

A Privacy Framework: Indistinguishable Privacy

Jinfei Liu
Emory University
Atlanta, USA
jinfei.liu@emory.edu

Li Xiong
Emory University
Atlanta, USA
lxiong@mathcs.emory.edu

Jun Luo
Shenzhen Institutes of
Advanced Technology, CAS
jun.luo@siat.ac.cn

ABSTRACT

In this paper we illustrate a privacy framework named Indistinguishable¹ Privacy. Indistinguishable privacy could be deemed as the formalization of the existing privacy definitions in privacy preserving data publishing as well as secure multi-party computation. We introduce three variants of the representative privacy notions in the literature, Bayes-optimal privacy for privacy preserving data publishing, differential privacy for statistical data release, and privacy w.r.t. semi-honest behavior in the secure multi-party computation setting, and prove they are equivalent. To the best of our knowledge, this is the first work that illustrates the relationships of these privacy definitions and unifies them through one framework.

Categories and Subject Descriptors

H.2.7 [Database Management]: Database Administration—*Security, integrity, and protection*

Keywords

K-Anonymity, Bayes-optimal Privacy, Differential Privacy, Secure Multi-party Computation, Privacy w.r.t. Semi-honest Behavior, Indistinguishable Privacy

1. INTRODUCTION

In an age where the details of our life are recorded and stored in databases, a clear paradox is emerging between the need to preserve the privacy of individuals and the need to use these collected data for research or other purposes. Privacy preserving data analysis and publishing (PPDP) has received considerable attention in recent years [10, 8]. The goal is for a data custodian to release some views or statistical computations of the original private data, so that the released data remains practically useful while individual privacy for the data subjects is preserved. A related problem is secure multi-party computation (SMC) where a given

¹We note that this notion came from [11] first and was developed further in [9].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EDBT/ICDT'13 Workshops, March 18-22, 2013, Genoa, Italy.
Copyright 2013 ACM 978-1-4503-1599-9/13/03 ...\$15.00.

number of participants wish to compute a public function on their private inputs without disclosing the private inputs to each other. With decades' development, there are three essential privacy principles in the literature on PPDP and SMC: Bayes-optimal privacy [24] in the PPDP setting, and differential privacy [6] in the PPDP setting, and privacy w.r.t. semi-honest behavior [12] in the SMC setting.

L. Sweeney [33, 32] introduced the well-known *k*-anonymity principle. It is a syntactic property on anonymized databases: if only certain attributes known as quasi-identifiers (QIDs) are considered, each tuple in the anonymized database should appear at least *k* times. However, it is generally considered too weak. Many other derivative privacy principles have been introduced since then, e.g., *l*-diversity [24] and *t*-closeness [21]. A. Machanavajjhala et al. [24] studied that *k*-anonymity is susceptible to homogeneity an background knowledge attacks. To address this problem, they show the notion of Bayes-optimal privacy that provides privacy even when the data publisher does not know what kind of knowledge is possessed by the adversary.

The notion of differential privacy was introduced by C. Dwork et al. [6, 9]. An “algorithm” (or, “mechanism”) \mathcal{A} satisfies ϵ -differential privacy if and only if for two datasets D_1 and D_2 which differ at only one tuple, the distributions of $\mathcal{A}(D_1)$ and $\mathcal{A}(D_2)$ differ at most by a multiplicative factor of e^ϵ . A relaxed version of ϵ -differential privacy, which we use (ϵ, δ) -differential privacy to denote, allows an error probability bounded by δ . Satisfying differential privacy ensures that even if the adversary has full knowledge of the values of a tuple t , as well as full knowledge of other tuples in the dataset, and is only uncertain about whether t is in the input dataset, the adversary is unlikely to have precise information about the tuple's presence in the database. Differential privacy is achieved by introducing randomness into query answers. The original algorithm for achieving differential privacy, commonly called the Laplace mechanism [9], returns the sum of the true answer and random noise drawn from a Laplace distribution. The scale of the distribution is determined by a property of the query called sensitivity: roughly the maximum possible change to the query answer induced by the addition or removal of one tuple. Higher sensitivity queries are more likely to reveal individual tuples and thus they should receive greater noise.

