An Alphabet of Leakage Measures

Atefeh Gilani, Gowtham R. Kurri, Oliver Kosut, Lalitha Sankar

2022 IEEE Information Theory Workshop (ITW) | 978-1-6654-8341-4/22/$31.00 ©2022 IEEE | DOI: 10.1109/ITW54588.2022.9965918

***Abstract*—We introduce a family of information leakage measures called *maximal* α, β*-leakage*, parameterized by real numbers α and β. The measure is formalized via an operational definition involving an adversary guessing an unknown function of the data given the released data. We obtain a simple, computable expression for the measure and show that it satisfies several basic properties such as monotonicity in β for a fixed α, non-negativity, data processing inequalities, and additivity over independent releases. Finally, we highlight the relevance of this family by showing that it bridges several known leakage measures, including maximal α-leakage ($\beta = 1$), maximal leakage ($\alpha = \infty, \beta = 1$), local differential privacy ($\alpha = \infty, \beta = \infty$), and local Rényi differential privacy ($\alpha = \beta$).**

I. Introduction

How much information does an observation released to an adversary reveal/leak about correlated sensitive data? This fundamental question arises in many secrecy and privacy problems whenever data about users is stored online (e.g., social networks and cloud-based services) and a certain level of information leakage is unavoidable in exchange for certain services. In an effort to quantify this leakage precisely, a variety of privacy measures have been proposed [1]–[9].

For any leakage measure, one of the key challenges is to associate an operational interpretation in terms of its definition. Only a few leakage measures possess such an operational meaning. For example, the works in [3], [5], [6], which pertain to the release of observation due to a side channel, measure privacy in terms of an adversary's gain in *guessing* the sensitive data after observing the released data. Issa *et al.* [5] introduce maximal leakage (MaxL), which quantifies the maximal logarithmic gain in the probability of correctly guessing any arbitrary function of the original data from the released data. Liao *et al.* [6] later generalized maximal leakage to a family of leakages, maximal α-leakage (Max-αL), that allows tuning the measure to specific applications. In particular, similar to MaxL, Max-αL quantifies the maximal logarithmic gain in a monotonically increasing power function (dependent on α) applied to the probability of correctly guessing.

Among leakage measures motivated by worst-case adversaries, differential privacy (DP) [7] has emerged as the gold standard. A differentially private algorithm guarantees that its outputs restrict the adversary from distinguishing between neighboring data entries. When privacy guarantees have to be provided in a distributed setting, local DP (LDP) [8], [9] provides such strong guarantees for every pair of (users) data entries. In the context of composing DP outputs sequentially, Rényi differential privacy (RDP) [4] has emerged as a meaningful variant to compute overall DP guarantees. Specifically, RDP relaxes DP based on the Rényi divergence.

No single measure of privacy/information leakages suits all the scenarios in practice. In this paper, we undertake the study of unifying various measures of information leakage so that the leakage measure can be tailored to different settings depending on the context. Motivated by [5], [6], in Section III, we introduce a leakage measure, *maximal* α, β*-leakage*, which is parameterized by two real numbers α and β. We obtain a simple computable expression for it and show that this family of measures encompasses a host of existing leakage measures: in particular, Max-αL ($\beta = 1$), MaxL ($\alpha = \infty, \beta = 1$), LDP ($\alpha = \beta = \infty$), and local Rényi differential privacy (LRDP) ($\alpha = \beta$) — a notion of RDP defined analogous to LDP (see Fig. 1). An important consequence of our result is an operational interpretation of LDP and LRDP (Section IV). We note that this subsumes an operational meaning of LDP given by Issa *et al.* [5] via maximal realizable leakage, a leakage measure definition concerned with worst-case analysis, akin to LDP. Interestingly, maximal α, β-leakage is defined in terms of average-case analysis (in the spirit of MaxL and Max-αL), and yet, it recovers the worst-case LDP by exploiting the interplay between the parameters α and β. We also show that this general leakage measure satisfies all the axiomatic properties of a measure of information leakage, including non-negativity, equality to zero if and only if the original data and the released data are independent of each other, and data-processing inequalities. This can be viewed as another proof that LDP satisfies both the post-processing and linkage inequalities unlike DP which does not satisfy the linkage inequality [10]. Due to space constraints, we have omitted some proofs; detailed proofs for all results are in [11].

