

Research Title:

Scalable Distributed Learning-based Approach for Joint Navigation in Local Environments

Primary Investigator:

Name: **Dr. Gilad Katz**

Faculty: **Department of Software and Information Systems Engineering**

Academic Institute: **Ben-Gurion University of the Negev**

This research introduces a scalable distributed learning-based approach for joint traffic navigation that addresses critical limitations in multi-agent vehicle coordination through a novel generative master architecture. Rather than relying on fixed, discrete command sets that scale poorly and create information silos between vehicles, the proposed system employs a generative master module that produces a fixed-size "proto-plan" embedding containing coordinated instructions for all vehicles. Individual vehicle agents then integrate this shared plan into their decision-making processes, enabling scalable coordination that adapts to variable numbers of vehicles, heterogeneous vehicle capabilities, and new action types without requiring architectural modifications or retraining. The project currently encompasses two parallel research directions with concrete implementations and experimental validations. In the first direction, we implemented a generative master module which communicates with individual agent modules. Through dynamic intersection scenarios, we demonstrated that the system successfully coordinates multiple DRL-controlled vehicles, first with two vehicles and then scaling to three vehicles without any architectural modifications. The experiments showed complete collision prevention after initial learning periods, with positive rewards indicating effective coordination, and the architecture exhibited strong generalization capabilities when handling additional vehicles. In the second direction, we implemented a Decision Transformer trained on diverse expert trajectories in realistic highway simulation environments. The evaluation showed that the Decision Transformer learns reward-conditioned policies where achieved returns match target values, with success rates reaching 46% at medium reward thresholds. While the approach is effective, the crash

rates observed in our experiment validates the necessity of master-level coordination. These implementations demonstrate the feasibility of the generative approach and provide concrete evidence for the system's ability to learn consistent, reward-driven policies across varying coordination requirements.