Pointwise Maximal Leakage on General Alphabets

Sara Saeidian, Giulia Cervia, Tobias J. Oechtering, Mikael Skoglund

Pointwise Maximal Leakage on General Alphabets

Sara Saeidian*, Giulia Cervia†, Tobias J. Oechtering*, and Mikael Skoglund*
*KTH Royal Institute of Technology, 100 44 Stockholm, Sweden, {saeidian, oech, skoglund}@kth.se
†IMT Nord Europe, Centre for Digital Systems, F-59000 Lille, France, giulia.cervia@imt-nord-europe.fr

***Abstract*—Pointwise maximal leakage (PML) is an operationally meaningful privacy measure that quantifies the amount of information leaking about a secret X to a single outcome of a related random variable Y. In this paper, we extend the notion of PML to random variables on arbitrary probability spaces. We develop two new definitions: First, we extend PML to countably infinite random variables by considering adversaries who aim to guess the value of discrete (finite or countably infinite) functions of X. Then, we consider adversaries who construct estimates of X that maximize the expected value of their corresponding gain functions. We use this latter setup to introduce a highly versatile form of PML that captures many scenarios of practical interest whose definition requires no assumptions about the underlying probability spaces.**

I. Introduction

Recently, the problem of private data analysis has attracted much attention from an information-theoretic perspective. A wide variety of privacy measures, for example, mutual information [1–5], measures based on f-divergences [6, 7], probability of correctly guessing [8], information privacy [9, 10], and log-lift [11, 12] are studied that aim to quantify the amount of information leaking about a (private) random variable X by disclosing a related random variable Y. (See [13] for a recent survey on information-theoretic privacy measures.) *Pointwise maximal leakage* (PML) [14, 15] is one such measure that is particularly robust and operationally meaningful. Introduced by Saeidian et al. [14], PML is a generalization of the pre-existing notion of *maximal leakage* [16–18]. While maximal leakage quantifies the information leaking from the average outcome of the random variable Y, PML can be used to measure the information leaking from each individual outcome $Y = y$. Hence, the pointwise definition sets up a more flexible framework in which information leakage is viewed as a random variable that can be bounded and controlled in different ways [14, Sec. III]. The original definition of maximal leakage is also retrieved as the expected value of the information leakage random variable.

To define PML, [14] exploits and unifies two seemingly different formalizations of maximal leakage: the *randomized function view* [18] and the *gain function view* [19]. The randomized function view of leakage assumes that an adversary attempts to guess the value of an arbitrary discrete (randomized) function of X, denoted by U. Then, PML is defined as (the logarithm of) the multiplicative increase in the probability of correctly guessing the value of U after observing an outcome $Y = y$, compared to the prior probability of correctly guessing the value of U. On the other hand, the gain function view of leakage considers an adversary who wishes to maximize the expected value of an arbitrary non-negative gain function. In this case, PML is defined as (the logarithm) of the multiplicative increase in the expected gain of an adversary who has observed $Y = y$, compared to the prior expected gain. In [14, Thm. 2], it is shown that when X takes values in a finite alphabet, the two definitions of leakage are in fact equivalent.

While PML is a robust and meaningful privacy measure, currently its definition is restricted to finite alphabet random variables. Our goal in this paper is to extend the notion of PML to random variables on general measurable spaces. In [18, Thm. 7], Issa et al. undertake a similar task and define a version of maximal leakage on general alphabets. Their definition, however, has certain limitations. Most importantly, while X is assumed to be a random variable on a measurable space, their setup still concerns an adversary who aims to guess the value of a finite alphabet random variable U, i.e., they consider the randomized function view of leakage. While this setup is suitable when X has a discrete (finite or countably infinite) alphabet, it results in a conceptually weaker definition in the general case as it is assumed that adversaries do not exploit the generalized structure of the space of X which is no longer restricted to be purely atomic. For example, if X is an absolutely continuous random variable (with respect to the Lebesgue measure), then an adversary can reasonably aim to construct a real-valued guess of X so as to maximize a gain function that is decreasing in the mean squared error. Such attack scenarios cannot be directly studied using the randomized function view of leakage.

In this paper, we generalize the definition of PML in two directions: First, in Section III, we extend the randomized function view of leakage to countable probability spaces (Thm. 2). Then, in Section IV, we use the gain function view of leakage to obtain a universal definition of PML that requires no assumptions about the underlying probability spaces (Thm. 3). This latter setup considers an adversary who is interested in maximizing the expected value of an arbitrary non-negative and measurable gain function, so the resulting notion of privacy is highly robust. We show that in both cases PML can be written as the Rényi divergence [20, 21] of order infinity of the posterior distribution of X from the

This work has been supported by the Strategic Research Agenda Program, Information and Communication Technology – The Next Generation (SRA ICT – TNG) funded by the Swedish Government and KTH Digital Futures center within the project DataLEASH.

prior distribution of X. Finally, in Corollary 1 we discuss the nuances of defining an information leakage random variable in the general setup and show how densities can be useful for deriving a form of PML that is a measurable function of Y. Throughout the paper, we give examples of common attack scenarios and evaluate PML for typical mechanisms, for example, adding independent Gaussian noise to a Gaussian random variable X (Example 4).

II. Notation and Background

A. Notation

We adopt the following notational conventions: $\mathbb{R}_+ = [0,\infty)$, $\bar{\mathbb{R}}_+ = [0,\infty]$, $\mathbb{Z} = \{\ldots,-1,0,1,\ldots\}$, $\mathbb{N} = \{0,1,\ldots\}$, $\mathbb{N}^* = \{1,2,\ldots\}$, and $[n] = \{1,\ldots,n\}$ with $n \in \mathbb{N}^*$. Sets are represented by uppercase letters, e.g., E, and $\mathbb{1}_E$ denotes the indicator function of E. Collections of sets are represented by calligraphic letters, e.g., $\mathcal{E}$. If E is a topological space then $\mathcal{B}_E$ denotes the Borel σ-algebra on E. Given a measurable space $(E,\mathcal{E})$, we use $\mathcal{E}_+$ to denote the set of all functions that are measurable relative to $\mathcal{E}$ and $\mathcal{B}_{\bar{\mathbb{R}}_+}$. Suppose μ and ν are measures on $(E,\mathcal{E})$ and assume that μ is σ-finite. If $\nu \ll \mu$, then we write $p = \frac{d\nu}{d\mu}$, or alternatively, $\nu(dx) = p(x)\,\mu(dx)$ to imply that $\int_E f(x)\ \nu(dx) = \int_E f(x)\,p(x)\ \mu(dx)$ for all $f \in \mathcal{E}_+$, where $p \in \mathcal{E}_+$ is the Radon-Nikodym derivative of ν with respect to μ.

