Reference Papers

Single-dimension Noise Injection

It's probably the oldest data perturbation technique. Rekash Agrawal's paper (2000) started the research area of privacy-preserving data mining. The basic ideas are injecting additive random noises to selected dimensions of a tabular dataset to hide the original values. However, these methods are subject to some known attacks, i.e., reconstruction attacks. It's mostly applied in the web model for users submitting perturbed data.

- "Privacy Preserving Data Mining", R. Agrawal and R. Srikant, Proc. ACM SIGMOD Conf. Management of Data, pp. 439-450, May 2000. [link]
- "On the Design and Quantification of Privacy Preserving Data Mining Algorithms,"
  D. Agrawal and C. C. Aggarwal, Proc. ACM SIGMOD, pp. 247-255, 2001. [link]
- "Deriving Private Information from Randomized Data," Z. Huang and W. Du and B. Chen, in Proceedings of the 2005 ACM SIGMOD Conference, pp. 37-48, Baltimroe, MD, 2005. [link]

Multiplicative Data Perturbation

Multiplicative data perturbation applies to multidimensional datasets (i.e., at least two dimensions) with a combination of linear transformation and random noise injection. The specific methods include rotation perturbation, random projection perturbation, geometric data perturbation (GDP), and random space perturbation (RASP). Some (e.g., RASP) are more resilient to *background-knowledge based attacks* than others. It preserves a certain multidimensional property of the perturbed dataset and thus many data mining models depending on that property can be applied directly to the perturbed dataset without much modification. However, without the perturbation keys, the mined models

cannot be used. Thus, multiplicative data perturbation is mostly used for outsourced computation (e.g., users export data and computation to curious cloud/service providers who are interested in users' private data). It cannot be applied in the web model for information is not shared with the curious party.

- Keke Chen and Ling Liu, " Geometric Data Perturbation for Privacy Preserving Outsourced Data Mining ", Journal of Knowledge and Information Systems (KAIS), 2011
- Huiqi Xu, Shumin Guo, and Keke Chen: "Building Confidential and Efficient Query Services in the Cloud with RASP Data Perturbation ", IEEE TKDE, 2014.
- Kun Liu, Chris Giannella and Hillol Kargupta, "An Attacker's View of Distance Preserving Maps for Privacy Preserving Data Mining," in Proceedings of the 10th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD'06), Berlin, Germany, September, 2006. [link]
- Deriving Private Information from Arbitrarily Projected Data, Songtao Guo and Xintao Wu, PKDD 2007

Randomized Response

Randomized response is mostly applied to perturb categorical values. It works in the web model. With a certain probability, the submitter flips the answer and then submits it. Some types of statistics of the submitted answers can be estimated with a certain level of accuracy. Later studies (the RAPPOR paper) show that it's equivalent to local differential privacy.

- Wenliang Du and Zhijun Zhan, "Using Randomized Response Techniques for Privacy-Preserving Data Mining," SIGKDD 2003. [link]

- RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response, Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova, ACM CCS 2014.

Data Anonymization

Data anonymization aims to protect persons' identities from published sensitive data such as patient records. It appears that removing the explicit identities such as names and SSNs are insufficient to protect identities. People can still be identified with *virtual identifiers*. The data anonymization research studies the possible attacks and the corresponding methods to protect persons from re-identification.

- L. Sweeney, "Achieving k-anonymity privacy protection using generalization and suppression," International Journal on Uncertainty, Fuzziness and Knowledge-based Systems, 10 (5), pp. 571-588, 2002. [link]
- Ashwin Machanavajjhala, Johannes Gehrke, Daniel Kifer, Muthuramakrishnan Venkitasubramaniam, "l -Diversity: Privacy Beyond k-Anonymity," p. 24, 22nd International Conference on Data Engineering (ICDE'06), 2006. [link]
- Ninghui Li, Tiancheng Li and Suresh Venkatasubramanian, t-Closeness: Privacy Beyond k-Anonymity and l-Diversity, ICDE 2007

Differential Privacy

Differential privacy was originally used to address the problem of query inference attack on databases. Database owners want to share information by exposing the query interface. Only certain types of queries, e.g., typically aggregate queries, are allowed to return statistics. Although statistical queries do not return individual record, query inference attack is possible by analyzing a series of carefully crafted queries to breach an individual or a small group of persons whose records are in the database. Differential privacy perturbs the query

results by adding noises designed with a certain mechanism such as the Laplace mechanism. It can be used to quantitatively define privacy and theoretically prove that privacy is protected, and thus is considered the golden standard in privacy-preserving computation. Differential privacy has been extended into different scenarios, such as non-interactive data release, local differential privacy, graph data release, differentially private data mining/deep learning, etc. It's a hot research area. We list some old but important works here.

