Maximal α-Leakage and its Properties

Jiachun Liao, Lalitha Sankar, Oliver Kosut
School of Electrical, Computer and Energy Engineering,
Arizona State University
Email: {jiachun.liao,lalithasankar,okosut}@asu.edu

Flavio P. Calmon
School of Engineering and Applied Sciences
Harvard University
Email: flavio@seas.harvard.edu

***Abstract*—Maximal α-leakage is a tunable measure of information leakage based on the quality of an adversary's belief about an arbitrary function of private data based on public data. The parameter α determines the loss function used to measure the quality of a belief, ranging from log-loss at $\alpha = 1$ to the probability of error at $\alpha = \infty$. We review its definition and main properties, including extensions to $\alpha < 1$, robustness to side information, and relationship to Rényi differential privacy.**

I. Introduction

In many applications it is important to understand how much information about one variable is "leaked" from another. Depending on how one defines the notion of information leakage, different measures emerge. Mutual information is a classic measure for quantifying information and often used to measure information secrecy [1] or leakage in data publishing settings [2], [3]. Recently, Issa *et al.* [4] introduced a measure, called *maximal leakage* (MaxL) [4], which considers an adversary's probability of guessing any (possibly random) function of the original dataset; the MaxL is the log of the ratio of the adversary's guessing probability with access to the released data to the probability without it. In the context of differential privacy (DP) [5] — along with its related measures approximate differential privacy and Rényi differential privacy (RDP) — have emerged as the most popular information measures.

The focus of this paper is the maximal α-leakage (MAL), originally introduced in [6]. MAL is a leakage measure with a tunable parameter α that ranges from 0 to ∞, thus creating a continuum of measures that includes both mutual information and MaxL as special cases. MAL is best understood through the lens of an adversary's ability to learn information about an unknown variable. Thus, it is quantified via a loss function that the adversary seeks to minimize. MAL is defined via a loss function called α-loss, which is again defined via the tunable parameter α. At $\alpha = 1$, this loss function becomes the log-loss [7]–[9], which leads to mutual information as the measure; at $\alpha = \infty$, this loss function becomes the 0-1 loss, which leads to maximal leakage as the measure. At other values of α, MAL turns out to be related to the Sibson and Arimoto variants of the mutual information; these two information measures can be viewed as different ways of extending the Rényi entropy to a "mutual information" quantity.

This material is based upon work supported by the National Science Foundation under Grant No. CCF-1901243.

We also consider MAL in the context of side information. An adversary's side information, which is generally unknown to the the data curator, can have a significant effect on the amount of information leaked to the adversary. One of the key advantages of DP is that it is robust to arbitrary side information [10]. Maximal leakage has also been investigated with respect to side information [11]. MAL has a natural "conditional" form, wherein the side information is explicitly modeled. It can be shown that the unconditional MAL upper bounds conditional MAL if the side information is conditionally independent of the released data given the original data. That is, MAL is robust to arbitrary side information that is not used in generating the released data from the original data. This surprising result provides further motivation for using MAL as a robust and tunable leakage metric.

This paper is organized as follows. Sec. II contains the primary definitions that will be of interest, including classical notions of Rényi information measures, as well as MAL and its variants. Sec. III presents the most important properties of MAL. Sec. IV provides a connection between MAL and RDP, showing that they are equivalent in the sense that if one is small, then the other is also small. Several proofs are given in the appendix. The primary innovations in the present paper beyond our prior work in [6], [12]–[14] are the extension of MAL to $\alpha < 1$, and the connection to RDP in Sec. IV.

II. Definitions

In this section we define the concepts we will be focused on throughout the paper. Many of these concepts are defined in terms of an order parameter α. We will often leave out a specific expression for the quantities for $\alpha = 1$ or $\alpha = \infty$; we adopt the convention that any quantity is implicitly defined by continuous extension for these values of α. We begin with the basic concepts of Rényi entropy and divergence [15].

Definition 1: Given a discrete distribution P_X, the Rényi entropy of order $\alpha \in (0,1) \cup (1,\infty)$ is defined as

$$H_\alpha(P_X) \triangleq \frac{1}{1-\alpha} \log \sum_x P_X(x)^\alpha. \tag{1}$$

Definition 2: Let P_X, Q_X be two discrete distributions over $\mathcal{X}$. The Rényi divergence between P_X and Q_X of order $\alpha \in (0,1) \cup (1,\infty)$ is defined as

$$D_\alpha(P_X \| Q_X) \triangleq \frac{1}{\alpha-1} \log \sum_x P_X(x)^\alpha \, Q_X(x)^{1-\alpha}. \tag{2}$$

978-1-7281-4760-4/20/$31.00 ©2020 IEEE

There are a number of ways to generalize Shannon's mutual information in a manner analogous to Rényi's generalizations of entropy and divergence [16]. Arimoto's approach [17] is based on the following conditional version of Rényi entropy.

