Department of Computer Science

Machine learning, security, and privacy: challenges and opportunities

Binghui Wang
Duke University

Hosted by Iulian Neamtiu

DATE: Monday, April 12, 2021
TIME: 2:30 PM - 3:30 PM
Webex Link: https://njit.webex.com/njit/j.php?MTID=mcf4f34868c528f064c084f57c2c98dd5

https://cs.njit.edu/seminars

Abstract: Machine learning provides more and more powerful tools for data analytics. On the other hand, security and privacy attacks increasingly involve data. Therefore, machine learning and security & privacy naturally intersect with each other as they both involve data, and there are many interesting questions at the intersections: i) How machine learning impacts security and privacy analytics design? ii) How security and privacy impact the deployment of machine learning systems?

In this talk, I will first talk about machine learning for security and privacy in social networks, particularly, graph-based collective classification to detect fake accounts in social networks. A long-standing challenge in collective classification is that existing methods cannot learn accurate edge weights, thus resulting in limited detection performance in practice. I address this challenge by developing a general weight learning framework. Our framework is applicable to many collective classification methods and significantly enhances detection performance for state-of-the-art methods, while with an acceptable computational overhead.

Second, I will present the work on stealing hyperparameters in machine learning. Hyperparameters are critical in machine learning and they are also deemed confidential because of their commercial value. I design a framework to accurately steal hyperparameters that are used to balance between the loss function and the regularization terms in the objective function of machine learning algorithms, and evaluate the attack theoretically. I also show that the hyperparameter stealing attack can be used by users to save economic costs when they use machine-learning-as-a-service platforms, e.g., Amazon machine learning, to train machine learning models. I further study countermeasures. The results highlight the need for new defenses against the hyperparameter stealing attack.

Finally, I will describe my future research plan on building machine learning systems that are provably secure and privacy-preserving and extending my research from cyber security and privacy to cyber-physical security and privacy.

Bio: Binghui Wang is currently a Postdoctoral Researcher in Electrical and Computer Engineering at Duke University. He obtained his Ph.D. from the Electrical and Computer Engineering Department at Iowa State University in July 2019. His research interests are trustworthy machine learning, data-driven security and privacy, and machine learning. His research has been published in top-tier security conferences such as IEEE S & P, CCS, NDSS, and AI/computer vision/data mining/networking/ conferences such as NeurIPS, ICLR, CVPR, WWW, AAAI, INFOCOM, and ICDM. His work has received several honors and awards, including 2020 DeepMind Best Extended Abstract Award, 2019 NDSS Distinguished Paper Award Honorable Mention, 2017 INFOCOM Paper for Fast Tracking, and has been also deployed by JD, the second-largest e-commerce company headquartered in China.