II. Preliminaries

We begin by reviewing the definitions of some existing information leakage measures, in particular, maximal α-leakage (which subsumes maximal leakage) and local (Rényi) DP.

Definition 1 (Maximal α-leakage [6])**.** *Let P_{XY} be a joint distribution, where X and Y represent the original data and the released data, respectively. The maximal α-leakage from X to Y, for $\alpha \in (1, \infty)$, is defined as*

$$\mathcal{L}_\alpha^{\max}(X \to Y) = \frac{\alpha}{\alpha - 1} \sup_{U - X - Y} \log \frac{\max_{P_{\hat{U}|Y}} \sum_{u,y} P_{UY}(u,y) P_{\hat{U}|Y}(u|y)^{\frac{\alpha-1}{\alpha}}}{\max_{P_{\hat{U}}} \sum_u P_U(u) P_{\hat{U}}(u)^{\frac{\alpha-1}{\alpha}}}, \quad (1)$$

The authors are with the School of Electrical, Computer and Energy Engineering at Arizona State University. Email: agilani2@asu.edu, gkurri@asu.edu, okosut@asu.edu, lalithasankar@asu.edu

This work is supported in part by NSF grants CIF-1901243, CIF-1815361, CIF-2007688, CIF-2134256, and SaTC-2031799.

where U represents any randomized function of X that the adversary is interested in guessing and takes values in an arbitrary finite alphabet. Moreover, $\hat{U}$ is an estimator of U with the same support as U.

Maximal α-leakage is a generalization of another measure of information leakage, the maximal leakage [5]. In particular, the latter recovers the former when $\alpha = \infty$. Liao *et al.* [6] showed that

$$\mathcal{L}_\alpha^{\max}(X \to Y) = \sup_{P_{\tilde{X}}} I_\alpha^{\mathrm{S}}(\tilde{X}; Y), \tag{2}$$

where the supremum is over all the probability distributions $P_{\tilde{X}}$ on the support of P_X and $I_\alpha^{\mathrm{S}}(\cdot;\cdot)$ is the Sibson mutual information of order α [12].

Definition 2 (Local differential privacy [8], [9])**.** *Given a conditional distribution $P_{Y|X}$, the local differential privacy (LDP) is defined as*

$$\mathcal{L}^{\mathrm{LDP}}(X \to Y) = \max_{\substack{y \in \mathcal{Y}, \\ x, x' \in \mathcal{X}}} \log \frac{P_{Y|X}(y|x)}{P_{Y|X}(y|x')}. \tag{3}$$

We may define local Rényi differential privacy as a generalization of local differential privacy based on the Rényi divergence [13].

Definition 3 (Local Rényi differential privacy)**.** *Given a conditional distribution $P_{Y|X}$, the local Rényi differential privacy (LRDP) is defined as*

$$\mathcal{L}^{\mathrm{LRDP}}(X \to Y) = \max_{x, x' \in \mathcal{X}} \frac{1}{\alpha - 1} \log \sum_y P_{Y|X}(y|x')^{1-\alpha} P_{Y|X}(y|x)^\alpha. \tag{4}$$

It can be verified using L'Hôpital's rule that LRDP simplifies to LDP as $\alpha \to \infty$.

III. Maximal α, β-Leakage

Motivated by the definitions of maximal leakage and maximal α-leakage, we introduce maximal α, β-leakage as follows.

Definition 4 (Maximal α, β-leakage)**.** *Given a conditional distribution $P_{Y|X}$, the maximal α, β-leakage from X to Y for $\alpha \in (1, \infty)$ and $\beta \in [1, \infty)$ is defined as*

$$\mathcal{L}_{\alpha,\beta}(X \to Y) := \sup_{P_X} \sup_{U \to X \to Y} \frac{\alpha}{\alpha - 1} \log \frac{\max\limits_{P_{\hat{U}|Y}} \left[\sum\limits_y P_Y(y) \left(\sum\limits_u P_{U|Y}(u|y) P_{\hat{U}|Y}(u|y)^{\frac{\alpha-1}{\alpha}} \right)^\beta \right]^{1/\beta}}{\max\limits_{P_{\hat{U}}} \sum\limits_u P_U(u) P_{\hat{U}}(u)^{\frac{\alpha-1}{\alpha}}}. \tag{5}$$

where $\hat{U}$ represents an estimator taking values from the same arbitrary finite alphabet as U. It is defined by continuous extension for $\alpha = \infty$ or $\beta = \infty$.