Definition 3: Given a joint distribution P_{XY}, the Arimoto conditional entropy of X given Y of order $\alpha \in (0,1) \cup (1,\infty)$ is defined as

$$H_\alpha(X|Y) \triangleq \frac{\alpha}{1-\alpha} \log \sum_y P_Y(y) \left(\sum_x P_{X|Y}(x|y)^\alpha \right)^{\frac{1}{\alpha}}. \quad (3)$$

Now we may define the Arimoto mutual information. The following also contains a definition for an Arimoto *conditional* mutual information, which is less common, but a natural definition nonetheless.

Definition 4: Given a joint distribution P_{XY}, the Arimoto mutual information of order $\alpha \in (0,1) \cup (1,\infty)$ is defined as

$$I_\alpha^{\mathrm{A}}(X;Y) \triangleq H_\alpha(X) - H_\alpha(X|Y). \quad (4)$$

Given a joint distribution P_{XYZ}, the Arimoto conditional mutual information is defined as

$$I_\alpha^{\mathrm{A}}(X;Y|Z) \triangleq H_\alpha(X|Z) - H_\alpha(X|Y,Z). \quad (5)$$

An alternative mutual information quantity, found by Sibson [18], is defined as follows.

Definition 5: Given a joint distribution P_{XY}, the Sibson mutual information of order $\alpha \in (0,1) \cup (1,\infty)$ is defined as

$$I_\alpha^{\mathrm{S}}(X;Y) \triangleq \inf_{Q_Y} D_\alpha(P_{XY} \| P_X \times Q_Y), \quad (6)$$

$$= \frac{\alpha}{\alpha-1} \log \sum_y \left(\sum_x P_X(x) P_{Y|X}(y|x)^\alpha \right)^{\frac{1}{\alpha}} \quad (7)$$

where the infimum in (6) is over all Q_Y over the same space as P_Y.

Next, we review the definition of maximal leakage [4], which has the same basic flavor as its generalization MAL.

Definition 6: Given a joint distribution P_{XY}, the *maximal leakage* from X to Y is

$$\mathcal{L}_{\mathrm{MaxL}}(X \to Y) \triangleq \sup_{U-X-Y} \log \frac{\max_{P_{\hat{U}|Y}} \mathbb{P}(\hat{U}=U)}{\max_{P_{\hat{U}}} \mathbb{P}(\hat{U}=U)} \quad (8)$$

where the supremum is over all random variables U with finite support satisfying the Markov chain condition. In the numerator, $\hat{U}$ is distributed according to the conditional distribution $P_{\hat{U}|Y}$, whereas in the denominator, $\hat{U}$ is distributed according to the unconditional distribution $P_{\hat{U}}$.

The intuition behind the definition for MaxL is that X represents the original dataset, whereas Y represents disclosed data that is available to an adversary. The variable U represents an arbitrary (possibly random) function of X that the adversary is interested in learning. The numerator is the best probability of the adversary correctly guessing U based on Y, whereas the denominator is the best probability of the adversary correctly guessing U *without* Y. Thus, the ratio of these quantities characterizes the usefulness of Y toward learning U. Taking a supremum over all possible functions U yields the *most* an adversary could learn about some aspect of X.

MAL generalizes MaxL by taking into account that an adversary's ability to learn U is not limited to its probability of correctly guessing U. In particular, an adversary's ability to learn a variable U based on information Y can be characterized via a loss function in the following manner. Given a loss function $\ell(p)$ for $p \in [0,1]$, we assume that the adversary finds the random estimator $P_{\hat{U}|Y}$ that minimizes

$$\mathbb{E}\left[\ell(P_{\hat{U}|Y}(U|Y)) \right]. \quad (9)$$

One can think of $P_{\hat{U}|Y}(u|y)$ as the adversary's *belief* that $U=u$ given its knowledge that $Y=y$. The loss function characterizes its cost for having imperfect information about u. Note that the optimal belief $P_{\hat{U}|Y}$ depends on the loss function itself. MAL is based on the so-called α-loss function, which is defined as follows.

Definition 7: Given a probability $p \in [0,1]$ and a parameter $\alpha > 0$, the α-loss is given by

$$\ell_\alpha(p) \triangleq \begin{cases} \frac{\alpha}{\alpha-1}(1-p^{\frac{\alpha-1}{\alpha}}), & \alpha \in (0,1) \cup (1,\infty), \\ \log \frac{1}{p}, & \alpha = 1, \\ 1-p, & \alpha = \infty. \end{cases} \quad (10)$$

It is easy to see that $\ell_\alpha(p)$ is continuous in α. Note also that at $\alpha = 1$, the α-loss is equivalent to the log-loss.

Prior to defining the MAL, we present a related metric called the α-leakage, defined as follows.

