Differential Privacy: A Historical Survey

MICHAEL HILTON
Cal Poly State University

The paper provides a historical survey of differential privacy. It begins with a discussion of the desideratum that the this work is based on. The original motivation for this was a concern for privacy in census data, as described by Dalenius. We will also discuss the landmark paper by Dwork which provides an initial description of differential privacy. There will be a discussion of two separate applications of differential privacy, location pattern mining and health data. Finally we will examine some cases which cause problems for the privacy guarantees afforded to us by differential privacy.

Categories and Subject Descriptors: H.2.8 [**Database Applications**]: Statistical Databases; K.4.1 [**Computers and Society**]: Privacy

General Terms: Databases, Security, Privacy

■

1. INTRODUCTION

The advent of large data sets and the computing power to quickly sift through them has brought privacy concerns to the forefront. As with almost everything in life, there is a tradeoff that must we wrestled with, between privacy and utility. Finding the correct balance between these two poles leads us to a point of tension, but one that may allow us to provide reasonable privacy guarantees while at the same time allowing users and society in general to benefit from the utility that these large datasets can provide.

2. STATISTICAL DISCLOSURE CONTROL

In 1977 Dalenius [Dalenius 1977] articulated a desideratum for statistical data sets. His paper does not talk about databases by name, he only talks about statistical results. In the paper he is concerned with disclosure control of data from a statistical data set. He offers the following definition of disclosure. Consider an object O_k, and a characteristic *D*, which is a survey characteristic. For the object O_k this characteristic assumes the value D_k. If the release of the statistics *S* makes it possible to determine the value of D_k more accurately than it is possible without access to S, a disclosure has taken place.
Dalenius then goes on to describe a typology for statistical disclosure. His typology makes use of 6 dimensions:

(1) kinds of statistics *S* released: micro-statics, or macro-statistics;
(2) the measurement scale we used to express *S*;
(3) accessibility of disclosure: direct of indirect disclosure;
(4) scope of disclosure: external or internal disclosure;
(5) the disclosing entities: *S*- of $S \times E$-disclosure

This typology results in at least 2^6 = 64 categories.

3. DIFFERENTIAL PRIVACY

In 2006 Cynthia Dwork wrote the first paper to define differential privacy[Dwork 2006]. The same author continued to work on the definition in this paper.[Dwork 2008] The paper defines differential privacy and then gives some ways to achieve that goal. As one would expect, there is a significant tonal shift that occurs between 1977 and 2006. Where the Dalenius paper talks about data in terms of census surveys and statistical data sets, Dwork talks about data in the context of a database.

3.1 Definitions

The paper begins with defining some terms. A statistic is a quantity computed from a specific sample. If a database is a representative sample of an underlying population, the goal of a privacy-preserving statistical database is to enable the user to learn properties of the population as a whole, while protecting the privacy of the individuals in the sample. As there will always be a trade off between privacy and utility, in this paper, privacy is paramount, so privacy goals are first defined, and then the utility that can be achieved given those privacy goals will be accepted.

3.1.1 *Auxiliary Information.* Dalenius articulated a desideratum that access to a statistical database should not enable one to learn anything about an individual that could not be learned without access. Dwork shows that this type of privacy cannot be achieved. The problem is in *auxiliary information*, which is defined as information available to the adversary other then from access to the statistical database. The example presented is the following. Suppose one's exact height was considered to be a highly sensitive piece of information, and that revealing it would be considered a breach of privacy. Assume that the database yields the average heights of women of different nationalities. An adversary who has access to the statistical database and the auxiliary information "Terry Gross is two inches shorter then the average Lithuanian woman" learns Terry Gross' height, while anyone learning only the auxiliary information, without access to the average heights, learns relatively little.

