CIS-TC Newsletter

Communications and Information Security Technical Committee

http://cis.committees.comsoc.org/

Issue 3 – December 2018

Contents

Message from The Chair

Welcome to the December issue of CIS-TC Newsletter!

This is the third issue of our newsletter. It is established as a mean to share important information with the members of CIS Technical Committee and to have some brief contact approximately every semester. We hope that it is a useful communication tool.

Contributions are open for everybody. Please feel free to send us any piece of news that you believe can be of interest to our technical committee.

Christmas is approaching and hence, holidays are coming. While you enjoy your time with family and friends, I hope you will be able to spend some time in reading our December Issue. It is an entirely fresh perspective and a new look and really exciting when reading through it. Thanks to our newsletter Editor-in-Chief, Dr. Francesco Chiti, for his effort and dedicated time to realize this issue. Also, you may notice that many parts of our newsletter are waiting for more input and updates from you to make our newsletter look better. On behalf of CIS-TC officers, best wishes to you, your families and friends for a healthy and joyful holiday season and a prosperous new year.

Sincerely,
Abderrahim Benslimane
Chair of CIS-TC

CISTC Outstanding Service Award 2018

The Committee decided to give the CIS-TC Outstanding Service Award 2018 to Professor Yi Qian at Department of Electrical and Computer Engineering, University of Nebraska-Lincoln, Nebraska (USA), for for his dedication and leadership to CIS-TC.

Prof. Qian has been awarded during the CIS-TC meeting held at IEEE Globecom 2018 on December 12th, 2018 in Abu Dhabi.

Yi Qian is a professor in the Department of Electrical and Computer Engineering, University of Nebraska-Lincoln (UNL). Prior to joining UNL, he worked in the telecommunications industry, academia, and the government. Some of his previous professional positions include serving as a senior member of scientific staff and a technical advisor at Nortel Networks, a senior systems engineer and a technical advisor at several start-up companies, an assistant professor at University of Puerto Rico at Mayaguez, and a senior researcher at National Institute of Standards and Technology.

His research interests include information assurance and network security, network design, network modeling, simulation and performance analysis for next generation wireless networks, wireless ad-hoc and sensor networks, vehicular networks, smart grid communication networks, broadband satellite networks, optical networks, high-speed networks and the Internet.

Prof. Yi Qian is a member of ACM and a Fellow member of IEEE. He was the Chair of IEEE Communications Society Technical Committee for Communications and Information Security from January 1, 2014 to December 31, 2015. He was the Technical Program Chair for IEEE International Conference on Communications (ICC) 2018. He is serving on the editorial boards for several international journals and magazines, including serving as the Editor-in-Chief for IEEE Wireless Communications Magazine. He is a Distinguished Lecturer for IEEE Vehicular Technology Society & IEEE Communications Society.

Forthcoming Meeting

The next IEEE ComSoc's Communication & Information Security TC (CISTC) meeting will be held at IEEE ICC 2019, 20-24 May 2019, Shanghai, China.

Featured Topics I

"The Accuracy-Privacy Tradeoff of Mobile Crowdsensing"

Zhu Han[1]

[1]University of Houston, Houston, Texas (USA)

The proliferation of mobile devices with built-in sensors has made mobile crowdsensing an efficient sensing paradigm especially in people-centric and Internet of Things (IoT) services. Crowdsensing users collect sensing data using their personal mobile devices, e.g., mobile phones and IoT gadgets. However, the development of crowdsensing services is impeded by many challenges, especially the criticism on the privacy protection of crowdsensing users. Service providers require true data, which is a key factor in optimizing data originated service. This introduces contradicting incentives of maximizing the privacy protection of users and the prediction accuracy of service providers. Most of the existing incentive models in the literature are monetary motivated with sole profit maximization objective, while the privacy incentive of users is neglected. Therefore, conventional monetary-based incentive models are inapplicable in privacy preserving crowdsensing systems, and new privacy-aware incentive models are required. Several major questions related to developing privacy-aware incentive models in mobile crowdsensing arise. First, how does the crowdsensing service define the contributions and payoff allocations of users with varying privacy levels? Second, do crowdsensing coalitions change the attained privacy of the cooperative users? Third, how do cooperative users divide the coalition payoff among themselves?