Throughout the paper, we assume that a probability space $(\Omega,\mathcal{H},\mathbb{P})$ is fixed in the background. Given $f \in \mathcal{H}_+$, the essential supremum of f with respect to $\mathbb{P}$ is $\operatorname{ess\,sup}_{\mathbb{P}} f = \sup\{c \in \mathbb{R}_+ \colon \mathbb{P}(f > c) > 0\}$. A mapping $X : \Omega \to E$ is called a random variable taking values in $(E,\mathcal{E})$ if X is measurable relative to $\mathcal{H}$ and $\mathcal{E}$. We use P_X to denote the distribution of X. A mapping $P_{Y|X} : E \times \mathcal{F} \to [0,1]$ is called a transition probability kernel (or simply kernel) from $(E,\mathcal{E})$ into $(F,\mathcal{F})$ if the mapping $x \mapsto P_{Y|X=x}(B)$ is in $\mathcal{E}_+$ for all $B \in \mathcal{F}$, and $P_{Y|X=x}(\cdot)$ is a probability measure on $(F,\mathcal{F})$ for all $x \in E$. Let P_{XY} be a probability measure on the product space $(E \times F, \mathcal{E} \otimes \mathcal{F})$ with marginals P_X and P_Y. Then we write $P_{XY}(dx,dy) = P_X(dx)\,P_{Y|X=x}(dy)$ to imply that $\mathbb{E}\,f = \int_{E\times F} f(x,y) P_{XY}(dx,dy) = \int_E P_X(dx) \int_F f(x,y)\ P_{Y|X=x}(dy)$ for all $f \in (\mathcal{E}\otimes\mathcal{F})_+$.

Suppose Y is a random variable taking values in $(F,\mathcal{F})$. Let σY denote the σ-algebra generated by Y on Ω. We use $\mathbb{E}[f \mid Y]$ to denote the conditional expectation of $f \in \mathcal{H}_+$ given σY. Since $\mathbb{E}[f \mid Y] \in (\sigma Y)_+$, then there exists $\phi \in \mathcal{F}_+$ such that $\mathbb{E}[f \mid Y] = \phi \circ Y$. Hence, we use the notation $\mathbb{E}[f \mid Y = y]$ to represent $\phi(y)$ for each $y \in F$.

B. PML for finite alphabet random variables

We begin by recalling the definition of Rényi divergence of order infinity [20, 21], which we later use to provide simplified expressions for PML.

Definition 1 (Rényi divergence of order ∞ [21, Thm. 6]): Let P and Q be probability measures on the measurable space $(\Omega,\mathcal{H})$. Let μ be a σ-finite measure satisfying $P \ll \mu$ and $Q \ll \mu$. The Rényi divergence of order ∞ of P from Q is defined as

$$D_\infty(P\|Q) = \log \sup_{A\in\mathcal{H}} \frac{P(A)}{Q(A)} = \log\left(\operatorname*{ess\,sup}_P \frac{p}{q}\right),$$

where $p = \frac{dP}{d\mu}$ and $q = \frac{dQ}{d\mu}$.
If $P \ll Q$ the divergence can also be expressed as

$$D_\infty(P\|Q) = \log\left(\operatorname*{ess\,sup}_P \frac{dP}{dQ}\right) = \log\left(\operatorname*{ess\,sup}_Q \frac{dP}{dQ}\right).$$

On the other hand, if $P \not\ll Q$, then $D_\infty(P\|Q) = \infty$.[1] When the sample space Ω is countable we may write (1) in the form

$$D_\infty(P\|Q) = \log\left(\sup_{\omega\in\Omega} \frac{P(\omega)}{Q(\omega)}\right).$$

In [14], Saeidian et al. use two different threat models to introduce PML on finite spaces: the randomized function view of leakage [14, Def. 1] based on the setup of Issa et al. [18], and the gain function view of leakage [14, Cor. 1] based on the setup of Alvim et al. [19]. The following definition is a generalization of [14, Def. 1] where the alphabets of X, Y, and U are no longer restricted to be finite but are allowed to be countable.

Definition 2 (Randomized function view of leakage): Suppose X is a random variable on the discrete set E, and Y is a random variable on the discrete set F. Let P_{XY} denote the joint distribution of X and Y. The pointwise maximal leakage from X to $y \in F$ is defined as[2]

$$\ell_{P_{XY}}(X \to y) \coloneqq \log \sup_{P_{U|X}} \frac{\sup_{P_{\hat{U}|Y}} \mathbb{P}\Big(U = \hat{U} \mid Y = y\Big)}{\max_{u\in G} P_U(u)},$$

where U is any random variable on a countable set G such that the Markov chain $U - X - Y$ holds.

Theorem 1 ([14, Thm. 1]): If X has a finite support, then the pointwise maximal leakage from X to $y \in F$, described by Definition 2, is given by

$$\ell_{P_{XY}}(X \to y) = D_\infty(P_{X|Y=y}\|P_X),$$

where $P_{X|Y=y}$ denotes the posterior distribution of X given $y \in F$.

III. PML for Discrete Random Variables: Randomized Function View

In this section, we show that the result of Theorem 1 can be extended to cases where the support of X is countably infinite. Note that the randomized function view described by Definition 2 is a special case of the more general gain function view presented in Definition 3 (see Example 2). Still, Definition 2 is of independent interest for us since it provides a strong operational meaning for PML on discrete probability spaces while its corresponding result in Theorem 2 can be proved without any measure theoretic intricacies.

Theorem 2: Let X and Y be random variables taking values in the countable sets E and F, respectively. The pointwise

[1] We use the conventions that $0/0 = 1$ and $x/0 = \infty$ if $x > 0$.

[2] To be able to define the leakage for all $y \in F$ when F is a countable set, we may assume that $\mathbb{P}(\cdot \mid Y = y) = \mathbb{P}(\cdot)$ if $P_Y(y) = 0$. That is, conditioning on events with probability zero equals no conditioning.

maximal leakage from X to $y \in F$, described by Definition 2, is given by

$$\ell_{P_{XY}}(X \to y) = D_\infty(P_{X|Y=y} \| P_X).$$

Theorem 2 is proved in Appendix A.

Example 1 (Geometric distribution): Suppose $X \sim \mathrm{Geom}(p)$ with $p \in (0,1)$. Let Y be a binary random variable defined through the kernel $P_{Y|X=x}(0) = 1 - P_{Y|X=x}(1) = q^x$ with $q \in (0,1)$ and $x \in \mathbb{N}^*$. Then, $P_Y(0) = 1 - P_Y(1) = \frac{pq}{1-q+pq}$, and

$$\ell_{P_{XY}}(X \to 0) = \log \sup_{x \in \mathbb{N}^*} \frac{P_{Y|X=x}(0)}{P_Y(0)} = \log \frac{1-q+pq}{p},$$

$$\ell_{P_{XY}}(X \to 1) = \log \sup_{x \in \mathbb{N}^*} \frac{P_{Y|X=x}(1)}{P_Y(1)} = \log \frac{1-q+pq}{1-q}.$$

IV. PML on Arbitrary Probability Spaces: Gain Function View

In this section, we present a universal definition of PML that not only showcases the robustness of PML as a privacy measure but also requires no assumptions about the underlying probability spaces. The setup is an extension of [19].

