

***Trusted Inference Engine: Preventing Neural Network Exfiltration
in Hardware Devices***

Michel A. Kinsy
Boston University

Tuesday, November 13, 2018

2:00 pm

EEB 132

Companies, in their push to incorporate artificial intelligence - in particular, machine learning - into their Internet of Things (IoT), system-on-chip (SoC), and automotive applications, will have to address a number of design challenges related to the secure deployment of artificial intelligence learning models and techniques. Machine learning (ML) models are often trained using private datasets that are very expensive to collect, or highly sensitive, using large amounts of computing power. The models are commonly exposed either through online APIs, or used in hardware devices deployed in the field or given to the end users. This gives incentives to adversaries to attempt to steal these ML models as a proxy for gathering datasets. While API-based model exfiltration has been studied before, the theft and protection of machine learning models on hardware devices have not been explored as of now. In this work, we examine this important aspect of the design and deployment of ML models. We illustrate how an attacker may acquire either the model or the model architecture through memory probing, side-channels, or crafted input attacks, and propose (1) power-efficient obfuscation as an alternative to encryption, and (2) timing side-channel countermeasures.



Michel A. Kinsy is an Assistant Professor in the Department of Electrical and Computer Engineering at Boston University (BU), where he directs the Adaptive and Secure Computing Systems (ASCS) Laboratory. He focuses his research on computer architecture, hardware-level security, neural network accelerator designs, and cyber-physical systems. Dr. Kinsy is an MIT Presidential Fellow, the 2018 MWSCAS Myril B. Reed Best Paper Award Recipient, DFT'17 Best Paper Award Finalist, and FPL'11 Tools and Open-Source Community Service Award Recipient. He earned his PhD in Electrical Engineering and Computer Science in 2013 from the Massachusetts Institute of Technology. His doctoral work in algorithms to emulate and control large-scale power systems at the microsecond resolution inspired further research by the MIT spin-off Typhoon HIL, Inc. Before joining the BU faculty, Dr. Kinsy was an assistant professor in the Department of Computer and Information Systems at the University of Oregon, where he directed the Computer Architecture and Embedded Systems (CAES) Laboratory. From 2013 to 2014, he was a Member of the Technical Staff at the MIT Lincoln Laboratory.