Definition 8: Given a joint distribution $P_{X,Y}$, the α-leakage from X to Y for $\alpha \in (0,1) \cup (1,\infty)$ is defined as

$$\mathcal{L}_\alpha(X \to Y) \triangleq \frac{\alpha}{\alpha-1} \log \frac{\frac{\alpha}{\alpha-1} - \min_{P_{\hat{X}|Y}} \mathbb{E}\left[\ell_\alpha(P_{\hat{X}|Y}(X|Y)) \right]}{\frac{\alpha}{\alpha-1} - \min_{P_{\hat{X}}} \mathbb{E}\left[\ell_\alpha(P_{\hat{X}}(X)) \right]}. \quad (11)$$

Note that the definition of α-leakage has a similar structure to that of MaxL, except it does not have a supremum over U. That is, it is concerned with an adversary that is specifically interested only in X, rather than a function of X. In addition, the numerator and denominator characterize the adversary's ability to learn X with or without Y via the expected α-loss.

We now define MAL, which is simply the α-leakage maximized over all possible functions of X.

Definition 9: Given a joint distribution $P_{X,Y}$ the maximal α-leakage from X to Y for $\alpha > 0$ is given by

$$\mathcal{L}_\alpha^{\max}(X \to Y) \triangleq \sup_{U \to X \to Y} \mathcal{L}_\alpha(U \to Y)$$

where the supremum is over all random variables U satisfying the Markov chain.

It is not hard to see that MAL for $\alpha = \infty$ is precisely MaxL. In the following section, we will present several results illustrating that MAL can be effectively computed given a distribution on P_{XY}, as well as several of its important properties.

To account for side information, we now define conditional versions of both α-leakage and MAL. The conditional α-leakage is defined by conditioning on the side information variable in both the numerator and denominator.

Definition 10: Given a joint distribution $P_{X,Y,Z}$, the conditional α-leakage for $\alpha \in (0,1) \cup (1,\infty)$ from X to Y given Z is defined as

$$\mathcal{L}_\alpha(X \to Y|Z) \triangleq \frac{\alpha}{\alpha-1} \log \frac{\frac{\alpha}{\alpha-1} - \min\limits_{P_{\hat{X}|Y,Z}} \mathbb{E}\left[\ell_\alpha(P_{\hat{X}|Y,Z}(X|Y,Z))\right]}{\frac{\alpha}{\alpha-1} - \min\limits_{P_{\hat{X}|Z}} \mathbb{E}\left[\ell_\alpha(P_{\hat{X}|Z}(X|Z))\right]}. \tag{12}$$

Finally, the conditional maximal α-leakage is defined by taking a supremum over functions U of X.

Definition 11: Given a joint distribution $P_{X,Y,Z}$, for $\alpha > 0$ the conditional maximal α-leakage from X to Y given Z is defined as

$$\mathcal{L}_\alpha^{\max}(X \to Y|Z) \triangleq \sup_{U-(X,Z)-Y} \mathcal{L}_\alpha(U \to Y|Z). \tag{13}$$

Note that the Markov chain condition allows U to be a (random) function of both the original dataset X as well as the side information Z. In other words, the quantity of interest to the adversary U may be related to its side information in an arbitrary manner, the dependence between U and (X,Z) must be independent of that between Y and (X,Z).

III. Properties of Maximal α-Leakage

We first review several properties of the leakage measures that have to do with simplifying them. That is, as given in the definitions it is not clear that these measures are computable. In fact, they are, and can be written as simple functions of the Arimoto or Sibson mutual informations. The following lemma, the proof of which requires a fairly simple calculation given in Appendix A, is our first tool in simplifying these measures.

Lemma 1: Given a distribution P_X, the minimal expected α-loss for an estimator $\hat{X}$ of X is

$$\min_{P_{\hat{X}}} \mathbb{E}\left[\ell_\alpha(P_{\hat{X}}(X))\right] = \begin{cases} \frac{\alpha}{\alpha-1}\left(1 - \exp\left\{\frac{1-\alpha}{\alpha} H_\alpha(X)\right\}\right), & \alpha \in (0,1)\cup(1,\infty), \\ H(X), & \alpha = 1, \\ 1 - \exp(-H_\infty(X)), & \alpha = \infty. \end{cases} \tag{14}$$

This lemma allows us to rewrite the numerator and denominator in the definition of the α-leakage. From there, we can show that the α-leakage is precisely the Arimoto mutual information.

Theorem 2: For any $\alpha > 0$, the α-leakage from X to Y simplifies to

$$\mathcal{L}_\alpha(X \to Y) = I_\alpha^A(X;Y). \tag{15}$$

Similarly, the conditional α-leakage is equal to the conditional Arimoto mutual information.

Theorem 3: For any $\alpha > 0$, the conditional α-leakage from X to Y given Z simplifies to

$$\mathcal{L}_\alpha(X \to Y|Z) = I_\alpha^A(X;Y|Z). \tag{16}$$

Thm. 3 is proved in Appendix B; note that Thm. 2 is a special case. Simplifying MAL and conditional MAL requires a significantly more complicated argument which precisely characterizes the optimal U. The resulting expressions are as follows. Appendix C contains the proof for the conditional version; again the unconditional version is a special case.