Definition 3 (Gain function view of leakage): Suppose X is a random variable taking values in $(E, \mathcal{E})$ with distribution P_X and Y is a random variable taking values in $(F, \mathcal{F})$. Let P_{XY} denote the joint distribution of X and Y. The pointwise maximal leakage from X to $y \in F$ is

$$\ell_{P_{XY}}(X \to y) \coloneqq \log \sup_{\substack{(D,\mathcal{D}),\\ g \in \Gamma}} \frac{\sup_{P_{W|Y}} \mathbb{E}[g(X,W) \mid Y = y]}{\sup_{w \in D} \mathbb{E}[g(X,w)]}, \quad (1)$$

where the supremum in the numerator is over all transition probability kernels $P_{W|Y}$ from $(F, \mathcal{F})$ into $(D, \mathcal{D})$, and Γ denotes the set of all gain functions defined as

$$\Gamma \coloneqq \{g \in (\mathcal{E} \otimes \mathcal{D})_+ \mid \sup_{w \in D} \mathbb{E}[g(X,w)] < \infty\}.$$

The above definition models an adversary who is interested in constructing an estimate of X, denoted by W, which would maximize the expected value of her gain function. W is a random variable taking values in $(D, \mathcal{D})$ and gain functions are picked from the set Γ. Then, to measure the amount of information leaking about X by disclosing an outcome $Y = y$, we evaluate the ratio of the adversary's expected gain given observation $y \in F$, and her expected gain without any observations. PML is then defined by taking the supremum of this ratio over all possible measurable spaces $(D, \mathcal{D})$ and all $g \in \Gamma$. Note that the requirement $\sup_{w \in D} \mathbb{E}[g(X,w)] < \infty$ implies that the adversary chooses a function such that her expected gain can potentially be improved upon observing $y \in F$.

Below, we provide two examples of gain functions that describe typical attack scenarios. The first one concerns an adversary who wishes to guess the value of a discrete function of X, denoted by U, which retrieves the setup of [18, Thm. 7]. This setup can be used to model a hypothesis-testing adversary. For example, we may take $U = \mathbb{1}_{A^*}(X)$ to model a binary hypothesis test for determining whether or not X is in the set $A^* \in \mathcal{E}$. The second example describes an adversary who aims to approximate the value of a random variable on a separable metric space.

Example 2 (Guessing a discrete function of X): Suppose U is a discrete random variable taking values in the set A and induced by the kernel $P_{U|X}$. To model an adversary who is interested in guessing the value of U we let $D = A$, define $\mathcal{D}$ to be the discrete σ-algebra on A (i.e., its power set), and express the gain function $g_\bullet$ as follows:

$$g_\bullet(x,w) = P_{U|X=x}(w), \quad x \in E, w \in D.$$

In this case, the denominator of (1) describes the prior probability of correctly guessing the value of U whereas the numerator represents the posterior probability of correctly guessing U given $Y = y$.

Example 3 (Approximate guessing in metric spaces): Let (A, ρ) be a complete and separable metric space. Suppose U is a random variable taking values in $(A, \mathcal{B}_A)$ induced by a kernel $P_{U|X}$. Fix $\varepsilon > 0$. Our goal is to model an adversary who attempts to guess the value of U within a radius of ε. Suppose D is a countable dense subset of A and $\mathcal{D}$ is the discrete σ-algebra on D. Let $B_\varepsilon(w) = \{a \in A : \rho(a,w) < \varepsilon\}$ denote the open ball of radius ε centered at $w \in D$. Consider the gain function $g_\sim$ defined as

$$g_\sim(x,w) = P_{U|X=x}(B_\varepsilon(w)), \quad x \in E, w \in D.$$

Note that the countability of D ensures that $g_\sim$ defined above is $\mathcal{E} \otimes \mathcal{D}$-measurable. Then, for fixed $w \in D$ we have

$$\mathbb{E}[g_\sim(X,w)] = \int P_{U|X=x}(B_\varepsilon(w))\, P_X(dx) = P_U(B_\varepsilon(w)) = \mathbb{P}[U \in B_\varepsilon(w)].$$

Hence, evaluating the denominator of (1) with $g_\sim$ yields the prior probability of approximately guessing U. Similarly, it can be verified that the numerator of (1) evaluated with $g_\sim$ describes the posterior probability of approximately guessing U given $Y = y$.

We are now in place to state the main result of our paper: that PML in the form described by Definition 3 can too be written as the Rényi divergence of the posterior distribution of X from the prior distribution of X. The proof is inspired by a result of van Erven and Harremoës [21, Thm. 2] where it is shown that the general expression for Rényi divergence can be written as the supremum of the divergence evaluated over all finite and measurable partitions of the underlying σ-algebra.

Theorem 3 requires a single assumption: that the joint distribution P_{XY} can be disintegrated into the marginal P_Y and a transition probability kernel $P_{X|Y}$ from $(F, \mathcal{F})$ into $(E, \mathcal{E})$. We discuss this assumption in Remark 1.

Theorem 3: Suppose there exists a transition probability kernel $P_{X|Y}$ from $(F, \mathcal{F})$ into $(E, \mathcal{E})$ such that $P_{XY}(dx, dy) =$

$P_Y(dy)P_{X|Y=y}(dx)$. Then, the pointwise maximal leakage from X to $y \in F$, described by Definition 3, is given by

$$\ell_{P_{XY}}(X \to y) = D_\infty(P_{X|Y=y}\|P_X). \tag{2}$$

Proof: Fix $y \in F$. To simplify the numerator of (1) we make use of the following lemma proved in Appendix B.

Lemma 1: Let W be a random variable induced by a transition probability kernel $P_{W|Y}$ from $(F,\mathcal{F})$ into $(D,\mathcal{D})$. Given $g \in \Gamma$, if there exists a transition probability kernel $P_{X|Y}$ from $(F,\mathcal{F})$ into $(E,\mathcal{E})$ such that $P_{XY}(dx,dy) = P_Y(dy)P_{X|Y=y}(dx)$, then

$$\sup_{P_{W|Y}} \mathbb{E}[g(X,W)\,|\,Y=y] = \sup_{w\in D}\int_E g(x,w)P_{X|Y=y}(dx), \tag{3}$$

where the supremum on the LHS is over all kernels from $(F,\mathcal{F})$ into $(D,\mathcal{D})$.

