|X}(a_v|x)/k.$

Then, observing Y, the probability of guessing V correctly with k guesses is:

$$\begin{aligned}
&\sup_{X-Y-(\hat{V}_i)_{i=1}^k} \mathbf{Pr}(V=\hat{V}_1 \vee \cdots \vee V=\hat{V}_k)\\
&= \sum_{y\in\mathcal{Y}} \max_{\substack{v_1,v_2,\dots,v_k\\ v_i\neq v_j, i\neq j}} \sum_{i=1}^{k}\sum_{x\in\mathcal{X}} P_X(x)P_{V|X}(v_i|x)P_{Y|X}(y|x)\\
&= \sum_{y\in\mathcal{Y}}\sum_{i=1}^{k} \max_{v_i\neq v_1,\dots,v_{i-1}} \sum_{x\in\mathcal{X}} P_X(x)P_{V|X}(v_i|x)P_{Y|X}(y|x)\\
&\overset{(a)}{=} \sum_{y\in\mathcal{Y}} \max_{u} \sum_{x\in\mathcal{X}} P_X(x)P_{U|X}(u|x)P_{Y|X}(y|x), \qquad (5)
\end{aligned}$$

where (a) follows by setting $v_i = (u^\star, i)$, where

$$u^\star = \operatorname*{argmax}_{u\in\mathcal{U}} \sum_{x\in\mathcal{X}} P_X(x)P_{U|X}(u|x)P_{Y|X}(y|x).$$

Now, note that (5) is simply the probability of guessing U correctly with a single guess after observing Y. A similar argument shows that, with no Y observation, the probability of guessing V correctly with k guesses is equal to the probability of guessing U correctly with a single guess, hence $\mathcal{L}^{(k)}(X\to Y)[V] = \mathcal{L}(X\to Y)[U]$, which establishes $\mathcal{L}^{(k)}(X\to Y) \geq \mathcal{L}(X\to Y)$.

We still need to show $\mathcal{L}(X\to Y) \geq \mathcal{L}^{(k)}(X\to Y)$. For any $P_{V|X}$, we construct $P_{U|X}$ such that $\mathcal{L}(X\to Y)[U] = \mathcal{L}^{(k)}(X\to Y)[V]$. So let $P_{V|X}$ be given, with associated alphabet $\mathcal{V}$, and let $\ell \triangleq |\mathcal{V}| \geq k$. Now, let

$$\begin{aligned}
\mathcal{U} &= \{S\subset\mathcal{V} : |S| = k\},\\
\text{and}\quad p_{U|X}(u|x) &= c\sum_{v\in u} p_{V|X}(v|x),
\end{aligned}$$

where $c = 1/\binom{\ell-1}{k-1}$. Then, observing Y, the probability of guessing U correctly with a single guess is

$$\begin{aligned}
&\sup_{X-Y-\hat{U}} \mathbf{Pr}(U=\hat{U})\\
&= \sum_{y\in\mathcal{Y}} \max_{u\in\mathcal{U}} \sum_{x\in\mathcal{X}} P_X(x)P_{U|X}(u|x)P_{Y|X}(y|x)\\
&= \sum_{y\in\mathcal{Y}} \max_{u\in\mathcal{U}} \sum_{x\in\mathcal{X}} P_X(x)\sum_{v\in u} P_{V|X}(v|x)P_{Y|X}(y|x)c\\
&= c\sum_{y\in\mathcal{Y}} \max_{\substack{v_1,v_2,\dots,v_k\\ v_i\neq v_j, i\neq j}} \sum_{x\in\mathcal{X}}\sum_{i=1}^{k} P_X(x)P_{V|X}(v_i|x)P_{Y|X}(y|x),
\end{aligned}$$

which is the probability, normalized by c, of guessing V correctly with k guesses after observing Y. A similar argument shows that, with no Y observation, the probability of guessing U correctly with a single guess is equal to the probability, normalized by c, of guessing V correctly with k guesses, hence $\mathcal{L}(X\to Y)[V] = \mathcal{L}^{(k)}(X\to Y)[U]$, which establishes $\mathcal{L}(X\to Y) \geq \mathcal{L}^{(k)}(X\to Y)$. ■

B. Maximal locational leakage

For locational leakage, the adversary only needs to generate a guess that is within a certain distance of the true function value. The term "locational" is motivated by the scenario in which the variable of interest U is a geographical location.

