NSF CREST MECIS – Spring 2026 Webinar Series – 5/1/2026

Presentation 1: **Stefan Cobeli (Doctoral Student), Dr. Bo Zou, University of Illinois Chicago (UIC)**

Title: Large Language Models for Transportation Mode Choice Prediction in Low-Resource Settings: Zero-Shot Inference Enhancement and Lightweight Adaptation

Abstract:

Our work investigates how large language models (LLMs) can be applied to transportation mode choice prediction using a large-scale, revealed-preference household travel survey from the Chicago metropolitan area. We propose three methodological contributions. First, a prompt augmentation strategy that incorporates external contextual information into natural-language trip descriptions. Second, an agentic inference approach in which the LLM refines its prediction sequentially across multiple prompting steps. Third, a lightweight adaptation method that uses frozen LLM embeddings with a trainable readout function. Our results show that the zero-shot agentic approach outperforms classical supervised baselines without requiring any labeled training data, and that the lightweight adaptation strategy achieves the strongest overall performance.

Presentation 2: **Fabian Hernandez (Master's Student), Dr. Qi Lu, University of Texas Rio Grande Valley (UTRGV)**

Title: Learning Adaptive Control for Safe Collaborative Autonomous Driving

Abstract:

Connected autonomous vehicles improve urban driving through collaborative perception via V2X communication. Frameworks such as V2Xverse use this collaboration for planning and perception, but their controllers rely on fixed parameters that cannot adapt to changing traffic. Control Barrier Functions (CBFs) enforce safety by constraining actions, yet fixed gains force a trade-off between overly cautious and unsafe behavior. This work introduces an adaptive CBF framework that learns state-dependent, class-specific gains for vehicles, pedestrians, and bicycles using constrained reinforcement learning. The learned policy adjusts a CBF-based quadratic program that modifies only longitudinal control. We evaluate 315 closed-loop trials across 105 routes in CARLA Town05. The adaptive method improves driving score to 75.81, reduces collisions by 80.3% versus baseline and 25.9% versus fixed CBF, and achieves higher route completion, showing that a simple safety layer can significantly improve performance without altering the upstream system.

Presentation 3: **Shaan Pakala (Doctoral Student), Dr. Evangelos Papalexakis, University of California Riverside (UCR)**

Title: Automating Data Science Pipelines with Tensor Completion

Abstract:

Hyperparameter optimization is an essential component in many data science pipelines, but it typically entails costly computations to exhaustively explore the design space. Similarly, neural architecture search (identifying the best design choices for a neural network) and query cardinality estimation (estimating the output cardinality of an SQL query) exhibit a similar combinatorial search

complexity. In this work, we abstract away these specific components of data science pipelines and instead model them as instances of tensor completion, where each variable of the search space corresponds to one mode of the tensor. Now in order to optimize these pipelines, we just identify the missing entries of the tensor (i.e. tensor completion). We thoroughly evaluate existing tensor methods, as well as introduce domain-inspired adaptations and an ensemble method, which achieve state-of-the-art performance. Furthermore, we release our generated datasets and code to encourage future work.

<u>Presentation 4:</u> Zhugang Liu (Doctoral Student), Dr. Jinghao Yang, University of Texas Rio Grande Valley (UTRGV)

Title: Bridging the Pretrain-to-Real Gap: Alignment Challenges in Deploying Generalist VLA Models for Additive Manufacturing

Abstract:

Generalist Vision-Language-Action (VLA) models mark a significant milestone in the Generative AI era. However, the prevailing reliance on simulated benchmarks obscures the severe physical and algorithmic domain shifts encountered during real-world hardware deployment. This research presents a comprehensive, high-performance framework that successfully deploys the 7B-parameter OpenVLA model directly onto a physical Franka Research 3 (FR3) robotic arm. Utilizing a meticulously optimized distributed cloud-edge architecture, we specifically target the highly dynamic workflows of additive manufacturing. Differentiating from simulation-only studies, we identify and decisively resolve the "system-level alignment challenges" intrinsic to this Pretrain-to-Real transition. Using the grasping of topologically complex 3D-printed parts as a rigorous evaluation metric, we detail our algorithmic mitigation strategies, including deterministic decoding enforcement, strict thresholding from continuous to binary, and geometric retargeting combined with Low-Rank Adaptation (LoRA). Our empirical results demonstrate that the aligned hardware system achieves a 98% reduction in spatial L2 error and ensures stable convergence of gripper states, providing a robust, field-tested taxonomy for deploying generative embodied AI in flexible industrial settings.

Please register for this webinar using the following QR Code: