

Secure Multiparty Computation for Privacy Preserving Data Mining

Yehida Lindell

Bar-Ilan University, Israel

INTRODUCTION

The increasing use of data-mining tools in both the public and private sectors raises concerns regarding the potentially sensitive nature of much of the data being mined. The utility to be gained from widespread data mining seems to come into direct conflict with an individual's need and right to privacy. Privacy-preserving data-mining solutions achieve the somewhat paradoxical property of enabling a data-mining algorithm to use data *without* ever actually seeing it. Thus, the benefits of data mining can be enjoyed without compromising the privacy of concerned individuals.

BACKGROUND

A classical example of a privacy-preserving data-mining problem is from the field of medical research. Consider the case that a number of different hospitals wish to jointly mine their patient data for the purpose of medical research. Furthermore, assume that privacy policy and law prevent these hospitals from ever pooling their data or revealing it to each other due to the confidentiality of patient records. In such a case, classical data-mining solutions cannot be used. Fortunately, privacy-preserving data-mining solutions enable the hospitals to compute the desired data-mining algorithm on the union of their databases without ever pooling or revealing their data. Indeed, the only information (provably) learned by the different hospitals is the output of the data-mining algorithm. This problem, whereby different organizations cannot directly share or pool their databases but must nevertheless carry out joint research via data mining, is quite common. For example, consider the interaction between different intelligence agencies in the USA. These agencies are suspicious of each other and do not freely share their data. Nevertheless, due to recent security needs, these agencies must run data-mining algorithms on their combined data. Another example relates to data that is held by governments. Until recently, the Canadian government held a vast federal database that pooled citizen data from a number of different government ministries (some called this database the “big brother” database). The Canadian government claimed that the data-

base was essential for research. However, due to privacy concerns and public outcry, the database was dismantled, thereby preventing that “essential research” from being carried out. This is another example of where privacy-preserving data mining could be used to find a balance between real privacy concerns and the need of governments to carry out important research.

Privacy-preserving data mining is actually a special case of a long-studied problem in cryptography: secure multiparty computation. This problem deals with a setting where parties with private inputs wish to jointly compute some function of their inputs. Loosely speaking, this joint computation should have the property that the parties learn the correct output and nothing else, even if some of the parties maliciously collude to obtain more information.

MAIN THRUST

In this article, I provide a succinct overview of secure multiparty computation and how it can be applied to the problem of privacy-preserving data mining. The main focus is on how security is formally defined, why this definitional approach is adopted, and what issues should be considered when defining security for privacy-preserving data-mining problems. Due to space constraints, the treatment in this chapter is both brief and informal. For more details, see Goldreich (2003) for a survey on cryptography and cryptographic protocols.

Security Definitions for Secure Computation

The aim of a secure multiparty computation task is for the participating parties to *securely* compute some function of their distributed and private inputs. But what does it mean for a computation to be *secure*? One way of approaching this question is to provide a list of *security properties* that should be preserved. The first such property that often comes to mind is that of *privacy* or confidentiality. A naïve attempt at formalizing privacy would be to require that each party learns nothing about the other parties' inputs, even if it behaves maliciously. However, such a definition is usually unat-

tainable, because the defined output of the computation itself typically reveals some information about the other parties' inputs. (For example, a decision tree computed on two distributed databases reveals some information about both databases.) Therefore, the privacy requirement is usually formalized by saying that the only information learned by the parties in the computation (again, even by those who behave maliciously) is that specified by the function output. Although privacy is a primary security property, it rarely suffices. Another important property is that of *correctness*; this states that the honest parties' outputs are correctly distributed even in the face of adversarial attack. A central question that arises in this process of defining security properties is "When is the list of properties complete?" This question is, of course, application dependent, which essentially means that for every new problem, the process of deciding which security properties are required must be reevaluated. I must stress that coming up with the right list of properties is often very difficult and it can take many years until one is convinced that a definition truly captures the necessary security requirements. Furthermore, an incomplete list of properties may easily lead to real security failures.