Yao first introduced the two party comparison problem (Yao's Millionaire Problem), where two millionaires want to find

out who is richer without disclosing their actual wealth, and developed a provably secure solution [37]. This was extended to multi-party computation by Goldreich et al. [13]. They developed a framework for secure multi-party computation (SMC), and proved that computing a function privately is equivalent to computing it securely in [12]. The protocol is *secure* if the parties learn only the result of the function and nothing else.

1.1 Contributions

While there are extensive literature based on each of the privacy definitions, few works studied their inter relationship. We believe it is important and significant to understand their relationships and the underlying principles that unify these definitions. The main contributions of this paper are as follows:

1. We present a privacy framework named indistinguishable privacy that could unify existing privacy definitions for privacy preserving data publishing and secure multi-party computation.
2. We define three variants of the representative privacy principles in the literature - Bayes-optimal privacy in privacy preserving data publishing setting and differential privacy in privacy preserving data publishing setting and privacy w.r.t. semi-honest behavior in the secure multi-party computation setting - and prove they are equivalent.
3. We also show that the indistinguishable Privacy framework satisfies the two privacy axioms.

The rest of the paper is organized as follows. Section 2 presents the new privacy framework, analyzes the relationship of the three definitions in the literature. We prove indistinguishable privacy satisfies privacy axioms in Section 3. We review the related work in Section 4. Section 5 concludes the paper.

2. INDISTINGUISHABLE PRIVACY

In this section, we illustrate our privacy framework *indistinguishable privacy*. We first discuss the underlying principle followed by the formal definition of the framework.

Let us briefly describe Shannon’s definition of “perfect secrecy” [29]. Consider an adversary with unlimited time and manpower available for analysis of intercepted cryptograms. Let the set of all possible messages be finite. These messages have a priori probability for each and are encoded and sent across the wire. When an adversary intercepts an encoded message, he can calculate its posterior probabilities for the various messages. Perfect secrecy is achieved if for all encoded messages the posterior probabilities are equal to the prior probabilities. Thus intercepting the message gives the adversary no information.

Compare with “perfect secrecy” that provably can not be broken even with unlimited time and manpower available, semantic security (computationally secure) is theoretically possible to break but it is infeasible to do so by any known practical means. Due to the use of cryptography knowledge,

SMC is based on semantic security. Hence, in this paper, our privacy framework definition also is based on semantic security.

Dalenius [4] defined a similar privacy goal for statistical database: access to the published data should not enable the attacker to learn anything extra about any target victim compared to no access to the database, even with the presence of any attacker’s background knowledge obtained from other sources. The literature on privacy preserving data publishing has attempted to formalize Dalenius’ goal or define similar notions for privacy. Machanavajjhala et al. defined a Bayes-optimal privacy notion for privacy preserving data publishing based on the uninformative principle (for published data).

Uninformative Principle [24]. The *published table* should provide the adversary with little additional information beyond the background knowledge. In other words, there should not be a large difference between the prior and posterior beliefs.

2.1 Notations and Privacy Principle

We define a formal indistinguishable privacy framework based on a generalized uninformative principle and show that it can be used to unify the privacy definitions for both PPDP and SMC. Given an operation of private data (e.g. a data release, a statistical computation, or a multi-party secure-computation), intercepting the operation knowledge should give an adversary no or little (ϵ) new information about sensitive knowledge in private data.

We first define the notations below.

Background Knowledge (BK) Background knowledge is what adversary obtained before one operation.

Sensitive Information (SI) Individual sensitive information is the information that needs to be protected, such as the salary or state of health of an individual.

Operation Knowledge (OK) Operation knowledge is the knowledge that a participant will obtain after an operation. For example, it can be the published data after the operation of data release.

We redefine the uninformative principle with a generalization from published data to the operation knowledge as follows:

Uninformative Principle for Operations. The *operation knowledge* should provide the adversary with little *sensitive information* beyond the *background knowledge*. In other words, the prior and posterior beliefs should be *indistinguishable*.

