Master Thesis Projects

Mina Alishahi
mina.sheikhalishahi@ou.nl, m.sheikhalishahi@tue.nl

1 Poisoning Attacks against LDP-based Federated Learning

Federated learning is a collaborative learning infrastructure in which the data owners do not need to share raw data with one another or rely on a single trusted entity. Instead, the data owners jointly train a Machine Learning model through executing the model locally on their own data and only share the model parameters with the aggregator. While the participants only share the updated parameters, still some private information about underlying data can be revealed from the shared parameters. To address this issue, Local Differential Privacy has been used as effective tool to protect information leakage over shared parameters in Federated Learning, say LDP-FED. However, it has not yet been investigated whether (and to what extent) the LDP-FED is resistant against data and model poisoning attacks. Also, if LDP-FED is not resistant against these attacks, how can we design a robust LDP-FL where its performance is negligibly affected by poisoning attacks.

This project aims to evaluate the resistance of the LDP-FED against poisoning attacks and to explore the possibilities of reducing the success rate of these attacks.
The following papers are suggested to be studied for this work:

1. Stacey Truex, Ling Liu, Ka-Ho Chow, Mehmet Emre Gursoy, Wenqi Wei; *LDP-Fed: Federated Learning with Local Differential Privacy*, CoRR, 2020.
2. Mohammad Naseri, Jamie Hayes, and Emiliano De Cristofaro; *Toward Robustness and Privacy in Federated Learning: Experimenting with Local and Central Differential Privacy*, CoRR, 2020.
3. Lingjuan Lyu, Han Yu, Xingjun Ma, Lichao Sun, Jun Zhao, Qiang Yan, Philip S. Yu, *Privacy and Robustness in Federated Learning: Attacks and Defenses*, arXiv, 2020.
4. Malhar Jere, Tyler Farnan, and Farinaz Koushanfar; *A Taxonomy of Attacks on Federated Learning*, IEEE Security & Privacy, 2021.
5. Xiaoyu Cao, Jinyuan Jia, Neil Zhenqiang Gong, *Data Poisoning Attacks to Local Differential Privacy Protocols*, CoRR, 2019.
6. Minghong Fang, Xiaoyu Cao, Jinyuan Jia, Neil Zhenqiang Gong; *Local Model Poisoning Attacks to Byzantine-Robust Federated Learning*, the 29th Usenix Security Symposium, 2020.

2 Privacy Preserving k-means/k-median Distributed Learning using Local Differential Privacy

The classical clustering algorithms were designed to be implemented in central server. However, in recent years data is generally located in distributed sites in different locations. Due to privacy concerns the data owners are unwilling to share their original data to public or even to each other. Several approaches have been proposed in the literature which protects the data owners' privacy while at the same time a clustering algorithm is shaped over protected data.

Given that the proposed solutions mainly need the presence of a trusted party, Local Differential Privacy (LDP) can be used as an effective solution that protects the data owner's data in her local device.

This study aims to investigate the application of LDP in the distributed learning of two well-known clustering algorithms, namely k-means and k-medians, in terms of utility loss and privacy leakage. Specifically, it explores the resistance of LDP-based k-means/k-medians clustering against poisoning attacks.
The following papers are suggested to be studied for this work:

1. Maria Florina Balcan, Steven Ehrlich, Yingyu Liang; *Distributed k-Means and k-Median Clustering on General Topologies*, NIPS, 2013.

2. Geetha Jagannathan, Rebecca N. Wright; *Privacy-Preserving Distributed k-Means Clustering over Arbitrarily Partitioned Data*, ACM SIGKDD, 2005.

3. Chang Xia, Jingyu Hua, Wei Tong, Sheng Zhong; *Distributed K-Means clustering guaranteeing local differential privacy*, Computer & Security journal, 2020.

