On the Relation Between Identifiability, Differential Privacy, and Mutual-Information Privacy

Weina Wang, *Student Member, IEEE,* Lei Ying, *Member, IEEE,* and Junshan Zhang, *Fellow, IEEE*

***Abstract*— This paper investigates the relation between three different notions of privacy: identifiability, differential privacy, and mutual-information privacy. Under a unified privacy-distortion framework, where the distortion is defined to be the expected Hamming distance between the input and output databases, we establish some fundamental connections between these three privacy notions. Given a maximum allowable distortion D, we define the privacy-distortion functions $\epsilon_{\mathrm{i}}^*(D)$, $\epsilon_{\mathrm{d}}^*(D)$, and $\epsilon_{\mathrm{m}}^*(D)$ to be the smallest (most private/best) identifiability level, differential privacy level, and mutual information between the input and the output, respectively. We characterize $\epsilon_{\mathrm{i}}^*(D)$ and $\epsilon_{\mathrm{d}}^*(D)$, and prove that $\epsilon_{\mathrm{i}}^*(D) - \epsilon_X \le \epsilon_{\mathrm{d}}^*(D) \le \epsilon_{\mathrm{i}}^*(D)$ for D within certain range, where ϵ_X is a constant determined by the prior distribution of the original database X, and diminishes to zero when X is uniformly distributed. Furthermore, we show that $\epsilon_{\mathrm{i}}^*(D)$ and $\epsilon_{\mathrm{m}}^*(D)$ can be achieved by the same mechanism for D within certain range, i.e., there is a mechanism that simultaneously minimizes the identifiability level and achieves the best mutual-information privacy. Based on these two connections, we prove that this mutual-information optimal mechanism satisfies ϵ-differential privacy with $\epsilon_{\mathrm{d}}^*(D) \le \epsilon \le \epsilon_{\mathrm{d}}^*(D)+2\epsilon_X$. The results in this paper reveal some consistency between two worst case notions of privacy, namely, identifiability and differential privacy, and an average notion of privacy, mutual-information privacy.**

***Index Terms*— Differential privacy, Hamming distance, identifiability, mutual information, rate-distortion.**

I. Introduction

PRIVACY has been an increasing concern in the emerging big data era, particularly with the growing use of personal data such as medical records or online activities for big data analysis. Analyzing these data results in new discoveries in science and engineering, but also puts individual's privacy at potential risks. Therefore, privacy-preserving data analysis, where the goal is to preserve the accuracy of data analysis while maintaining individual's privacy, has become one of the main challenges of this big data era. The basic idea of privacy-preserving data analysis is to add randomness in the released information to guarantee that an individual's information cannot be inferred. Intuitively, the higher the randomness is, the better privacy protection individual users get, but the less accurate (useful) the output statistical information is. While randomization seems to be inevitable, for the privacy-preserving data analysis it is of great interest to quantitatively define the notion of privacy. Specifically, we need to understand the amount of randomness needed to protect privacy while preserving usefulness of the data. To this end, we consider three different notions: identifiability, differential privacy and mutual-information privacy, where identifiability is concerned with the posteriors of recovering the original data from the released data, differential privacy is concerned with additional disclosures of an individual's information due to the release of the data, and mutual information measures the average amount of information about the original database contained in the released data.

While these three different privacy notions are defined from different perspectives, they are fundamentally related. The focus of this paper is to investigate the fundamental connections between these three different privacy notions in the following setting:

- We consider a non-interactive database releasing approach for privacy-preserving data analysis, where a synthetic database is released to the public. The synthetic database is a sanitized version of the original database, on which queries and operations can be carried out as if it was the original database. It is then natural to assume that the synthetic database and the original database are in the same "universe" so the entries have the same interpretation. Therefore we focus on mechanisms that map an input database to an output synthetic database in the same universe. Specifically, we consider a database consisting of n rows, each of which takes values from a finite domain $\mathcal{D}$ of size m. In this paper, the database is modeled as a discrete random variable X drawn from $\mathcal{D}^n$ with prior distribution p_X. A mechanism $\mathcal{M}$ takes a database X as input and outputs a database Y, which is also a random variable with alphabet $\mathcal{D}^n$.
- We define the *distortion* between the output database and the input database to be the expected Hamming distance. When the input and output are in the same universe, the Hamming distance measures the number of rows two databases differ in, which directly points to the number of rows that need to be modified in order to guarantee a given privacy level.

In this paper, we use a unified *privacy–distortion* framework to understand the relation between the three privacy notions. Given a maximum allowable distortion D, we define the privacy–distortion functions $\epsilon_{\mathrm{i}}^*(D)$, $\epsilon_{\mathrm{d}}^*(D)$, and $\epsilon_{\mathrm{m}}^*(D)$ to be

Manuscript received September 6, 2014; revised September 22, 2015; accepted June 7, 2016. Date of publication June 23, 2016; date of current version August 16, 2016. This work was supported by the National Science Foundation through the Division of Electrical, Communications and Cyber Systems under Grant ECCS-1255425 and Grant ECCS-1547294. This paper was presented at the 2014 Annual Allerton Conference on Communication, Control, and Computing.

The authors are with the School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ 85281 USA (e-mail: weina.wang@asu.edu; lei.ying.2@asu.edu; junshan.zhang@asu.edu).

Communicated by S. Jaggi, Associate Editor for Coding Techniques.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIT.2016.2584610

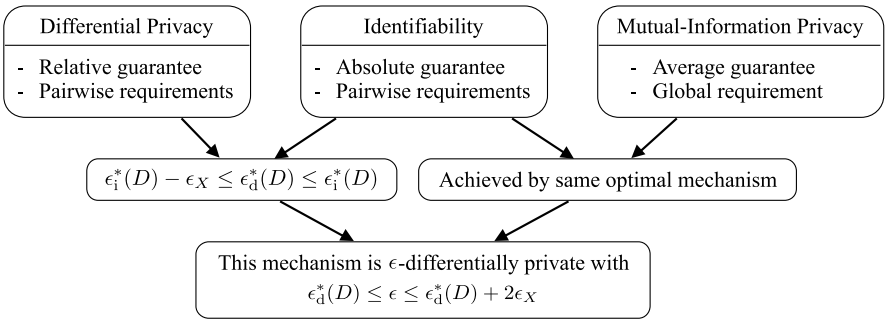

Fig. 1. Relation between identifiability, differential privacy and mutual-information privacy.

the smallest identifiability level, differential privacy level, and mutual information between the input and output, respectively. Then we have the following main results, which are also summarized in Fig. 1.

- We derive the exact form of the privacy–distortion function $\epsilon_{\mathrm{i}}^*(D)$ under the notion of identifiability, for certain range of the distortion values, by showing that $\epsilon_{\mathrm{i}}^*(D) = h^{-1}(D)$ regardless of the prior distribution, where

$$h^{-1}(D) = \ln\Big(\frac{n}{D} - 1\Big) + \ln(m-1).$$

  We further show that for the privacy–distortion function $\epsilon_{\mathrm{d}}^*(D)$ under the notion of differential privacy,

$$\epsilon_{\mathrm{i}}^*(D) - \epsilon_X \le \epsilon_{\mathrm{d}}^*(D) \le \epsilon_{\mathrm{i}}^*(D).$$

  The constant ϵ_X is determined by the prior distribution p_X only, given by

$$\epsilon_X = \max_{x, x' \in \mathcal{D}^n : x \sim x'} \ln \frac{p_X(x)}{p_X(x')},$$

  where $x \sim x'$ denotes that x and x' differ in exactly one row. When the input database has a uniform distribution, we have that $\epsilon_{\mathrm{i}}^* = \epsilon_{\mathrm{d}}^*$, i.e., differential privacy is equivalent to identifiability. Note that for ϵ_X to be finite, the prior p_X needs to have full a support on $\mathcal{D}^n$, i.e., $p_X(x) > 0$ for any $x \in \mathcal{D}^n$. When ϵ_X is large, differential privacy provides only weak guarantee on identifiability. In other words, when ϵ_X is large, it is possible to identify some entries of the database with non-trivial accuracy even if the differential privacy is satisfied. This is because differential privacy provides a *relative* guarantee about disclosures, which ensures that limited *additional* information of an individual is leaked in the released data in addition to the knowledge that an adversary has known. Identifiability, on the other hand, requires an *absolute* guarantee about disclosures when individuals' data is being inferred from the output database assuming that the prior p_X and the mechanism are both known to the adversary.
- The privacy–distortion functions $\epsilon_{\mathrm{i}}^*(D)$ and $\epsilon_{\mathrm{m}}^*(D)$ under the notions of identifiability and mutual-information privacy, respectively, can be achieved by the same mechanism for D within certain range, i.e., there is a mechanism that simultaneously minimizes the identifiability level and the mutual information between X and Y. We further prove that this mutual-information optimal mechanism satisfies ϵ-differential privacy that is within a constant difference from the optimal differential privacy level for the given maximum allowable distortion:

$$\epsilon_{\mathrm{d}}^*(D) \le \epsilon \le \epsilon_{\mathrm{d}}^*(D) + 2\epsilon_X.$$

  These results reveal certain consistency between identifiability and mutual-information privacy, and between differential privacy and mutual-information privacy when the prior p_X is uniform, although identifiability and differential privacy are defined based on "pairwise" requirements on distinguishability and can be viewed as "worst-case" guarantee of privacy, while mutual-information privacy is defined by "global" requirements and is considered to be an "average" notion of privacy. The value of $\epsilon_{\mathrm{m}}^*(D)$ is in bits and thus is not directly comparable with $\epsilon_{\mathrm{i}}^*(D)$ and $\epsilon_{\mathrm{d}}^*(D)$, but the fact that identifiability and mutual-information privacy can be optimized simultaneously in the setting studied in this paper reveals the fundamental connections between these three privacy notions.

A. Related Work

Differential privacy, as an emerging analytical foundation for privacy-preserving data analysis, was developed by a line of work [1]–[3], and since then both interactive model (e.g., [1], [4]–[9]) and non-interactive model (e.g., [8], [10]–[15]) have been studied in the literature. There is a vast and growing body of work on differential privacy, which we do not attempt to survey but refer interested readers to a thorough introduction by Dwork and Roth [16].

The privacy guarantee of differential privacy does not depend on the prior distribution of the original database, since it captures the additional disclosure caused by an information releasing mechanism on top of any given disclosure. With the prior taken into account, privacy notions based on the posterior have also been proposed. The seminal work of differential privacy [1] also proposed a semantically flavored definition of privacy, named semantic security, and showed its equivalence to differential privacy. This definition measures privacy by the difference between an adversary's prior knowledge of the database and the posterior belief given the output of the mechanism. Differential identifiability [17] and membership privacy [18] assume that a database entry can be traced back to the identify of an individual, and they quantify the leakage

of the information on whether an individual participates the database or not. Differential identifiability is defined to be the posterior probability for any individual to be the only unknown participant of a database given the entries of all the known participants and the output of the mechanism. This probability cannot be directly translated to a differential privacy level. Membership privacy is defined based on the difference between the prior and the posterior probability for an entity to be included in the database. Choosing appropriate prior distribution families makes differential privacy and differential identifiability instantiations of membership privacy under their database model. In this paper, the notion of identifiability is defined based on the indistinguishability between the posterior probabilities of neighboring databases given the output of the mechanism, which measures the hardness of identifying the data content of a database entry rather than the identity of the individual who contributes the data.

