Keynote: AI Canonical Architecture and Robust AI

David R. Martinez[*]

Abstract

This presentation addresses an AI canonical architecture suitable for a number of different classes of applications. Several examples will be shown focused on cyber security and potential vulnerabilities to adversarial attacks. One critical element of the end-to-end AI architecture is the need for robust AI. Significant advances have been made in AI algorithms and high performance computing. However, additional advancements in science and technology (S & T) are needed to validate the performance of AI systems. This performance assessment is very critical because AI systems are very brittle to adversarial modifications to the system. The AI canonical architecture starts with data conditioning, followed by classes of machine learning algorithms, human-machine teaming, modern computing, and robust AI. We will briefly address each of these areas. The presentation concludes with a summary of S & T challenges and recommendations.

[*] D. Martinez is with the Lincoln Laboratories, Massachusetts Institute of Technology, MA, USA. e-mail: dmartinez@ll.mit.edu
In: Joseph Collins, Prithviraj Dasgupta, Ranjeev Mittu (eds.): Proceedings of the AAAI Fall 2018 Symposium on Adversary-Aware Learning Techniques and Trends in Cybersecurity, Arlington, VA, USA, 18-19 October, 2018, published at http://ceur-ws.org