4. Pathum Chamikara, Mahawaga Arachchige, Peter Bertok, Ibrahim Khalil, Dongxi Liu, Seyit Camtepe, Mohammed Atiquzzaman; *Local Differential Privacy for Deep Learning*, IEEE Internet of Things Journal, 2020.

5. Malhar Jere, Tyler Farnan, and Farinaz Koushanfar; *A Taxonomy of Attacks on Federated Learning*, IEEE Security & Privacy, 2021.

3 Scalable Blockchain-based framework in Internet of Things (IoT)

Internet of Things (IoT) is an ever-increasing technology in which many of our daily objects are connected via Internet (or other networks) and transfer data for analysis or performing certain tasks. This property makes the IoT vulnerable to security threats needing to be addressed for building trust among clients.

Blockchain, a technology born with cryptocurrency, has shown its effectiveness and robustness against some security threats when integrated with IoT. Against its capability, the main issue of integrating blockchain with IoT is its scalability and efficiency with a large-scale network like IoT.

In this project we aim to explore the existing solutions addressing the scalability of blockchain in IoT and to investigate the possibilities of improving the existing ones by proposing new solutions.
The following papers are suggested to be studied for this work:

1. Hong-Ning Dai, Zibin Zheng, Yan Zhang; *Blockchain for Internet of Things: A Survey*, IEEE Internet of Things Journal, VOL. 6, NO. 5, 2019.

2. Hany F. Atlam, Muhammad Ajmal Azad, Ahmed G. Alzahrani, Gary Wills; *A Review of Blockchain in Internet of Things and AI*, Big Data and Cognitive Computing MDPI, 2020.

3. Tiago M. Fernandez-Carames, Paula Fraga-Lamas; *A Review on the Use of Blockchain for the Internet of Things*, IEEE Access, 2020.

4 Deep Learning for Partial Image Encryption

Face recognition has increasingly gained importance for a variety of applications, *e.g.*, surveillance in public places, access control in organizations, in tagging photos in social network, and border control at airports. The widespread application of face recognition, however, raises privacy risks as the individuals' biometric information can be used to profile and track people against their desire.

A typical solution to this problem is the application of Homomorphic Encryption, where an encrypted image is searched to be check in a list of images for a possible match. However, this solution is heavy in terms of both computation and communication costs as it requires all image's pixels to be encrypted. This is while all the pixels of an image do not contain privacy-sensitive information.

In this project, we plan to investigate the application of Deep Learning in detecting the users' identifiable pixels (instead of all pixels) for partial encryption of an image.

1. Zekeriya Erkin, Martin Franz, Jorge Guajardo, Stefan Katzenbeisser, Inald Lagendijk, Tomas Toft; *Privacy-Preserving Face Recognition*, International Symposium on Privacy Enhancing Technologies Symposium, 2009.

2. Peiyang He, Charlie Griffin, Krzysztof Kacprzyk, Artjom Joosen, Michael Collyer, Aleksandar Shtedritski, Yuki M. Asano; *Privacy-preserving Object Detection*, arXive, 2021.

5 Deep Learning for Detecting Network Traffic Attack

The expansion of new communication technologies and services, along with an increasing number of interconnected network devices, web users, services, and applications, contributes to making computer networks ever larger and more complex as systems. However, network anomalies pose significant challenges to many on-line services, which their performance is highly dependent to network performance. For instance, a faulty airport network caused nine hours delay in all of its fights in 2007. To address the issues related to network anomalies, the security solutions need to analyze, detect, and stop such attacks in real time. Although there is a significant amount of technical and scientific literature on anomaly detection methods for network traffic, still 1) a new generated (simulated) dataset that contains a wide range of network attacks (detectable through network traffic monitoring) is missing; 2) the valuable step of feature selection is often underrepresented and treated inattentively in the literature; and 3) the detection techniques suffer from considerable false error rate.

The aim of this project is to address these issues by analyzing network traffic using Deep Learning.
The following papers are suggested to be studied for this project:

1. A. Kind, M. P. Stoecklin, and X. Dimitropoulos; *Histogram-based traffic anomaly detection*, IEEE Transactions on Network and Service Management, vol. 6, no. 2, 2009.