Theorem 4: For $\alpha > 0$, the maximal α-leakage simplifies to

$$\mathcal{L}_\alpha^{\max}(X \to Y) = \begin{cases} \sup\limits_{P_{\tilde{X}} \ll P_X} I_\alpha^{\mathrm{S}}(\tilde{X};Y) = \sup\limits_{P_{\tilde{X}} \ll P_X} I_\alpha^{\mathrm{A}}(\tilde{X};Y), & \alpha \neq 1, \\ I(X;Y), & \alpha = 1. \end{cases} \tag{17}$$

where $P_{\tilde{X}} \ll P_X$ means that $P_{\tilde{X}}$ is absolutely continuous with respect to P_X; i.e., the support of $P_{\tilde{X}}$ is contained within that of P_X.

Theorem 5: For $\alpha > 0$, conditional maximal α-leakage simplifies to

$$\mathcal{L}_\alpha^{\max}(X \to Y|Z) = \begin{cases} \sup\limits_{z \in \mathrm{supp}(Z)} \sup\limits_{\substack{P_{\tilde{X}|Z=z} \\ \ll P_{X|Z=z}}} I_\alpha^{\mathrm{S}}(\tilde{X};Y|Z=z), & \alpha \neq 1, \\ I(X;Y|Z), & \alpha = 1 \end{cases} \tag{18}$$

where $\mathrm{supp}(Z)$ indicates the support of Z, and $I_\alpha^{\mathrm{S}}(\tilde{X};Y|Z=z)$ is the Sibson mutual information for the distribution $P_{\tilde{X},Y|Z=z}$.

The following theorem gives some other properties of MAL. Most of these are derived from properties of the Sibson mutual information.

Theorem 6: For $\alpha > 0$, maximal α-leakage

1. is quasi-convex in $P_{Y|X}$;
2. is monotonically non-decreasing in α for $\alpha \neq 1$;
3. satisfies data processing inequalities: let random variables X, Y, Z form a Markov chain, i.e., $X - Y - Z$, then
$$\mathcal{L}_\alpha^{\max}(X \to Z) \leq \mathcal{L}_\alpha^{\max}(X \to Y) \tag{19a}$$
$$\mathcal{L}_\alpha^{\max}(X \to Z) \leq \mathcal{L}_\alpha^{\max}(Y \to Z). \tag{19b}$$
4. satisfies
$$\mathcal{L}_\alpha^{\max}(X \to Y) \geq 0 \tag{20}$$
with equality if and only if X is independent of Y, and
$$\mathcal{L}_\alpha^{\max}(X \to Y) \leq \begin{cases} \log|\mathcal{X}| & \alpha \neq 1 \\ H(P_X) & \alpha = 1 \end{cases} \tag{21}$$

with equality if X is a deterministic function of Y.

Consider two disclosed versions Y_1 and Y_2 of X. The following theorem upper bounds the maximal α-leakage to an adversary who has access to both Y_1 and Y_2 simultaneously. This composition result allows composing multiple releases under a total leakage constraint.

Theorem 7: If X, Y_1, Y_2 satisfies the Markov chain $Y_1 - X - Y_2$, then

$$\mathcal{L}_\alpha^{\max}(X \to Y_1, Y_2) \leq \mathcal{L}_\alpha^{\max}(X \to Y_1) + \mathcal{L}_\alpha^{\max}(X \to Y_2). \tag{22}$$

The following theorem shows that in many scenarios of interest, the conditional MAL is upper bounded by MAL. This suggests that limiting the unconditional MAL also limits the amount an adversary can learn about the private data X, even if the adversary has access to side information, the details of which are completely unknown to the data curator.

Theorem 8: If the Markov chain $Z - X - Y$ holds, then

$$\mathcal{L}_\alpha^{\max}(X \to Y|Z) \leq \mathcal{L}_\alpha^{\max}(X \to Y). \tag{23}$$

IV. Relationship between Maximal α-Leakage and Rényi Differential Privacy

The maximal α-leakage depends on the statistics of the underlying dataset, and therefore, is regarded as a context-aware metric. On the contrary, differential privacy (DP) quantifies the worst case information leakage; as such, it is a context-free metric which is independent of the statistical information of data. Several variants of DP have been proposed for the sake of preserving utility, including Rényi differential privacy (RDP) [19]. RDP is superficially similar to MAL in that they both make use of variations of Rényi's information measures, but are in many ways different metrics. Even so, we show in this section that the two measures are equivalent in the sense of capturing the same collection of mechanisms which provide a specified level of privacy protection. This result implies that MAL can reach out to context-free metrics, and therefore, extends the scope of information leakage measures that can be linked to MAL.

RDP is based on the notion of adjacency, wherein two datasets are adjacent if they differ only in exactly one element [19]. To properly compare against MAL, which is defined without a notion of adjacent datasets, we extend RDP to the local privacy context [20] and formally define *local Rényi differential privacy* (LRDP) as follows.

Definition 12: Given a mechanism $P_{Y|X}$, the local Rényi differential privacy of order $\alpha > 0$ is given by

$$\mathcal{L}_\alpha^{\text{LRDP}}(X \to Y) \triangleq \sup_{x,x' \in \mathcal{X}} D_\alpha(P_{Y|X=x} \| P_{Y|X=x'}). \tag{24}$$

where $P_{Y|X=x}$ and $P_{Y|X=x'}$ are the two conditional probabilities of Y given $X = x$ and $X = x'$, respectively.

