COMPUTER SCIENCE & ENGINEERING
UNIVERSITY OF MICHIGAN

DISSERTATION DEFENSE

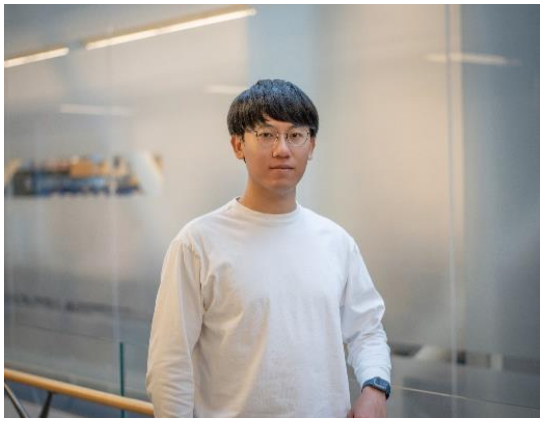

Jiachen Sun

On Improving Robustness of Deep Neural Networks for Computer Vision

Monday, October 30, 2023
9:00am – 11:00am
3725 Beyster
Hybrid – Zoom
Passcode: 405537

ABSTRACT: Over the past decade, deep learning has spearheaded a revolution in a myriad of applications, achieving remarkable success. This is particularly evident in the field of computer vision, which has undergone a transformative evolution, fueled by deep neural networks and the availability of extensive datasets. However, despite the monumental advancements in computer vision research, significant vulnerabilities persist. Numerous studies have illustrated that deep networks are fundamentally susceptible to alterations, including subtle perturbations and style changes, in input data that maintain consistent conceptual integrity in human perception. This includes phenomena such as common corruptions, adversarial perturbations, and domain shifts. To address the identified challenges within computer vision, my dissertation embarked on a systematic exploration and resolution of the robustness dilemma inherent in both 2D and 3D realms through an array of approaches encompassing comprehensive benchmarking, architectural refinements, train-time strategies, test-time adaptation, and system-level methodologies. My dissertation is articulated over four segments. In the initial segment, benchmarking and enhancement of adversarial and general robustness in 2D image classification are attended to. We introduced FourierMix as a novel data augmentation technique and VPA as a test-time prompting method to ameliorate certified adversarial and general robustness, respectively. The ensuing segment revolves around benchmarking and improving the robustness of 3D point cloud recognition against adversarial attacks and common corruptions. We first proposed a novel dataset, ModelNet40-C, accompanied by systematic benchmarking results. Subsequently, we devised a new global pooling operation, DeepSym, and harnessed self-supervised learning to heighten the efficacy of adversarial training in 3D point cloud recognition. Furthermore, we proposed PointDP, a diffusion-driven purifier, as a defense against adversarial point clouds. Transitioning to a more complex task, the third segment is dedicated to boosting the robustness of LiDAR-based 3D object detection. A universal vulnerability within LiDAR-based perception was identified, followed by the proposal of both system- and model-level countermeasures to mitigate the vulnerability. In the final part, we introduced CALICO, aiming at elevating 2D and 3D multimodal perception efficiency and robustness via constrastive pretraining strategies.

CHAIR: Prof. Z. Morley Mao