2.2 Indistinguishable Privacy Definition

Consider an adversary Alice, we denote her prior belief about the sensitive information as:

$$\alpha_{(BK,SI)} = P(SI|BK)$$

After an operation is executed, Alice’s posterior belief is:

$$\beta_{(BK,SI,OK)} = P(SI|BK,OK)$$

Table 1: Summary of Notations and Mapping with Existing Privacy Definitions

Notation \ Definition	Bayes-optimal Privacy	Differential Privacy	Privacy w.r.t Semi-honest Behavior in SMC
Background Knowledge (BK)	Dataset that can be linked with the quasi-identifying attributes in the original dataset	Query result from dataset D_1 (all but one record in the database)	own input and output
Sensitive Information (SI)	Each individual's sensitive attributes	Each individual's presence	Other party's private input and output
Operation Knowledge (OK)	Anonymized data	Query result from dataset D_2 which differs only in one tuple from Dataset D_1	Intermediate result received during communication with the other party

Definition 1. (Indistinguishable Privacy). An operation satisfies indistinguishable privacy if the following holds for an adversary with polynomially bounded resources available:

$$\alpha_{(BK,SI)} \leq e^\epsilon \beta_{(BK,SI,OK)} \quad (1)$$

The Inequation (1) indicates that before and after the operation, the difference between the prior and posterior beliefs of the adversary about the *sensitive information* (not the complete information) differs by no more than a very small fraction ϵ . ϵ can be considered as the disclosure. When ϵ is small, $\ln(1 + \epsilon) \approx \epsilon$.

Given the above definition, we can show that each of the three existing privacy definitions can be mapped to the indistinguishable privacy framework. Table 1 shows a summary of the mapping.

EXAMPLE 1. *Bayes-optimal Privacy: Alice has background knowledge (BK), e.g. publicly available voter registration list, that can be used to link the quasi-identifying attributes in the dataset and identify the individuals in order to compute the sensitive information (SI). After she obtains the published dataset (OK), she could compute the sensitive information (SI) based on both the background knowledge and the published data (BK, OK).*

EXAMPLE 2. *Differential Privacy: Alice has background knowledge (BK), in the worst case, about all but one record in the database, which is considered as D_1 . She can compute sensitive information (SI) based on the query result on D_1 . After she obtains the query answer (OK) from the real dataset D_2 which only differs from D_1 by one record, she could compute the sensitive information (SI) based on D_1 's query result and D_2 's query result (BK, OK).*

EXAMPLE 3. *Privacy w.r.t. Semi-honest Behavior in SMC: Alice has some knowledge (BK) to compute sensitive information (SI). After she receives intermediate result (OK) from the other party during the execution of the protocol, she could compute the sensitive information (SI) based on her own background knowledge and the intermediate result (BK, OK).*

2.3 Relationships of Existing Privacy Definitions

We present three variants of the existing privacy definitions and show their equality in this subsection.

2.3.1 Variant of Bayes-Optimal Privacy

We present a variant of Bayes-optimal privacy [24] that bounds the differences between the prior and posterior beliefs.

Consider an original table T which contains a tuple t about Bob. Let S denote the set of all sensitive attributes and Q denote the set of all nonsensitive attributes. Alice's prior belief about Bob's sensitive information is denoted as $\alpha_{(q,s)}$, where s is Bob's sensitive attribute values and q is his nonsensitive attribute values:

$$\alpha_{(q,s)} = P(t[S] = s | t[Q] = q)$$

After the anonymization operation, Alice observes the anonymized table T^* , her posterior belief about Bob's sensitive attribute is denoted as $\beta_{(q,s,T^*)}$:

$$\beta_{(q,s,T^*)} = P(t[S] = s | t[Q] = q \wedge \exists t^* \in T^*, t \rightarrow t^*).$$

Definition 2. (ϵ -Bayes-optimal privacy for anonymized data). The anonymized table T^* satisfies ϵ -adversarial privacy if the following holds:

$$\alpha_{(q,s)} \leq e^\epsilon \beta_{(q,s,T^*)}.$$

2.3.2 Variant of Differential Privacy

We present a variant of differential privacy based on the notion of free-lunch privacy [17].

Definition 3. (ϵ -free-lunch privacy). A randomized algorithm \mathcal{A} satisfies ϵ -free-lunch privacy if for any pair of database instances D_1, D_2 (not just differing at one tuple) and for any set S ,

$$P(\mathcal{A}(D_1) \in S) \leq e^\epsilon P(\mathcal{A}(D_2) \in S)$$

2.3.3 Variant of Privacy w.r.t. Semi-honest Behavior in SMC

We present a variant of the privacy definition w.r.t. Semi-honest Behavior in SMC based on [11].