We remark that the definition of maximal α, β-leakage in (5) nearly recovers the definition of maximal α-leakage from (1) (and thus maximal leakage also) when $\beta = 1$. The reason is that maximal α-leakage depends on the distribution of X only through its support, as shown in (2), so including the supremum over P_X does not change the value. Moreover, for $\alpha = 1$, maximal α, β-leakage recovers Shannon channel capacity. We have included the supremum in the definition of maximal α, β-leakage in order to recover some of the worst-case measures such as LDP and LRDP as special cases. One can view the introduction of β into the summation in the numerator in (5) as allowing a continuous transition from a simple average over y (at $\beta = 1$) to a maximum over y (at $\beta = \infty$). This maximum over y is present in the definition of maximal *realizable* leakage [5, Definition 8], which corresponds to $\alpha = \beta = \infty$, and has been shown to be equal to LDP. This allows us to view maximal leakage and maximal realizable leakage as two corner points of the inner optimization problem in (5) for $\alpha = \infty$ and $\beta = 1$ and $\beta = \infty$, respectively (see also Fig. 1).

The following theorem simplifies the expression of maximal α, β-leakage in (5).

Theorem 1. *Maximal α, β-leakage defined in* (5) *simplifies to*

$$\mathcal{L}_{\alpha,\beta}(X \to Y) = \max_{x'} \sup_{P_{\tilde{X}}} \frac{\alpha}{(\alpha - 1)\beta} \log \sum_y P_{Y|X}(y|x')^{1-\beta} \left(\sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)^\alpha \right)^{\beta/\alpha}, \tag{6}$$

where $P_{\tilde{X}}$ is a probability distribution on the support of P_X.

A detailed proof for Theorem 1 is given in Section V-A. For $\beta \le \alpha$, the quantity inside the log in (6) is concave in $P_{\tilde{X}}$; thus the supremum over $P_{\tilde{X}}$ can be efficiently solved using convex optimization techniques. As we will show in Section IV, for $\beta \ge \alpha$, the supremum over $P_{\tilde{X}}$ can be replaced by a maximum over $x \in \mathcal{X}$. Thus, in either case the quantity in (6) can be efficiently computed for finite alphabets.

Like other leakage measures, maximal α, β-leakage satisfies several basic properties such as non-negativity, data processing inequalities and additivity, as shown in the following theorem.

Theorem 2. *For $\alpha \in (1, \infty)$ and $\beta \in [1, \infty)$, maximal α, β-leakage*

1) *is monotonically non-decreasing in β for a fixed α;*
2) *satisfies data processing inequalities, i.e., for the Markov chain $X - Y - Z$:*

$$\mathcal{L}_{\alpha,\beta}(X \to Z) \le \mathcal{L}_{\alpha,\beta}(X \to Y) \tag{7a}$$

$$\mathcal{L}_{\alpha,\beta}(X \to Z) \le \mathcal{L}_{\alpha,\beta}(Y \to Z). \tag{7b}$$

3) *is non-negative, i.e.,*

$$\mathcal{L}_{\alpha,\beta}(X \to Y) \ge 0 \tag{8}$$

with equality if and only if X and Y are independent.

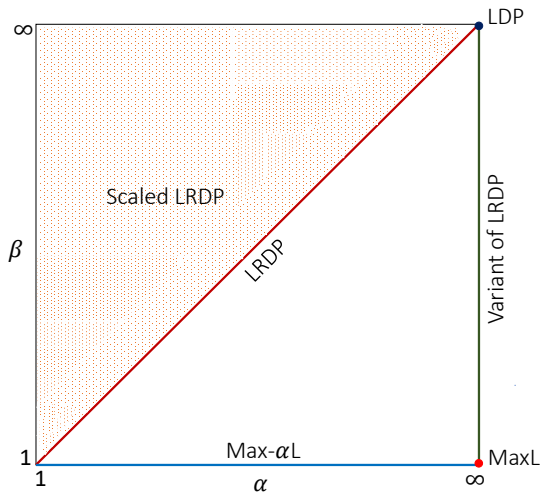

Fig. 1. Relationship between maximal α, β-leakage and other leakage measures as a function of α and β. For $\alpha \leq \beta$, we obtain a scaled LRDP of order β. For $\alpha = \beta = \infty$, we obtain LDP. For $\beta = 1$, we obtain maximal α-leakage which simplifies to maximal leakage for $\alpha = \infty$. Finally, for $\alpha = \infty$ and arbitrary β, we obtain a variant of LRDP.

