STEVENS
INSTITUTE OF TECHNOLOGY
1870

Ph.D. DISSERTATION DEFENSE

Candidate:	Ayomide Akinsanya
Degree:	Doctor of Philosophy
School/Department:	Charles V. Schaefer, Jr. School of Engineering and Science (SES) / Computer Science
Date:	Wednesday, August 19, 2026
Time/Location:	1:00PM / Gateway North Room 303
Title:	Timing Channels and Efficiency Attacks in Adaptive Neural Networks
Chairperson:	Dr. Tegan Brennan, Department of Computer Science, Stevens Institute of Technology
Committee Members:	Dr. Hui Wang, Department of Computer Science, Stevens Institute of Technology Dr. Shaoyi Huang, Department of Computer Science, Stevens Institute of Technology Dr. Koduvayur Subbalakshmi, Department of Computer Engineering, Stevens Institute of Technology

ABSTRACT

Adaptive neural networks (AdNNs) have enjoyed wide adoption as an alternative to traditional deep neural networks (DNNs) due to their ability to dynamically adjust computation during inference. By leveraging optimization mechanisms such as early-exits and other adaptive inference strategies, AdNNs can reduce latency, computational cost, memory requirements, and energy consumption of traditional DNNs without significantly sacrificing their accuracy, making them particularly attractive for deployment in resource-constrained environments such as mobile devices, IoT systems, and real-time low-latency applications. However, little is known about the security and privacy implications of this adaptive behavior. It is therefore critical to investigate whether the very mechanisms that enable efficiency in AdNNs may inadvertently introduce novel attack surfaces. In this dissertation, we present the first systematic study of timing side channels in AdNNs, demonstrating how adversaries can exploit the adaptive behavior of AdNNs to infer sensitive information about user inputs at inference time. We then extend this threat to federated learning environments, showing that a malicious participant in a collaborative training procedure can deliberately craft timing channels into the jointly learned model and subsequently exploit them during inference. Leveraging the same underlying intuition, we additionally propose FedSlowdown, an efficiency attack that intentionally degrades AdNN computational efficiency, forcing adaptive models to behave like traditional DNNs and negating their key advantages in latency and energy consumption. On the defensive side, we introduce the Composite Inter-Attribute Variance (CIAV), a novel metric for quantifying inter-attribute confidence asymmetry in AdNNs, and propose Dual Attribute-Aware Temperature Scaling (DAATS), a novel training-time defense that eliminates timing channel vulnerability at its source by directly targeting the asymmetric confidence distributions that drive timing leakage. Finally, we additionally develop a runtime monitoring framework for early detection of timing-based information leakage in deployed AdNN models. These contributions aim to establish a comprehensive understanding of

both the privacy and efficiency risks facing AdNNs, to expose the environments in which they are most vulnerable, and to provide principled and practical countermeasures that ensure AdNNs remain both secure and efficient in practice.