Michigan AI Symposium

RESPONSIBLE AI

October 17 | 9am-5pm

North Campus, Beyster Building 1670 & Tishman Hall

AI LAB

Causal Inference for Robust and Responsible NLP

Lightning Talk **| Zhijing Jin**

Abstract: Despite the remarkable progress in large language models (LLMs), it is well-known that natural language processing (NLP) models tend to fit for spurious correlations, which can lead to unstable behavior under domain shifts or adversarial attacks. In my PhD research, I have developed various adversarial attacks and stress tests for NLP, while also integrating causal inference as a systematic framework for enhancing robustness throughout the entire process, from data collection and model training to evaluation. Through these endeavors, I aim to advance the development of robust and responsible NLP systems by leveraging causal inference.

Neural Caches for Monte Carlo Partial Differential Equation Solver

Lightning Talk **| Guandao Yang**

Abstract: This paper presents a method that uses neural networks as a caching mechanism to reduce the variance of Monte Carlo Partial Differential Equation solvers, such as the Walk-on-Spheres algorithm [Sawhney and Crane 2020]. While these Monte Carlo PDE solvers have the merits of being unbiased and discretization-free, their high variance often hinders real-time applications. On the other hand, neural networks can approximate the PDE solution, and evaluating these networks at inference time can be very fast. However, neural-network-based solutions may suffer from convergence difficulties and high bias. Our hybrid system aims to combine these two potentially complementary solutions by training a neural field to approximate the PDE solution using supervision from a WoS solver. This neural field is then used as a cache in the WoS solver to reduce variance during inference. We demonstrate that our neural field training procedure is better than the commonly used self-supervised objectives in the literature. We also show that our hybrid solver exhibits lower variance than WoS with the same computational budget: it is significantly better for small compute budgets and provides smaller improvements for larger budgets, reaching the same performance as WoS in the limit.

Personalized Preference-Bandits: Learning to make customized user predictions through preference elicitation

***Lightning Talk* | Aadirupa Saha**

Abstract: Customer statistics collected in several real-world systems have reflected that users often prefer eliciting their liking for a given pair of items, say (A,B), in terms of relative queries like: "Do you prefer Item A over B?", rather than their absolute counterparts: ``How much do you score items A and B on a scale of [0-10]?". Drawing inspirations, in the search for a more effective feedback collection mechanism, led to the famous formulation of Dueling Bandits (DB), which is a widely studied online learning framework for efficient information aggregation from relative/comparative feedback. However despite the novel objective, unfortunately, most of the existing DB techniques were limited only to simpler settings of finite decision spaces, and stochastic environments, which are unrealistic in practice. We will see the basic problem formulations for DB and familiarize ourselves with some of the breakthrough results. Following this, will dive deep into a more practical framework of contextual dueling bandits (C-DB) where the goal of the learner is to make customized predictions based on the user contexts: We will see a new algorithmic approach that can efficiently achieve the optimal O(\sqrt T) regret performance for this problem, resolving an open problem from Dudík et al. [COLT, 2015]. We will conclude with some interesting open problems.

Discovering the Right Things to Design with Artificial Intelligence

***Lightning Talk* | Nur Yildirim**

Abstract: Advances in artificial intelligence have enabled unprecedented technical capabilities, yet making these advances useful in the real world remains challenging. Today, most AI projects fail, often because product teams select suboptimal places to apply AI. I argue that the current process for designing AI products and services is broken. In this talk, I present a new innovation process that helps teams identify low risk, high value use cases where moderate AI performance can create value for people. Through case studies, I highlight the vital role human-computer interaction research plays in finding applications where humans can benefit and thrive with AI.

Deep Learning Tools for Large-Scale Multi-modal Neural Data.
***Lightning Talk* | Lu Mi**

Abstract: Recent years have marked notable advancements in neuroscience, largely attributable to the development of advanced tools enabling more detailed studies of the brain. However, despite such advancements, our accessibility and comprehension of the brain's intricacies are still in nascent stages. My research primarily aims to expand our abilities to collect and interpret large-scale, multi-modal neural data - including anatomical structure, functional activity, transcriptomics, behaviors - by leveraging cutting-edge computer vision and machine learning methods. By enhancing the accuracy, efficiency, and scalability of acquisition and analysis workflows for neuroscience studies, we aspire to expedite scientific discoveries in unraveling of the mysteries of coding, computation, and learning processes in the brain. Furthermore, this research paves the way for developing innovative brain-inspired AI frameworks, potentially closing the loop between artificial and natural intelligence.

Towards 4D Capture of Objects and Scenes from Casual Videos
***Lightning Talk* | Gengshan Yang**

Abstract: Prior methods for 4D capture often require specialized sensors or body templates, making them less applicable to the diverse objects and scenarios one may encounter in everyday life. To address this, we build 4D models utilizing casually-taken monocular videos, aiming to more faithfully capture the diversity of the world around us. Although the problem is challenging due to its under-constrained nature, recent advances in inverse graphics and vision foundation models allow us to approach it as an analysis-by-synthesis task, where the goal is to find a physically plausible representation that synthesizes videos consistent with the vision models. Building upon this framework, we present methods to capture deformable objects and their surrounding scenes from in-the-wild video footage.