- Calibrating Noise to Sensitivity in Private Data Analysis Cynthia Dwork, Frank McSherry, Kobbi Nissim, Adam Smith, in Third Theory of Cryptography Conference (TCC 2006), Springer, March 1, 2006,
- Privacy Integrated Queries, Frank McSherry, ACM SIGMOD 2009.
- Mechanism Design via Differential Privacy, Frank McSherry, Kunal Talwar, IEEE FOCS, 2007
- Data Mining with Differential Privacy, Arik Friedman and Assaf Schuster, ACM KDD 2010

Private information retrieval and oblivious RAM

PIR and ORAM have the same purpose: protecting access patterns from curious service providers, in the setting that users run programs or access databases that are hosted by a curious service provider. It is used in public database (e.g., patent database) providers or outsourced computation. The basic idea is to hide the real accessed items among fake random accesses.

- Private Information Retrieval, Benny Chor, Oded Goldreich, Eval Kushilevitz and Madhu Sudan, Journal of ACM, 1999
- A Simple ORAM, Kai-min Chung and Rafael Pass, https://eprint.iacr.org/2013/243

- Path ORAM: an Extremely simple ORAM, Emil Stefanov, et al., https://eprint.iacr.org/2013/280.pdf

Garbled Circuits and Randomized Secret Sharing

Garbled circuits are mostly for secure two-party computation, although expensive extension might also be used for multiple parties. In the two-party setting, each party holds a part of the input of a function. The purpose is to evaluate the function and one of the parties learns only the function output, nothing else. It can be used in collaborative computation and outsourced computation. SMC handles >=3 parties collaboratively computing a function. The most important technique is secret sharing.

- Protocols for Secure Computations, A. Yao, Annual Symposium on Foundations of Computer Science, 1982
- Faster Secure Two-Party Computation Using Garbled Circuits, Yan Huang, David Evans, Jonathan Katz, and Lior Malka, USENIX Security Symposium, 2011
  - SecureML: A System for Scalable Privacy-Preserving Machine Learning, Payman Mohassel and Yupeng Zhang, IEEE Symposium on Security and Privacy, 2017

Homomorphic encryption

Partially homomorphic encryption allows to compute the encrypted sum or product of two encrypted values without decrypting them. Fully homomorphic encryption implements both operations. Theoretically, all the functions can be implemented based on the two homomorphic operations: addition and multiplication. Homomorphic encryption can

be used in outsourced computation. It protects both data and models mined from the data from curious service providers.

- A Survey on Homomorphic Encryption Schemes: Theory and Implementation, ABBAS ACAR, HIDAYET AKSU, and A. SELCUK ULUAGAC, MAURO CONTI, https://arxiv.org/pdf/1704.03578.pdf

- (Leveled) Fully Homomorphic Encryption Without Bootstrapping, Z. Brakerski, C. Gentry, and V. Vaikuntanathan, Innovations in Theoretical Computer Science Conference (ITSC), 2012,

Location privacy

Location-based services (LBS) become popular with the wide use of smart phones. The convenience of LBS comes with the privacy concern. Learning one's locations can seriously breach privacy. The purpose of preserving location privacy is to preserve location privacy from LBS providers, while preserving the utility of LBS as much as possible.

- Location Privacy in Pervasive Computing. A. Beresford and F. Stajano. IEEE Pervasive Computing, 2003
- SpaceTwist: Managing the Trade-Offs Among Location Privacy, Query Performance, and Query Accuracy in Mobile Services. Man Lung Yiu, Christian S. Jensen, Xuegang Huang, Hua Lu, ICDE08
- The New Casper: Query Processing for Location Services without Compromising Privacy. Mohamed Mokbel, Chi-Yin Chow, Walid Aref. VLDB 2006.

Privacy in online social networks

Privacy in online social networks involves many aspects. Among them one is the issue with privacy settings. Users have to set up tens of privacy settings without the knowledge of privacy implication of each item. The other is sharing social graph data with curious parties while preserving individual user's privacy. It studies the attacks utilizing nodes, edges, or subgraph structures and the methods countering the attacks.