4) *satisfies additivity: i.e., if* (X_i, Y_i) *for* $i = 1, 2, \ldots, n$ *are independent, then*

$$\mathcal{L}_{\alpha,\beta}(X_1, \ldots, X_n \to Y_1, \ldots, Y_n) = \sum_{i=1}^{n} \mathcal{L}_{\alpha,\beta}(X_i \to Y_i). \tag{9}$$

Proofs for most of the claims of Theorem 2 are in Section V-B.

Remark 1. Although maximal α, β-leakage is monotonic in only one of its orders, if we consider a reparameterization in which $\tau \in [0, 1]$ and $\beta = \frac{\alpha}{1-\tau(1-\alpha)}$, the new leakage measure is non-increasing in τ for a fixed α, and non-decreasing in α for a fixed τ. A detailed proof is provided in [11].

IV. Relationship between maximal α, β-leakage, and other measures

As mentioned earlier, maximal α, β-leakage recovers maximal α-leakage (and thus maximal leakage) when $\beta = 1$. The choices of α and β help to recover other leakage measures such as a scaled LRDP for $\alpha \leq \beta$, LRDP for $\alpha = \beta$, LDP for $(\alpha = \infty, \beta = \infty)$, and a variant of LRDP for $\alpha = \infty$ and arbitrary β, as shown in Fig. 1. We now present these in detail.

When $\alpha \leq \beta$, we have

$$\begin{aligned}
&\mathcal{L}_{\alpha,\beta}(X \to Y) \\
&= \max_{x'} \sup_{P_{\tilde{X}}} \frac{\alpha}{(\alpha-1)\beta} \log \sum_y P_{Y|X}(y|x')^{1-\beta} \\
&\quad \left(\sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)^{\alpha} \right)^{\beta/\alpha} \qquad (10) \\
&= \max_{x'} \max_{x} \frac{\alpha}{(\alpha-1)\beta} \log \sum_y P_{Y|X}(y|x')^{1-\beta} P_{Y|X}(y|x)^{\beta} \qquad (11)
\end{aligned}$$

where (11) follows because the argument of the logarithm in (10) is convex in $P_{\tilde{X}}$ and so the supremum is attained at an extreme point. This quantity is a scaled LRDP of order β. Furthermore, if $\alpha = \beta$, the expression in (11) reduces to

$$\begin{aligned}
&\mathcal{L}_{\beta,\beta}(X \to Y) \\
&= \max_{x'} \max_{x} \frac{1}{\beta-1} \log \sum_y P_{Y|X}(y|x')^{1-\beta} P_{Y|X}(y|x)^{\beta}, \qquad (12)
\end{aligned}$$

which is exactly LRDP of order β. Taking a limit as $\beta \to \infty$ in (12) gives

$$\begin{aligned}
\mathcal{L}_{\infty,\infty}(X \to Y) &= \max_{x'} \max_{x} \log \left(\max_y \frac{P_{Y|X}(y|x)}{P_{Y|X}(y|x')} \right) \qquad (13) \\
&= \max_{x,x',y} \log \frac{P_{Y|X}(y|x)}{P_{Y|X}(y|x')} \qquad (14)
\end{aligned}$$

which is LDP. So $\mathcal{L}_{\infty,\beta}(X \to Y)$ passes from maximal leakage at $\beta = 1$ to LDP at $\beta = \infty$.

From Theorem 1, for $\alpha = \infty$ and arbitrary β, we obtain

$$\begin{aligned}
&\mathcal{L}_{\infty,\beta}(X \to Y) \\
&= \max_{x'} \frac{1}{\beta} \log \sum_y P_{Y|X}(y|x')^{1-\beta} \max_x P_{Y|X}(y|x)^{\beta}. \qquad (15)
\end{aligned}$$

This quantity is a variant of LRDP, and it differs from LRDP of order β only in that the max over x is inside the summation over y rather than outside. As far as we know, this quantity has not appeared before.

Finally, we propose an extension of maximal α, β-leakage so as to include the data of multiple users, shown in (16). Here, $X^n = (X_1, X_2, \ldots, X_n)$ is a dataset with n entries, X_{-i} represents all entries except the ith, and as usual Y is the released data. Thus, the measure in (16) characterizes a situation in which an adversary may have side information access to all entries of the dataset except one. It is clear that

$$\sup_{P_{X^n}} \max_i \sup_{U \to X^n \to Y} \frac{\alpha}{(\alpha-1)\beta} \log \frac{\sup_{P_{\hat{U}|X_{-i},Y}} \sum_{x_{-i},y} P_{X_{-i},Y}(x_{-i},y) \left[\sum_u P_{U|X_{-i},Y}(u|x_{-i},y) P_{\hat{U}|X_{-i},Y}(u|x_{-i},y)^{\frac{\alpha-1}{\alpha}} \right]^{\beta}}{\sup_{P_{\hat{U}|X_{-i}}} \sum_{x_{-i}} P_{X_{-i}}(x_{-i}) \left[\sum_u P_{U|X_{-i}}(u|x_{-i}) P_{\hat{U}|X_{-i}}(u|x_{-i})^{\frac{\alpha-1}{\alpha}} \right]^{\beta}}. \tag{16}$$

this measure collapses to maximal α,β-leakage for $n=1$. We conjecture that, for $n>1$, it recovers (non-local) DP [7] ($\alpha=\infty,\beta=\infty$) and RDP [4] ($\alpha=\beta$) for multi-user data across neighbouring databases.

V. Proofs

Due to space limitations, some proofs have been omitted. Complete proofs can be found in [11].

A. Proof of Theorem 1

For $\alpha\in(1,\infty)$ and $\beta\in[1,\infty)$, we first bound $\mathcal{L}_{\alpha,\beta}(X\to Y)$ from above and then, give an achievable scheme.

Upper Bound: Consider the denominator of (5):

$$\max_{P_{\hat{U}}}\sum_u P_U(u)P_{\hat{U}}(u)^{\frac{\alpha-1}{\alpha}}. \tag{17}$$

This is solved by $P_U(u)P_{\hat{U}}(u)^{-1/\alpha}=\nu$, for some constant ν. So we have

$$P_{\hat{U}}(u)=\frac{P_U(u)^\alpha}{\sum_{u'}P_U(u')^\alpha}. \tag{18}$$

Thus the denominator becomes

$$\sum_u P_U(u)\left(\frac{P_U(u)^\alpha}{\sum_{u'}P_U(u')^\alpha}\right)^{\frac{\alpha-1}{\alpha}}=\left(\sum_u P_U(u)^\alpha\right)^{\frac{1}{\alpha}}. \tag{19}$$

Similarly, the numerator becomes

$$\left[\sum_y P_Y(y)\left(\sum_u P_{U|Y}(u|y)^\alpha\right)^{\beta/\alpha}\right]^{1/\beta}. \tag{20}$$

Thus, the logarithmic term in (5) reduces to

$$\log\frac{\left[\sum_y P_Y(y)\left(\sum_u P_{U|Y}(u|y)^\alpha\right)^{\beta/\alpha}\right]^{1/\beta}}{\left(\sum_u P_U(u)^\alpha\right)^{1/\alpha}} \tag{21}$$

$$=\log\frac{\left[\sum_y P_Y(y)^{1-\beta}\left(\sum_u P_{U,Y}(u,y)^\alpha\right)^{\beta/\alpha}\right]^{1/\beta}}{\left(\sum_u P_U(u)^\alpha\right)^{1/\alpha}} \tag{22}$$

$$=\frac{1}{\beta}\log\sum_y P_Y(y)^{1-\beta}\left[\frac{\sum_u P_U(u)^\alpha P_{Y|U}(y|u)^\alpha}{\sum_u P_U(u)^\alpha}\right]^{\frac{\beta}{\alpha}}. \tag{23}$$