2. R. Chapaneri and S. Shah; *A comprehensive survey of machine learning based network intrusion detection*, in Smart Intelligent Computing and Applications, S. C. Satapathy, V. Bhateja, and S. Das, Eds. Springer, 2019.

6 Local Differential Privacy (LDP) in Protecting the Privacy of Internet of Things (IoT)

The Internet of things (IoT) include physical objects with sensors, software, some processing technologies, which connect and exchange data with other devices and systems over the Internet or other communication network. One of the main challenges in Internet of Things is the users' privacy. While several approaches in the literature have been proposed to protect the information privacy in IoT, still a thorough analysis of the application of Local Differential Privacy (LDP) in this setting is missing. LDP offers a strong level of privacy in which the individuals perturb their data locally before sending them to the third party (named aggregator). This manes that the LDP eliminates the need of a trusted party in the middle. In this project, we aim to investigate the application of LDP in protecting the privacy of IoT data, while still the results of some statistical analyses over protected data is practically useful.
The following papers are suggested to be studied for this project:

1. Chao Li, Balaji Palanisamy; *Privacy in Internet of Things: From Principles to Technologies*, IEEE Internet Of Things Journal, VOL. 6, NO. 1, 2019.

2. Diego Mendez, Ioannis Papapanagiotou, Baijian Yang; *Internet of Things: Survey on Security and Privacy*, https://arxiv.org/pdf/1707.01879.pdf, 2017.

3. Mengmeng Yang, Lingjuan Lyu, Jun Zhao, Tianqing Zhu, Kwok-Yan Lam; *Local Differential Privacy and Its Applications: A Comprehensive Survey*, https://arxiv.org/pdf/2008.03686.pdf, 2015.

7 Game Theory Meets Privacy-preserving Distributed Learning

Companies, organization, and even individuals find mutual benefits in sharing their own information to make better decisions or to increase their revenues. However, generally for privacy concerns the data holders are unwilling to share their own table of data but are interested in getting information from other parties' data. Thence, it is an essential task to define a platform in which several aspects of data-sharing come under consideration and through a game theoretic approach all parties relax their privacy requirements as much as possible to have a more effective output.

In this project, we plan to define data sharing as a game in which several aspects are considered as: 1) the value of shared data (freshness, size,...), 2) privacy gain (in terms of anonymization, differential privacy, etc.), 3) trust o reputation, and 4) the utility of result. The output of the game is setting the Nash Equilibrium in a way that the best balance in terms of utility and privacy is obtained.
The following papers are suggested to be studied for this project:

1. Ningning Ding, Zhixuan Fang, Jianwei Huang; *Incentive Mechanism Design for Federated Learning with Multi-Dimensional Private Information*, 18th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT), 2020.

2. Yufeng Zhan, Jie Zhang, Zicong Hong, Leijie Wu, Peng Li, Song Guo; *A Survey of Incentive Mechanism Design for Federated Learning*, IEEE Transactions on Emerging Topics in Computing, 2021.

3. Ningning Ding; Zhixuan Fang; Lingjie Duan; Jianwei Huang; *Incentive Mechanism Design for Distributed Coded Machine Learning*, IEEE Conference on Computer Communications (InfoComm), 2021.

8 The Output Privacy of Collaborative Classifiers' Learning

Privacy preserving data mining has focused on obtaining valid result when the input data is private. For example, secure multi-party computation techniques are utilized to construct a data-mining algorithm on whole distributed data, without revealing the original data. However, these approaches still might leave potential privacy breaches, *e.g.*, by looking at the structure of a decision tree constructed on the protected shared data.