3.1.2 *Utility.* While it is true that a mechanism that always outputs an empty string, or a purely random string, clearly preserves privacy, we need utility beyond that. In order for a mechanism to be *useful* its output should not be predictable by the user, but the unpredictability must not stem only from random choices made by the mechanism. Intuitively there should be a vector of questions whose answers *should* be learnable by a user, but whose answers are not known in advance. The paper then posits a *utility vector*. This is a binary vector of some fixed length. We can think of the utility vector as answers to questions about the data.

3.1.3 *Differential Privacy.* The paper offers a formal definition of differential privacy.

DEFINITION 1. *A randomized function κ gives ϵ-differential privacy if for all data sets D1 and D2 differing on at most one element, and all S $\leq$ Range(κ),*

$$\Pr[\kappa(\text{D1}) \in \text{S}] \leq \exp(\epsilon) \times \Pr[\kappa(\text{D2}) \in \text{S}]$$

ϵ is the statistical distance we use to define the strength of the privacy. A mechanism κ satisfying this definition addresses concerns that any participant might have about the leakage of her personal information *x*: even if the participant removed her data from the

data set, no outputs (and thus consequences of outputs) would become significantly more or less likely. For example, the presence or absence of an individual in a database should not significantly affect their chance of receiving insurance coverage.
This definition can be extended to discuss group privacy as well. A collection of c participants can be assured the same outcomes where the privacy dilation is bound by $\exp(\epsilon c)$. While this value may be tolerable for small values of c, the specific aim of data is to disclose aggregate information about large groups, so we should expect the privacy bounds to disintegrate with increasing group size.
As previously discussed, differential privacy is not an absolute guarantee of privacy. However, society has decided that benefits of certain databases outweigh the costs, and differential privacy ensures that the additional risk incurred by participating in the social beneficial databases is limited.

3.2 Implementation of Differential Privacy

The paper proceeds to describe a concrete interactive privacy mechanism achieving ϵ-differential privacy. The mechanism works by adding appropriately chosen random noise to the answer $a = f(X)$, where f is the *query function* and X is the database; thus the query functions may operate on the entire database as once. The paper states that they can be simple -eg, "Count the number of rows in the database satisfying a given predicate" - or complex - e.g., "Compute the median value for each column; if the Column 1 median exceeds the Column 2 median, then output a histogram of the numbers of points in the the S of orthants, else provide a histogram of the numbers of points in a different set T of orthants."

3.2.1 *Exponential Noise.* The paper states that they achieve ϵ-differential privacy by the addition of random noise whose magnitude is choses as a function of the largest change a single participant could have on the output to the query function; this quantity is referred to as the *sensitivity* of the function.

DEFINITION 2. *For $f : \mathcal{D} \rightarrow R^d$, the L1-sensitivity of f is*

$$\Delta f = \max_{D_1, D_2} ||f(D_1) - f(D_2)||_1$$

for all D_1, D_2 differing in at most one element.

For many types of queries Δf will be quite small. For example, simple counting queries such as "How many rows have the property P?" have $\Delta f = 1$. The author states that the technics they describe work best, introduce the least noise, when Δf is small. This sensitivity is a property of the function alone, and is independent of the database. The sensitivity captures how great a difference (between the value of f on two databases differing in a single element) must be hidden by the additive noise generated by the curator.

4. LOCATION PATTERN MINING

Differential privacy has been applied to a wide range disparate data sets and applications. One way that it has been applied is for privacy preservation during location pattern mining[Ho and Ruan 2011]. With the proliferation of location acquisition technology on mobile devises, there is a trend in collecting human movement data for behavior analysis and pattern mining. To obtain meaningful patterns, a large and diverse amount of individual location history records has to be collected for analysis or mining. One of the main concerns for individuals who participate in such a data collection is the disclosure of their location when a user queries for the analysis or mining results. This dataset is an ideal candidate for differential privacy, because following the definition previously described, there should not be anything that can be learned about the user that is different had they not participated in the location mapping system.
If user can be guaranteed that their privacy is protected, then they will be more likely to participate in such a system. The authors state that this increase in privacy will lead to increased participation which in turn will lead to pattern mining capabilities that will have more "social utility".