Using Jensen's inequality and the Markov chain $U-X-Y$, we have

$$P_{Y|U}(y|u)^\alpha=\left(\sum_x P_{X|U}(x|u)P_{Y|X}(y|x)\right)^\alpha \tag{24}$$

$$\le\sum_x P_{X|U}(x|u)P_{Y|X}(y|x)^\alpha. \tag{25}$$

So maximal α,β-leakage may be bounded from above by

$$\sup_{P_X}\sup_{U\to X\to Y}\frac{\alpha}{(\alpha-1)\beta}\log\sum_y P_Y(y)^{1-\beta}\left[\frac{\sum_{u,x}P_U(u)^\alpha P_{X|U}(x|u)P_{Y|X}(y|x)^\alpha}{\sum_u P_U(u)^\alpha}\right]^{\frac{\beta}{\alpha}} \tag{26}$$

$$\le\sup_{P_X}\sup_{P_{\tilde{X}}}\frac{\alpha}{(\alpha-1)\beta}\log\sum_y P_Y(y)^{1-\beta}\left[\sum_x P_{\tilde{X}}(x)P_{Y|X}(y|x)^\alpha\right]^{\frac{\beta}{\alpha}} \tag{27}$$

where

$$P_{\tilde{X}}(x)=\frac{\sum_u P_U(u)^\alpha P_{X|U}(x|u)}{\sum_u P_U(u)^\alpha}. \tag{28}$$

Lower Bound: The proof is based on the expression in (23) as well as "shattering" method. Consider a random variable U such that $U\to X\to Y$ form a Markov chain and $H(X|U)=0$. For each x, let $\mathcal{U}_x$ be a finite set such that $U=u\in\mathcal{U}_x$ if and only if $X=x$ and $\mathcal{U}=\bigcup_{x\in\mathcal{X}}\mathcal{U}_x$. Moreover, given $X=x$ let U be uniformly distributed on $\mathcal{U}_x$. That is,

$$P_{U|X}(u|x)=\begin{cases}\frac{1}{|\mathcal{U}_x|} & \text{for all } u\in\mathcal{U}_x\\ 0 & \text{otherwise,}\end{cases} \tag{29}$$

and so

$$P_{Y|U}(y|u)=\begin{cases}P_{Y|X}(y|x) & \text{for all } u\in\mathcal{U}_x\\ 0 & \text{otherwise.}\end{cases} \tag{30}$$

Therefore, we have

$$\frac{\sum_u P_U(u)^\alpha P_{Y|U}(y|u)^\alpha}{\sum_u P_U(u)^\alpha} \tag{31}$$

$$=\frac{\sum_{x\in\mathcal{X}}\sum_{u\in\mathcal{U}_x}\left(\frac{P_X(x)P_{U|X}(u|x)}{P_{X|U}(x|u)}\right)^\alpha P_{Y|U}(y|u)^\alpha}{\sum_{x\in\mathcal{X}}\sum_{u\in\mathcal{U}_x}\left(\frac{P_X(x)P_{U|X}(u|x)}{P_{X|U}(x|u)}\right)^\alpha} \tag{32}$$

$$=\frac{\sum_x|\mathcal{U}_x|^{1-\alpha}P_X(x)^\alpha P_{Y|X}(y|x)^\alpha}{\sum_x|\mathcal{U}_x|^{1-\alpha}P_X(x)^\alpha}. \tag{33}$$

So we may bound maximal α,β-leakage from below by

$$\sup_{P_X}\sup_{\mathcal{U}_x}\frac{\alpha}{(\alpha-1)\beta}\log\sum_y P_Y(y)^{1-\beta}\left(\frac{\sum_x|\mathcal{U}_x|^{1-\alpha}P_X(x)^\alpha P_{Y|X}(y|x)^\alpha}{\sum_x|\mathcal{U}_x|^{1-\alpha}P_X(x)^\alpha}\right)^{\frac{\beta}{\alpha}} \tag{34}$$

$$=\sup_{\substack{P_X,\\ P_{\tilde{X}}}}\frac{\alpha}{(\alpha-1)\beta}\log\sum_y P_Y(y)^{1-\beta}\Big(\sum_x P_{\tilde{X}}(x)P_{Y|X}(y|x)^\alpha\Big)^{\frac{\beta}{\alpha}} \tag{35}$$

where here $P_{\tilde{X}}(x)=\frac{|\mathcal{U}_x|^{1-\alpha}P_X(x)^\alpha}{\sum_x|\mathcal{U}_x|^{1-\alpha}P_X(x)^\alpha}$, and we have used the fact that any distribution $P_{\tilde{X}}(x)$ can be reached with appropriate choice of $|\mathcal{U}_x|$, assuming $P_X(x)>0$ for all x; this condition can be assumed because any P_X is arbitrarily close to a distribution with full support. Thus, combining (27) and (35), we have

$$\mathcal{L}_{\alpha,\beta}(X\to Y)=\sup_{P_X}\sup_{P_{\tilde{X}}}\frac{\alpha}{(\alpha-1)\beta}$$

$$\log \sum_y P_Y(y)^{1-\beta} \left(\sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)^\alpha \right)^{\frac{\beta}{\alpha}}. \quad (36)$$