Information-theoretic privacy measures including mutual information, min-entropy, equivocation, etc, are relatively classical and have a rich history (e.g., [19]–[31]). When mutual information is used as the privacy notion, the problem of finding the optimal tradeoff between privacy and distortion can usually be formulated as a rate–distortion problem in the field of information theory (see [32] for an introduction) [25], [27]–[31]. In this paper, we also utilize results from the celebrated rate–distortion theory to characterize the optimal privacy–distortion tradeoff. However, we are more interested in the relation between the optimal privacy–distortion tradeoffs with different privacy notions: mutual information, differential privacy, and identifiability, and we quantify the impact of the prior explicitly. The work of du Pin Calmon and Fawaz [27] and Makhdoumi and Fawaz [28] showed that when a mechanism satisfies ϵ-information privacy (defined based on the difference between the prior of the database and the posterior given the output), it is 2ϵ-differentially private, and the mutual information between the database and the output is upper bounded by $\epsilon/\ln 2$. But differential privacy alone does not imply a bound on the mutual information if the possible values and sizes of the database and the output and the prior can be chosen freely. McGregor et al. [33] and De [34] showed that ϵ-differential privacy implies upper bounds on the mutual information in the order of $O(\epsilon n)$ and $O(\epsilon d)$, respectively, where n is the size of the database and d is the dimension of the data entry. Alvim et al. [26] showed that differential privacy implies a bound on the min-entropy leakage. The above relations between information-theoretic privacy notions and differential privacy, however, are not for the optimal privacy with distortion constraint, although they can contribute to building relations between the optimal tradeoffs. Sarwate and Sankar [31] showed that the result in [33] indicates a one direction bound between the optimal differential privacy and the optimal mutual information given the same distortion constraint. Mir [29] pointed out that the mechanism that achieves the optimal rate–distortion also guarantees a certain level of differential privacy. However, whether this differential privacy level is optimal or how far it is from optimal was not answered.

II. Model

Consider a database consisting of n rows, each of which corresponds to the data of a single individual. Each individual's data contains some sensitive information such as the individual's health status. Suppose that each row takes values from a domain $\mathcal{D}$. Then $\mathcal{D}^n$ is the set of all possible values of a database. Two databases, denoted by $x, x' \in \mathcal{D}^n$, are said to be *neighbors* if they differ in exactly one row. Let $x \sim x'$ denote the neighboring relation. In this paper, we assume that the domain $\mathcal{D}$ is a finite set and model a database as a discrete random variable X with alphabet $\mathcal{D}^n$ and probability mass function (pmf) p_X. Suppose $|\mathcal{D}| = m$, where m is an integer and $m \geq 2$. A (randomized) mechanism $\mathcal{M}$ takes a database x as the input, and outputs a random variable $\mathcal{M}(x)$.

Definition 1 (Mechanism): A *mechanism* $\mathcal{M}$ is specified by an *associated mapping* $\phi_{\mathcal{M}}\colon \mathcal{D}^n \to \mathcal{F}$, where $\mathcal{F}$ is the set of multivariate cdf's on some range $\mathcal{R}$. Taking database X as the input, the mechanism $\mathcal{M}$ outputs a $\mathcal{R}$-valued random variable Y with $\phi_{\mathcal{M}}(x)$ as the multivariate conditional cdf of Y given $X = x$.

In this paper, we focus on mechanisms for which the range is the same as the alphabet of X, i.e., $\mathcal{R} = \mathcal{D}^n$. Then the output Y is also a discrete random variable with alphabet $\mathcal{D}^n$, which can be interpreted as a synthetic database. Denote the conditional pmf of Y given $X = x$ defined by the cdf $\phi_{\mathcal{M}}(x)$ as $p_{Y|X}(\cdot \mid x)$. Then a mechanism in this setting is fully specified by $p_{Y|X}$. When using this mechanism, the database curator samples from $p_{Y|X}(\cdot \mid x)$ to generate a synthetic database Y. The form of the mechanism is assumed to be public since it may be of interest to data analysts.

Throughout this paper we use the following basic notation. We denote the set of real numbers by $\mathbb{R}$, the set of nonnegative real numbers by $\mathbb{R}^+$, and the set of nonnegative integers by $\mathbb{N}$. Let $\overline{\mathbb{R}}^+ = \mathbb{R}^+ \cup \{+\infty\}$.

A. Different Notions of Privacy

In addition to the output database Y, we assume that the adversary also knows the prior distribution p_X, which represents the side information the adversary has, and the privacy-preserving mechanism $\mathcal{M}$. The three notions of privacy studied in this paper are defined next.

Definition 2 (Identifiability): A mechanism $\mathcal{M}$ satisfies *ϵ-identifiability* for some $\epsilon \in \overline{\mathbb{R}}^+$ if for any pair of neighboring elements $x, x' \in \mathcal{D}^n$ and any $y \in \mathcal{D}^n$,

$$p_{X|Y}(x \mid y) \leq e^{\epsilon}\, p_{X|Y}(x' \mid y). \tag{1}$$

The notion of identifiability is defined based on the indistinguishability between any two neighboring databases from a Bayesian view. When a mechanism satisfies ϵ-identifiability for a small ϵ, two close (neighboring) databases cannot be distinguished from the posterior probabilities after observing the output database, which makes any individual's data hard to identify. To see the semantic implications of identifiability, we consider the following "worst-case" type of adversaries, who are called *informed adversaries* [1]. An adversary of this type knows $n-1$ database entries and tries to identify the value of the remaining one. The notation of identifiability

is defined based on neighboring databases to reflect this worst-case scenario. Consider adversaries who know X_{-i}, i.e., all the database entries except X_i. The requirement (1) of ϵ-identifiability indicates that for any $x_i, x_i' \in \mathcal{D}$, any $x_{-i} \in \mathcal{D}^{n-1}$ and any $y \in \mathcal{D}^n$,

$$\Pr\{X_i = x_i \mid X_{-i} = x_{-i}, Y = y\} \leq e^{\epsilon} \Pr\{X_i = x_i' \mid X_{-i} = x_{-i}, Y = y\}.$$

Therefore, when ϵ-identifiability is satisfied, even for such a worst-case adversary, the probability of correctly identifying the value of X_i is still no greater than $\frac{1}{1+(m-1)e^{-\epsilon}}$, which is close to randomly guessing when ϵ is small. We say that identifiability provides an *absolute* guarantee about disclosures since when it is satisfied, the probability of correctly identifying some individual's data is limited, and thus no bad disclosure can occur. This will become more clear when we discuss the relative guarantee provided by differential privacy.

We remark that in some cases, not all values of ϵ are achievable for ϵ-identifiability. The smallest achievable identifiability level is constrained by the prior p_X, since an adversary can always identify the values of the database entries based on the prior. When the prior itself is very disclosive, no mechanism can make the database entries less identifiable. To illustrate, we give the following example.

Example 1: Consider a database X with a single binary entry, i.e., $\mathcal{D} = \{0, 1\}$ and $n = 1$. Suppose the prior is given by $p_X(0) = 0.55$ and $p_X(1) = 0.45$. Consider the mechanism $\mathcal{M}$ specified by

$$p_{Y|X}(0 \mid 0) = p_{Y|X}(1 \mid 1) = 0.6,$$
$$p_{Y|X}(1 \mid 0) = p_{Y|X}(0 \mid 1) = 0.4.$$

Then the mechanism $\mathcal{M}$ satisfies ϵ-identifiability for $\epsilon \approx 0.6$. Therefore, the probability of correctly identifying X is guaranteed to be no greater than $\frac{1}{1+e^{-\epsilon}} \approx 0.65$. The smallest identifiability level that can be achieved for this prior is $\epsilon = \ln(0.55/0.45) \approx 0.2$. Now consider another prior that is given by $p_X(0) = 0.9$ and $p_X(1) = 0.1$. Then the mechanism $\mathcal{M}$ satisfies ϵ-identifiability for $\epsilon \approx 2.6$. In this case, no matter what mechanism is used, guessing that $X = 0$ yields a probability of correctness that is no less than 0.9. For an adversary with this prior, which indicates that the adversary has very good knowledge about the entry, no mechanism can achieve ϵ-identifiability for $\epsilon < \ln(0.9/0.1) \approx 2.2$.

Definition 3 (Differential Privacy [1], [2]): A mechanism $\mathcal{M}$ satisfies *ϵ-differential privacy* for some $\epsilon \in \overline{\mathbb{R}}^+$ if for any pair of neighboring elements $x, x' \in \mathcal{D}^n$ and any $y \in \mathcal{D}^n$,

$$p_{Y|X}(y \mid x) \leq e^{\epsilon} p_{Y|X}(y \mid x'). \tag{2}$$

Note that Definition 3 is equivalent to the definition of differential privacy in the seminal work [1], [2] under the model in this paper, although the languages used are slightly different. The differential privacy property of a mechanism is only determined by the associated mapping represented by $p_{Y|X}$ and does not depend on the prior.

In contrast to identifiability, differential privacy provides a *relative* guarantee about disclosures [2]. If some privacy disclosure were to happen with certain probability, differential privacy guarantees that this disclosure probability increases by at most a multiplicative factor after the output of the mechanism is published. So only limited *additional* risk will be caused by the mechanism. To illustrate, we give the following example.

Example 2: We again consider the database X and the mechanism $\mathcal{M}$ in Example 1. The mechanism $\mathcal{M}$ satisfies ϵ-differential privacy for $\epsilon = \ln(0.6/0.4) \approx 0.4$ regardless of the prior p_X. If the prior is given by $p_X(0) = 0.9$ and $p_X(1) = 0.1$, then before seeing the output Y, the probability of correctly identifying X is 0.9. Suppose that the adversary observes an output $Y = 0$. Then the probability of correctly identifying X becomes $\Pr(X = 0 \mid Y = 0) \approx 0.93$, which improves by a factor of approximately $e^{0.03}$. Differential privacy guarantees that this multiplicative factor is at most e^{ϵ}. Note that differential privacy is not intended to reverse the previous disclosure. In this example, the adversary is able to identify X with probability 0.9 using her prior information. After observing Y, the adversary is still able to identify X with a high probability (≈ 0.93), but only a small multiplicative factor is caused by the mechanism $\mathcal{M}$.

Definition 4 (Mutual-Information Privacy): A mechanism $\mathcal{M}$ satisfies *ϵ-mutual-information privacy* for some $\epsilon \in \overline{\mathbb{R}}^+$ if the mutual information between X and Y satisfies $I(X; Y) \leq \epsilon$, where

$$I(X; Y) = \sum_{x, y \in \mathcal{D}^n} p_{X,Y}(x, y) \log \frac{p_{X,Y}(x, y)}{p_X(x) p_Y(y)}.$$

The notion of mutual information is an information-theoretic notion of privacy, which measures the *average* amount of information about X contained in Y. The mutual information is minimized and equal to 0 when X and Y are independent, and it is maximized and equal to $H(X)$ when $Y = X$.

B. Distortion

In this paper, we measure the usefulness of a mechanism by the distortion between the input database X and the output Y, where smaller distortion corresponds to greater usefulness. Consider the (generalized) Hamming distance $d: \mathcal{D}^n \times \mathcal{D}^n \to \mathbb{N}$, where the distance $d(x, x')$ between any two elements $x, x' \in \mathcal{D}^n$ is the number of rows they differ in. We define the distortion between X and Y to be the expected Hamming distance

$$\mathbb{E}[d(X, Y)] = \sum_{x \in \mathcal{D}^n} \sum_{y \in \mathcal{D}^n} p_X(x) p_{Y|X}(y \mid x) d(x, y).$$

The Hamming distance also characterizes the neighboring relation on $\mathcal{D}^n$. Two elements $x, x' \in \mathcal{D}^n$ are neighbors if and only if $d(x, x') = 1$.

C. Privacy–Distortion Function

A privacy–distortion pair (ϵ, D) is said to be *achievable* if there exists a mechanism $\mathcal{M}$ with output Y such that $\mathcal{M}$ satisfies ϵ-privacy level and $\mathbb{E}[d(X, Y)] \leq D$. The *privacy–distortion function* $\epsilon^*: \mathbb{R}^+ \to \overline{\mathbb{R}}^+$ is defined by

$$\epsilon^*(D) = \inf\{\epsilon: \ (\epsilon, D) \text{ is achievable}\},$$

which is the smallest privacy level given the distortion constraint $\mathbb{E}[d(X,Y)] \leq D$. We are only interested in the range $[0,n]$ for D since this is the meaningful range for distortion. The privacy–distortion function depends on the prior p_X, which reflects the impact of the prior on the privacy–distortion tradeoff. To characterize the privacy–distortion function, we also consider the *distortion–privacy function* $D^*: \overline{\mathbb{R}}^+ \to \mathbb{R}^+$ defined by

$$D^*(\epsilon) = \inf\{D: \ (\epsilon, D) \text{ is achievable}\},$$

which is the smallest achievable distortion given privacy level ϵ.

In this paper we consider three different notions of privacy: identifiability, differential privacy and mutual-information privacy, so we denote the privacy–distortion functions under these three notions by ϵ_{i}^*, ϵ_{d}^* and ϵ_{m}^*, respectively.

III. Identifiability Versus Differential Privacy

In this section, we establish a fundamental connection between identifiability and differential privacy. We characterize their privacy–distortion functions through studying the distortion–privacy functions. Given privacy level ϵ_{i} and ϵ_{d}, the minimum distortion level is the solution to the following optimization problems.

The Privacy–Distortion Problem Under Identifiability (PD-I):

$$\begin{aligned}
\min_{p_{X|Y}, p_Y} \ & \sum_{x\in\mathcal{D}^n}\sum_{y\in\mathcal{D}^n} p_Y(y)p_{X|Y}(x \mid y)d(x,y) \\
\text{subject to } & p_{X|Y}(x \mid y) \leq e^{\epsilon_{\mathrm{i}}} p_{X|Y}(x' \mid y), \\
& \qquad \forall x, \quad x' \in \mathcal{D}^n : x \sim x', \ \ y \in \mathcal{D}^n, && (3)\\
& \qquad \sum_{x\in\mathcal{D}^n} p_{X|Y}(x \mid y) = 1, \quad \forall y \in \mathcal{D}^n, && (4)\\
& \qquad p_{X|Y}(x \mid y) \geq 0, \quad \forall x, \ \ y \in \mathcal{D}^n, && (5)\\
& \qquad \sum_{y\in\mathcal{D}^n} p_{X|Y}(x \mid y)p_Y(y) = p_X(x), \\
& \qquad \forall x \in \mathcal{D}^n, && (6)\\
& \qquad p_Y(y) \geq 0, \quad \forall y \in \mathcal{D}^n. && (7)
\end{aligned}$$

The Privacy–Distortion Problem Under Differential Privacy (PD-DP):

$$\begin{aligned}
\min_{p_{Y|X}} \ & \sum_{x\in\mathcal{D}^n}\sum_{y\in\mathcal{D}^n} p_X(x)p_{Y|X}(y \mid x)d(x,y) \\
\text{subject to } & p_{Y|X}(y \mid x) \leq e^{\epsilon_{\mathrm{d}}} p_{Y|X}(y \mid x'), \\
& \qquad \forall x, \quad x' \in \mathcal{D}^n : x \sim x', \ \ y \in \mathcal{D}^n, && (8)\\
& \qquad \sum_{y\in\mathcal{D}^n} p_{Y|X}(y \mid x) = 1, \quad \forall x \in \mathcal{D}^n, && (9)\\
& \qquad p_{Y|X}(y \mid x) \geq 0, \quad \forall x, y \in \mathcal{D}^n. && (10)
\end{aligned}$$

Note that to obtain the distortion–privacy functions, we need to find a mechanism $p_{Y|X}$ to minimize the distortion subject to privacy constraints. However, for identifiability, since it is defined based on $p_{X|Y}$, we change the optimization variable from $p_{Y|X}$ to $(p_{X|Y}, p_Y)$ in PD-I, and the constraints (4)–(7) ensure that PD-I is equivalent to the original distortion–privacy problem.

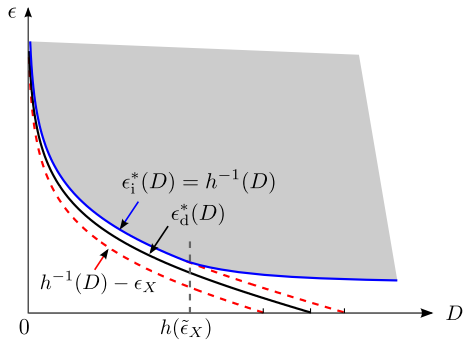

Fig. 2. The privacy–distortion functions ϵ_{i}^* under identifiability and ϵ_{d}^* under differential privacy satisfy $\epsilon_{\mathrm{i}}^*(D) - \epsilon_X \leq \epsilon_{\mathrm{d}}^*(D) \leq \epsilon_{\mathrm{i}}^*(D)$ for D within certain range.

For convenience, we first define two constants ϵ_X and $\widetilde{\epsilon}_X$ that are determined by the prior p_X. Let

$$\epsilon_X = \max_{x,x'\in\mathcal{D}^n : x\sim x'} \ln \frac{p_X(x)}{p_X(x')}, \tag{11}$$

which is the maximum prior probability difference between two neighboring databases. For ϵ_X to be finite, the prior distribution p_X needs to have full support on $\mathcal{D}^n$, i.e., $p_X(x) > 0$ for any $x \in \mathcal{D}^n$. To define $\widetilde{\epsilon}_X$, note that the prior p_X puts constraints on the posterior probabilities, as given by the constraint (6) in PD-I. We say $\{p_{X|Y}(x \mid y), y, \in \mathcal{D}^n\}$ is *feasible* if there exists a pmf p_Y such that it is the marginal pmf of Y. Let $\widetilde{\epsilon}_X$ be the smallest ϵ such that the following posterior probabilities are feasible:

$$p_{X|Y}(x \mid y) = \frac{e^{-\epsilon d(x,y)}}{\left(1 + (m-1)e^{-\epsilon}\right)^n}, \quad x, y \in \mathcal{D}^n.$$

We will see that the $p_{X|Y}$ in the above form plays an important role in solving PD-I. For any p_X, $\widetilde{\epsilon}_X$ is finite since when $\epsilon \to +\infty$, the pmf $p_Y = p_X$ is the marginal pmf of Y. Finally we define the function

$$h^{-1}(D) = \ln\left(\frac{n}{D} - 1\right) + \ln(m-1).$$

Recall that $\epsilon_{\mathrm{i}}^*(D)$ and $\epsilon_{\mathrm{d}}^*(D)$ denote the minimum identifiability level and minimum differential privacy level for a maximum allowable distortion D. The connection between the privacy–distortion functions ϵ_{i}^* and ϵ_{d}^* is established in the following theorem. See Fig. 2 for an illustration.

Theorem 1: For identifiability, the privacy–distortion function ϵ_{i}^* of a database X with $\epsilon_X < +\infty$ satisfies

$$\begin{cases} \epsilon_{\mathrm{i}}^*(D) = h^{-1}(D), & 0 \leq D \leq h(\widetilde{\epsilon}_X), \\ \epsilon_{\mathrm{i}}^*(D) \geq \max\{h^{-1}(D), \epsilon_X\}, & h(\widetilde{\epsilon}_X) < D \leq n. \end{cases} \tag{12}$$

For differential privacy, the privacy–distortion function ϵ_{d}^* of a database X satisfies the following bounds for any D with $0 \leq D \leq n$:

$$\max\{h^{-1}(D) - \epsilon_X, 0\} \leq \epsilon_{\mathrm{d}}^*(D) \leq \max\{h^{-1}(D), 0\}. \tag{13}$$

From the theorem above, we can see that $0 \leq \epsilon_{\mathrm{i}}^*(D) - \epsilon_{\mathrm{d}}^*(D) \leq \epsilon_X$ when $0 \leq D \leq h(\widetilde{\epsilon}_X)$. The lemmas needed in the proof of this theorem can be found in the appendix. Here

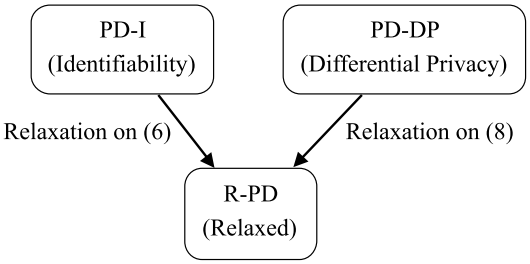

Fig. 3. Both PD-I and PD-DP boil down to R-PD through different relaxations.

we give a sketch of the proof, which consists of the following key steps:

- The first key step is to show that both PD-I and PD-DP, through (respective) relaxations as shown in Fig. 3, boil down to the same optimization problem.
  Relaxed Privacy–Distortion (R-PD):

$$\min_{p_{X|Y},p_Y} \sum_{x\in\mathcal{D}^n}\sum_{y\in\mathcal{D}^n} p_Y(y)p_{X|Y}(x\mid y)d(x,y)$$
$$\text{subject to } p_{X|Y}(x\mid y)\le e^{\epsilon}p_{X|Y}(x'\mid y),$$
$$\forall x,\quad x'\in\mathcal{D}^n : x\sim x',\ y\in\mathcal{D}^n, \tag{14}$$
$$\sum_{x\in\mathcal{D}^n} p_{X|Y}(x\mid y)=1,\quad \forall y\in\mathcal{D}^n, \tag{15}$$
$$p_{X|Y}(x\mid y)\ge 0,\quad \forall x,y\in\mathcal{D}^n, \tag{16}$$
$$\sum_{y\in\mathcal{D}^n} p_Y(y)=1, \tag{17}$$
$$p_Y(y)\ge 0,\quad \forall y\in\mathcal{D}^n. \tag{18}$$

  Relaxing the constraint (6) in PD-I to the constraint (17) gives R-PD. Now consider PD-DP. For any neighboring $x, x'\in\mathcal{D}^n$, $p_X(x)\le e^{\epsilon_X}p_X(x')$ according to the definition of ϵ_X, and a necessary condition for the constraint (8) to be satisfied is

$$p_X(x)p_{Y|X}(y\mid x)\le e^{\epsilon_{\mathrm{d}}+\epsilon_X}p_X(x')p_{Y|X}(y\mid x'). \tag{19}$$

  Therefore, replacing constraint (8) with (19) and letting $\epsilon=\epsilon_{\mathrm{d}}+\epsilon_X$, we obtain R-PD. So R-PD can be regarded as a relaxation of both PD-I and PD-DP.
- To solve R-PD, it suffices to solve the following optimization problem for any fixed $y\in\mathcal{D}^n$:

$$\min_{p_{X|Y}} \sum_{x\in\mathcal{D}^n} p_{X|Y}(x\mid y)d(x,y)$$
$$\text{subject to } p_{X|Y}(x\mid y)\le e^{\epsilon}p_{X|Y}(x'\mid y),$$
$$\forall x,\quad x'\in\mathcal{D}^n : x\sim x',$$
$$\sum_{x\in\mathcal{D}^n} p_{X|Y}(x\mid y)=1,$$
$$p_{X|Y}(x\mid y)\ge 0,\quad \forall x\in\mathcal{D}^n.$$

  Intuitively, to minimize the objective function, which is the average distortion between X and y, we should assign larger probability to $p_{X|Y}(x\mid y)$ with smaller $d(x,y)$, and smaller probability to $p_{X|Y}(x\mid y)$ with larger $d(x,y)$. For the x such that $x=y$, we should assign the largest value to $p_{X|Y}(x\mid y)$ since $d(x,y)=0$, and as x goes far way from y, we should assign smaller and smaller values to $p_{X|Y}(x\mid y)$. However, the privacy constraint limits the decreasing rate we can use as x goes far away from y due to the neighboring relations. In Lemma 1, we prove that the optimal solution is given by

$$p_{X|Y}(x\mid y)=\frac{e^{-\epsilon d(x,y)}}{\left(1+(m-1)e^{-\epsilon}\right)^n},\quad x,y\in\mathcal{D}^n, \tag{20}$$

  where the probability $p_{X|Y}(x\mid y)$ decreases with rate e^{ϵ} as $d(x,y)$ increases. This is the fastest possible decreasing rate with the privacy constraint, so this solution gives the smallest distortion.
- By Lemma 1, the minimum distortion of R-PD is $D^*_{\mathrm{relaxed}}(\epsilon)=h(\epsilon)$, which gives lower bounds on the distortion–privacy functions under identifiability and under differential privacy. By the connection between distortion–privacy function and privacy–distortion function, Lemma 2 shows that $\epsilon^*_{\mathrm{i}}(D)\ge h^{-1}(D)$ and $\epsilon^*_{\mathrm{d}}(D)\ge h^{-1}(D)-\epsilon_X$ for any D with $0\le D\le n$. Lemma 3 shows another lower bound on ϵ^*_{i}, combining which with the lower bound in Lemma 2 gives the lower bound in Theorem 1.
  Next we design achievable mechanisms to prove the upper bounds in Theorem 1. Notice that when the posterior probabilities given by the solution $p_{X|Y}$ in (20) is feasible, the mechanism that corresponds to this $p_{X|Y}$ satisfies ϵ-identifiability. Therefore, the lower bound for identifiability is achievable in this case. Consider the mechanism $\mathcal{E}^{\epsilon}_{\mathrm{i}}$ specified by

$$p_{Y|X}(y\mid x)=\frac{p_Y(y)e^{-\epsilon d(x,y)}}{p_X(x)\left(1+(m-1)e^{-\epsilon}\right)^n},\quad x,y\in\mathcal{D}^n, \tag{21}$$

  where $\epsilon\ge\widetilde{\epsilon}_X$ and p_Y is the corresponding pmf of Y. The mechanism $\mathcal{E}^{\epsilon}_{\mathrm{i}}$ corresponds to the posterior distributions given by $p_{X|Y}$ in (20). Lemma 4 shows that the mechanism $\mathcal{E}^{\epsilon}_{\mathrm{i}}$ guarantees an identifiability level of ϵ with distortion $h(\epsilon)$ when $\epsilon\ge\widetilde{\epsilon}_X$, which yields the equality in (12) when combining with the lower bound above.
- For differential privacy, consider the mechanism $\mathcal{E}^{\epsilon}_{\mathrm{d}}$ specified by the conditional probabilities

$$p_{Y|X}(y\mid x)=\frac{e^{-\epsilon d(x,y)}}{\left(1+(m-1)e^{-\epsilon}\right)^n},\quad x,y\in\mathcal{D}^n, \tag{22}$$

  where $\epsilon\ge 0$. Note that in contrast with the mechanism $\mathcal{E}^{\epsilon}_{\mathrm{i}}$, the mechanism $\mathcal{E}^{\epsilon}_{\mathrm{d}}$ itself has the same form as the solution $p_{X|Y}$ in (20). Lemma 5 shows that the mechanism $\mathcal{E}^{\epsilon}_{\mathrm{d}}$ satisfies ϵ-differential privacy with distortion $h(\epsilon)$, which provides the upper bound in (13). We remark that the mechanism $\mathcal{E}^{\epsilon}_{\mathrm{d}}$ has the same form as an exponential mechanism with score function $q=-d$ [35], where the score function has a sensitivity $\Delta q=1$. In general, an exponential mechanism with parameter ϵ is $2\epsilon\Delta q$-differentially private. However, the mechanism $\mathcal{E}^{\epsilon}_{\mathrm{d}}$ is ϵ-differentially private without the factor 2 since the normalizing term in the denominator of (22) does not depend on x. The mechanism $\mathcal{E}^{\epsilon}_{\mathrm{d}}$ can also be cast as a randomized response [36].

IV. Mutual-Information Privacy Versus Identifiability and Differential Privacy

In this section, we first discuss the relation between mutual-information privacy and identifiability. Then based on this relation and the relation between identifiability and differential privacy derived in the last section, we further establish a connection between mutual-information privacy and differential privacy.

Theorem 2: For any D with $0 \le D \le h(\widetilde{\epsilon}_X)$, the identifiability optimal mechanism $\mathcal{E}_{\mathrm{i}}^{\epsilon}$ with $\epsilon = h^{-1}(D)$ is also mutual-information optimal.

By this theorem, the privacy–distortion functions $\epsilon_{\mathrm{i}}^*(D)$ and $\epsilon_{\mathrm{m}}^*(D)$ under the notions of identifiability and mutual-information privacy, respectively, can be achieved by the same mechanism for D within certain range. This theorem indicates a consistency between identifiability and mutual-information privacy under the privacy–distortion framework since they can be optimized simultaneously.

Recall that given a maximum allowable distortion D, the privacy–distortion function $\epsilon_{\mathrm{m}}^*(D)$ under mutual-information privacy for an input database X with prior p_X is given by the optimal value of the following convex optimization problem.

The Privacy and Distortion Problem Under Mutual-Information Privacy (PD-MIP):

$$\min_{p_{Y|X}} \ I(X;Y)$$

$$\text{subject to} \sum_{x\in\mathcal{D}^n}\sum_{y\in\mathcal{D}^n} p_X(x)p_{Y|X}(y\mid x)d(x,y) \le D, \quad (23)$$

$$\sum_{y\in\mathcal{D}^n} p_{Y|X}(y\mid x) = 1, \quad \forall x \in \mathcal{D}^n, \quad (24)$$

$$p_{Y|X}(y\mid x) \ge 0, \quad \forall x, y \in \mathcal{D}^n. \quad (25)$$

Note that this formulation has the same form as the formulation in the celebrated rate–distortion theory (e.g., see [32]), and thus the privacy–distortion function under mutual-information privacy is identical to the rate–distortion function in this setting. Studies on the rate–distortion function [32], [37] have revealed the structure of an optimal solution of PD-MIP using Karush-Kuhn-Tucker (KKT) conditions [38]. We utilize these results to prove Theorem 2.

Proof of Theorem 2: By the KKT conditions for PD-MIP, the mutual information is minimized by

$$p_{Y|X}(y\mid x) = \frac{p_Y(y)e^{-\lambda d(x,y)}}{\sum_{y'\in\mathcal{D}^n} p_Y(y')e^{-\lambda d(x,y')}}, \quad x, y \in \mathcal{D}^n,$$

if there exists a pmf p_Y of Y and $\lambda \ge 0$ such that

$$\sum_{x\in\mathcal{D}^n} \frac{p_X(x)e^{-\lambda d(x,y)}}{\sum_{y'\in\mathcal{D}^n} p_Y(y')e^{-\lambda d(x,y')}} = 1, \quad \text{if } p_Y(y) > 0, \quad (26)$$

$$\sum_{x\in\mathcal{D}^n} \frac{p_X(x)e^{-\lambda d(x,y)}}{\sum_{y'\in\mathcal{D}^n} p_Y(y')e^{-\lambda d(x,y')}} \le 1, \quad \text{if } p_Y(y) = 0, \quad (27)$$

$$\lambda\left(\sum_{x\in\mathcal{D}^n}\sum_{y\in\mathcal{D}^n} \frac{p_X(x)p_Y(y)e^{-\lambda d(x,y)}}{\sum_{y'\in\mathcal{D}^n} p_Y(y')e^{-\lambda d(x,y')}} d(x,y) - D\right) = 0, \quad (28)$$

where λ is the Lagrange multiplier for the distortion constraint (23). This optimal solution has an exponential form. Recall that the identifiability optimal mechanism $\mathcal{E}_{\mathrm{i}}^{\epsilon}$ in (21) also has an exponential form. In what follows we prove that for properly chosen λ, the conditions (26)–(28) are satisfied under $\mathcal{E}_{\mathrm{i}}^{\epsilon}$.

For any $0 \le D \le h(\widetilde{\epsilon}_X)$, consider the mechanism $\mathcal{E}_{\mathrm{i}}^{\epsilon}$ with $\epsilon = h^{-1}(D)$. Let $\lambda = \epsilon$. Recall that under $\mathcal{E}_{\mathrm{i}}^{\epsilon}$,

$$p_{Y|X}(y\mid x) = \frac{p_Y(y)e^{-\epsilon d(x,y)}}{p_X(x)\big(1+(m-1)e^{-\epsilon}\big)^n}, \quad x, y \in \mathcal{D}^n.$$

Since $p_{Y|X}$ satisfies that

$$\sum_{y'\in\mathcal{D}^n} p_{Y|X}(y'\mid x) = 1,$$

we have

$$\sum_{y'\in\mathcal{D}^n} p_Y(y')e^{-\epsilon d(x,y')} = p_X(x)\big(1+(m-1)e^{-\epsilon}\big)^n.$$

Then for any $y \in \mathcal{D}^n$,

$$\begin{aligned}
&\sum_{x\in\mathcal{D}^n} \frac{p_X(x)e^{-\epsilon d(x,y)}}{\sum_{y'\in\mathcal{D}^n} p_Y(y')e^{-\epsilon d(x,y')}} \\
&\quad= \sum_{x\in\mathcal{D}^n} \frac{p_X(x)e^{-\epsilon d(x,y)}}{p_X(x)\big(1+(m-1)e^{-\epsilon}\big)^n} \\
&\quad= 1,
\end{aligned}$$

which indicates that (26) and (27) are satisfied. We can verify that

$$\begin{aligned}
&\sum_{x\in\mathcal{D}^n}\sum_{y\in\mathcal{D}^n} \frac{p_X(x)p_Y(y)e^{-\epsilon d(x,y)}}{\sum_{y'\in\mathcal{D}^n} p_Y(y')e^{-\epsilon d(x,y')}} d(x,y) \\
&\quad= \sum_{y\in\mathcal{D}^n} p_Y(y) \sum_{x\in\mathcal{D}^n} \frac{p_X(x)e^{-\epsilon d(x,y)}d(x,y)}{p_X(x)\big(1+(m-1)e^{-\epsilon}\big)^n} \\
&\quad= h(\epsilon) \\
&\quad= D,
\end{aligned}$$

which indicates that (28) is satisfied. Therefore, the mechanism $\mathcal{E}_{\mathrm{i}}^{\epsilon}$ with $\epsilon = h^{-1}(D)$ gives an optimal solution of PD-MIP, which completes the proof. ■

Next, we establish a connection between differential privacy and mutual-information privacy based on Theorem 2 and Theorem 1.

Corollary 1: For any D with $0 \le D \le h(\widetilde{\epsilon}_X)$, the mutual-information optimal mechanism $\mathcal{E}_{\mathrm{i}}^{\epsilon}$ with $\epsilon = h^{-1}(D)$ is ϵ_{d}-differentially private with $\epsilon_{\mathrm{d}}^*(D) \le \epsilon_{\mathrm{d}} \le \epsilon_{\mathrm{d}}^*(D) + 2\epsilon_X$.

It has been pointed out in [29] that a mechanism that achieves the optimal rate–distortion also guarantees a certain level of differential privacy. However, whether this differential privacy level is optimal or how far it is from optimal was not answered. Our result in Corollary 1 further shows that the gap between the differential privacy level of the mutual-information optimal mechanism $\mathcal{E}_{\mathrm{i}}^{\epsilon}$ and the optimal differential privacy level is no greater than $2\epsilon_X$, which is a constant determined by the prior p_X. Therefore, given a distortion constraint,

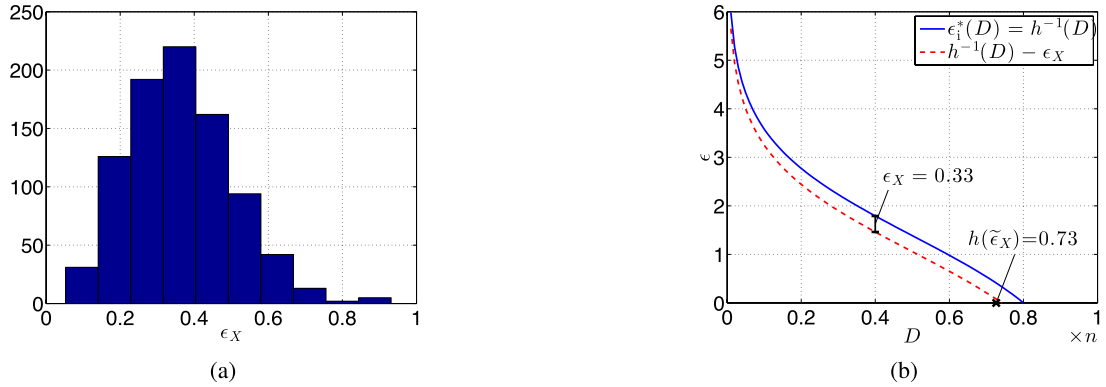

Fig. 4. Illustration of the characterizations of the privacy–distortion functions. (a) Histogram of ϵ_X for 887 databases. (b) The privacy–distortion function under identifiability is given by $\epsilon_{\mathrm{i}}^*(D) = h^{-1}(D)$ for $0 \le D \le h(\widetilde{\epsilon}_X)$, where $h(\widetilde{\epsilon}_X) = 0.73n$. The identifiability optimal mechanism that achieves this curve is also mutual-information optimal. The privacy–distortion function under differential privacy, $\epsilon_{\mathrm{d}}^*(D)$, lies between $\epsilon_{\mathrm{i}}^*(D) = h^{-1}(D)$ and $h^{-1}(D) - \epsilon_X$, where $\epsilon_X = 0.33$.

optimizing for mutual information leads to a differentially private mechanism whose privacy level is close to the optimal differential privacy level. When the prior is uniform, this mutual-information optimal mechanism achieves exactly the optimal differential privacy level. Note that differential privacy can be viewed as providing "worst-case" privacy guarantee since it is agnostic to prior distributions. Although mutual-information privacy is an "average" notion of privacy, when the prior is uniform, it exhibits certain consistency with differential privacy.

Proof of Corollary 1: By Theorem 2, the mechanism $\mathcal{E}_{\mathrm{i}}^{\epsilon}$ with $\epsilon = h^{-1}(D)$ is mutual-information optimal. According to its form, we can verify that $\mathcal{E}_{\mathrm{i}}^{\epsilon}$ with $\epsilon = h^{-1}(D)$ is ϵ_{d}-differentially private with $\epsilon_{\mathrm{d}} = h^{-1}(D) + \epsilon_X$. Since $\epsilon_{\mathrm{d}}^*(D)$ is the minimum differential privacy level with distortion constraint given by D, we have $\epsilon_{\mathrm{d}} \ge \epsilon_{\mathrm{d}}^*(D)$. By Theorem 1, $h^{-1}(D) \le \epsilon_{\mathrm{d}}^*(D) + \epsilon_X$. Thus $\epsilon_{\mathrm{d}} \le \epsilon_{\mathrm{d}}^*(D) + 2\epsilon_X$, which completes the proof. ■

An Illustration: We demonstrate the characterizations of the privacy–distortion functions in Theorem 1 and 2 using prior distributions based on a Netflix dataset [39]. The dataset consists of movie ratings from users, with each rating on a scale from 1 to 5 (integer) stars. We view the ratings of a movie from active users as a database and generate ratings uniformly at random for missing entries. We first calculate the corresponding ϵ_X, assuming that entries of a database are drawn i.i.d. from a distribution. The constant ϵ_X bounds the gap between the upper and lower bounds on $\epsilon_{\mathrm{d}}^*(D)$, and also bounds $\epsilon_{\mathrm{i}}^*(D) - \epsilon_{\mathrm{d}}^*(D)$. In Fig. 4, we show the histogram of ϵ_X for 887 most reviewed movies (databases). Next, we pick a database whose prior distribution of each entry is given by

$$p_{X_i}(1) = 0.2533, \quad p_{X_i}(2) = 0.1821, \quad p_{X_i}(3) = 0.1821,$$
$$p_{X_i}(4) = 0.1873, \quad p_{X_i}(5) = 0.1953.$$

For this prior, we have $\epsilon_X = 0.33$ and $\widetilde{\epsilon}_X = 0.41$. In Fig. 4, we draw the privacy–distortion function $\epsilon_{\mathrm{i}}^*(D) = h^{-1}(D)$ under identifiability for $0 \le D \le h(\widetilde{\epsilon}_X)$, where the value $h(\widetilde{\epsilon}_X) = 0.73n$ is displayed in the figure. The identifiability optimal mechanism that achieves this curve is also mutual-information optimal. The curve $\epsilon_{\mathrm{i}}^*(D) = h^{-1}(D)$ gives an upper bound on the privacy–distortion function $\epsilon_{\mathrm{d}}^*(D)$ under differential privacy. We also draw the curve $\max\{h^{-1}(D) - \epsilon_X, 0\}$, which is a lower bound on $\epsilon_{\mathrm{d}}^*(D)$.

V. Conclusions

In this paper, we investigated the relation between three different notions of privacy: identifiability, differential privacy and mutual-information privacy, where identifiability guarantees indistinguishability between posterior probabilities, differential privacy guarantees limited additional disclosures, and mutual information is an information-theoretic notion. Under a unified privacy–distortion framework, where the distortion is defined to be the expected Hamming distance between the input and output databases, we established some fundamental connections between these three privacy notions. Given a maximum allowable distortion D within certain range, the smallest identifiability level $\epsilon_{\mathrm{i}}^*(D)$ and the smallest differential privacy level $\epsilon_{\mathrm{d}}^*(D)$ are proved to satisfy $\epsilon_{\mathrm{i}}^*(D) - \epsilon_X \le \epsilon_{\mathrm{d}}^*(D) \le \epsilon_{\mathrm{i}}^*(D)$, where ϵ_X is a constant determined by the prior of the original database, and diminishes to zero when the prior is uniform. Next, we showed that there is a mechanism that simultaneously minimizes the identifiability level and the mutual information given the same maximum allowable distortion within certain range. We further showed that this mechanism satisfies ϵ-differential privacy with $\epsilon_{\mathrm{d}}^*(D) \le \epsilon \le \epsilon_{\mathrm{d}}^*(D) + 2\epsilon_X$.

Our findings in this study reveal some fundamental connections between the three notions of privacy. With these three notions of privacy being defined, many interesting issues deserve further attention. The connections we have established in this work are based on the distortion measure of Hamming distance, which is closely tied with the neighboring relations, and we assume that the output synthetic database and the original database are in the same universe. It would be of great interest to study the connections of these privacy notions under other common distortion measures and other output formats. We remark that our results for Hamming distance can be used to prove lower bounds on the distortion of a differentially private mechanism when the distortion is measured by the distortion at the worst-case query in a query class [40]. Some other interesting directions are as follows.

In some cases, the prior p_X is imperfect. Then for privacy notions depending on the prior such as identifiability and mutual-information privacy, it is natural to ask how we can protect privacy with robustness over the prior distribution. Identifiability and differential privacy impose requirements on neighboring databases to protect an individual's privacy. Then are there any practical scenarios that we would desire to generalize this "pairwise" privacy to "group" privacy? The connections between membership privacy and these three notions of privacy also need to be explored, since membership privacy has been proposed as a unifying framework for privacy definitions.

Appendix
Proof of Theorem 1

Lemma 1: The minimum distortion $D^*_{\text{relaxed}}(\epsilon)$ of the relaxed optimization problem R-PD satisfies

$$D^*_{\text{relaxed}}(\epsilon) = h(\epsilon), \tag{A.1}$$

where

$$h(\epsilon) = \frac{n}{1+\frac{e^\epsilon}{m-1}}.$$

Proof: We first prove the following claim, which gives a lower bound on the minimum distortion $D^*_{\text{relaxed}}(\epsilon)$.

Claim: Any feasible solution $\{p_{X|Y}(x \mid y), x, y \in \mathcal{D}^n\}$ of R-PD satisfies

$$\sum_{x\in\mathcal{D}^n} p_{X|Y}(x \mid y) d(x,y) \geq h(\epsilon).$$

Proof of the Claim: Consider any feasible $\{p_{X|Y}(x \mid y), y \in \mathcal{D}^n\}$. For any $y \in \mathcal{D}^n$ and any integer l with $0 \leq l \leq n$, let $\mathcal{N}_l(y)$ be the set of elements with distance l to y, i.e.,

$$\mathcal{N}_l(y) = \{v \in \mathcal{D}^n : d(v,y) = l\}. \tag{A.2}$$

Denote $P_l = \Pr\{X \in \mathcal{N}_l(y) \mid Y = y\}$. Then

$$\sum_{x\in\mathcal{D}^n} p_{X|Y}(x \mid y) d(x,y) = \sum_{l=0}^{n} l P_l.$$

We first derive a lower bound on P_n. For any $u \in \mathcal{N}_{l-1}(y)$, $\mathcal{N}_1(u) \cap \mathcal{N}_l(y)$ consists of the neighbors of u that are in $\mathcal{N}_l(y)$. By the constraint (14), for any $v \in \mathcal{N}_1(u) \cap \mathcal{N}_l(y)$,

$$p_{X|Y}(u \mid y) \leq e^\epsilon p_{X|Y}(v \mid y). \tag{A.3}$$

Each $u \in \mathcal{N}_{l-1}(y)$ has $n-(l-1)$ rows that are the same with the corresponding rows of y. Each neighbor of u in $\mathcal{N}_l(y)$ can be obtained by changing one of these $n-(l-1)$ rows to a different element in $\mathcal{D}$, which is left with $m-1$ choices. Therefore, each $u \in \mathcal{N}_{l-1}(y)$ has $(n-l+1)(m-1)$ neighbors in $\mathcal{N}_l(y)$. By similar arguments, each $v \in \mathcal{N}_l(y)$ has l neighbors in $\mathcal{N}_{l-1}(y)$. Taking summation of (A.3) over $u \in \mathcal{N}_{l-1}(y)$, $v \in \mathcal{N}_l(y)$ with $u \sim v$ yields

$$\begin{aligned}&\sum_{u\in\mathcal{N}_{l-1}(y)} (n-l+1)(m-1) p_{X|Y}(u \mid y)\\ &\quad\leq e^\epsilon \sum_{u\in\mathcal{N}_{l-1}(y)} \sum_{v\in\mathcal{N}_1(u)\cap\mathcal{N}_l(y)} p_{X|Y}(v \mid y).\end{aligned}$$

Thus

$$\begin{aligned}&(n-l+1)(m-1)P_{l-1}\\ &\quad\leq e^\epsilon \sum_{v\in\mathcal{N}_l(y)} \sum_{u\in\mathcal{N}_1(v)\cap\mathcal{N}_{l-1}(y)} p_{X|Y}(v \mid y) &\text{(A.4)}\\ &\quad= e^\epsilon l P_l. &\text{(A.5)}\end{aligned}$$

Recall that $N_l \triangleq |\mathcal{N}_l(x)| = \binom{n}{l}(m-1)^l$. Then by (A.5) we obtain that, for any l with $1 \leq l \leq n$,

$$\frac{P_{l-1}}{N_{l-1}} \leq \frac{P_l}{N_l} e^\epsilon.$$

As a consequence, for any l with $0 \leq l \leq n$,

$$P_l \leq \frac{N_l}{N_n} e^{(n-l)\epsilon} P_n. \tag{A.6}$$

Since $\sum_{l=0}^n P_l = 1$, taking summation over l in (A.6) yields

$$\begin{aligned}1 &\leq P_n \frac{1}{N_n e^{-n\epsilon}} \sum_{l=0}^{n} N_l e^{-l\epsilon}\\ &= P_n \frac{\left(1+(m-1)e^{-\epsilon}\right)^n}{N_n e^{-n\epsilon}},\end{aligned}$$

i.e.,

$$P_n \geq \frac{N_n e^{-n\epsilon}}{\left(1+(m-1)e^{-\epsilon}\right)^n}.$$

This lower bound on P_n gives the following lower bound:

$$\begin{aligned}\sum_{l=0}^{n} l P_l \geq \sum_{l=0}^{n} l\left(P_l + a\frac{N_l e^{-l\epsilon}}{\sum_{k=0}^{n-1} N_k e^{-k\epsilon}}\right)\\ + \frac{n N_n e^{-n\epsilon}}{\left(1+(m-1)e^{-\epsilon}\right)^n},\end{aligned}$$

where $a = P_n - \frac{N_n e^{-n\epsilon}}{\left(1+(m-1)e^{-\epsilon}\right)^n}$.

Consider the following optimization problem:

$$\begin{aligned}\min\ &\sum_{l=0}^{n-1} l Q_l\\ \text{subject to}\ &Q_l \geq 0, \quad l = 0,1,\ldots,n-1,\\ &\frac{Q_{l-1}}{N_{l-1}} \leq \frac{Q_l}{N_l} e^\epsilon, \quad l = 1,2,\ldots,n-1,\\ &\sum_{l=0}^{n-1} Q_l = 1 - \frac{N_n e^{-n\epsilon}}{\left(1+(m-1)e^{-\epsilon}\right)^n}.\end{aligned}$$

Suppose the optimal solution of this problem is $\{Q_0^*, Q_1^*, \ldots, Q_{n-1}^*\}$. Then

$$\sum_{l=0}^{n-1} l\left(P_l + a\frac{N_l e^{-l\epsilon}}{\sum_{k=0}^{n-1} N_k e^{-k\epsilon}}\right) \geq \sum_{l=0}^{n-1} l Q_l^*$$

as $\left\{P_l + a\frac{N_l e^{-l\epsilon}}{\sum_{k=0}^{n-1} N_k e^{-k\epsilon}}, l = 0,1,\ldots,n-1\right\}$ is a feasible solution. Therefore,

$$\sum_{l=0}^{n} l P_l \geq \sum_{l=0}^{n-1} l Q_l^* + \frac{n N_n e^{-n\epsilon}}{\left(1+(m-1)e^{-\epsilon}\right)^n}.$$

Similar to $\{P_l, l = 0, \ldots, n\}$, $\{Q_l^*, l = 0, \ldots, n-1\}$ satisfies

$$Q_l^* \le \frac{N_l}{N_{n-1}} e^{(n-1-l)\epsilon} Q_{n-1}^*. \tag{A.7}$$

Since $\sum_{l=0}^{n-1} Q_l^* = 1 - \frac{N_n e^{-n\epsilon}}{\left(1+(m-1)e^{-\epsilon}\right)^n}$, taking summation over l in (A.7) yields

$$Q_{n-1}^* \ge \frac{N_{n-1} e^{-(n-1)\epsilon}}{\left(1+(m-1)e^{-\epsilon}\right)^n}.$$

Using similar arguments we have

$$\sum_{l=0}^{n-1} l Q_l^* \ge \sum_{l=0}^{n-2} l C_l^* + \frac{(n-1)N_{n-1} e^{-(n-1)\epsilon}}{\left(1+(m-1)e^{-\epsilon}\right)^n},$$

where $\{C_l^*, l = 0, \ldots, n-2\}$ is the optimal solution of

$$\begin{aligned}
\min \ & \sum_{l=0}^{n-2} l C_l \\
\text{subject to } & C_l \ge 0, \quad l = 0, 1, \ldots, n-2, \\
& \frac{C_{l-1}}{N_{l-1}} \le \frac{C_l}{N_l} e^{\epsilon}, \quad l = 1, 2, \ldots, n-2, \\
& \sum_{l=0}^{n-2} C_l = 1 - \frac{N_{n-1} e^{-(n-1)\epsilon}}{\left(1+(m-1)e^{-\epsilon}\right)^n} \\
& \qquad - \frac{N_n e^{-n\epsilon}}{\left(1+(m-1)e^{-\epsilon}\right)^n}.
\end{aligned}$$

Continue this procedure we obtain

$$\sum_{l=0}^{n} l P_l \ge \sum_{l=0}^{n} \frac{l N_l e^{-(n-l)\epsilon}}{\left(1+(m-1)e^{-\epsilon}\right)^n} = \frac{n}{1+\frac{e^{\epsilon}}{m-1}} = h(\epsilon).$$

Therefore, for any feasible $\{p_{X|Y}(x \mid y), x, y \in \mathcal{D}^n\}$,

$$\sum_{x \in \mathcal{D}^n} p_{X|Y}(x \mid y) d(x, y) = \sum_{l=0}^{n} l P_l \ge h(\epsilon),$$

which completes the proof of the claim.

By this claim, any feasible solution satisfies

$$\sum_{x \in \mathcal{D}^n} \sum_{y \in \mathcal{D}^n} p_Y(y) p_{X|Y}(x \mid y) d(x, y) \ge h(\epsilon).$$

Therefore

$$D_{\text{relaxed}}^*(\epsilon) \ge h(\epsilon). \tag{A.8}$$

Next we prove the following claim, which gives an upper bound on the minimum distortion $D_{\text{relaxed}}^*(\epsilon)$.

Claim: Consider

$$p_{X|Y}(x \mid y) = \frac{e^{-\epsilon d(x,y)}}{\left(1+(m-1)e^{-\epsilon}\right)^n}, \quad x, y \in \mathcal{D}^n,$$

and any $\{p_Y(y), y \in \mathcal{D}^n\}$ *with*

$$\sum_{y \in \mathcal{D}^n} p_Y(y) = 1, \qquad p_Y(y) \ge 0, \quad \forall y \in \mathcal{D}^n.$$

Then $\{p_{X|Y}(x \mid y), x, y \in \mathcal{D}^n\}$ *and* $\{p_Y(y), y \in \mathcal{D}^n\}$ *form a feasible solution of R-PD, and*

$$\sum_{x \in \mathcal{D}^n} \sum_{y \in \mathcal{D}^n} p_Y(y) p_{X|Y}(x \mid y) d(x, y) = h(\epsilon).$$

Proof of the Claim: Obviously the considered $\{p_{X|Y}(x \mid y), x, y \in \mathcal{D}^n\}$ and $\{p_Y(y), y \in \mathcal{D}^n\}$ satisfy constraints (16)–(18). Therefore to prove the feasibility, we are left with constraint (14) and (15). We first verify constraint (14). Consider any pair of neighboring elements $x, x' \in \mathcal{D}^n$ and any $y \in \mathcal{D}^n$. Then by the triangle inequality,

$$d(x, y) \le d(x', y) - d(x', x) = d(x', y) - 1.$$

Therefore,

$$\begin{aligned}
p_{X|Y}(x \mid y) &= \frac{e^{-\epsilon d(x,y)}}{\left(1+(m-1)e^{-\epsilon}\right)^n} \\
&\le \frac{e^{-\epsilon(d(x',y)-1)}}{\left(1+(m-1)e^{-\epsilon}\right)^n} \\
&= e^{\epsilon} p_{X|Y}(x' \mid y).
\end{aligned}$$

Next we verify constraint (15). For any $y \in \mathcal{D}^n$ and any integer l with $0 \le l \le n$, let $\mathcal{N}_l(x)$ be the set of elements with distance l to y as defined in (A.2). Then it is easy to see that $N_l \triangleq |\mathcal{N}_l(y)| = \binom{n}{l}(m-1)^l$, and for any $y \in \mathcal{D}^n$,

$$\mathcal{D}^n = \bigcup_{l=0}^{n} \mathcal{N}_l(y).$$

Therefore, for any $y \in \mathcal{D}^n$,

$$\begin{aligned}
\sum_{x \in \mathcal{D}^n} p_{X|Y}(x \mid y) &= \sum_{x \in \mathcal{D}^n} \frac{e^{-\epsilon d(x,y)}}{\left(1+(m-1)e^{-\epsilon}\right)^n} \\
&= \frac{1}{\left(1+(m-1)e^{-\epsilon}\right)^n} \sum_{l=0}^{n} \sum_{x \in \mathcal{N}_l(y)} e^{-\epsilon d(x,y)} \\
&= \frac{1}{\left(1+(m-1)e^{-\epsilon}\right)^n} \sum_{l=0}^{n} \binom{n}{l} (m-1)^l e^{-\epsilon l} \\
&= 1.
\end{aligned}$$

With feasibility verified, we can proceed to calculate the distortion. Let $g_\epsilon = 1 + (m-1)e^{-\epsilon}$. Then

$$\begin{aligned}
&\sum_{x \in \mathcal{D}^n} \sum_{y \in \mathcal{D}^n} p_Y(y) p_{X|Y}(x \mid y) d(x, y) \\
&\quad = \frac{1}{(g_\epsilon)^n} \sum_{y \in \mathcal{D}^n} p_Y(y) \sum_{l=0}^{n} \sum_{x \in \mathcal{N}_l(y)} e^{-\epsilon d(x,y)} d(x, y) \\
&\quad = \frac{1}{(g_\epsilon)^n} \sum_{y \in \mathcal{D}^n} p_Y(y) \sum_{l=0}^{n} \binom{n}{l} (m-1)^l e^{-\epsilon l} l \\
&\quad = \frac{n(m-1)e^{-\epsilon}\left(1+(m-1)e^{-\epsilon}\right)^{n-1}}{(g_\epsilon)^n} \sum_{y \in \mathcal{D}^n} p_Y(y) \\
&\quad = \frac{n}{1+\frac{e^{\epsilon}}{m-1}} \\
&\quad = h(\epsilon),
\end{aligned}$$

which completes the proof of the claim.

By this claim, there exists a feasible solution such that

$$\sum_{x \in \mathcal{D}^n} \sum_{y \in \mathcal{D}^n} p_Y(y) p_{X|Y}(x \mid y) d(x, y) = h(\epsilon),$$

which implies

$$D^*_{\text{relaxed}}(\epsilon) \le h(\epsilon).$$

Combining this upper bound with the lower bound (A.8) gives

$$D^*_{\text{relaxed}}(\epsilon) = h(\epsilon).$$

□

Lemma 2: The optimal value $D^*_{\text{relaxed}}(\epsilon) = h(\epsilon)$ of R-PD implies the following lower bounds for any D with $0 \le D \le n$:

$$\epsilon^*_{\text{i}}(D) \ge h^{-1}(D), \tag{A.9}$$

$$\epsilon^*_{\text{d}}(D) \ge \max\{h^{-1}(D) - \epsilon_X, 0\}. \tag{A.10}$$

Proof: First we derive the lower bound on $\epsilon^*_{\text{i}}(D)$. Let δ be an arbitrary positive number. For any D with $0 \le D \le n$, let $\epsilon_{D,\delta} = \epsilon^*_{\text{i}}(D) + \delta$. Then by the definition of ϵ^*_{i}, we have that $(\epsilon_{D,\delta}, D)$ is achievable under identifiability. Therefore

$$D \ge D^*_{\text{i}}(\epsilon_{D,\delta}) \ge D^*_{\text{relaxed}}(\epsilon_{D,\delta}) = h(\epsilon_{D,\delta}),$$

where $D^*_{\text{i}}(\cdot)$ is the optimal value of PD-I. Since h is a decreasing function, this implies that $\epsilon_{D,\delta} \ge h^{-1}(D)$. Therefore

$$\epsilon^*_{\text{i}}(D) \ge h^{-1}(D) - \delta.$$

Letting $\delta \to 0$ yields

$$\epsilon^*_{\text{i}}(D) \ge h^{-1}(D).$$

Next we derive the lower bound on $\epsilon^*_{\text{d}}(D)$ using arguments similar to those in the proof of the lower bound on $\epsilon^*_{\text{i}}(D)$. Let δ be an arbitrary positive number. For any D with $0 \le D \le n$, let $\epsilon_{D,\delta} = \epsilon^*_{\text{d}}(D) + \delta$. Then by the definition of ϵ^*_{d}, we have that $(\epsilon_{D,\delta}, D)$ is achievable under differential privacy. Therefore

$$D \ge D^*_{\text{d}}(\epsilon_{D,\delta}) \ge D^*_{\text{relaxed}}(\epsilon_{D,\delta} + \epsilon_X) = h(\epsilon_{D,\delta} + \epsilon_X),$$

where $D^*_{\text{d}}(\cdot)$ is the optimal value of PD-DP. Since h is a decreasing function, this implies that $\epsilon_{D,\delta} + \epsilon_X \ge h^{-1}(D)$. Therefore

$$\epsilon^*_{\text{d}}(D) \ge h^{-1}(D) - \epsilon_X - \delta.$$

Letting $\delta \to 0$ yields

$$\epsilon^*_{\text{d}}(D) \ge h^{-1}(D) - \epsilon_X.$$

Since the privacy level is nonnegative, we obtain the lower bound in (A.10). □

Lemma 3: The privacy–distortion function ϵ^*_{i} of a database X is bounded from below as

$$\epsilon^*_{\text{i}}(D) \ge \epsilon_X$$

for any D with $0 \le D \le n$, where ϵ_X is the constant defined in (11)*.*

Proof: Suppose by contradiction that there exists a D with $0 \le D \le n$ such that $\epsilon^*_{\text{i}}(D) < \epsilon_X$. Let δ be an arbitrary positive number with $0 < \delta < \epsilon_X - \epsilon^*_{\text{i}}(D)$, and let $\epsilon = \epsilon^*_{\text{i}}(D) + \delta$. Then $\epsilon < \epsilon_X$ and (ϵ, D) is achievable under identifiability. Consider the mechanism that achieves (ϵ, D). Then by the requirement of identifiability, for any neighboring $x, x' \in \mathcal{D}^n$ and any $y \in \mathcal{D}^n$,

$$p_{X|Y}(x \mid y) \le e^{\epsilon} p_{X|Y}(x' \mid y). \tag{A.11}$$

Let $p_Y(\cdot)$ be the pmf of the output Y. Then $p_Y(y) \ge 0$ for any $y \in \mathcal{D}^n$. Therefore, multiplying both sides of (A.11) by $p_Y(y)$ and taking summation over $y \in \mathcal{D}^n$ yield

$$\sum_{y \in \mathcal{D}^n} p_{X|Y}(x \mid y) p_Y(y) \le \sum_{y \in \mathcal{D}^n} e^{\epsilon} p_{X|Y}(x' \mid y) p_Y(y),$$

which implies

$$p_X(x) \le e^{\epsilon} p_X(x').$$

Then there do not exist neighboring $x, x' \in \mathcal{D}^n$ with $p_X(x) = e^{\epsilon_X} p_X(x')$ since $\epsilon < \epsilon_X$, which contradicts with the definition of ϵ_X in (11). □

Lemma 4: For $\epsilon \ge \widetilde{\epsilon}_X$, the mechanism $\mathcal{E}^{\epsilon}_{\text{i}}$ defined in (21) *satisfies ϵ-identifiability, and the distortion of $\mathcal{E}^{\epsilon}_{\text{i}}$ is given by $\mathbb{E}[d(X, Y)] = h(\epsilon)$.*

Proof: Consider any $\epsilon \ge \widetilde{\epsilon}_X$. Then under the mechanism $\mathcal{E}^{\epsilon}_{\text{i}}$, the posterior probability for any $x, y \in \mathcal{D}^n$ is given by

$$p_{X|Y}(x \mid y) = \frac{p_{Y|X}(y \mid x) p_X(x)}{p_Y(y)} = \frac{e^{-\epsilon d(x,y)}}{\left(1 + (m-1)e^{-\epsilon}\right)^n}.$$

As shown in the proof of Lemma 1, this $\{p_{X|Y}(x \mid y), x, y \in \mathcal{D}^n\}$ and the corresponding $\{p_Y(y), y \in \mathcal{D}^n\}$ form an optimal solution of the relaxed optimization problem R-PD. Following the same arguments as in the proof of Lemma 1 we can conclude that $\mathcal{E}^{\epsilon}_{\text{i}}$ satisfies ϵ-identifiability, and the distortion of $\mathcal{E}^{\epsilon}_{\text{i}}$ is given by $\mathbb{E}[d(X, Y)] = h(\epsilon)$. □

Lemma 5: The mechanism $\mathcal{E}^{\epsilon}_{\text{d}}$ defined in (22) *satisfies ϵ-differential privacy, and the distortion of $\mathcal{E}^{\epsilon}_{\text{d}}$ is given by $\mathbb{E}[d(X, Y)] = h(\epsilon)$.*

Proof: Under mechanism $\mathcal{E}^{\epsilon}_{\text{d}}$, $\{p_{Y|X}(y \mid x), x, y \in \mathcal{D}^n\}$ has the same form as the posteriors under mechanism $\mathcal{E}^{\epsilon}_{\text{i}}$. Therefore still by similar arguments as in the proof of Lemma 1, $\mathcal{E}^{\epsilon}_{\text{d}}$ satisfies ϵ-differential privacy, and the distortion of $\mathcal{E}^{\epsilon}_{\text{d}}$ is given by $\mathbb{E}[d(X, Y)] = h(\epsilon)$. □

REFERENCES

[1] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proc. 3rd Conf. Theory Cryptograph. (TCC)*, New York, NY, USA, 2006, pp. 265–284.

[2] C. Dwork, "Differential privacy," in *Proc. 33rd Int. Conf. Automata, Lang. Program. (ICALP)*, Venice, Italy, 2006, pp. 1–12.

[3] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: Privacy via distributed noise generation," in *Proc. 24th Annu. Int. Conf. Theory Appl. Cryptograph. Techn. (EUROCRYPT)*, St. Petersburg, Russia, 2006, pp. 486–503.

[4] K. Nissim, S. Raskhodnikova, and A. Smith, "Smooth sensitivity and sampling in private data analysis," in *Proc. 39th Ann. ACM Symp. Theory Comput. (STOC)*, San Diego, CA, USA, 2007, pp. 75–84.

[5] A. Ghosh, T. Roughgarden, and M. Sundararajan, "Universally utility-maximizing privacy mechanisms," in *Proc. 41st Ann. ACM Symp. Theory Comput. (STOC)*, Bethesda, MD, USA, 2009, pp. 351–360.

[6] A. Roth and T. Roughgarden, "Interactive privacy via the median mechanism," in *Proc. 42nd Ann. ACM Symp. Theory Comput. (STOC)*, Cambridge, MA, USA, 2010, pp. 765–774.

[7] M. Hardt and G. N. Rothblum, "A multiplicative weights mechanism for privacy-preserving data analysis," in *Proc. 51st Ann. IEEE Symp. Found. Comput. Sci. (FOCS)*, Las Vegas, NV, USA, Oct. 2010, pp. 61–70.

[8] A. Gupta, A. Roth, and J. Ullman, "Iterative constructions and private data release," in *Proc. 9th Int. Conf. Theory Cryptograph. (TCC)*, Sicily, Italy, 2012, pp. 339–356.

[9] S. Muthukrishnan and A. Nikolov, "Optimal private halfspace counting via discrepancy," in *Proc. 44th Annu. ACM Symp. Theory Comput. (STOC)*, New York, NY, USA, 2012, pp. 1285–1292.

[10] A. Blum, K. Ligett, and A. Roth, "A learning theory approach to non-interactive database privacy," in *Proc. 14th Annu. ACM Symp. Theory Comput. (STOC)*, Victoria, BC, Canada, 2008, pp. 609–618.
[11] C. Dwork, M. Naor, O. Reingold, G. N. Rothblum, and S. Vadhan, "On the complexity of differentially private data release: Efficient algorithms and hardness results," in *Proc. 41st Annu. ACM Symp. Theory Comput. (STOC)*, Bethesda, MD, USA, 2009, pp. 381–390.
[12] S. P. Kasiviswanathan, M. Rudelson, A. Smith, and J. Ullman, "The price of privately releasing contingency tables and the spectra of random matrices with correlated rows," in *Proc. 42nd Annu. ACM Symp. Theory Comput. (STOC)*, Cambridge, MA, USA, 2010, pp. 775–784.
[13] J. Ullman and S. Vadhan, "PCPs and the hardness of generating private synthetic data," in *Proc. 8th Conf. Theory Cryptograph. (TCC)*, Providence, RI, USA, 2011, pp. 400–416.
[14] M. Hardt, K. Ligett, and F. McSherry, "A simple and practical algorithm for differentially private data release," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, Lake Tahoe, NV, USA, Dec. 2012, pp. 2339–2347.
[15] M. Bun, J. Ullman, and S. Vadhan, "Fingerprinting codes and the price of approximate differential privacy," in *Proc. 46th Annu. ACM Symp. Theory Comput. (STOC)*, New York, NY, USA, 2014, pp. 1–10.
[16] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Found. Trends Theoretical Comput. Sci.*, vol. 9, nos. 3–4, pp. 211–407, Aug. 2014.
[17] J. Lee and C. Clifton, "Differential identifiability," in *Proc. 18th Annu. ACM SIGKDD Conf. Knowl. Discovery Data Mining (KDD)*, Beijing, China, 2012, pp. 1041–1049.
[18] N. Li, W. Qardaji, D. Su, Y. Wu, and W. Yang, "Membership privacy: A unifying framework for privacy definitions," in *Proc. ACM SIGSAC Conf. Comput. Commun. Security (CCS)*, Berlin, Germany, 2013, pp. 889–900.
[19] D. Agrawal and C. C. Aggarwal, "On the design and quantification of privacy preserving data mining algorithms," in *Proc. Symp. Principles Database Syst. (PODS)*, Santa Barbara, CA, USA, 2001, pp. 247–255.
[20] D. Clark, S. Hunt, and P. Malacaria, "Quantitative information flow, relations and polymorphic types," *J. Logic Comput.*, vol. 15, no. 2, pp. 181–199, Apr. 2005.
[21] G. Smith, "On the foundations of quantitative information flow," in *Proc. 12th Int. Conf. Found. Softw. Sci. Comput. Struct. (FSSACS)*, York, U.K., 2009, pp. 288–302.
[22] Y. Zhu and R. Bettati, "Anonymity vs. information leakage in anonymity systems," in *Proc. 25th IEEE Int. Conf. Distrib. Comput. Syst. (ICDCS)*, Columbus, OH, USA, Jun. 2005, pp. 514–524.
[23] K. Chatzikokolakis, C. Palamidessi, and P. Panangaden, "Anonymity protocols as noisy channels," in *Proc. 2nd Int. Conf. Trustworthy Global Comput. (TGC)*, Lucca, Italy, 2007, pp. 281–300.
[24] K. Chatzikokolakis, T. Chothia, and A. Guha, "Statistical measurement of information leakage," in *Proc. 16th Int. Conf. Tools Algorithms Construction Anal. Syst. (TACAS)*, Paphos, Cyprus, 2010, pp. 390–404.
[25] D. Rebollo-Monedero, J. Forné, and J. Domingo-Ferrer, "From t-closeness-like privacy to postrandomization via information theory," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 11, pp. 1623–1636, Nov. 2010.
[26] M. S. Alvim, M. E. Andrés, K. Chatzikokolakis, P. Degano, and C. Palamidessi, "Differential privacy: On the trade-off between utility and information leakage," in *Formal Aspects Security Trust* (Lecture Notes in Computer Science), vol. 7140, Berlin, Germany: Springer, 2012, pp. 39–54.
[27] F. du Pin Calmon and N. Fawaz, "Privacy against statistical inference," in *Proc. 50th Annu. Allerton Conf. Commun., Control Comput.*, Monticello, IL, USA, Oct. 2012, pp. 1401–1408.
[28] A. Makhdoumi and N. Fawaz, "Privacy-utility tradeoff under statistical uncertainty," in *Proc. 51st Annu. Allerton Conf. Commun., Control Comput.*, Monticello, IL, USA, Oct. 2013, pp. 1627–1634.
[29] D. J. Mir, "Information-theoretic foundations of differential privacy," in *Foundations and Practice of Security* (Lecture Notes in Computer Science), vol. 7743, Berlin, Germany: Springer, 2013, pp. 374–381.
[30] L. Sankar, S. R. Rajagopalan, and H. V. Poor, "Utility-privacy tradeoffs in databases: An information-theoretic approach," *IEEE Trans. Inf. Forensics Security*, vol. 8, no. 6, pp. 838–852, Jun. 2013.
[31] A. D. Sarwate and L. Sankar, "A rate-disortion perspective on local differential privacy," in *Proc. 52nd Annu. Allerton Conf. Commun., Control Comput.*, Monticello, IL, USA, Sep./Oct. 2014, pp. 903–908.
[32] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Hoboken, NJ, USA: Wiley, 2006.
[33] A. McGregor, I. Mironov, T. Pitassi, O. Reingold, K. Talwar, and S. Vadhan, "The limits of two-party differential privacy," in *Proc. 51st Annu. IEEE Symp. Found. Comput. Sci. (FOCS)*, Las Vegas, NV, USA, Oct. 2010, pp. 81–90.
[34] A. De, "Lower bounds in differential privacy," in *Theory of Cryptography* (Lecture Notes in Computer Science), vol. 7194, Berlin, Germany: Springer, 2012, pp. 321–338.
[35] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *Proc. 48th Annu. IEEE Symp. Found. Comput. Sci. (FOCS)*, Providence, RI, USA, Oct. 2007, pp. 94–103.
[36] S. L. Warner, "Randomized response: A survey technique for eliminating evasive answer bias," *J. Amer. Statist. Assoc.*, vol. 60, no. 309, pp. 63–69, Mar. 1965.
[37] R. E. Blahut, "Computation of channel capacity and rate-distortion functions," *IEEE Trans. Inf. Theory*, vol. 18, no. 4, pp. 460–473, Jul. 1972.
[38] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge Univ. Press, 2004.
[39] *Netflix Prize*. [Online]. Available: http://www.netflixprize.com
[40] W. Wang, L. Ying, and J. Zhang, "A minimax distortion view of differentially private query release," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Pacific Grove, CA, USA, Nov. 2015, pp. 1046–1050.

Weina Wang (S'13) received her B.E. degree in Electronic Engineering from Tsinghua University, Beijing, China, in 2009. She is currently pursuing a Ph.D. degree in the School of Electrical, Computer and Energy Engineering at Arizona State University, Tempe, AZ. Her research interests include resource allocation in stochastic networks, data privacy and game theory. She won the Joseph A. Barkson Fellowship for the 2015-16 academic year from Arizona State University. She received the Kenneth C. Sevcik Outstanding Student Paper Award at ACM SIGMETRICS 2016.

Lei Ying (M'08) received his B.E. degree from Tsinghua University, Beijing, China, and his M.S. and Ph.D. in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign. He currently is an Associate Professor at the School of Electrical, Computer and Energy Engineering at Arizona State University, and an Associate Editor of the IEEE/ACM Transactions on Networking.

His research interest is broadly in the area of stochastic networks, including data privacy, social networks, cloud computing, and communication networks. He is coauthor with R. Srikant of the book Communication Networks: An Optimization, Control and Stochastic Networks Perspective, Cambridge University Press, 2014.

He won the Young Investigator Award from the Defense Threat Reduction Agency (DTRA) in 2009 and NSF CAREER Award in 2010. He was the Northrop Grumman Assistant Professor in the Department of Electrical and Computer Engineering at Iowa State University from 2010 to 2012. He received the best paper award at IEEE INFOCOM 2015.

Junshan Zhang (M'01–SM'05–F'12) received his Ph.D. degree from the School of ECE at Purdue University in 2000. He joined the School of ECEE at Arizona State University in August 2000, where he has been Fulton Chair Professor since 2015. His research interests fall in the general field of information networks and its intersections with social networks and power networks. His current research focuses on fundamental problems in information networks and energy networks, including modeling and optimization for cyber-physical systems, optimization/control of mobile social networks and cognitive radio networks, and privacy/security in information networks.

Prof. Zhang is a fellow of the IEEE, and a recipient of the ONR Young Investigator Award in 2005 and the NSF CAREER award in 2003. He received the Outstanding Research Award from the IEEE Phoenix Section in 2003. He co-authored two papers that won the Best Paper Runner-up Award of IEEE INFOCOM 2009 and IEEE INFOCOM 2014, and a paper that won IEEE ICC 2008 Best Paper Award. He was TPC co-chair for a number of major conferences in communication networks, including IEEE INFOCOM 2012 and ACM MOBIHOC 2015. He was the general chair for WiOpt 2016 and IEEE Communication Theory Workshop 2007. He was an Associate Editor for IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS, an editor for the *Computer Network Journal* and an editor *IEEE Wireless Communication Magazine*. He was a Distinguished Lecturer of the IEEE Communications Society. He is currently serving as an editor-at-large for IEEE/ACM TRANSACTIONS ON NETWORKING and an editor for *IEEE Network Magazine*.