First, we assume that $P_{X|Y=y} \ll P_X$. For notational convenience, let $f(x) \coloneqq \frac{dP_{X|Y=y}}{dP_X}(x)$ denote the Radon-Nikodym derivative of $P_{X|Y=y}$ with respect to P_X. We begin by showing that $\ell_{P_{XY}}(X \to y) \le D_\infty(P_{X|Y=y}\|P_X)$. Fix an arbitrary measurable space $(D,\mathcal{D})$, and a gain function $g \in \Gamma$. We can write

$$\begin{aligned}
\frac{\sup_{P_{W|Y}} \mathbb{E}[g(X,W)\mid Y=y]}{\sup_{w\in D}\mathbb{E}[g(X,w)]} &= \frac{\sup_{w\in D}\int_E g(x,w)\,P_{X|Y=y}(dx)}{\sup_{w\in D}\int_E g(x,w)\,P_X(dx)}\\
&\le \sup_{w\in D}\frac{\int_E g(x,w)\,P_{X|Y=y}(dx)}{\int_E g(x,w)\,P_X(dx)}\\
&= \sup_{w\in D}\frac{\int_E g(x,w)\,f(x)\,P_X(dx)}{\int_E g(x,w)\,P_X(dx)}\\
&\le \operatorname*{ess\,sup}_{P_X} f\\
&= \exp\Big(D_\infty(P_{X|Y=y}\|P_X)\Big).
\end{aligned}$$

Thus, $\ell_{P_{XY}}(X \to y) \le D_\infty(P_{X|Y=y}\|P_X)$.

Now, we show that $\ell_{P_{XY}}(X \to y) \ge D_\infty(P_{X|Y=y}\|P_X)$. Since $P_{X|Y=y} \ll P_X$ we may, without loss of generality, assume that $f(x) < \infty$ for all $x \in E$. Let $D = \mathbb{Z}\cup\{-\infty\}$, and suppose $\mathcal{D}$ is the discrete σ-algebra on D. Fix $\varepsilon > 0$ and consider the following (countable and disjoint) partition of E:

$$B_w^\varepsilon = \{x \in E : e^{w\varepsilon} \le f(x) < e^{(w+1)\varepsilon}\}, \quad w \in D, \tag{4}$$

which is indexed by D. Note that since $f(x)$ is $\mathcal{E}$-measurable, $B_w^\varepsilon \in \mathcal{E}$ for all $w \in D$. Let us define the gain function $g^* : E\times D \to \mathbb{R}_+$ as follows:

$$g^*(x,w) = \begin{cases} \frac{1}{P_X(B_w^\varepsilon)}\,\mathbb{1}_{B_w^\varepsilon}(x) & \text{if } P_X(B_w^\varepsilon) > 0,\\ 0 & \text{if } P_X(B_w^\varepsilon) = 0.\end{cases}$$

Then, we can write

$$\begin{aligned}
\exp\Big(\ell_{P_{XY}}(X\to y)\Big) &\ge \frac{\sup_{P_{W|Y}}\mathbb{E}[g^*(X,W)\mid Y=y]}{\sup_{w\in D}\mathbb{E}[g^*(X,w)]}\\
&= \frac{\sup_{w\in D}\int_E g^*(x,w)\,P_{X|Y=y}(dx)}{\sup_{w\in D}\int_E g^*(x,w)\,P_X(dx)}\\
&= \frac{\sup_{w\in D:P_X(B_w^\varepsilon)>0}\int_E \frac{1}{P_X(B_w^\varepsilon)}\mathbb{1}_{B_w^\varepsilon}(x)\,P_{X|Y=y}(dx)}{\sup_{w\in D:P_X(B_w^\varepsilon)>0}\int_E \frac{1}{P_X(B_w^\varepsilon)}\mathbb{1}_{B_w^\varepsilon}(x)\,P_X(dx)}\\
&= \frac{\sup_{w\in D:P_X(B_w^\varepsilon)>0}\frac{P_{X|Y=y}(B_w^\varepsilon)}{P_X(B_w^\varepsilon)}}{\sup_{w\in D:P_X(B_w^\varepsilon)>0}\frac{P_X(B_w^\varepsilon)}{P_X(B_w^\varepsilon)}}\\
&= \sup_{w\in D:P_X(B_w^\varepsilon)>0}\frac{P_{X|Y=y}(B_w^\varepsilon)}{P_X(B_w^\varepsilon)} \qquad (5a)\\
&= \operatorname*{ess\,sup}_{P_X}\bar{f}, \qquad (5b)
\end{aligned}$$

where

$$\bar{f}(x) \coloneqq \sum_{w\in D}\frac{P_{X|Y=y}(B_w^\varepsilon)}{P_X(B_w^\varepsilon)}\,\mathbb{1}_{B_w^\varepsilon}(x), \quad x\in E.$$

In (5b), we have written (5a) as a function of x. We have replaced the supremum over w in (5a) with the essential supremum over x in (5b) because $\bar{f}$ is constant on each set B_w^ε. Note that $\bar{f}(x) < \infty$ for all $x \in E$ even if there exists $w \in D$ such that $P_X(B_w^\varepsilon) = 0$. This is because $P_{X|Y=y} \ll P_X$ and we use the convention that $0/0 = 1$.

Let $\mathcal{G} \coloneqq \sigma\{B_w^\varepsilon\}$ denote the σ-algebra on E generated by the collection of sets $\{B_w^\varepsilon\}$. We argue that $\bar{f}$ is (a version of) the conditional expectation of f given $\mathcal{G}$, that is, $\bar{f} = \mathbb{E}[f\mid\mathcal{G}]$. Clearly, $\bar{f}$ is $\mathcal{G}$-measurable, so we should verify that $\int_A f\,dP_X = \int_A \bar{f}\,dP_X$ for all $A \in \mathcal{G}$. It is, however, sufficient to verify this equality for $A = B_w^\varepsilon$ because each non-empty set in $\mathcal{G}$ can be written as a countable union of sets in $\{B_w^\varepsilon\}$ and the monotone convergence theorem ensures that $\int_{\cup_i C_i} f\,dP_X = \sum_i \int_{C_i} f\,dP_X$ for each countable collection of disjoint sets $\{C_i\}$ in $\mathcal{G}$. Thus, by noting that

$$\int_{B_w^\varepsilon}\bar{f}\,dP_X = P_{X|Y=y}(B_w^\varepsilon) = \int_{B_w^\varepsilon} f\,dP_X,$$

for all $w \in D$ we conclude that $\bar{f} = \mathbb{E}[f\mid\mathcal{G}]$.

Finally, we can write

$$\begin{aligned}
\ell_{P_{XY}}(X\to y) &\ge \log\operatorname*{ess\,sup}_{P_X}\mathbb{E}[f\mid\mathcal{G}]\\
&\ge \log\left(\left(\operatorname*{ess\,sup}_{P_X} f\right)e^{-\varepsilon}\right) \qquad (6a)\\
&= \log\operatorname*{ess\,sup}_{P_X} f - \varepsilon\\
&= D_\infty(P_{X|Y=y}\|P_X) - \varepsilon,
\end{aligned}$$

where (6a) is due to the fact that by the definition of the sets $\{B_w^\varepsilon\}$ in (4), $\mathbb{E}[f\mid\mathcal{G}]$ never differs from f by more than a factor of e^ε. Then, letting $\varepsilon \to 0$, we obtain $\ell(X\to y) \ge D_\infty(P_{X|Y=y}\|P_X)$, which completes the proof for the case $P_{X|Y=y} \ll P_X$.