- Wherefore Art Thou R3579X? Anonymized Social Networks, Hidden Patterns, and Structural Steganography, Lars Backstrom, Cynthia Dwork, Jon Kleinberg, WWW2007
- Resisting Structural Identification in Anonymized Social Networks Michael Hay, Gerome Miklau, David Jensen, Don Towsley, and Philipp Weis, Conference on Very Large Databases (VLDB) 2008
- Shumin Guo, and Keke Chen. "Mining Privacy Settings to Find Optimal Privacy-Utility Tradeoffs for Social Network Services", ASE/IEEE International Conference on Privacy, Security, Risk and Trust (PASSAT), 2012.

Privacy-preserving data mining

Privacy-preserving data mining has two settings. One is to share the private data with curious parties for mining models while preserving individual's privacy. The other is outsourcing data to a curious cloud service provider for mining data without breaching data privacy. Often, solutions will use a novel hybrid of homomorphic encryption, secure multiparty computation, and data perturbation.

- Y. Lindell and B. Pinkas, "Privacy Preserving Data Mining," Advances in Cryptology (CRYPTO'00), pp. 36--53, 2000. (use SMC)

- J. Vaidya and C. Clifton, "Privacy-Preserving K-Means Clustering over Vertically Partitioned Data," in Proceedings of the ninth ACM

SIGKDD international conference on Knowledge discovery and data mining, 2003. [link] (use random secret sharing + partial HE)

- K. Chen and L. Liu. Privacy preserving data classification with rotation perturbation. In Proceedings of the Fifth IEEE International Conference on Data Mining (ICDM05), Houston, TX, November 2005.
- Sagar Sharma and **Keke Chen** "Confidential Boosting with Random Linear Classifiers for Outsourced User-generated Data", European Symposium on Research in Computer Security (ESORICS), 2019 [pdf] (use PHE + garbled circuits)

Outsourced database services

Outsourcing database services is popular due to the wide use of cloud and services computing. Users submit queries to the cloud or other service provider, which confidentially process. The purpose is to protect data and model privacy from curious service providers, while finishing query processing.

- Curtmola, Reza and Garay, Juan and Kamara, Seny and Ostrovsky, Rafail, Searchable symmetric encryption: improved definitions and efficient constructions, ACM Computer and Communications Security, 2006
- Secure kNN computation on encrypted databases, Wong, Wai Kit and Cheung, David Wai-lok and Kao, Ben and Mamoulis, Nikos, SIGMOD 2009
- Huiqi Xu, Shumin Guo, and Keke Chen: "Building Confidential and Efficient Query Services in the Cloud with RASP Data Perturbation ", IEEE TKDE, 2014.
- Raluca Ada Popa, Catherine M. S. Redfield, Nickolai Zeldovich, and Hari Balakrishnan.

CryptDB: Protecting Confidentiality with Encrypted Query Processing.
In *Proceedings of the 23rd ACM Symposium on Operating Systems Principles (SOSP)*, Cascais, Portugal, October 2011

Model Inversion Attack

This line of research studies the attacks that can use published statistical/machine learning models to recover private information

- Matt Fredrikson, Somesh Jha, and Thomas Ristenpart, *Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures*, ACM Conference on Computer and Communications Security, 2015
- Matthew Fredrikson, Eric Lantz, Somesh Jha, Simon Lin, David Page, and Thomas Ristenpart. Privacy in Pharmacogenetics: An End-to-End Case Study of Personalized Warfarin Dosing, 23rd USENIX Security Symposium USENIX Security,

Membership Inference Attack

This line of research studies the attacks that can use published statistical/machine learning models to infer the likelihood of a targeted record being used in training the model.

- Membership Inference Attacks against Machine Learning Models
  Reza Shokri, Marco Stronati, Congzheng Song, Vitaly Shmatikov

Privacy-Preserving Deep Learning

Applying differential privacy in deep learning or federated learning.

- Deep Learning with Differential Privacy by Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, Li Zhang, ACM CCS 2016
- Hitaj, Briland and Ateniese, Giuseppe and Perez-Cruz, Fernando, Deep Models Under the GAN: Information Leakage from Collaborative Deep Learning, ACM SIGSAC Conference on Computer and Communications Security, 2017

Trusted Execution Environments (TEE)

This line of research studies the hardware-assisted TEE. CPUs supporting TEEs use special instructions to create and access a secure enclave that cannot be breached by compromised OS or hypervisor. It's a promising solution for confidential computing in the cloud or collaborative computing scenarios.

- Intel SGX explained, by Victor Costan and Srinivas Devadas
- SCONE: Secure Linux Containers with Intel SGX, OSDI2016
- Graphene-SGX: A Practical Library OS for Unmodified Applications on SGX, ATC17