An alternative manner of defining LRDP, more in line with DP conventions, is to state that a mechanism $P_{Y|X}$ satisfies (α, γ)-LRDP if $\mathcal{L}_\alpha^{\text{LRDP}}(X \to Y) \leq \gamma$. Here, we find it more useful to define the measure as in (24), as it can be more easily compared to MAL.

We present the connection of privacy captured by MAL and LRDP in the following theorem, which is proved in Appendix D.

Theorem 9: For any mechanism $P_{Y|X}$ and any $\alpha > 0$,

$$\mathcal{L}_\alpha^{\max}(X \to Y) \leq \mathcal{L}_\alpha^{\text{LRDP}}(X \to Y). \tag{25}$$

Let $\tau = \min_{x,y} P_{Y|X}(y|x)$. For any $\alpha > 1$,

$$\mathcal{L}_\alpha^{\text{LRDP}}(X \to Y) \leq \log\left(1 + \frac{1}{\tau}\sqrt{\frac{2\mathcal{L}_\alpha^{\max}(X \to Y)}{\log e}}\right). \tag{26}$$

Note that the result in Theorem 9 is tight for perfect privacy, i.e., $\mathcal{L}_\alpha^{\max}(X \to Y) = 0$ iff $\mathcal{L}_\alpha^{\text{LRDP}}(X \to Y) = 0$.

From Thm. 9, we conclude that LRDP and MAL are equivalent in the sense that a mechanism has small MAL if and only if it has small LRDP.

Appendix A
Proof of Lemma 1

For $\alpha \in (0,1) \cup (1, \infty)$, the minimal expected α-loss is given by

$$\min_{P_{\hat{X}}} \mathbb{E}\left[\ell_\alpha(P_{\hat{X}}(X))\right] \tag{27}$$

$$= \min_{P_{\hat{X}}} \frac{\alpha}{\alpha - 1}\left(1 - \sum_x P_X(x) P_{\hat{X}}(x)^{\frac{\alpha-1}{\alpha}}\right). \tag{28}$$

Note that for $\alpha > 1$, $P_{\hat{X}}(x)^{\frac{\alpha-1}{\alpha}}$ is a concave function of $P_{\hat{X}}$, meaning the overall function in (28) is convex. For $\alpha < 1$, $P_{\hat{X}}(x)^{\frac{\alpha-1}{\alpha}}$ is a convex function of $P_{\hat{X}}$, but since $\frac{\alpha}{\alpha-1}$ is negative, the overall function is again convex. Either way, the minimization in (28) amounts to a convex optimization problem subject to the constraint that $P_{\hat{X}}$ is in the simplex. Incorporating the constraint that $\sum_x P_{\hat{X}}(x) = 1$, the Lagrangian is given by

$$L(P_{\hat{X}}, \nu) \triangleq \frac{\alpha}{\alpha - 1}\left(1 - \sum_x P_X(x) P_{\hat{X}}(x)^{\frac{\alpha-1}{\alpha}}\right) + \nu\left(\sum_x P_{\hat{X}}(x) - 1\right) \tag{29}$$

where ν is a Lagrange multiplier. Thus, the KKT condition is

$$0 = \frac{\partial}{\partial P_{\hat{X}}(x)} L(P_{\hat{X}}, \nu) = -P_X(x) P_{\hat{X}}(x)^{-\frac{1}{\alpha}} + \nu. \tag{30}$$

Solving for $P_{\hat{X}}$, and finding the correct value of ν to put $P_{\hat{X}}$ in the simplex, we find that the optimal $P_{\hat{X}}$ is

$$P_{\hat{X}}(x) = \frac{P_X(x)^\alpha}{\sum_{x'} P_X(x')^\alpha}. \tag{31}$$

Plugging this into the objective function, we find

$$\min_{P_{\hat{X}}} \mathbb{E}\left[\ell_\alpha(P_{\hat{X}}(X))\right] = \frac{\alpha}{\alpha - 1}\left[1 - \left(\sum_x P_X(x)^\alpha\right)^{\frac{1}{\alpha}}\right] \tag{32}$$

$$= \frac{\alpha}{\alpha-1}\left[1-\exp\left\{\frac{1-\alpha}{\alpha}H_\alpha(X)\right\}\right]. \tag{33}$$