The Ideal/Real Model Paradigm

Due to these difficulties, the standard definitions of secure computation today follow an alternative approach called the *ideal/real model paradigm*. This has been the dominant paradigm in the investigation of secure computation in the last 15 years; see Canetti (2000) for the formal definition and references therein for related definitional work. Loosely speaking, this paradigm defines the security of a real protocol by comparing it to an *ideal computing scenario*, in which the parties interact with an external trusted and incorruptible party. In this ideal execution, the parties all send their inputs to the trusted party (via ideally secure communication lines). The trusted party then computes the function on these inputs and sends each party its specified output. Such a computation embodies the goal of secure computation, and it is easy to see that the properties of privacy and correctness hold in the ideal model. In addition to the fact that these and other security properties are preserved in an ideal execution, the simplicity

of the ideal model provides an intuitively convincing security guarantee. For example, notice that the only message a party sends in an ideal execution is its input. Therefore, the only power that a corrupted party has is to choose its input as it wishes (behavior that is typically legitimate anyway).

So far, I have defined an ideal execution in an ideal world. However, in the real world, the parties run a protocol without any trusted help. Despite this fact, a secure real protocol should somehow emulate an ideal execution. That is, a real protocol that is run by the parties (in a world where no trusted party exists) is *secure* if no adversary can do more harm in a real execution than in an execution that takes place in the ideal world. Stated differently, for any adversary carrying out a successful attack on a real protocol, there exists an adversary that successfully carries out the same attack in the ideal world. This suffices because, as I have shown, no successful attacks can be carried out in an ideal execution. Thus, no successful attacks can be carried out on the real protocol, implying that it is secure. See Figure 1 for a diagram of the real and ideal models.


Note that security is required to hold for *every* adversary carrying out *any feasible attack* (within the parameters defined for the adversary, as discussed next).

Defining the Model

The preceding informal description of the ideal/real model paradigm expresses the intuition that a real execution should behave just like an ideal execution. In order to obtain a complete and formal definition, it is crucial that both the ideal and real models are fully defined. Among other things, this involves defining the real network model and the adversary's power, including any assumptions on its behavior. A secure protocol only provides real-world security guarantees if the mathematical definition of the real computation and adversarial models accurately reflects the real network and adversarial threats that exist.

I now briefly discuss a number of parameters that are considered when defining the network model and the adversary; this list is far from comprehensive. Two central considerations that arise when defining the net-

Figure 1.



3 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/chapter/secure-multiparty-computation-privacy-preserving/10743

Related Content

Internet Data Mining Using Statistical Techniques

Kuldeep Kumar (2008). *Data Warehousing and Mining: Concepts, Methodologies, Tools, and Applications* (pp. 1446-1453).

www.irma-international.org/chapter/internet-data-mining-using-statistical/7708

View Selection and Materialization

Zohra Bellahsene (2010). *Data Warehousing Design and Advanced Engineering Applications: Methods for Complex Construction* (pp. 114-130).

www.irma-international.org/chapter/view-selection-materialization/36611

Data Mining and Homeland Security

Jeffrey W. Seifert (2008). *Data Warehousing and Mining: Concepts, Methodologies, Tools, and Applications* (pp. 3630-3638).

www.irma-international.org/chapter/data-mining-homeland-security/7853

A Haplotype Analysis System for Genes Discovery of Common Diseases

Takashi Kido (2008). *Data Warehousing and Mining: Concepts, Methodologies, Tools, and Applications* (pp. 1674-1687).

www.irma-international.org/chapter/haplotype-analysis-system-genes-discovery/7723

Management of Data Streams for Large-Scale Data Mining

Jon R. Wright, Gregg T. Vesonder and Tamraparni Dasu (2008). *Data Warehousing and Mining: Concepts, Methodologies, Tools, and Applications* (pp. 2644-2658).

www.irma-international.org/chapter/management-data-streams-large-scale/7789