From Hardware to Software: Defending the Next Generation of Artificial Intelligence and Machine Learning Applications

by

For the Ph.D. degree in Computer Science and Engineering

The growing energy costs of artificial intelligence (AI) and machine learning (ML) workloads have motivated research efforts into low-power acceleration platforms. One popular platform is field-programmable gate arrays (FPGAs), due to their low-power and in-field reconfigurability. As the use of these specialized hardware platforms becomes more prevalent, concerns around the security of these systems have intensified. This has led to a significant body of research in adversarial machine learning, aimed at securing both the software and hardware components of ML applications. This dissertation focuses on hardware security in the context of FPGA-based ML systems. I explore the unique security risks associated with deploying ML applications on FPGAs and present novel methods for securing these systems against various cyberattacks, including IP theft. To ensure the overall security and integrity of FPGA-based ML systems, this work addresses security concerns at every