4.1 Privacy Policies

There are two specific privacy policies that the authors of this paper would like to control: (i) the precision of the discovered locations and (ii) the output counts for queries on these locations. Differential privacy ensures that the "ability of an adversary should be essentially the same, independent of whether any individual opts in to, or opts out of, the dataset". If a specific user decides to opt out from a location history database, and an adversary can detect that, an interesting change can occur in the data, from Region A to Region A'.
The authors of the paper claim that conventional privacy mechanism that adds Laplace noise to (count) data may not work (well) in practice even though it provides theoretical privacy guarantees. One has to be cautious when this privacy preserving mechanism is applied to new algorithms or a new problem settings.

4.1.1 *Location Pattern Mining Differential Privacy.* The authors of the paper present an algorithm for differential privacy for location pattern mining data sets. They start out by identifying two issues that need to be overcome to make differential privacy guarantee practical for the outputs of a location pattern mining algorithm. First, the ability to control magnitude of sensitivity is critical for Laplace noise perturbation. The second is to distribute the desired level for differential privacy to different steps to to achieve practical differential privacy. In order to address these issue, the authors split the differential privacy level ϵ between two steps: (i) spatial decomposition or the preprocessing step and the (ii) pattern mining algorithm outputs.
The problem with the Laplace noise perturbation privacy mechanism is the magnitude of the (global) sensitivity, Δf. The criterion for an interesting location point is the number (or count) of stay points in a specific region. If one was to define Δf from the whole spatial domain, Δf can be a very large number. The paper presents an example. Suppose John's house is an interesting location which is frequently visited by John who stays there. This "global" sensitivity has little relevance to the other interesting locations. The solution the authors propose is to use spatial decomposition to break the global problem into smaller local problems where smaller local sensitivity is used in the privacy mechanism. This strategy provides better localized output accuracy at a fixed differential privacy level. This decomposition is accomplished with a region quad tree. The local sensitivity is then computed locally for each region in the quad tree.
The Density-Based Spatial Clustering of Applications with Noise (DBSCAN) clustering algorithm is applied to the outputs of the quad tree. The authors do not go into detail on the DBSCAN as it is an established technique.

5. HEALTH DATA DIFFERENTIAL PRIVACY

Health Data is another area where there is a push for differential privacy[Dankar and El Emam 2012]. The authors of the paper out-

line a number of characteristics that need to be considered in an practical mechanism used to preserve privacy. While some of their concerns are technological, and some are more social, these concerns must all be considered before a practical health data differential privacy system is developed and used to actually protect the general public's private data.

5.1 Data types

Health data contains categorial data such as diagnosis codes, procedure codes, drugs dispensed, laboratory test ordered, and geographical data about the patient and provider. There is also numeric data such as age, length of stay in hospital, and time since last visit. Both types of variables need to be addressed by any practical solution. It has been shown that Laplace noise added to numeric data can distort the values significantly.

5.2 User Acceptance

While it is not a technical concern, the paper points out some significant concerns about user acceptance of new differential privacy approach. Health data is often disclosed by professionals that have a set method and established code. Since it will be challenging to convince users to abandon their established code, at least in the short term, a non-interactive mechanism would be most suited for this community.

5.3 A Priori Knowledge

This non-interactive mechanism would only work for releasing statistics for computations that are known a priori. While this would work for surveillance and reporting services, there would be utility for differentially private statistics. Beyond such applications, many other data uses would require actual data publishing to meet the needs of the analyst community, according to the authors of this paper.

5.4 Legal Ramifications

Another important consideration is the law. There are many specific privacy laws in many jurisdictions. Current health privacy statues in the US, Canada, and Europe do not specify the acceptable risk and often use the "reasonableness" standard. In practice, one relies on precedent to justify the risk thresholds that are used. The current privacy models have been used for more then two decades. During that time, acceptable levels of risk have been defined by guidelines, policies, court cases and regulatory orders. In differential privacy, important parameters such as ϵ have no intrinsic meaning, and almost no precedent of actual health data releases to justify the choice of any value.