Since the choice of P_X only impacts P_Y, and the supremum of a convex function is attained at an extreme point, we may simplify (36) as follows.

$$\mathcal{L}_{\alpha,\beta}(X \to Y) = \max_{x'} \sup_{P_{\tilde{X}}} \frac{\alpha}{(\alpha-1)\beta} \log \sum_y P_{Y|X}(y|x')^{1-\beta} \left(\sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)^\alpha \right)^{\beta/\alpha}.$$

B. Proof of Theorem 2

Monotonicity in β: For $\alpha \in (1,\infty)$, $\beta_1, \beta_2 \in [1,\infty)$ and $\beta_2 > \beta_1$, consider the argument of the logarithm in (6):

$$\sum_y P_{Y|X}(y|x')^{1-\beta_1} \left[\sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)^\alpha \right]^{\frac{\beta_1}{\alpha}} \quad (37)$$

$$= \sum_y P_{Y|X}(y|x') \left[P_{Y|X}(y|x')^{-\alpha} \sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)^\alpha \right]^{\frac{\beta_2\beta_1}{\alpha\beta_2}} \quad (38)$$

$$\le \left[\sum_y P_{Y|X}(y|x') \left(P_{Y|X}(y|x')^{-\alpha} \sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)^\alpha \right)^{\frac{\beta_2}{\alpha}} \right]^{\frac{\beta_1}{\beta_2}} \quad (39)$$

$$= \left[\sum_y P_{Y|X}(y|x')^{1-\beta_2} \left(\sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)^\alpha \right)^{\frac{\beta_2}{\alpha}} \right]^{\frac{\beta_1}{\beta_2}} \quad (40)$$

where the inequality results from applying Jensen's inequality to the concave function $f: x \to x^p$ $(x \ge 0,\ p < 1)$. For $\alpha \in (1,\infty)$ and $\beta \in [1,\infty)$, the function $f: t \to \frac{\alpha}{(\alpha-1)\beta} \log t$ is increasing in $t > 0$. Therefore, we have

$$\frac{\alpha}{(\alpha-1)\beta_1} \log \sum_y P_{Y|X}(y|x')^{1-\beta_1} \left[\sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)^\alpha \right]^{\frac{\beta_1}{\alpha}} \quad (41)$$

$$\le \frac{\alpha}{(\alpha-1)\beta_2} \log \sum_y P_{Y|X}(y|x')^{1-\beta_2} \left[\sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)^\alpha \right]^{\frac{\beta_2}{\alpha}}. \quad (42)$$

Taking the max over x' and supremum over $P_{\tilde{X}}$ completes the proof. Another way to prove this property is to consider the numerator in (5) as the β-norm of a random variable. Since the β-norm of a random variable is non-decreasing in β, maximal α,β-leakage is non-decreasing in β.

Data processing inequalities: Let random variables X, Y, Z form a Markov chain, i.e., $X - Y - Z$. First consider the post-processing inequality, that is,

$$\mathcal{L}_{\alpha,\beta}(X \to Z) \le \mathcal{L}_{\alpha,\beta}(X \to Y). \quad (43)$$

The proof of this, contained in the extended version [11], is based on the expression for maximal α,β-leakage in (6), and proper use of Jensen's inequality.

We now prove the linkage inequality, that is

$$\mathcal{L}_{\alpha,\beta}(X \to Z) \le \mathcal{L}_{\alpha,\beta}(Y \to Z), \quad (44)$$

using the definition of maximal α,β-leakage in (5). Let $f(P_{UY})$ be the quantity inside the suprema in (5), which depends only on the joint distribution of U and Y. Similarly, we write $f(P_{UZ})$ for the same quantity when Y is replaced by Z. For the Markov chain $X - Y - Z$, we have

$$\mathcal{L}_{\alpha,\beta}(X \to Z) = \sup_{P_X} \sup_{U \to X \to Z} f(P_{UZ}) \quad (45)$$

$$= \sup_{P_X} \sup_{U \to X \to Y \to Z} f(P_{UZ}) \quad (46)$$

$$\le \sup_{P_X} \sup_{U \to Y \to Z} f(P_{UZ}) \quad (47)$$

$$\le \sup_{P_Y} \sup_{U \to Y \to Z} f(P_{UZ}) \quad (48)$$

$$= \mathcal{L}_{\alpha,\beta}(Y \to Z)$$

where (46) follows because P_{UZ} are the same under the Markov chains $U - X - Z$ and $U - X - Y - Z$, and (48) follows from the fact that a subset of all distributions P_Y is reachable from the distribution P_X.