On the other hand, if $P_{X|Y=y} \not\ll P_X$ then there exists $A_0 \in \mathcal{E}$ such that $P_X(A_0) = 0$ and $P_{X|Y=y}(A_0) > 0$. Let $(D,\mathcal{D})$ be

an arbitrary measurable space, and consider the gain function $g(x,w) = \mathbb{1}_{A_0}(x)$ for all $w \in D$. Then, it is easy to see that $\mathbb{E}[g(X,W) \mid Y=y] = P_{X|Y=y}(A_0) > 0$ for all kernels $P_{W|Y}$ while $\sup_{w\in D} \mathbb{E}[g(X,w)] = 0$. Hence, $\ell(X \to y) = D_\infty(P_{X|Y=y}\|P_X) = \infty$, as desired. ∎

Remark 1: Theorem 3 assumes that the joint distribution P_{XY} can be disintegrated into the marginal P_Y and a kernel $P_{X|Y}$. This can be achieved in different ways. For example, we may start with a distribution P_Y and a kernel $P_{X|Y}$ and construct P_{XY} such that it satisfies $P_{XY}(dx,dy) = P_Y(dy)P_{X|Y=y}(dx)$. Otherwise, we may assume that $(E,\mathcal{E})$ is a Borel space [22, Def. 8.35] in which case the existence of a regular version of the conditional probability $\mathbb{P}[\,\cdot \mid Y]$ restricted to $\sigma X \subset \mathcal{H}$ is guaranteed [23, Thm. IV.2.18]. In this latter case, $P_{X|Y}$ is any kernel satisfying $\mathbb{P}[X \in A \mid Y](\omega) = P_{X|Y=Y(\omega)}(A)$ for $\mathbb{P}$-almost all $\omega \in \Omega$ and all $A \in \mathcal{E}$. Hence, Theorem 3 requires that P_{XY} can be disintegrated into a marginal distribution and a kernel, though it is immaterial how this is actually achieved. For a detailed discussion on disintegration theorems and the existence of regular conditional probabilities see [24].

Equipped with Theorem 3, we can calculate the information leaking from an observation $y \in F$. However, this result alone is insufficient for obtaining an information leakage random variable $\ell_{P_{XY}}(X \to Y)$. The difficulty is that the mapping $y \mapsto \ell_{P_{XY}}(X \to y)$ must be $\mathcal{F}$-measurable and there are certain nuances associated with this task. For example, we need to ensure that if $P_{X|Y=y} \ll P_X$, then the Radon-Nikodym derivative $\frac{dP_{X|Y=y}}{dP_X}$ is jointly measurable in (x,y), or that the set $\{y \in F : P_{X|Y=y} \ll P_X\}$ is measurable.

To obtain a measurable version of $\ell_{P_{XY}}(X \to y)$ we use the pragmatic assumption that the joint distribution P_{XY} is absolutely continuous with respect to the product of two σ-finite measures on $(E,\mathcal{E})$ and $(F,\mathcal{F})$ [25, Sec. 2.6]. This assumption also has the advantage of guaranteeing that $P_{XY}(dx,dy) = P_Y(dy)P_{X|Y=y}(dx)$ holds.

Corollary 1 (Privacy leakage random variable): Suppose P_{XY} is a probability measure on the product space $(E\times F, \mathcal{E}\otimes\mathcal{F})$ satisfying

$$P_{XY}(dx,dy) = p(x,y)\,\mu(dx)\,\nu(dy), \quad x \in E, y \in F, \quad (7)$$

where μ and ν are σ-finite measures on $(E,\mathcal{E})$ and $(F,\mathcal{F})$ respectively, and $p \in (\mathcal{E}\otimes\mathcal{F})_+$. Then, there exists a transition probability kernel $P_{X|Y}$ from $(F,\mathcal{F})$ into $(E,\mathcal{E})$ such that

1) $P_{XY}(dx,dy) = P_Y(dy)\,P_{X|Y=y}(dx)$;
2) $P_{X|Y=y} \ll P_X$ for ν-almost all $y \in F$; and
3) the mapping $y \mapsto \ell_{P_{XY}}(X \to y)$ is $\mathcal{F}$-measurable.

Corollary 1 is proved in Appendix C.

Remark 2: The assumption of Corollary 1 guarantees that $P_{XY} \ll P_X \otimes P_Y$, where $P_X \otimes P_Y$ denotes the product of P_X and P_Y. This assumption also allows us to write PML in different forms using densities:

$$\begin{aligned}\ell_{P_{XY}}(X \to y) &= \operatorname*{ess\,sup}_{P_X} i(X;y)\\ &= \log\left(\operatorname*{ess\,sup}_{P_X} \frac{f_{X|Y}(X,y)}{f_X(X)}\right)\\ &= \log\left(\operatorname*{ess\,sup}_{P_X} \frac{f_{Y|X}(y,X)}{f_Y(y)}\right),\end{aligned}$$

where $i(X;Y) = \log\frac{dP_{XY}}{dP_X\otimes P_Y}(X,Y)$ denotes the information density (the expectation of which is mutual information), $f_{X|Y} \in (\mathcal{E}\otimes\mathcal{F})_+$ denotes the density of $P_{X|Y}$ with respect to μ, and $f_X \in \mathcal{E}_+$ denotes the density of P_X with respect to μ. Densities $f_{Y|X}$ and f_Y are defined similarly.

Below, we calculate PML when X has Gaussian distribution and Y is obtained by adding independent Gaussian noise to X. Further examples are provided in Appendix E. Here, densities are defined with respect to the Lebesgue measure.

Example 4 (Additive Gaussian Noise): Suppose $Y = X+N$ where $X \sim \mathcal{N}(0,\sigma_X^2)$, $N \sim \mathcal{N}(0,\sigma_N^2)$, and X and N are independent. The PML from X to $y \in \mathbb{R}$ is given by

$$\begin{aligned}\ell_{P_{XY}}(X \to y) &= \log \sup_{x\in\mathbb{R}} \frac{f_{Y|X}(y,x)}{f_Y(y)}\\ &= \frac{1}{2}\log\left(1+\frac{\sigma_X^2}{\sigma_N^2}\right) + \frac{y^2}{2(\sigma_X^2+\sigma_N^2)}.\end{aligned}$$

As expected, for fixed $y \in \mathbb{R}$ and σ_X^2, taking $\sigma_N^2 \to \infty$ implies $\ell_{P_{XY}}(X \to y) \to 0$.

V. Conclusions

In this paper, we have extended the notion of PML to random variables with general alphabets. Theorem 2 describes a direct extension of the randomized function view of leakage from finite to countably infinite random variables. On the other hand, Theorem 3 illustrates that the gain function view of leakage can be used to define PML on arbitrary probability spaces.

The following points are worth emphasizing:

Properties of PML. We have shown that in all cases PML can be written as the Rényi divergence of order ∞ of the posterior distribution of X from the prior distribution of X. Hence, some properties of PML, such as non-negativity and satisfying a data-processing inequality, are directly inherited from the Rényi divergence [21]. We leave the detailed development of the properties of PML as future work.