Appendix B
Proof of Theorem 3

Assume $\alpha \in (0,1)\cup(1,\infty)$. Lemma 1 applies for an unconditional distribution P_X, but we can apply it to evaluate the denominator of the definition of conditional α-leakage in (12), where there is a conditional estimator $P_{\hat{X}|Z}$, as follows:

$$\min_{P_{\hat{X}|Z}} \mathbb{E}\left[\ell_\alpha(P_{\hat{X}|Z}(X|Z))\right] \tag{34}$$

$$= \min_{P_{\hat{X}|Z}} \sum_z P_Z(z)\mathbb{E}\left[\ell_\alpha(P_{\hat{X}|Z}(X|Z))\Big| Z=z\right] \tag{35}$$

$$= \sum_z P_Z(z) \min_{P_{\hat{X}|Z=z}} \mathbb{E}\left[\ell_\alpha(P_{\hat{X}|Z}(X|Z))\Big| Z=z\right] \tag{36}$$

$$= \sum_z P_Z(z)\frac{\alpha}{\alpha-1}\left(1-\exp\left\{\frac{1-\alpha}{\alpha}H_\alpha(X|Z=z)\right\}\right) \tag{37}$$

$$= \frac{\alpha}{\alpha-1}\left(1-\sum_z P_Z(z)\left(\sum_x P_{X|Z}(x|z)^\alpha\right)^{\frac{1}{\alpha}}\right) \tag{38}$$

$$= \frac{\alpha}{\alpha-1}\left(1-\exp\left\{\frac{1-\alpha}{\alpha}H_\alpha(X|Z)\right\}\right). \tag{39}$$

Similarly, the expression in the numerator of (12) is given by

$$\min_{P_{\hat{X}|Y,Z}} \mathbb{E}\left[\ell_\alpha(P_{\hat{X}|Y,Z}(X|Y,Z))\right] = \frac{\alpha}{\alpha-1}\left(1-\exp\left\{\frac{1-\alpha}{\alpha}H_\alpha(X|Y,Z)\right\}\right). \tag{40}$$

Therefore

$$\mathcal{L}_\alpha(X\to Y|Z) = \frac{\alpha}{\alpha-1}\log\frac{\exp\left\{\frac{1-\alpha}{\alpha}H_\alpha(X|Y,Z)\right\}}{\exp\left\{\frac{1-\alpha}{\alpha}H_\alpha(X|Z)\right\}} \tag{41}$$

$$= H_\alpha(X|Z) - H_\alpha(X|Y,Z) \tag{42}$$

$$= I_\alpha^A(X;Y|Z). \tag{43}$$

Appendix C
Proof of Theorem 5

Let $\alpha \in (0,1)\cup(1,\infty)$. The case of $\alpha=1$ is relatively simple and is addressed in [14]. From Thm. 3, we have

$$\mathcal{L}_\alpha^{\max}(X\to Y|Z) = \sup_{U-(X,Z)-Y} I_\alpha^{\mathrm{A}}(U;Y|Z). \tag{44}$$

Given any U satisfying $U-(X,Z)-Y$, we may upper bound

$$I_\alpha^{\mathrm{A}}(U;Y|Z) = \frac{\alpha}{\alpha-1}\log\frac{\sum_{y,z}\left(\sum_u P_{U,Y,Z}(u,y,z)^\alpha\right)^{\frac{1}{\alpha}}}{\sum_z\left(\sum_u P_{U,Z}(u,z)^\alpha\right)^{\frac{1}{\alpha}}} \tag{45}$$

$$\le \sup_{z\in\mathrm{supp}(Z)} I_\alpha^{\mathrm{A}}(U;Y|Z=z) \tag{46}$$

$$\le \sup_{z\in\mathrm{supp}(Z)}\ \sup_{P_{\hat{X}|\tilde{U}}:P_{\hat{X}|\tilde{U}}\ll P_{X|Z=z}}\ \sup_{P_{\tilde{U}}} I_\alpha^{\mathrm{A}}(\tilde{U};Y|Z=z) \tag{47}$$

$$= \sup_{z\in\mathrm{supp}(Z)}\ \sup_{P_{\hat{X}|\tilde{U}}:P_{\hat{X}|\tilde{U}}\ll P_{X|Z=z}}\ \sup_{P_{\tilde{U}}} I_\alpha^{\mathrm{S}}(\tilde{U};Y|Z=z) \tag{48}$$

$$\le \sup_{z\in\mathrm{supp}(Z)}\ \sup_{P_{\tilde{X}}\ll P_{X|Z=z}} I_\alpha^{\mathrm{S}}(\tilde{X};Y|Z=z) \tag{49}$$

where

- (46) follows from the fact that for nonnegative a_i, b_i, $\frac{\sum_i a_i}{\sum_i b_i} \le \max_i \frac{a_i}{b_i}$,
- (48) follows because Arimoto and Sibson MIs have the same supremum over the input distribution,
- (49) follows from the facts that Sibson MI satisfies the data processing inequality, and $\tilde{U}-\tilde{X}-Y|Z=z$ forms a Markov chain.