Non-negativity: Consider the logarithmic term in (6):

$$\log \sum_y P_{Y|X}(y|x')^{1-\beta} \left(\sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)^\alpha \right)^{\frac{\beta}{\alpha}} \quad (49)$$

$$\ge \log \sum_y P_{Y|X}(y|x')^{1-\beta} \left(\sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x) \right)^{\beta} \quad (50)$$

$$= \log \sum_y P_{Y|X}(y|x') \left(\frac{\sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)}{P_{Y|X}(y|x')} \right)^{\beta} \quad (51)$$

$$\ge \log \left(\sum_y P_{Y|X}(y|x') \ \frac{\sum_x P_{\tilde{X}}(x) P_{Y|X}(y|x)}{P_{Y|X}(y|x')} \right)^{\beta} \quad (52)$$

$$= \log \left(\sum_{x,y} P_{\tilde{X}}(x) P_{Y|X}(y|x) \right)^{\beta} = \log 1 = 0 \quad (53)$$

where both inequalities follow from applying Jensen's inequality to the convex function $f: x \to x^p$ $(x \ge 0,\ p \ge 1)$ and the fact that logarithmic functions are increasing. Equality holds in the first inequality if and only if for any $y \in \mathcal{Y}$, $P_{Y|X}(y|x)$ are the same for all $x \in \mathcal{X}$. Thus, we have

$$P_{Y|X}(y|x) = P_Y(y) \quad x \in \mathcal{X}, y \in \mathcal{Y} \quad (54)$$

which means X and Y are independent. This condition is also sufficient for equality in the second inequality.

Additivity: Due to space limitations, the proof is relegated to [11].

REFERENCES

[1] B. Rassouli and D. Gündüz, "Optimal utility-privacy trade-off with total variation distance as a privacy measure," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 594–603, 2020.

[2] S. Asoodeh, F. Alajaji, and T. Linder, "On maximal correlation, mutual information and data privacy," in *IEEE 14th Canadian workshop on information theory (CWIT)*, 2015, pp. 27–31.

[3] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, "Privacy-aware guessing efficiency," in *IEEE International Symposium on Information Theory (ISIT)*, 2017, pp. 754–758.

[4] I. Mironov, "Rényi differential privacy," in *IEEE 30th computer security foundations symposium (CSF)*, 2017, pp. 263–275.

[5] I. Issa, A. B. Wagner, and S. Kamath, "An operational approach to information leakage," *IEEE Transactions on Information Theory*, vol. 66, no. 3, pp. 1625–1657, 2020.

[6] J. Liao, O. Kosut, L. Sankar, and F. P. Calmon, "Tunable measures for information leakage and applications to privacy-utility tradeoffs," *IEEE Transactions on Information Theory*, vol. 65, no. 12, pp. 8043–8066, 2019.

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

[8] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith, "What can we learn privately?" *SIAM Journal on Computing*, vol. 40, no. 3, pp. 793–826, 2011.

[9] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Local privacy and statistical minimax rates," in *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. IEEE, 2013, pp. 429–438.

[10] Y. O. Basciftci, Y. Wang, and P. Ishwar, "On privacy-utility tradeoffs for constrained data release mechanisms," in *2016 Information Theory and Applications Workshop (ITA)*, 2016, pp. 1–6.

[11] A. Gilani, G. R. Kurri, O. Kosut, and L. Sankar, "An alphabet of leakage measures," May 2022, github.com/SankarLab/Alphabet-Leakage-Measures/blob/main/AnAlphabetofLeakageMeasures.pdf.

[12] R. Sibson, "Information radius," *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, vol. 14, no. 2, pp. 149–160, 1969.

[13] A. Rényi, "On measures of entropy and information," in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, 1961, pp. 547–561.