Mechanism design via PML. In Examples 4 and 7, we have calculated PML in two typical setups involving the Gaussian distribution. These examples signify the advantage of PML over the average-case notion of maximal leakage: In many situations, including these Gaussian examples and even the discrete case of Example 5, $\exp\left(\ell_{P_{XY}}(X \to Y)\right)$ is not P_Y-integrable (i.e., its expectation is infinite). However, one can still characterize privacy using probability bounds on PML. Such bounds can be useful for designing privacy-preserving mechanisms, where simple mechanisms are conceived by adding noise with suitable parameters to X. We shall explore mechanism design with PML in future works.

References

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

[2] W. Wang, L. Ying, and J. Zhang, "On the relation between identifiability, differential privacy, and mutual-information privacy," *IEEE Transactions on Information Theory*, vol. 62, no. 9, pp. 5018–5029, 2016.

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

[4] B. Rassouli and D. Gündüz, "On perfect privacy," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 177–191, 2021.

[5] 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.

[6] M. Diaz, H. Wang, F. P. Calmon, and L. Sankar, "On the robustness of information-theoretic privacy measures and mechanisms," *IEEE Transactions on Information Theory*, vol. 66, no. 4, pp. 1949–1978, 2019.

[7] 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, 2019.

[8] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, "Estimation efficiency under privacy constraints," *IEEE Transactions on Information Theory*, vol. 65, no. 3, pp. 1512–1534, 2018.

[9] B. Jiang, M. Seif, R. Tandon, and M. Li, "Context-aware local information privacy," *IEEE Transactions on Information Forensics and Security*, 2021.

[10] F. P. Calmon and N. Fawaz, "Privacy against statistical inference," in *2012 50th annual Allerton conference on communication, control, and computing (Allerton)*. IEEE, 2012, pp. 1401–1408.

[11] H. Hsu, S. Asoodeh, and F. P. Calmon, "Information-theoretic privacy watchdogs," in *2019 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2019, pp. 552–556.

[12] P. Sadeghi, N. Ding, and T. Rakotoarivelo, "On properties and optimization of information-theoretic privacy watchdog," in *2020 IEEE Information Theory Workshop (ITW)*. IEEE, 2021, pp. 1–5.

[13] M. Bloch, O. Günlü, A. Yener, F. Oggier, H. V. Poor, L. Sankar, and R. F. Schaefer, "An overview of information-theoretic security and privacy: Metrics, limits and applications," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 5–22, 2021.

[14] S. Saeidian, G. Cervia, T. J. Oechtering, and M. Skoglund, "Pointwise maximal leakage," *arXiv preprint arXiv:2205.04935*, 2022.

[15] ——, "Pointwise maximal leakage," in *2022 IEEE International Symposium on Information Theory (ISIT)*, 2022, pp. 626–631.

[16] G. Smith, "On the foundations of quantitative information flow," in *International Conference on Foundations of Software Science and Computational Structures*. Springer, 2009, pp. 288–302.

[17] 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.

[18] 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, 2019.

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

[20] 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*. University of California Press, 1961, pp. 547–561.

[21] T. van Erven and P. Harremoës, "Rényi divergence and Kullback-Leibler divergence," *IEEE Transactions on Information Theory*, vol. 60, no. 7, pp. 3797–3820, 2014.

[22] A. Klenke, *Probability theory: A comprehensive course*. Springer Science & Business Media, 2013.

[23] E. Çinlar, *Probability and stochastics*. Springer, 2011, vol. 261.

[24] A. M. Faden, "The existence of regular conditional probabilities: Necessary and sufficient conditions," *The Annals of Probability*, pp. 288–298, 1985.

[25] Y. Polyanskiy and Y. Wu, "Lecture notes on information theory," 2019.

APPENDIX

A. Proof of Theorem 2

The finite case has been proved in [14], so here we assume that $E = \{x_1, x_2, \ldots\}$ is countably infinite. Fix $y \in F$. The argument for showing that $\ell_{P_{XY}}(X \to y) \leq D_\infty(P_{X|Y=y}\|P_X)$ is identical to the finite case laid out in [14, Thm. 1], so we only prove that $\ell_{P_{XY}}(X \to y) \geq D_\infty(P_{X|Y=y}\|P_X)$. Fix a small constant $\delta > 0$ and let k be an integer satisfying $\sum_k^\infty P_X(x_k) = P_X(A_\delta) < \delta$, where $A_\delta \coloneqq \{x_k, x_{k+1}, \ldots\}$. For notational convenience, we will use the shorthands $\xi(B) \coloneqq \dfrac{\sum_{x\in B} P_{X|Y=y}(x)}{\sum_{x\in B} P_X(x)}$ with $B \subset E$, and

$$\ell_U(X \to y) \coloneqq \log \frac{\sup_{P_{\hat{U}|Y}} \mathbb{P}\Big(U = \hat{U} \mid Y = y\Big)}{\max_{u\in G} P_U(u)},$$

for a given random variable U induced by a kernel $P_{U|X}$. We prove the inequality for the following two cases: First, we assume that there exists $x^* \in E$ such that $\log \xi(x^*) = D_\infty(P_{X|Y=y}\|P_X)$. Then, we consider the case where the supremum in the expression of the Rényi divergence is not attained by any $x \in E$.

Suppose $\log \xi(x^*) = D_\infty(P_{X|Y=y}\|P_X)$. Assume δ is sufficiently small so that x^* is not in the set A_δ. Define the finite random variable W with alphabet $D = \{1, \ldots, k\}$ according to the kernel $P_{W|X=x_i}(i) = 1$ for $i \in [k-1]$ and $P_{W|X=x_i}(k) = 1$ for $i \geq k$. Then, $P_W(i) = P_X(x_i)$ for $i \in [k-1]$ and $P_W(k) = P_X(A_\delta)$. Now, consider the random variable $U_\lightning$ induced by the *shattering channel* $P_{U_\lightning|W}$ defined in [18, Thm. 1], which yields

$$\begin{aligned}
\ell_{P_{XY}}(X \to y) &= \sup_{P_{U|X}} \ell_U(X \to y) \\
&\geq \ell_{U_\lightning}(X \to y) \\
&= \log \max_{i\in[k]} \frac{P_{W|Y=y}(i)}{P_W(i)} \qquad (8a) \\
&= \log \max\{\xi(x_1), \ldots, \xi(x^*), \ldots, \xi(x_{k-1}), \xi(A_\delta)\} \\
&\geq \log \xi(x^*) = D_\infty(P_{X|Y=y}\|P_X),
\end{aligned}$$

where (8a) is due to the definition of the shattering channel. Hence, $\ell_{P_{XY}}(X \to y) \geq D_\infty(P_{X|Y=y}\|P_X)$ holds in the first case.