Definition 4. (ϵ -privacy w.r.t. Semi-honest Behavior for SMC). For a function f , we say that \prod privately computes

one operation f if there exist probabilistic polynomial time algorithms, denoted S_1 and S_2 for any ϕ and φ , such that

$$\begin{aligned} & P(\{S_1(x, f_1(x, y)) = \phi\}_{x, y \in \{0, 1\}^*}) \\ & \leq e^\epsilon P(\{view_1^\Pi(x, y) = \phi\}_{x, y \in \{0, 1\}^*}), \end{aligned} \quad (2)$$

$$\begin{aligned} & P(\{S_2(y, f_2(x, y)) = \varphi\}_{x, y \in \{0, 1\}^*}) \\ & \leq e^\epsilon P(\{view_2^\Pi(x, y) = \varphi\}_{x, y \in \{0, 1\}^*}), \end{aligned} \quad (3)$$

where x and y are the input, $f(x, y)$ is the output. $f_1(x, y)$ (resp., $f_2(x, y)$) denotes the first (resp., second) element of $f(x, y)$. The view of the first (resp., second) party during an execution of Π on (x, y) , denoted $view_1^\Pi(x, y)$ (resp., $view_2^\Pi(x, y)$). “ $\leq e^\epsilon$ ” is equivalent to “ $\stackrel{\epsilon}{\equiv}$ ” in [11].

2.3.4 Equivalence

THEOREM 1. *The three definitions: ϵ -Bayes-optimal privacy for anonymized data, ϵ -free-lunch privacy for a randomized algorithm, and ϵ -privacy w.r.t. semi-honest behavior for SMC, are equivalent.*

PROOF. From Bayes-optimal privacy’s viewpoint:

$$\alpha_{(q,s)} \leq e^\epsilon \beta_{(q,s,T^*)}$$

$$\Leftrightarrow P(t[S] = s | t[Q] = q) \leq e^\epsilon P(t[S] = s | t[Q] = q \wedge \exists t^* \in T^*, t \rightarrow t^*)$$

$$\Leftrightarrow \alpha_{(BK,SI)} \leq e^\epsilon \beta_{(BK,SI,OK)}$$

where $(t[S] = s) = SI$, $(t[Q] = q) = BK$ and $(\exists t^* \in T^*, t \rightarrow t^*) = OK$.

From differential privacy’s viewpoint:

$$P(\mathcal{A}(D_1) \in S) \leq e^\epsilon P(\mathcal{A}(D_2) \in S)$$

$$\Leftrightarrow \alpha_{(BK,SI)} \leq e^\epsilon \beta_{(BK,SI,OK)}$$

where *whether the certain tuple existing* = SI , *query result from $D_1 = BK$* and *query result from $D_2 = OK$* .

From privacy w.r.t semi-honest behavior’s viewpoint:

$$\begin{aligned} & P(\{S_1(x, f_1(x, y)) = \phi\}_{x, y \in \{0, 1\}^*}) \\ & \leq e^\epsilon P(\{view_1^\Pi(x, y) = \phi\}_{x, y \in \{0, 1\}^*}) \\ & \Leftrightarrow \alpha_{(BK,SI)} \leq e^\epsilon \beta_{(BK,SI,OK)} \end{aligned}$$

where $(y, f_2(x, y)) = SI$, $(S_1(x, f_1(x, y))) = BK$ and $(view_1^\Pi(x, y)) = OK$. Similarly, Inequation (3) satisfies the constraint.

Summing up the above, the theorem is correct. \square

3. PRIVACY AXIOMS

Modern design guidelines for privacy definitions include two fundamental axioms known as transformation invariance and convexity [16]. While the ideas contained in the axioms have been accepted by the privacy community for a long time, only recently there has been an insistence that privacy definitions should satisfy them.

In this section we show indistinguishable privacy framework satisfies both fundamental axioms, thus ensuring that it satisfies modern design guidelines.