Definition 3 (Maximal Locational Leakage): Given a joint distribution P_{XY} on finite alphabets $\mathcal{X}$ and $\mathcal{Y}$, and a metric space $\mathcal{U}$ (with its associated Borel σ-field), the maximal locational leakage from X to Y is defined as

$$\mathcal{L}_{\mathcal{U}}(X\to Y) = \sup_{\substack{U:U-X-Y\\ \exists u:\mathbf{Pr}(U\in B(u))>0}} \log\frac{\sup_{\hat{u}(.)}\mathbf{Pr}(U\in B(\hat{u}(Y)))}{\sup_{\hat{u}}\mathbf{Pr}(U\in B(\hat{u}))}, \qquad (6)$$

where $B(u)$ denotes the closed unit ball centered at $u\in\mathcal{U}$.

Lemma 2: For any joint distribution P_{XY} on finite alphabets $\mathcal{X}$ and $\mathcal{Y}$, and any metric space $\mathcal{U}$,

$$\mathcal{L}_{\mathcal{U}}(X\to Y) \leq \mathcal{L}(X\to Y),$$

with equality if $\mathcal{U}$ has an infinitely countable subset S, such that no pair of its elements can be contained in a single unit ball.

Proof:

Consider any U and $\hat{u}(Y)$ in the maximization of (6):

$$\begin{aligned}
&\mathbf{Pr}(U\in B(\hat{u}(Y)))\\
&\leq \sum_{y\in\mathcal{Y}}\sup_{u\in\mathcal{U}} P(U\in B(u), Y=y)\\
&= \sum_{y\in\mathcal{Y}}\sup_{u\in\mathcal{U}}\sum_{x\in\mathcal{X}} P(U\in B(u), X=x, Y=y)\\
&= \sum_{y\in\mathcal{Y}}\sup_{u\in\mathcal{U}}\sum_{x\in\mathcal{X}} P(U\in B(u))\cdot\\
&\qquad\qquad P(X=x|U\in B(u))P_{Y|X}(y|x)\\
&\leq \sum_{y\in\mathcal{Y}}\sup_{u\in\mathcal{U}} P(U\in B(u))\sup_{x\in\mathcal{X}} p_{Y|X}(y|x)\\
&= \left[\sum_{y\in\mathcal{Y}}\sup_{x\in\mathcal{X}} p_{Y|X}(y|x)\right]\sup_{u\in\mathcal{U}} P(U\in B(u)).
\end{aligned}$$

Therefore,

$$\mathcal{L}_{\mathcal{U}}(X\to Y) \leq \log\sum_{y\in\mathcal{Y}}\sup_{x\in\mathcal{X}} P_{Y|X}(y|x) = \mathcal{L}(X\to Y).$$

If $\mathcal{U}$ satisfies the lemma condition (e.g., $\mathcal{U}$ is unbounded), then exact guessing of discrete functions can be simulated by choosing S to be the support of U. Hence $\mathcal{L}_{\mathcal{U}}(X \to Y) \geq \mathcal{L}(X \to Y)$, which implies the equality. ■

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] N. Li, T. Li, and S. Venkatasubramanian, "t-closeness: Privacy beyond k-anonymity and l-diversity," in *Data Engineering, 2007. ICDE 2007. IEEE 23rd International Conference on*, April 2007, pp. 106–115.

[3] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *Foundations of Computer Science, 2007. FOCS '07. 48th Annual IEEE Symposium on*, Oct 2007, pp. 94–103.