Next, suppose the supremum in the expression of the Rényi divergence is not attained by any $x \in E$. Let $M = D_\infty(P_{X|Y=y}\|P_X)$, where $M \in (0, \infty]$. Since the supremum is not attained by any x, then we must have $\limsup_{n\to\infty} \log \xi(x_n) = M$. Fix an arbitrary $0 < m < M$, and define the set $B_m = \{x \in E : \xi(x) > m\}$. Clearly, for each $0 < m < M$ and $0 < \delta < 1$ the two sets A_δ and B_m have non-empty intersection. Define $E_{m,\delta} \coloneqq A_\delta \cap B_m$ and $E'_{m,\delta} \coloneqq A_\delta \setminus B_m$. Similarly to the previous case, we define a finite random variable W' with alphabet $D' = \{1, \ldots, k+1\}$ according to the kernel $P_{W'|X=x_i}(i) = 1$ for $i \in [k-1]$, $P_{W'|X=x}(k) = 1$ for $x \in E_{m,\delta}$, and $P_{W'|X=x}(k+1) = 1$ for $x \in E'_{m,\delta}$. Now, consider the random variable $U'_\lightning$ induced by the shattering channel $P_{U'_\lightning|W'}$. We obtain

$$\begin{aligned}
&\ell_{P_{XY}}(X \to y) \geq \ell_{U'_\lightning}(X \to y) \\
&= \log \max_{i\in[k+1]} \frac{P_{W'|Y=y}(i)}{P_{W'}(i)} \\
&= \log \max\big\{\xi(x_1), \ldots, \xi(x_{k-1}), \xi(E_{m,\delta}), \xi(E'_{m,\delta})\big\} \\
&\geq \log \xi(E_{m,\delta}) = \log \frac{\sum_{x\in E_{m,\delta}} P_{X|Y=y}(x)}{\sum_{x\in E_{m,\delta}} P_X(x)} \\
&\geq \log \inf_{x\in E_{m,\delta}} \frac{P_{X|Y=y}(x)}{P_X(x)} \geq m,
\end{aligned}$$

where the last inequality follows by the definitions of the sets $E_{m,\delta}$ and B_m. Finally, taking $m \to M$ yields

$$\ell_{P_{XY}}(X \to y) \geq M = D_\infty(P_{X|Y=y}\|P_X),$$

as desired.

B. Proof of Lemma 1

To show that the LHS of (3) lower bounds the RHS, fix an arbitrary kernel $P_{W|Y}$. We have

$$\begin{aligned}
&\mathbb{E}[g(X,W) \mid Y = y] = \int_{E\times D} g(x,w)\ P_{XW|Y=y}(dx, dw) \\
&= \int_D P_{W|Y=y}(dw) \int_E g(x,w)\ P_{X|Y=y}(dx) \\
&\leq \int_D P_{W|Y=y}(dw) \Big(\sup_{w\in D} \int_E g(x,w)\ P_{X|Y=y}(dx)\Big) \\
&= \sup_{w\in D} \int_E g(x,w)\ P_{X|Y=y}(dx).
\end{aligned}$$

Taking the supremum over all kernels $P_{W|Y}$ we obtain

$$\sup_{P_{W|Y}} \mathbb{E}[g(X,W) \mid Y = y] \leq \sup_{w\in D} \int_E g(x,w)\ P_{X|Y=y}(dx).$$

To show that the LHS of (3) upper bounds the RHS, fix an arbitrary $a < \sup_{w\in D} \int_E g(x,w)\ P_{X|Y=y}(dx)$. Then, there exists $w' \in D$ such that $\int_E g(x,w')\ P_{X|Y=y}(dx) \geq a$. Let δ_w denote the Dirac measure defined by

$$\delta_w(A) = \begin{cases} 1 & \text{if } w \in A, \\ 0 & \text{if } w \notin A, \end{cases}$$

for each $A \in \mathcal{D}$. We can write

$$\begin{aligned}
&\sup_{P_{W|Y}} \mathbb{E}[g(X,W) \mid Y = y] \\
&\qquad \geq \int_E P_{X|Y=y}(dx) \int_D g(x,w)\ \delta_{w'}(dw) \\
&\qquad = \int_E g(x,w')\ P_{X|Y=y}(dx)
\end{aligned}$$

$$\geq a.$$

Then, letting $a \to \sup_{w\in D} \int_E g(x,w)\ P_{X|Y=y}(dx)$ we obtain the desired inequality.

C. Proof of Corollary 1

Define the functions

$$q(y) \coloneqq \int_E p(x,y)\ \mu(dx), \quad y \in F,$$

$$r(x) \coloneqq \int_F p(x,y)\ \nu(dy), \quad x \in E,$$

$$k(x,y) \coloneqq \begin{cases} \frac{p(x,y)}{q(y)} & \text{if } q(y) > 0, \\ r(x) & \text{if } q(y) = 0, \end{cases} \quad x \in E, y \in F.$$

Let

$$P_{X|Y=y}(A) \coloneqq \int_A k(x,y)\ \mu(dx), \quad A \in \mathcal{E}, y \in F.$$

It can easily be checked that $P_{X|Y}$ defined above is a transition probability kernel from $(F,\mathcal{F})$ into $(E,\mathcal{E})$.

First, we show that $P_{X|Y=y} \ll P_X$ holds ν-almost everywhere. Suppose $A_0 \in \mathcal{E}$ satisfies $P_X(A_0) = 0$. Noting that $P_X(dx) = r(x)\mu(dx)$, we have

$$P_X(A_0) = \int_{A_0} r(x)\ \mu(dx) = \int_E \mathbb{1}_{A_0}(x)\, r(x)\ \mu(dx) = 0,$$

i.e., $\mathbb{1}_{A_0}(x)\, r(x) = 0$ μ-almost everywhere. In other words,

$$\mu(A_0 \cap \{x \in E : r(x) > 0\}) = 0. \tag{9}$$

Now, if $q(y) = 0$, then $P_{X|Y=y}(A_0) = P_X(A_0) = 0$ by construction. So, suppose $q(y) > 0$. In this case, we have

$$\begin{aligned} P_{X|Y=y}(A_0) &= \frac{1}{q(y)} \int_{A_0} p(x,y)\ \mu(dx) \\ &= \frac{1}{q(y)} \left(\int_{A_0\cap\{r>0\}} p(x,y)\ \mu(dx) \right. \\ &\qquad \left. + \int_{A_0\cap\{r=0\}} p(x,y)\ \mu(dx) \right). \end{aligned}$$

The first integral is zero due to (9). Moreover, for each $x \in E$, $r(x) = 0$ implies that $p(x,y) = 0$ ν-almost everywhere; thus, the second integral is also zero ν-almost everywhere. We conclude that $P_{X|Y=y} \ll P_X$ for ν-almost all $y \in F$.

