

Title: Formal Verification of Deep Neural Networks: Challenges and Recent Advances

Abstract: Neural networks have become a crucial element in modern artificial intelligence. When applying neural networks to mission-critical tasks such as those in autonomous systems and medical devices, it is often desirable to formally verify their trustworthiness such as safety, robustness and correctness. In this talk, I will first introduce the problem of neural network verification and the challenges of guaranteeing the behavior of a neural network given input specifications. Then, I will discuss the bound-propagation-based algorithms (e.g., CROWN and beta-CROWN), which are efficient, scalable and powerful techniques for formal verification of neural networks and can also be generalizable to computational graphs beyond neural networks. My talk will highlight state-of-the-art verification techniques used in our α, β -CROWN (alpha-beta-CROWN) verifier that won the 2nd and 3rd International Verification of Neural Networks Competition (VNN-COMP 2021 and 2022), as well as recent advances and open challenges in the field of neural network verification.

Bio: Huan Zhang is a postdoctoral researcher at CMU, supervised by Prof. Zico Kolter. He received his Ph.D. degree at UCLA in 2020. Huan's research focuses on the trustworthiness of artificial intelligence, especially on developing formal verification methods to guarantee the robustness and safety of machine learning. Huan was awarded an IBM Ph.D. fellowship and he led a multi-institutional team that won the 2021 and 2022 International Verification of Neural Networks Competition. Huan received the 2021 AdvML Rising Star Award sponsored by MIT-IBM Watson AI Lab.

Photo: https://huan-zhang.com/images/Huan_Zhang_photo.jpg