Armon Barton

University of Texas, Arlington

Towards Defending Deep Neural Networks Against Adversarial Examples, and Predicting Efficient and Anonymous Tor Circuits

**** GE-104, 1500-1600, 4 March 2019 ****

Abstract: Deep learning is becoming a technology central to the safety of cars, the security of networks, and the correct functioning of many other types of systems. Unfortunately, attackers can create adversarial examples, small perturbations to inputs that trick deep neural networks into making a misclassification. While the original images classify as the correct digit, the adversarial examples all classify as the number '9'. To solve this problem, I proposed PadNet, a stacked defense against adversarial examples that does not require knowledge of the attack techniques used by the attacker. PadNet combines two novel techniques: Defensive Padding and Targeted Gradient Minimizing (TGM). I will show results indicating that PadNet significantly increases robustness against adversarial examples compared to Adversarial Logit Pairing. Against a state-of-the art attack, PadNet improved robustness by 92% compared to ALP. I also observed that PadNet was adaptable to various attacks without knowing the attacker's techniques, and therefore allows the training cost to be fixed unlike Adversarial Logit Pairing.

My anonymity research was devoted toward improving the usability and security of the popular anonymous communication system called Tor. The Tor anonymity system provides online privacy for millions of users, but it is slower than typical web browsing. To improve Tor performance, I proposed PredicTor, a path selection technique that uses a Random Forest classifier trained on recent measurements of Tor to predict the performance of a proposed path. If the path is predicted to be fast, the client then builds a circuit using those relays. I implemented PredicTor in the Tor source code and showed through live Tor experiments and Shadow simulations that PredicTor improves Tor network performance by 11% to 23% compared to Vanilla Tor and by 7% to 13% compared to the previous state-of-the-art scheme. My experiments showed that PredicTor is the first path selection algorithm to dynamically avoid highly congested nodes during times of high congestion and avoid long-distance paths during times of low congestion.

Biography: Armon Barton earned his Ph.D. degree in computer science and engineering in 2018 from the University of Texas at Arlington, and his B.S. degree in mechanical engineering in 2006 from Texas Tech University. His research interests lie at the intersection of cybersecurity and machine learning and are focused on secure machine learning and anonymous communications.