Next, we argue that $P_{XY}(dx,dy) = P_Y(dy)\, P_{X|Y=y}(dx)$. Noting that $P_Y(dy) = q(y)\,\nu(dy)$ and $P_{X|Y=y}(dx) = k(x,y)\mu(dx)$ we write

$$\begin{aligned} P_Y(dy)\, P_{X|Y=y}(dx) &= q(y)\, k(x,y)\, \mu(dx)\, \nu(dy) \\ &= \begin{cases} p(x,y)\, \mu(dx)\, \nu(dy) & \text{if } q(y) > 0, \\ 0 & \text{if } q(y) = 0. \end{cases} \\ &= p(x,y)\, \mu(dx)\, \nu(dy) \qquad (10a) \\ &= P_{XY}(dx,dy), \end{aligned}$$

where (10a) is due to the fact that for each $y \in F$, $q(y) = 0$ implies $p(x,y) = 0$ μ-almost everywhere and a Radon-Nikodym derivative is specified uniquely up to almost everywhere equivalence.

Finally, we show that the mapping $y \mapsto \ell_{P_{XY}}(X \to y)$ is $\mathcal{F}$-measurable. Define the set $B_0 = \{y \in F : \int_{\{r=0\}} k(x,y)\, \mu(dx) = 0\}$ which is guaranteed to be in $\mathcal{F}$ by Fubini's theorem. The leakage $\ell_{P_{XY}}(X \to y)$ can be expressed as

$$\ell_{P_{XY}}(X \to y) = \begin{cases} \operatorname{ess\,sup}_{P_X} \left(\frac{k(x,y)}{r(x)} \right) & \text{if } y \in B_0, \\ \infty & \text{if } y \notin B_0. \end{cases}$$

Note that $\dfrac{k(x,y)}{r(x)}$, which is an $(\mathcal{E}\otimes\mathcal{F})$-measurable function, is used as the Radon-Nikodym derivative $\frac{dP_{X|Y=y}}{dP_X}$. It remains to show that the essential supremum of a jointly measurable function is measurable. We state this in the form of a lemma, proved in Appendix D.

Lemma 2: Given measurable spaces $(E,\mathcal{E})$ and $(F,\mathcal{F})$, suppose $s \in (\mathcal{E}\otimes\mathcal{F})_+$. Let P_X be a probability measure on $(E,\mathcal{E})$. Then, the function $t : F \to \bar{\mathbb{R}}_+$ defined as $t(y) = \operatorname{ess\,sup}_{P_X} s(x,y)$ is $\mathcal{F}$-measurable.

Equipped with Lemma 2, we conclude that the mapping $y \mapsto \ell_{P_{XY}}(X \to y)$ is $\mathcal{F}$-measurable, as desired.

D. Proof of Lemma 2

To show that t is $\mathcal{F}$-measurable it suffices to show that the inverse image $t^{-1}(c,\infty]$ is in $\mathcal{F}$ for each $c \in \mathbb{R}_+$. Fix an arbitrary $c \in \mathbb{R}_+$. Given $y \in F$, define the set $C_y = \{x \in E : s(x,y) > c\}$ which is in $\mathcal{E}$ by the measurability of the mapping $x \mapsto s(x,y)$ for fixed $y \in F$. Now, we write

$$\begin{aligned} t^{-1}(c,\infty] &= \{y \in F : t(y) > c\} \\ &= \{y \in F : P_X(\{x \in E : s(x,y) > c\}) > 0\} \\ &= \{y \in F : P_X(C_y) > 0\} \\ &= \left\{ y \in F : \left(\int_E \mathbb{1}_{C_y}(x)\ P_X(dx) \right) > 0 \right\}. \end{aligned}$$

The mapping $(x,y) \mapsto \mathbb{1}_{C_y}(x)$ is $\mathcal{E}\otimes\mathcal{F}$-measurable since $\{(x,y) \in E \times F : \mathbb{1}_{C_y}(x) = 1\} = \{(x,y) \in E \times F : s(x,y) > c\} \in \mathcal{E}\otimes\mathcal{F}$. Then, Fubini's theorem ensures that $y \mapsto \int_E \mathbb{1}_{C_y}(x)\ P_X(dx)$ is $\mathcal{F}$-measurable, which in turn, implies that $t^{-1}(c,\infty]$ belongs to $\mathcal{F}$.

E. Other examples

Example 5 (Poisson and Binomial distributions): Suppose $X \sim \mathrm{Pois}(\lambda p)$, where $\lambda > 1$, $p \in (0,1)$, and $\lambda(1-p) < 1$. Assume Y is defined through the kernel

$$P_{Y|X=x}(y) = \begin{cases} \dfrac{\big(\lambda(1-p)\big)^{y-x} e^{-\lambda(1-p)}}{(y-x)!} & \text{if } y \geq x, \\ 0 & \text{if } y < x, \end{cases}$$

where $x \in \mathbb{N}$. It can be easily verified that $X \mid Y = y \sim \mathrm{Binom}(y, p)$. Hence, the PML from X to $y \in \mathbb{N}$ is given by

$$\begin{aligned} \ell_{P_{XY}}(X \to y) &= \log \sup_{x \in \mathbb{N}} \frac{P_{X|Y=y}(x)}{P_X(x)} \\ &= \log \max_{x \in \{0,\ldots,y\}} \frac{P_{X|Y=y}(x)}{P_X(x)} \\ &= \log \big(e^{\lambda p}\, \lambda^{-y}\, y!\big). \end{aligned}$$

Example 6 (Gaussian mixtures): Suppose $X \sim \mathrm{Ber}(\frac{1}{2})$ is an equiprobable Bernoulli random variable, and $Y \mid X = x \sim \mathcal{N}(x, \sigma^2)$ has Gaussian distribution with mean $x \in \{0,1\}$ and variance σ^2. The PML from X to each $y \in \mathbb{R}$ can be computed as

$$\begin{aligned} \ell_{P_{XY}}(X \to y) &= \log \max_{x \in \{0,1\}} \frac{f_{Y|X}(y,x)}{f_Y(y)} \\ &= \log \frac{2}{\exp\Big(-\frac{|y-\frac{1}{2}|}{\sigma^2}\Big) + 1}. \end{aligned}$$

Specifically, $\ell_{P_{XY}}(X \to \frac{1}{2}) = 0$ and $\lim_{y\to\infty} \ell_{P_{XY}}(X \to y) = \lim_{y\to-\infty} \ell_{P_{XY}}(X \to y) = \log 2$.

Example 7 (Bivariate Gaussian): Suppose X and Y are zero-mean jointly Gaussian random variables with variances σ_X^2 and σ_Y^2, respectively, and correlation coefficient $\rho \in (-1,1)$. Then, $Y \mid X = x \sim \mathcal{N}(\frac{\sigma_Y}{\sigma_X}\rho x, (1-\rho^2)\sigma_Y^2))$, and the PML from X to $y \in \mathbb{R}$ is

$$\ell_{P_{XY}}(X \to y) = \begin{cases} \frac{y^2}{2\sigma_Y^2} - \frac{1}{2}\log(1-\rho^2) & \text{if } \rho \neq 0, \\ 0 & \text{if } \rho = 0. \end{cases}$$