Our research in [1] provides answers for the aforementioned questions by presenting a novel incentive framework for privacy preservation and accuracy maximization in mobile crowdsensing. The sensing users select their preferred data anonymization levels without knowing the privacy preferences of the other users. The data anonymization is inversely proportional to the accuracy of data analytics of the service provider. Accordingly, the users are paid based on their marginal contributions to the service accuracy. The users can be also penalized with a negative payoff if they cause a marginal harm to the service accuracy, e.g., an outlier providing misleading data. Moreover, a set of k cooperative users can jointly work by forming a crowdsensing coalition, increasing the anonymity privacy protection measured by the k-anonymity metric. The total coalition payoff is then divided among the cooperative users based on their marginal contributions to the coalition's data quality. Our experiments on a real-world dataset of crowdsensing activity recognition system show that the payoff allocation of a particular user does not directly depend on the contributed data size but on the data quality. Likewise, the payoff allocation is found to decrease as the privacy level increases.

References

[1] M. A. Alsheikh, Y. Jiao, D. Niyato, P. Wang, D. Leong, and Z. Han, "The Accuracy-Privacy Tradeoff of Mobile Crowdsensing," *IEEE Communications Magazines, special issue on Sustainable Incentive Mechanisms For Mobile Crowdsensing*, vol. 55, no. 6, pp. 132-139, August 2017.

FEATURED TOPICS II

"The Broken Promise of Decentralized Deep Learning"

Briland Hitaj, Giuseppe Ateniese[1]

[1]Department of Computer Science, Stevens Institute of Technology (USA)

In the past, many attempted a myriad of often-complicated processes to obtain gold from base metals. To some extent, Deep Learning (DL) can be seen as the modern version of the ancient science of Alchemy. Deep Neural Networks (DNNs), the technology residing at the core of DL, are capable of processing tremendous amounts of data, extracting relevant information (features) with little to no human intervention. Through this "alchemical" process, DNNs have significantly transformed various areas of computer science, substantially outperforming previous machine learning (ML) techniques.

Despite its indisputable performance, Deep Learning's undeniable success can be distilled down to two main components: 1) massive amounts of (training) data and 2) powerful computational resources. As a result, DL can prove to be an infeasible task for entities lacking these requirements. Note, that small amounts of training data can lead to models that overfit or memorize the data, thus impractical. As a viable solution, data can be pooled in centralized datasets residing in third-party servers, managed by an entity that satisfies the requirements to train a good DL model, i.e., computational power and resources. The third-party (entity) trains the DNN and typically provides query access to the users. However, this *centralized learning* scheme can be ground for several privacy violations:

- Provided data can lie on third-party's servers indefinitely,
- End-users have no control over the ways their data is utilized,
- End-users have no control over what else could be learned from their data.

For instance, while training a model on face detection/ classification, inspecting the image background for useful features, the model can inadvertently learn information relevant to the location where such images were taken, thus constituting a potential privacy violation. Furthermore, pooling all the data in a centralized data center might be impractical for entities (users) dealing with sensitive information, such as governmental agencies, banks or hospitals.

Decentralized deep learning was introduced as an alternative that allows the participants to jointly train a deep learning model on a specific task without explicitly revealing their training data. Collaborative deep learning by Shokri & Shmatikov [19] or federated learning proposed by Google researchers [11], [12] are a few of the well-known instances of decentralized deep learning. In this setting, the users train a local model on their own data and then share selected updates with other participants, (note that each of the participants has a replica of the model). This feat enables the participants to indirectly influence the learning process of other members while training a model on the given task with no access to the actual training data.

However, is decentralized deep learning actually privacy preserving? Can a malicious participant utilize the *indirect influence* present in the learning process to obtain sensitive information on target participant's training data? In our ACM CCS'17 [10] paper, we are the first to question whether decentralized deep learning *indeed* preserves the privacy of participant's training data. While parameter sharing in a decentralized learning process might sound beneficial at first glance; these parameters are rich with information learned from the actual training data of participant/s. In short, they can be seen as a compressing mechanism to the input data.

