

Differential Privacy

Cynthia Dwork

Microsoft Research
dwork@microsoft.com

Abstract. In 1977 Dalenius articulated a desideratum for statistical databases: nothing about an individual should be learnable from the database that cannot be learned without access to the database. We give a general impossibility result showing that a formalization of Dalenius' goal along the lines of semantic security cannot be achieved. Contrary to intuition, a variant of the result threatens the privacy even of someone not in the database. This state of affairs suggests a new measure, *differential privacy*, which, intuitively, captures the increased risk to one's privacy incurred by participating in a database. The techniques developed in a sequence of papers [8, 13, 3], culminating in those described in [12], can achieve any desired level of privacy under this measure. In many cases, extremely accurate information about the database can be provided while simultaneously ensuring very high levels of privacy.

1 Introduction

A statistic is a quantity computed from a 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. The work discussed herein was originally motivated by exactly this problem: how to reveal useful information about the underlying population, as represented by the database, while preserving the privacy of individuals. Fortunately, the techniques developed in [8, 13, 3] and particularly in [12] are so powerful as to broaden the scope of private data analysis beyond this original “representative” motivation, permitting privacy-preserving analysis of an object that is itself of intrinsic interest. For instance, the database may describe a concrete interconnection network – not a sample subnetwork – and we wish to reveal certain properties of the network without releasing information about individual edges or nodes. We therefore treat the more general problem of *privacy-preserving analysis of data*.

A rigorous treatment of privacy requires definitions: What constitutes a failure to preserve privacy? What is the power of the adversary whose goal it is to compromise privacy? What auxiliary information is available to the adversary (newspapers, medical studies, labor statistics) even without access to the database in question? Of course, utility also requires formal treatment, as releasing no information or only random noise clearly does not compromise privacy; we

will return to this point later. However, in this work privacy is paramount: we will first define our privacy goals and then explore what utility can be achieved given that the privacy goals will be satisfied¹.

A 1977 paper of Dalenius [6] articulated a desideratum that foreshadows for databases the notion of semantic security defined five years later by Goldwasser and Micali for cryptosystems [15]: access to a statistical database should not enable one to learn anything about an individual that could not be learned without access². We show this type of privacy cannot be achieved. The obstacle is in *auxiliary information*, that is, information available to the adversary other than from access to the statistical database, and the intuition behind the proof of impossibility is captured by the following example. Suppose one's exact height were considered a highly sensitive piece of information, and that revealing the exact height of an individual were a privacy breach. 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 than the average Lithuanian woman" learns Terry Gross' height, while anyone learning only the auxiliary information, without access to the average heights, learns relatively little.

There are two remarkable aspects to the impossibility result: (1) it applies regardless of whether or not Terry Gross is in the database and (2) Dalenius' goal, formalized as a relaxed version of semantic security, cannot be achieved, while semantic security for cryptosystems can be achieved. The first of these leads naturally to a new approach to formulating privacy goals: the risk to one's privacy, or in general, any type of risk, such as the risk of being denied automobile insurance, should not substantially increase as a result of participating in a statistical database. This is captured by *differential privacy*.

The discrepancy between the possibility of achieving (something like) semantic security in our setting and in the cryptographic one arises from the utility requirement. Our adversary is analogous to the eavesdropper, while our user is analogous to the message recipient, and yet there is no decryption key to set them apart, they are one and the same. Very roughly, the database is designed to convey certain information. An auxiliary information generator knowing the data therefore knows much about what the user will learn from the database. This can be used to establish a shared secret with the adversary/user that is unavailable to anyone not having access to the database. In contrast, consider a cryptosystem and a pair of candidate messages, say, $\{0, 1\}$. Knowing which message is to be encrypted gives one no information about the ciphertext; intuitively, the auxiliary information generator has "no idea" what ciphertext the eavesdropper will see. This is because by definition the ciphertext must have no utility to the eavesdropper.

¹ In this respect the work on privacy diverges from the literature on secure function evaluation, where privacy is ensured only modulo the function to be computed: if the function is inherently disclosive then privacy is abandoned.

² Semantic security against an eavesdropper says that nothing can be learned about a plaintext from the ciphertext that could not be learned without seeing the ciphertext.