The axioms are:

AXIOM 1. (*Transformation Invariance [16]*). *If an algorithm \mathcal{M} satisfies a privacy definition and \mathcal{A} is any algorithm such that (1) its domain contains the range of \mathcal{M} and (2) its random bits (if any) are statistically independent from the random bits (if any) of \mathcal{M} , then the algorithm $\mathcal{A} \circ \mathcal{M}$, which first runs \mathcal{M} on the data and then runs \mathcal{A} on the output should also satisfy the same privacy definition.*

AXIOM 2. (*Convexity [16]*). *If \mathcal{M}_1 and \mathcal{M}_2 both satisfy a privacy definition, and $p \in [0, 1]$, then the algorithm \mathcal{M}^p which runs \mathcal{M}_1 with probability p and \mathcal{M}_2 with probability $1 - p$ should also satisfy the privacy definition.*

The following theorem confirms that the indistinguishable privacy framework satisfies modern privacy design guidelines.

THEOREM 2. *Indistinguishable privacy satisfies the axioms of convexity and transformation invariance.*

PROOF. Let \mathcal{M} be an algorithm that satisfies indistinguishable privacy, for any s we have:

$$P(SI = s | BK) \leq e^\epsilon P(SI = s | BK, OK(\mathcal{M}))$$

where $OK(\mathcal{M})$, the operation knowledge of \mathcal{M} , is what you obtain after the operation \mathcal{M} .

Step 1: First we prove it satisfies the axiom of transformation invariance. Let \mathcal{A} be any algorithm whose domain contains the range of \mathcal{M} and whose random bits are independent from those of \mathcal{M} .

$$\begin{aligned} & P(SI = s | BK) \leq e^\epsilon P(SI = s | BK, OK(\mathcal{M})) \\ & = e^\epsilon \frac{P(SI = s, BK, OK(\mathcal{M}))}{P(BK, OK(\mathcal{M}))} \\ & = e^\epsilon \frac{P(SI = s, BK, OK(\mathcal{M})) \times P(OK(\mathcal{A} \circ \mathcal{M}))}{P(BK, OK(\mathcal{M})) \times P(OK(\mathcal{A} \circ \mathcal{M}))} \\ & = e^\epsilon \frac{P(SI = s, BK, OK(\mathcal{A} \circ \mathcal{M}), OK(\mathcal{M}))}{P(BK, OK(\mathcal{A} \circ \mathcal{M}), OK(\mathcal{M}))} \\ & = e^\epsilon P(SI = s | BK, OK(\mathcal{A} \circ \mathcal{M}), OK(\mathcal{M})) \end{aligned}$$

Thus $\mathcal{A} \circ \mathcal{M}$ also satisfies the privacy definition.

Step 2: Now we prove that it satisfies that axiom of convexity. Let \mathcal{M}_1 and \mathcal{M}_2 be algorithms that satisfy the axiom of convexity. Let \mathcal{M}_1 and \mathcal{M}_2 be algorithms that satisfy the privacy definition indistinguishable privacy, let $p \in [0, 1]$ and let \mathcal{M}^p be the algorithm that runs \mathcal{M}_1 with probability p and \mathcal{M}_2 with probability $1 - p$.

$$P(SI = s | BK)$$

$$\begin{aligned}
&= pP(SI = s|BK) + (1 - p)P(SI = s|BK) \\
&\leq e^\epsilon pP(SI = s|BK, OK(\mathcal{M}_1)) + e^\epsilon (1 - p)P(BK, OK(\mathcal{M}_2)) \\
&= e^\epsilon P(SI = s|BK, OK(\mathcal{M}^p))
\end{aligned}$$

Thus \mathcal{M}^p also satisfies the privacy definition. \square

4. RELATED WORK

In a k -anonymized dataset, each record is *indistinguishable* from at least $k - 1$ other records with respect to certain identifying attributes. Some follow-up notions include l -diversity [24], which requires that the distribution of a sensitive attribute in each equivalence class has at least l “well represented” values. l -diversity is based on the Bayes-optimal privacy which involves modeling background knowledge as a probability distribution over the attributes and uses Bayesian inference techniques to reason about privacy. t -closeness [21], which requires that the distribution of a sensitive attribute in any equivalent class is close to the distribution of the attribute in the overall table. A series [20, 30, 1, 31, 35] of papers based on k -anonymity are presented. For a detailed survey, please see [10].

The principle of differential privacy was developed in a series of works [6, 9, 2, 19, 17, 5, 3, 36, 14]. Differential privacy requires that computations be insensitive to changes in any particular individual’s record, thereby restricting data leaks through the results. It represents a major breakthrough in privacy preserving data analysis. In an attempt to make differential privacy more amenable to more sensitive queries, several relaxations have been developed, including (ϵ, δ) -differential privacy [2, 9]. Three basic general approaches to achieve differential privacy add Laplace noise proportional to the query’s global sensitivity [6, 9], adding noise related to the smooth bound of the query’s local sensitivity [27], and the exponential mechanism to select a result among all possible results [26]. A survey on these results can be found in [7].

Secure two party computation was first investigated by Yao [37], and was later generalized to multi-party computation in [11, 12]. The basic problem is Yao’s Millionaires’ problem that two millionaires would like to know who is richer, but neither wants to reveal their real worth. Abstractly, the problem is simply comparing two numbers, each held by one party, however, without either party revealing its number to the other. Yao presented a solution for any efficiently computable function restricted to two parties and semi-honest adversaries. Goldreich first [12] introduced the privacy w.r.t semi-honest behavior. Due to the inefficiency of generic protocols, some research has focused on finding efficient protocols for specific problem of secure computation. See privacy preserving SVM [23] and privacy preserving k -means [34, 15] for just a few examples.

The work closest to ours are V. Rastogi et al. [28], A. McGregor et al. [25], and N. Li et al. [22]. [28] has the original equivalent result on adversarial privacy and differential privacy, which is of independent interest: an algorithm is ϵ -indistinguishable if and only if it is private for a particular class of adversaries. [25] demonstrates a connection between differential privacy and deterministic extraction from Santha-Vazirani sources. [22] uses the random sampling step

to bridge the gap between k -anonymity and differential privacy. D. Kifer et al. [18] introduced a rigorous and customizable framework for privacy which is the most closest work to ours. The main difference between our paper and others is that we consider the privacy w.r.t. semi-honest behavior while others’ work like D. Kifer et al. only consider the k -anonymity and differential privacy. We prove that Bayes-optimal privacy, differential privacy are similar with privacy w.r.t. semi-honest behavior which is the original idea in some extent (variants of three definitions in our paper).

5. CONCLUSIONS AND DISCUSSIONS

In this paper, we illustrate a new privacy framework named indistinguishable privacy and define three variants of Bayes-optimal privacy, differential privacy and privacy w.r.t. semi-honest behavior, and prove they are equivalent. This is the first work to unify the existing definitions: Bayes-optimal privacy, differential privacy and privacy w.r.t. semi-honest behavior into one framework.

We also note that this is only a first step towards unifying the distinct privacy notions and many challenging issues remain to be solved. While the framework allows the framing of similar sets of equations for each of the privacy definition, it is also important to devise ways to capture numerically the various parameters. For instance, finding expressions to represent the background knowledge BK in the context of k -anonymity or syntactic privacy is itself a challenging issue. We hope there will be more work in this direction from the community that help us understand the foundations of the privacy definitions.

Acknowledgement

This research is partially supported by the National Science Foundation under grant CNS-1117763 and NSF of China under project 11271351. The authors would also like to thank the anonymous reviewers for their valuable comments that helped improve the final version of the paper.

6. REFERENCES

- [1] M. M. Baig, J. Li, J. Liu, and H. Wang. Cloning for privacy protection in multiple independent data publications. In *CIKM*, pages 885–894, 2011.
- [2] A. Blum, C. Dwork, F. McSherry, and K. Nissim. Practical privacy: the *sqlq* framework. In *PODS*, pages 128–138, 2005.
- [3] R. Chen, N. Mohammed, B. C. M. Fung, B. C. Desai, and L. Xiong. Publishing set-valued data via differential privacy. *PVLDB*, 4(11):1087–1098, 2011.
- [4] T. Dalenius. Towards a methodology for statistical disclosure control. *Statistik Tidskrift*, (15), 1977.
- [5] B. Ding, M. Winslett, J. Han, and Z. Li. Differentially private data cubes: optimizing noise sources and consistency. In *SIGMOD Conference*, pages 217–228, 2011.
- [6] C. Dwork. Differential privacy. In *ICALP (2)*, pages 1–12, 2006.
- [7] C. Dwork. Differential privacy: A survey of results. In *TAMC*, pages 1–19, 2008.
- [8] C. Dwork. A firm foundation for private data analysis. *Commun. ACM*, 54(1):86–95, Jan. 2011.