We now lower bound $\mathcal{L}_\alpha^{\max}(X;Y|Z)$ by constructing a specific U satisfying $U-(X,Z)-Y$. We will define a variable U with alphabet consisting of disjoint subsets $\mathcal{U}_{x,z}$ for each x,z. Let $n_{x,z}=|\mathcal{U}_{x,z}|$ to be determined later. Define

$$P_{U|X,Z}(u|x,z) = \begin{cases}\frac{1}{n_{x,z}}, & u\in\mathcal{U}_{x,z},\\ 0, & \text{otherwise.}\end{cases} \tag{50}$$

With some hindsight, we define random variables $\tilde{Z},\tilde{X}$ with joint distribution given by

$$P_{\tilde{Z}}(z) \sim \left(\sum_x n_{x,z}^{1-\alpha}P_{X,Z}(x,z)^\alpha\right)^{1/\alpha}, \tag{51}$$

$$P_{\tilde{X}|\tilde{Z}}(x|z) \sim n_{x,z}^{1-\alpha}P_{X,Z}(x,z)^\alpha \tag{52}$$

where $\sim$ indicates that the distribution is proportional to the RHS expression. Note that for any $\alpha\neq 1$, there exists choices for $n_{x,z}$ that make $P_{\tilde{X},\tilde{Z}}$ to be any distribution with the same support as $P_{X,Z}$. For U as constructed above, we can evaluate the conditional Arimoto mutual information as

$$\exp\left\{\frac{\alpha-1}{\alpha}I_\alpha^{\mathrm{A}}(U;Y|Z)\right\} \tag{53}$$

$$= \frac{\sum_{y,z}\left(\sum_u P_{U,Y,Z}(u,y,z)^\alpha\right)^{\frac{1}{\alpha}}}{\sum_z\left(\sum_u P_{U,Z}(u,z)^\alpha\right)^{\frac{1}{\alpha}}} \tag{54}$$

$$= \frac{\sum_{y,z}\left(\sum_x n_{x,z}^{1-\alpha}P_{X,Y,Z}(x,y,z)^\alpha\right)^{\frac{1}{\alpha}}}{\sum_z\left(\sum_x n_{x,z}^{1-\alpha}P_{X,Z}(x,z)^\alpha\right)^{\frac{1}{\alpha}}} \tag{55}$$

$$= \sum_z P_{\tilde{Z}}(z)\sum_y\left(\frac{\sum_x n_{x,z}^{1-\alpha}P_{X,Y,Z}(x,y,z)^\alpha}{\sum_x n_{x,z}^{1-\alpha}P_{X,Z}(x,z)^\alpha}\right)^{\frac{1}{\alpha}} \tag{56}$$

$$= \sum_z P_{\tilde{Z}}(z)\sum_y\left(\sum_x P_{\tilde{X}|\tilde{Z}}(x|z)P_{Y|X,Z}(y|x,z)^\alpha\right)^{\frac{1}{\alpha}} \tag{57}$$

$$= \sum_z P_{\tilde{Z}}(z)\exp\left\{\frac{\alpha-1}{\alpha}I_\alpha^S(\tilde{X};Y|\tilde{Z}=z)\right\}. \tag{58}$$

We may maximize over the choice $P_{\tilde{X},\tilde{Z}}$ (implicitly choosing $n_{x,z}$) to lower bound the conditional MAL by

$$\mathcal{L}_\alpha^{\max}(X \to Y|Z) \tag{59}$$

$$\geq \sup_{\substack{P_{\tilde{X},\tilde{Z}} \\ \ll P_{X,Z}}} \frac{\alpha}{\alpha-1} \log \sum_z P_{\tilde{Z}}(z) \exp\left\{\frac{\alpha-1}{\alpha} I_\alpha^S(\tilde{X};Y|\tilde{Z}=z)\right\} \tag{60}$$

$$= \sup_{z \in \mathrm{supp}(Z)} \sup_{\substack{P_{\tilde{X}|Z=z} \\ \ll P_{X|Z=z}}} I_\alpha^S(\tilde{X};Y|Z=z). \tag{61}$$

Appendix D
Proof of Theorem 9

We will make use of the following lemma, which bounds the Rényi divergence in terms of total variational distance.

Lemma 10: ([21, (1), (6)]) Let P and Q be two arbitrary probability distributions of the random variable X. For $\alpha > 1$,

$$\frac{1}{2}|P-Q|_{\mathrm{TV}}^2 \log e \leq D_\alpha(P\|Q), \tag{62}$$

$$D_\alpha(P\|Q) \leq \log\left(1 + \frac{|P-Q|_{\mathrm{TV}}}{2\min_x Q(x)}\right) \tag{63}$$

where $|P-Q|_{\mathrm{TV}} = \sum_x |P(x)-Q(x)|$ is the total variational distance.

Given a mechanism $P_{Y|X}$, we may upper bound the MAL by LRDP as follows:

$$\mathcal{L}_\alpha^{\max}(X \to Y) = \sup_{P_{\tilde{X}}} \inf_{Q_Y} D_\alpha(P_{\tilde{X}Y}\|P_{\tilde{X}} \times Q_Y) \tag{64}$$

$$\leq \inf_{Q_Y} \max_x D_\alpha(P_{Y|X=x}\|Q_Y) \tag{65}$$

$$\leq \max_{x,x'} D_\alpha(P_{Y|X=x}\|P_{Y|X=x'}) \tag{66}$$

$$= \mathcal{L}_\alpha^{\mathrm{LRDP}}(X \to Y) \tag{67}$$

where (64) follows from Thm. 4 and the definition of Sibson mutual information in (6).