[4] D. X. Song, D. Wagner, and X. Tian, "Timing analysis of keystrokes and timing attacks on SSH," in *Proceedings of the 10th USENIX Security Symposium - Volume 10*. Berkeley, CA, USA: USENIX Association, 2001.

[5] P. Venkitasubramaniam, T. He, and L. Tong, "Anonymous networking amidst eavesdroppers," *Information Theory, IEEE Transactions on*, vol. 54, no. 6, pp. 2770–2784, June 2008.

[6] Y. Wang and G. E. Suh, "Efficient timing channel protection for on-chip networks," in *Networks on Chip (NoCS), 2012 Sixth IEEE/ACM International Symposium on*, May 2012, pp. 142–151.

[7] J. R. Rao, P. Rohatgi, H. Scherzer, and S. Tinguely, "Partitioning attacks: or how to rapidly clone some GSM cards," in *IEEE Proc. Symposium on Security and Privacy*. IEEE, 2002, pp. 31–41.

[8] A. D. Wyner, "The wire-tap channel," *Bell System Technical Journal*, vol. 54, no. 8, pp. 1355–1387, Oct 1975.

[9] S. Leung-Yan-Cheong and M. Hellman, "The Gaussian wire-tap channel," *Information Theory, IEEE Transactions on*, vol. 24, no. 4, pp. 451–456, Jul 1978.

[10] C. E. Shannon, "Communication theory of secrecy systems," *Bell system technical journal*, vol. 28, no. 4, pp. 656–715, 1949.

[11] P. K. Gopala, L. Lai, and H. El Gamal, "On the secrecy capacity of fading channels," *Information Theory, IEEE Transactions on*, vol. 54, no. 10, pp. 4687–4698, Oct 2008.

[12] L. Lai and H. El Gamal, "The relay-eavesdropper channel: cooperation for secrecy," *Information Theory, IEEE Transactions on*, vol. 54, no. 9, pp. 4005–4019, Sept 2008.

[13] I. Csiszár and J. Körner, "Broadcast channels with confidential messages," *Information Theory, IEEE Transactions on*, vol. 24, no. 3, pp. 339–348, May 1978.

[14] D. Gunduz, E. Erkip, and H. V. Poor, "Lossless compression with security constraints," in *Information Theory, 2008. ISIT 2008. IEEE International Symposium on*, July 2008, pp. 111–115.

[15] L. Sankar, S. Rajagopalan, and H. Poor, "Utility-privacy tradeoffs in databases: An information-theoretic approach," *Information Forensics and Security, IEEE Transactions on*, vol. 8, no. 6, pp. 838–852, June 2013.

[16] F. Calmon, A. Makhdoumi, and M. Médard, "Fundamental limits of perfect privacy," in *Information Theory (ISIT), 2015 IEEE International Symposium on*, June 2015, pp. 1796–1800.

[17] 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.

[18] V. Anantharam, A. Gohari, S. Kamath, and C. Nair, "On hypercontractivity and a data processing inequality," in *Information Theory (ISIT), 2014 IEEE International Symposium on*, June 2014, pp. 3022–3026.

[19] S. Verdú, "α-mutual information," in *Information Theory and Applications Workshop (ITA), 2015*, Feb 2015, pp. 1–6.

[20] R. Sibson, "Information radius," *Zeitschrift fr Wahrscheinlichkeitstheorie und Verwandte Gebiete*, vol. 14, no. 2, pp. 149–160, 1969. [Online]. Available: http://dx.doi.org/10.1007/BF00537520

[21] G. Smith, "On the foundations of quantitative information flow," in *Foundations of Software Science and Computational Structures*, ser. Lecture Notes in Computer Science, L. de Alfaro, Ed. Springer Berlin Heidelberg, 2009, vol. 5504, pp. 288–302.

[22] C. Braun, K. Chatzikokolakis, and C. Palamidessi, "Quantitative notions of leakage for one-try attacks," *Electronic Notes in Theoretical Computer Science*, vol. 249, pp. 75–91, 2009.