In our threat model, the adversary exploits the real-time nature of the decentralized learning process to influence the target victim into revealing more sensitive information on their respective private training data than intended (or needed). Moreover, we devise our attack based on Generative Adversarial Networks (GANs) [8], which to the best of our knowledge is the first use of GANs as an attack mechanism. Furthermore, our attack works even when differential privacy (DP) as proposed in [1], [19], is deployed in a decentralized setting. We emphasize however that the problem is not associated with DP but rather with its incorrect use, see also [17].

GANs attempt to learn the training data distribution [8]. Therefore, even though the inferred samples might not be the exact training data records, they do have a striking resemblance to the actual training data. Consider for instance a hypothetical scenario where Apple deploys a decentralized learning mechanism to improve its face-recognition mechanism (FaceID) introduced with the iPhone X series. In such a scenario, the prototypical samples obtained via GAN could potentially constitute a severe privacy violation. Even in scenarios where the target class encompasses a diverse population of training records, the results obtained via a *properly* deployed GAN-attack will capture the variance. For instance, in a collaboratively trained gender classifier where the target is the 'female' class, the GAN will *indeed* produce samples that look like the female records in the training data which might not constitute a privacy violation per se. However, given that the GAN is attempting to learn the training data distribution [8] of the 'female' class, our attack [10] will allow the adversary to infer potentially sensitive information from the generated samples, such as the presence of glasses, race, and more. This is in line with state-of-the-art works in the domain [2], [5], [7].

As is often the case in computer security studies, once a flaw in a system is found other researchers extend it in a short time. Recently, Nasr et al. [15] and Melis et al. [14] show evidence of a form of

membership/property inference on collaborative learning that exploits the active nature of this learning process (i.e., model updates). Their type of inference (seen as a decisional problem) is limited in this context since the adversary needs auxiliary training data correctly labeled and, for example, will not work in the FaceID scenario above. Moreover, DP typically serves as a suitable defense mechanism against membership inference attacks. However, this type of inference can be applied by a passive attacker, thus invalidating all security claims of collaborative learning [19].

Pursuant to our work, new attacks have been proposed considering the privacy implications of decentralized deep learning [4], [3], [20]; and recently the works by Phong et al. [16], Hayes et al. [9], Fung et al. [6], Shayan et al. [18], and more, make the first attempts at devising effective countermeasures for such attacks. Moreover, McMahan et al. [13] propose the use of participant-level differential privacy as an effective countermeasure. However, they do so in a scenario with thousands of participants, and it is unclear whether it is effective with fewer participants, as also pointed out in [4]. Therefore, we believe that there is still considerable ground for further improvements of a decentralized deep learning mechanism for real-life applications.