5.5 Convincing the Public

Since the privacy of health data is a public concern, any new mechanism to protect that data will need to be explained in a manner that can convince the general public that their privacy is in fact being protected. In the context of differential privacy, it is quite challenging to explain to a patient the meaning of ϵ, and how it is used to disclose or provide access to their data, for example.

6. NO FREE LUNCH WITH DIFFERENTIAL PRIVACY

While differential privacy has many great applications, it is not a silver bullet that solves all problems. The rise of differential privacy has led to a critical re-examination of it, such as the one by Kifer and Machanavajjhala.[Kifer and Machanavajjhala 2011] The authors use a no-free-lunch theorem which defines non-privacy as a game. The authors argue that it is not possible to provide privacy and utility without making assumptions about how the data is generated. They argue that the privacy of an individual is preserved when it is possible to limit the inference of an attacker about the participation of the individual in the data generating process.

6.1 Social Networks

One application of differential privacy that they authors of this paper look at is the application to social networks. Preserving privacy in a social network is a challenge because of the effect that nodes have on each other. While the differential privacy algorithm gives us assurances about the effect of the removal of a certain node from the database, there is the matter of influence over other nodes. For example, if a tuple contains information about a user Bob, deleting Bob's tuple would not remove the *influence* of this tuple on the data. This *influence* causes there to be *evidence* of of Bob's tuple even after Bob's tuple has been removed from the data.

6.2 Tabular Data with Previously Released Statistics

Another specific case sited by the authors of this paper is how to preserve privacy when answering queries over a table for which deterministic statistics have already been released. The example that the paper presents is the U.S. Census Bureau. As one of the largest consumers of privacy technology, the Census Bureau both collects and disseminates data about the U.S. population. Many of the datasets that they release have been perturbed, masked, or otherwise modified to protect the confidentiality of individuals in the data.

While privacy is generally a driving concern, in some cases, utility is more important then privacy. One example of that is the release of population counts which are used for allocating seats in the House of Representatives. Since this is a highly-charged political issue, these counts must be as accurate as possible. There are other queries that may be returned with differential privacy protections after the initial deterministic statistical release. However, the fact that an adversary has hard data about that table, can cause the privacy guarantees to degrade.

7. CONCLUSION

Differential Privacy is a relatively new privacy ensuring mechanism, but as the number and volume of databases with private data continues to grow, this will continue to be a powerful and important tool. In this paper we have shown the history of differential privacy starting with Dalenius' initial mathematical methodology, through Dwork's landmark paper defining differential privacy, and then examined some applications of differential privacy, as well as looking at a critical examination of some of its drawbacks.

The conclusion of this author is that differential privacy is an effective and powerful tool, but like all tools it must be used properly and one must understand what they are doing in order to fully take advantage of the privacy guarantees that differential privacy affords.

REFERENCES

DALENIUS, T. 1977. Towards a methodology for statistical disclosure control. *Statistik Tidskrift 15,* 429-444, 2–1.

DANKAR, F. K. AND EL EMAM, K. 2012. The application of differential privacy to health data. In *Proceedings of the 2012 Joint EDBT/ICDT Workshops*. EDBT-ICDT '12. ACM, New York, NY, USA, 158–166.

DWORK, C. 2006. Differential privacy. *Automata, languages and programming*, 1–12.

DWORK, C. 2008. Differential privacy: A survey of results. *Theory and Applications of Models of Computation*, 1–19.

HO, S.-S. AND RUAN, S. 2011. Differential privacy for location pattern mining. In *Proceedings of the 4th ACM SIGSPATIAL International Workshop on Security and Privacy in GIS and LBS*. SPRINGL '11. ACM, New York, NY, USA, 17–24.

KIFER, D. AND MACHANAVAJJHALA, A. 2011. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*. SIGMOD '11. ACM, New York, NY, USA, 193–204.