We now prove the bound in the opposite direction. For any mechanism $P_{Y|X}$, we have

$$\mathcal{L}_\alpha^{\mathrm{LRDP}}(X \to Y) \tag{68}$$

$$= \max_{x',x} D_\alpha(P_{Y|X=x}\|P_{Y|X=x'}) \tag{69}$$

$$\leq \max_{x',x} \log\left(1 + \frac{|P_{Y|X=x} - P_{Y|X=x'}|_{\mathrm{TV}}}{2\min_{x,y} P(y|x)}\right) \tag{70}$$

$$\leq \inf_{Q_Y} \max_{x',x} \log\left(1 + \frac{|P_{Y|X=x} - Q_Y|_{\mathrm{TV}} + |Q_Y - P_{Y|X=x'}|_{\mathrm{TV}}}{2\tau}\right) \tag{71}$$

$$\leq \inf_{Q_Y} \max_x \log\left(1 + \frac{|P_{Y|X=x} - Q_Y|_{\mathrm{TV}}}{\tau}\right) \tag{72}$$

$$\leq \inf_{Q_Y} \max_x \log\left(1 + \frac{1}{\tau}\sqrt{\frac{2D_\alpha(P_{Y|X=x}\|Q_Y)}{\log e}}\right) \tag{73}$$

$$= \log\left(1 + \frac{1}{\tau}\sqrt{\frac{2\inf_{Q_Y} \max_x D_\alpha(P_{Y|X=x}\|Q_Y)}{\log e}}\right) \tag{74}$$

$$= \log\left(1 + \frac{1}{\tau}\sqrt{\frac{2\mathcal{L}_\alpha^{\max}(X \to Y)}{\log e}}\right) \tag{75}$$

where

- (70) follows by applying (63),
- (71) is due to the triangle inequality for total variation,
- (73) follows by applying (62),
- (75) follows from the equivalent form of MAL in Thm. 4.

References

[1] C. E. Shannon, "Communication theory of secrecy systems," *The Bell System Technical Journal*, vol. 28, no. 4, pp. 656–715, Oct. 1949.

[2] L. Sankar, S. R. Rajagopalan, and H. V. Poor, "Utility-privacy tradeoffs in databases: An information-theoretic approach," *IEEE Trans. on Inform. For. and Sec.*, vol. 8, no. 6, pp. 838–852, 2013.

[3] F. P. Calmon, M. Varia, and M. Médard, "On information-theoretic metrics for symmetric-key encryption and privacy," in *Proc. 52nd Annual Allerton Conf. on Commun., Control, and Comput.*, 2014.

[4] I. Issa, S. Kamath, and A. B. Wagner, "An operational measure of information leakage," in *2016 Annual Conference on Information Science and Systems (CISS)*, 2016, pp. 234–239.

[5] C. Dwork, "Differential privacy: A survey of results," in *Lecture Notes in Computer Science*. New York:Springer, Apr. 2008.

[6] J. Liao, O. Kosut, L. Sankar, and F. P. Calmon, "A tunable measure for information leakage," in *IEEE ISIT*, June 2018, pp. 701–705.

[7] N. Merhav and M. Feder, "Universal prediction," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2124–2147, Oct 1998.

[8] T. A. Courtade and R. D. Wesel, "Multiterminal source coding with an entropy-based distortion measure," in *IEEE International Symposium on Information Theory Proceedings*, July 2011, pp. 2040–2044.

[9] F. du Pin Calmon and N. Fawaz, "Privacy against statistical inference," in *50th Annual Allerton Conference on Communication, Control, and Computing*, 2012.

[10] S. P. Kasiviswanathan and A. D. Smith, "A note on differential privacy: Defining resistance to arbitrary side information," *arXiv:0803.3946v3*, 2015.

[11] 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.

[12] J. Liao, O. Kosut, L. Sankar, and F. P. Calmon, "Privacy under hard distortion constraints," in *2018 IEEE Information Theory Workshop (ITW)*, 2018, pp. 1–5.

[13] J. Liao, O. Kosut, L. Sankar, and F. du Pin 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.

[14] J. Liao, L. Sankar, O. Kosut, and F. P. Calmon, "Robustness of maximal α-leakage to side information," in *2019 IEEE International Symposium on Information Theory (ISIT)*, 2019, pp. 642–646.

[15] A. Rényi, "On measures of entropy and information," in *4th Berkeley Symp. Math. Stat. Prob.*, 1961, pp. 547–561.

[16] S. Verdú, "α-mutual information," in *2015 Information Theory and Applications Workshop (ITA)*, 2015.

[17] S. Arimoto, "Information measures and capacity of order α for discrete memoryless channels," in *Coll. Math. Soc.*, Hungary, 1975, pp. 41–52.

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

[19] I. Mironov, "Rényi differential privacy," in *Proceedings of 30th IEEE Computer Security Foundations Symposium (CSF)*, 2017, pp. 263–275.

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

[21] I. Sason and S. Verdú, "Upper bounds on the relative entropy and rényi divergence as a function of total variation distance for finite alphabets," in *2015 IEEE Information Theory Workshop-Fall (ITW)*. IEEE, 2015, pp. 214–218.