References

[1] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, “Deep learning with differential privacy,” in Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, ser. CCS ’16. New York, NY, USA: ACM, 2016, pp. 308–318.
[2] G. Ateniese, L. V. Mancini, A. Spognardi, A. Villani, D. Vitali, and G. Felici, “Hacking smart machines with smarter ones: How to extract meaningful data from machine learning classifiers,” International Journal of Security and Networks, vol. 10, no. 3, pp. 137–150, 2015.
[3] H. Bae, J. Jang, D. Jung, H. Jang, H. Ha, and S. Yoon, “Security and privacy issues in deep learning,” CoRR, vol. abs/1807.11655, 2018.
[4] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, “How to backdoor federated learning,” CoRR, vol. abs/1807.00459, 2018.
[5] M. Fredrikson, S. Jha, and T. Ristenpart, “Model inversion attacks that exploit confidence information and basic countermeasures,” in Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. ACM, 2015, pp. 1322–1333.
[6] C. Fung, J. Koerner, S. Grant, and I. Beschastnikh, “Dancing in the dark: private multi-party machine learning in an untrusted setting,” arXiv preprint arXiv:1811.09712, 2018.
[7] K. Ganju, Q. Wang, W. Yang, C. A. Gunter, and N. Borisov, “Property inference attacks on fully connected neural networks using permutation invariant representations,” in Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, ser. CCS ’18. ACM, 2018, pp. 619–633.
[8] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in Advances in neural information processing systems, 2014, pp. 2672–2680.
[9] J. Hayes and O. Ohrimenko, “Contamination attacks and mitigation in multi-party machine learning,” in Advances in Neural Information Processing Systems, 2018, pp. 6602–6614.
[10] B. Hitaj, G. Ateniese, and F. P´erez-Cruz, “Deep models under the GAN: Information leakage from collaborative deep learning,” in Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, ser. CCS ’17. New York, NY, USA: ACM, 2017, pp. 603–618.
[11] B. McMahan and D. Ramage. (2017) Federated learning: Collaborative machine learning without centralized training data. https://research.googleblog.com/2017/04/federated-learning-collaborative.html.
[12] H. B. McMahan, E. Moore, D. Ramage, and B. A. y Arcas, “Federated learning of deep networks using model averaging,” CoRR, vol. abs/1602.05629, 2016.
[13] H. B. McMahan, D. Ramage, K. Talwar, and L. Zhang, “Learning differentially private language models without losing accuracy,” arXiv preprint arXiv:1710.06963, 2017.
[14] L. Melis, C. Song, E. D. Cristofaro, and V. Shmatikov, “Exploiting unintended feature leakage in collaborative learning,” in IEEE S&P, 2018.
[15] M. Nasr, R. Shokri, and A. Houmansadr, “Comprehensive privacy analysis of deep learning: Stand-alone and federated learning under passive and active white-box inference attacks,” arXiv preprint arXiv:1812.00910., 2018.
[16] L. T. Phong, Y. Aono, T. Hayashi, L. Wang, and S. Moriai, “Privacypreserving deep learning via additively homomorphic encryption,” IEEE Transactions on Information Forensics and Security, vol. 13, pp. 1333–1345, 2017.
[17] M. A. Rahman, T. Rahman, R. Laganiere, N. Mohammed, and Y. Wang, “Membership nference attack against differentially private deep learning model”, Transactions on Data Privacy, vol. 11, no. 1, pp. 61–79, 2018.
[18] M. Shayan, C. Fung, C. J. Yoon, and I. Beschastnikh, “Biscotti: A ledger for private and secure peer-to-peer machine learning,” arXiv preprint arXiv:1811.09904, 2018.
[19] R. Shokri and V. Shmatikov,“Privacy-preserving deep learning,” in Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security, ser. CCS ’15. New York, NY, USA: ACM, 2015, pp. 1310–1321.
[20] Z. Wang, M. Song, Z. Zhang, Y. Song, Q. Wang, and H. Qi, “Beyond inferring class representatives: User-level privacy leakage from federated learning,” arXiv preprint arXiv:1812.00535., 2018.

CIS-TC Organized Symposia

- *Communications and Information System Security Symposium* – **ICC2018**, 20-24 May 2018 Kansas City, MO, USA (http://icc2018.ieee-icc.org)
- *Communications and Information System Security Symposium* – **Globecom2018,** 9-13 December 2018, Abu Dhabi, UAE (http://globecom2018.ieee-globecom.org)

CIS-TC Affiliate Conferences

- **ICACT 2018**
  - *20th IEEE International Conference on Advanced Communications Technology*
  - 11-14 February, 2018 Elysian Ganchon Ski Resort, GW, Korea
  - http://www.icact.org

- **WTS 2018**
  - *Wireless Telecommunications Symposium*
  - 18-20 April, 2018 Phoenix (Chandler), Arizona, USA
  - http://www.cpp.edu/~wtsi

- **CNS 2018**
  - *Conference on Communications and Network Security*
  - 30 May-1 June 2018, Beijing, China
  - http://cns2018.ieee-cns.org

- **MoWNet 2018**
  - *International conference on selected topics in Mobile and Wireless Networking (@ International 5G Summit)*
  - 20-22 June 2018, Tangier, Morocco
  - http://mownet.org

- **WiMob 2018**
  - *14th IEEE International Conference on Wireless and Mobile Computing, Networking and Communications*
  - 15-17 Oct, 2018, Limassol, Cyprus
  - http://www.wimob.org/wimob2018