

COLLOQUIUM

Sponsored by
The University of Georgia
Department of Computer Science,
Institute for Artificial Intelligence, & Faculty of Robotics

Using Symbolic Artificial Intelligence to Explain Statistical Machine Learning

Dr. Arthur Choi

University of California at Los Angeles

Friday, November 6, 2020

11:00 AM - 12:15 PM

Zoom - <https://zoom.us/j/7065830827?from=addon>

Abstract:

Recent and rapid advances in Artificial Intelligence (AI), particularly in the form of deep neural networks, has opened many new possibilities, but it has also brought with it many new challenges. In particular, it has become increasingly apparent that while deep neural networks are highly performant, they can also be opaque and brittle. We do not have enough understanding of why and when they work well, and why they may fail completely when faced with new situations not seen in the training data. In this talk, we propose a symbolic approach to explaining the behavior and verifying the properties of machine learning models, which is based on sustained advances in logical and probabilistic reasoning. We show how our approach facilitates the analysis of a neural network, helping us to understand its behavior, and in turn, providing directions towards learning better and more robust models.

Biography:

Arthur Choi is a research scientist in the Computer Science Department at the University of California at Los Angeles (UCLA). He obtained his Ph.D. from UCLA in 2010. His research interests are in logical and probabilistic approaches to reasoning, machine learning, tractable knowledge representations, and most recently in explainable artificial intelligence (XAI). He has led medal-winning teams in international probabilistic inference competitions. He has served as an SPC/PC member at conferences such as IJCAI, AAAI, UAI, NeurIPS, and ICML.