[23] M. Alvim, K. Chatzikokolakis, C. Palamidessi, and G. Smith, "Measuring information leakage using generalized gain functions," in *Computer Security Foundations Symposium (CSF), 2012 IEEE 25th*, June 2012, pp. 265–279.

[24] B. Espinoza and G. Smith, "Min-entropy as a resource," *Information and Computation*, vol. 226, pp. 57 – 75, 2013, special Issue: Information Security as a Resource.

[25] M. S. Alvim, M. E. Andrés, K. Chatzikokolakis, and C. Palamidessi, "On the relation between differential privacy and quantitative information flow," in *Automata, Languages and Programming*. Springer, 2011, pp. 60–76.

[26] G. Barthe and B. Kopf, "Information-theoretic bounds for differentially private mechanisms," in *Computer Security Foundations Symposium (CSF), 2011 IEEE 24th*, June 2011, pp. 191–204.

[27] C. Dwork, "Differential privacy: A survey of results," in *Theory and applications of models of computation*. Springer, 2008, pp. 1–19.

[28] S. Goldwasser and S. Micali, "Probabilistic encryption," *Journal of Computer and System Sciences*, vol. 28, no. 2, pp. 270 – 299, 1984.

[29] A. Russell and H. Wang, "How to fool an unbounded adversary with a short key," *Information Theory, IEEE Transactions on*, vol. 52, no. 3, pp. 1130–1140, March 2006.

[30] Y. Dodis and A. Smith, "Entropic security and the encryption of high entropy messages," in *Theory of Cryptography*, ser. Lecture Notes in Computer Science, J. Kilian, Ed. Springer Berlin Heidelberg, 2005, vol. 3378, pp. 556–577.

[31] M. Bellare, S. Tessaro, and A. Vardy, "Semantic security for the wiretap channel," in *Advances in Cryptology–CRYPTO 2012*. Springer, 2012, pp. 294–311.

[32] F. Calmon, M. Varia, M. Médard, M. Christiansen, K. Duffy, and S. Tessaro, "Bounds on inference," in *Communication, Control, and Computing (Allerton), 2013 51st Annual Allerton Conference on*, Oct 2013, pp. 567–574.

[33] C. T. Li and A. E. Gamal, "Maximal correlation secrecy," *CoRR*, vol. abs/1412.5374, 2014. [Online]. Available: http://arxiv.org/abs/1412.5374

[34] M. Greenacre, *Correspondence analysis in practice*. CRC press, 2007.

[35] H. Yamamoto, "Rate-distortion theory for the Shannon cipher system," *Information Theory, IEEE Transactions on*, vol. 43, no. 3, pp. 827–835, 1997.

[36] C. Schieler and P. Cuff, "The henchman problem: measuring secrecy by the minimum distortion in a list," in *Information Theory (ISIT), 2014 IEEE International Symposium on*. IEEE, 2014, pp. 596–600.

[37] ——, "The henchman problem: measuring secrecy by the minimum distortion in a list," *CoRR*, vol. abs/1410.2881, 2014. [Online]. Available: http://arxiv.org/abs/1410.2881

[38] N. Weinberger and N. Merhav, "A large deviations approach to secure lossy compression," *arXiv preprint arXiv:1504.05756*, 2015.

[39] I. Issa and A. B. Wagner, "Measuring secrecy by the probability of a successful guess," *CoRR*, vol. abs/1507.02342, 2015. [Online]. Available: http://arxiv.org/abs/1507.02342

[40] I. Wagner and D. Eckhoff, "Technical privacy metrics: a systematic survey," *arXiv preprint arXiv:1512.00327*, 2015.

[41] M. M. Prabhakaran and V. M. Prabhakaran, "Rényi information complexity and an information theoretic characterization of the partition bound," *CoRR*, vol. abs/1511.07949, 2015. [Online]. Available: http://arxiv.org/abs/1511.07949