- [9] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, pages 265–284, 2006.
- [10] B. C. M. Fung, K. Wang, R. Chen, and P. S. Yu. Privacy-preserving data publishing: A survey of recent developments. *ACM Comput. Surv.*, 42(4), 2010.
- [11] O. Goldreich. *Secure Multi-party Computation (working draft)*. 1998.
- [12] O. Goldreich. *The Foundations of Cryptography - Volume 2, Basic Applications*. Cambridge University Press, 2004.
- [13] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *STOC*, pages 218–229, 1987.
- [14] A. Inan, M. Kantarcioglu, G. Ghinita, and E. Bertino. Private record matching using differential privacy. In *EDBT*, pages 123–134, 2010.
- [15] G. Jagannathan and R. N. Wright. Privacy-preserving distributed k -means clustering over arbitrarily partitioned data. In *KDD*, pages 593–599, 2005.
- [16] D. Kifer and B. Lin. An axiomatic view of statistical privacy and utility. *Journal of Privacy and Confidentiality*, 4(1):2, 2012.
- [17] D. Kifer and A. Machanavajjhala. No free lunch in data privacy. In *SIGMOD Conference*, pages 193–204, 2011.
- [18] D. Kifer and A. Machanavajjhala. A rigorous and customizable framework for privacy. In *Proceedings of the 31st symposium on Principles of Database Systems*, pages 77–88. ACM, 2012.
- [19] C. Li, M. Hay, V. Rastogi, G. Miklau, and A. McGregor. Optimizing linear counting queries under differential privacy. In *PODS*, pages 123–134, 2010.
- [20] J. Li, R. C.-W. Wong, A. W.-C. Fu, and J. Pei. Anonymization by local recoding in data with attribute hierarchical taxonomies. *IEEE Trans. Knowl. Data Eng.*, 20(9):1181–1194, 2008.
- [21] N. Li, T. Li, and S. Venkatasubramanian. t -closeness: Privacy beyond k -anonymity and l -diversity. In *ICDE*, pages 106–115, 2007.
- [22] N. Li, W. H. Qardaji, and D. Su. Provably private data anonymization: Or, k -anonymity meets differential privacy. *CoRR*, abs/1101.2604, 2011.
- [23] K.-P. Lin and M.-S. Chen. Privacy-preserving outsourcing support vector machines with random transformation. In *KDD*, pages 363–372, 2010.
- [24] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkatasubramanian. l -diversity: Privacy beyond k -anonymity. *TKDD*, 1(1), 2007.
- [25] A. McGregor, I. Mironov, T. Pitassi, O. Reingold, K. Talwar, and S. P. Vadhan. The limits of two-party differential privacy. In *FOCS*, pages 81–90, 2010.
- [26] F. McSherry and K. Talwar. Mechanism design via differential privacy. In *FOCS*, pages 94–103, 2007.
- [27] K. Nissim, S. Raskhodnikova, and A. Smith. Smooth sensitivity and sampling in private data analysis. In *STOC*, pages 75–84, 2007.
- [28] V. Rastogi, M. Hay, G. Miklau, and D. Suciu. Relationship privacy: output perturbation for queries with joins. In *PODS*, pages 107–116, 2009.
- [29] C. Shannon. Communication theory of secrecy systems. *Bell system technical journal*, 28(4):656–715, 1949.
- [30] P. Shi, L. Xiong, and B. C. M. Fung. Anonymizing data with quasi-sensitive attribute values. In *CIKM*, pages 1389–1392, 2010.
- [31] X. Sun, H. Wang, J. Li, and J. Pei. Publishing anonymous survey rating data. *Data Min. Knowl. Discov.*, 23(3):379–406, 2011.
- [32] L. Sweeney. Achieving k -anonymity privacy protection using generalization and suppression. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5):571–588, 2002.
- [33] L. Sweeney. k -anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5):557–570, 2002.
- [34] J. Vaidya and C. Clifton. Privacy-preserving k -means clustering over vertically partitioned data. In *KDD*, pages 206–215, 2003.
- [35] R. C.-W. Wong, J. Li, A. W.-C. Fu, and K. Wang. (α, k) -anonymous data publishing. *J. Intell. Inf. Syst.*, 33(2):209–234, 2009.
- [36] X. Xiao, G. Wang, and J. Gehrke. Differential privacy via wavelet transforms. In *ICDE*, pages 225–236, 2010.
- [37] A. C.-C. Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167, 1986.