The aim of this project is to investigate how the output of a classifier constructed collaboratively over private data violates the input data privacy. We then plan to propose solutions to reduce the privacy leakage in this setting. The following papers are suggested to be studied for this project:

1. Qi Jia, Linke Guo, Zhanpeng Jin, Yuguang Fang; *Preserving Model Privacy for Machine Learning in Distributed Systems*, IEEE Transactions on Parallel and Distributed Systems, 2018.

2. Reza Shokri, Marco Stronati, Congzheng Song, Vitaly Shmatikov; *Membership Inference Attacks Against Machine Learning Models*, IEEE Symposium on Security and Privacy, 2017.

3. Ting Wang, Ling Liu, Output Privacy in Data Mining, ACM Transactions on Database Systems, 2011.

4. Radhika Kotecha, Sanjay Garg; Preserving output-privacy in data stream classification, Progress in Artificial Intelligence, June 2017, Volume 6, Issue 2, pp 87–10.

9 On the Trade-off between Utility Loss and Privacy Gain in LDP-based Distributed Learning

Local Differential Privacy (LDP) is a notion of privacy that provides a very strong privacy guarantee by protecting confidential information on users' sides. In this setting generally the users employ a randomization mechanism to perturb their data on their devices following the rules of a mechanism properly. The collected data when aggregated preserves some statistical properties, *e.g.*, mean value can be computed out of perturbed data. This interesting property of LDP has lead to its wide application in many real-world scenarios. In particular, it has been used as an effective tool in privacy preserving distributed machine learning. However, a thorough analysis on finding the trade-off of between the utility loss and privacy gain on LDP-based distributed learning is missing.

In this project we plan to investigate the utility-privacy trade-offs in learning some well-known classifiers when they are trained on distributed data respecting LDP.
The following papers are suggested to be studied for this project:

1. Emre Yilmaz, Mohammad Al-Rubaie, Morris Chang; *Locally Differentially Private Naive Bayes Classification*, https://arxiv.org/pdf/1905.01039.pdf, 2019.

2. Mengmeng Yang, Lingjuan Lyu, Jun Zhao, Tianqing Zhu, Kwok-Yan Lam; *Local Differential Privacy and Its Applications: A Comprehensive Survey*, https://arxiv.org/pdf/2008.03686.pdf, 2015.

3. Mario S. Alvim, Miguel E. Andres, Konstantinos Chatzikokolakis, Pierpaolo Degano, Catuscia Palamidessi; *Differential Privacy: on the trade-off between Utility and Information Leakage*, 2011.

10 Deep Learning for Private Text Generation

The recent development of Deep Learning has led to its success in tasks related to text processing. In particular, Recurrent Neural Network (specifically LSTM) has served as effective tool in next-word prediction. However, the application of Deep Learning in 1) generating a text which respects some privacy constrains and 2) predicting the next word in a sentence in such a way protects confidential information is currently missing.

In this project, we plan to employ Deep Learning as a useful tool to detect the words and sentences that cause privacy violation through uniquely identifying a person (or other confidential information linked to a person) in a text and replace them with meaningful substitute words. Also, we plan to design LSTM that suggests the next-words by considering the text privacy protection.

1. Shervin Minaee, Nal Kalchbrenner, Erik Cambria, Narjes Nikzad, Meysam Chenaghlu, Jianfeng Gao; *Deep Learning–based Text Classification: A Comprehensive Review*, ACM Computing Surveys, 2021.

2. Ankush Chatterjee, Umang Gupta, Manoj Kumar Chinnakotla, Radhakrishnan Srikanth, Michel Galley, Puneet Agrawal; *Understanding Emotions in Text Using Deep Learning and Big Data*; Computers in Human Behavior, 2019.

3. Hong Liang, Xiao Sun, Yunlei Sun & Yuan Gao; *Text feature extraction based on deep learning: a review*, EURASIP Journal on Wireless Communications and Networking volume, 2017.

4. Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Francoise Beaufays, Sean Augenstein, Hubert Eichner; *Federated Learning for Mobile Keyboard Prediction*, 2019.