

MotionToken-MAPF: Coordination in Mixed Autonomy via Language-Based Motion Prediction

Emanuel Munoz

Introduction and Problem Statement

Context:

Autonomous vehicles are becoming essential for the future of logistics and urban mobility (robotaxis, trucks).

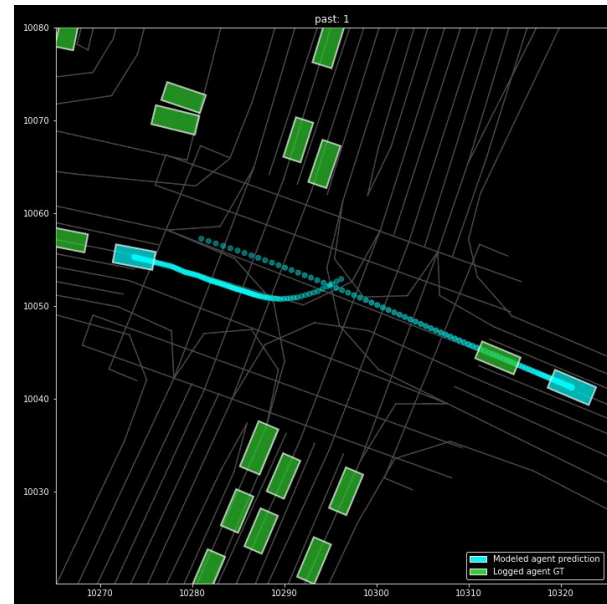
We must manage a mixed-autonomy coexistence between Coordinated Autonomous Vehicles (CAV) and Human-Driven Vehicles (HDV).

The Problem:

CAVs must predict HDV behavior accurately even when they interact between each other.

Objective:

To develop a multi-agent cooperative motion framework for specific high-conflict scenarios to avoid collisions and find optimal trajectories.



Background

Multi-agent planning coordination

- BK-PBS [2] proposes to use PBS and a G-CNN for conditional behavior prediction.
- Limitation 1: Only considers a single primitive (lane changing).
- Limitation 2: It does not forecast a full trajectory (shortsighted and optimistic).

Behavior conditional prediction


- MotionLM [1]: Forecasting as an NLP paradigm, treating agent trajectories as sequences of tokens.
- Generates trajectories autoregressively, allowing for a physically consistent and temporally dense future state, enforcing causality.

Proposal: MotionToken-MAPF

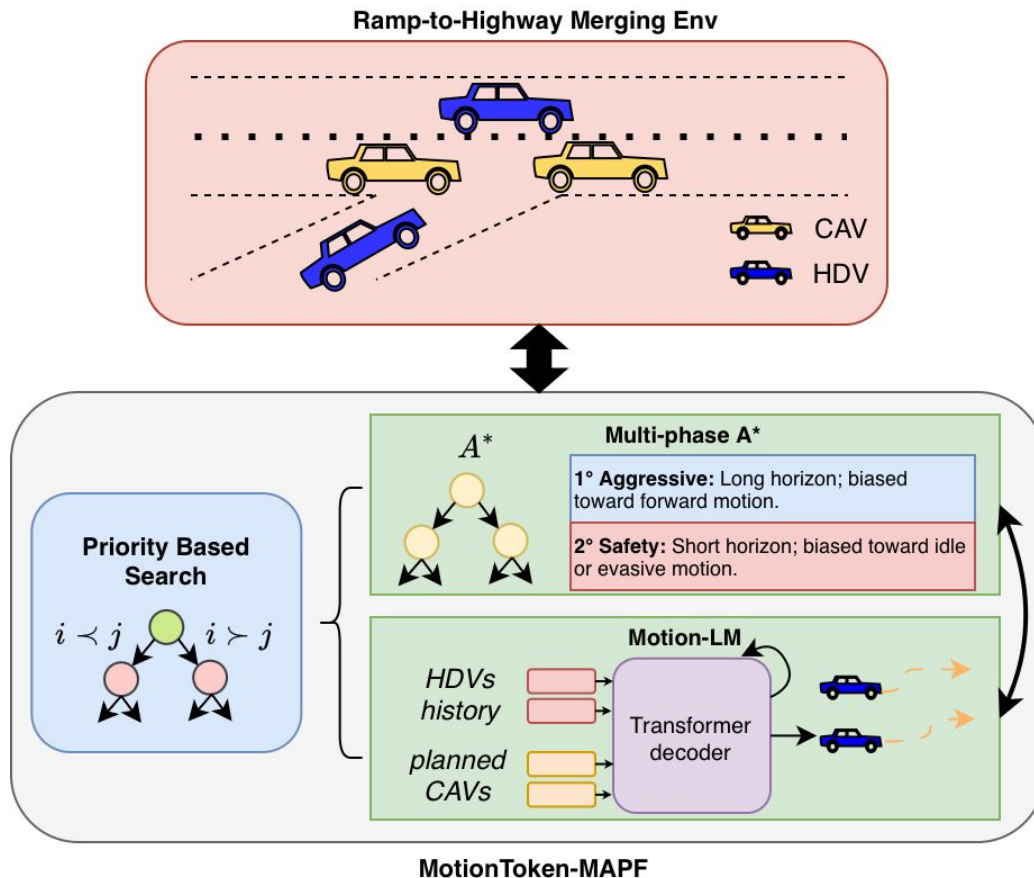
- Use BK-PBS backbone, but use MotionLM instead.
- Handle braking, idling, and emergency maneuvers through trajectory forecasting.
- Powerful architecture for scaling: heterogeneous agents and complex environments.

[1] A. Seff et al., "MotionLM: Multi-Agent Motion Forecasting as Language Modeling," en Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2023.

[2] H. Zheng, Z. Yan, y C. Wu, "Multi-Agent Path Finding for Mixed Autonomy Traffic Coordination," en 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2024.



Proposed pipeline



PBS and Behavior prediction

Hierarchical Coordination Logic

- **High-Level (PBS):** Resolves multi-agent conflicts by constructing a Priority DAG.
- **Mid-Level (Multi-phase A):*** Implements sequential decision-making with two behaviors:
 - **Phase 1 (Aggressive):** 70m planning horizon biased toward forward progress.
 - **Phase 2 (Safety Fallback):** 30m horizon focused on deceleration, idling, or evasive maneuvers.
- **Low-Level (A*):** Computes individual trajectories constrained by the plans of higher-priority agents.
- **Primitive Level (Action Space):** Discrete search across six motion primitives:
 - *Accelerate, Idle, Decelerate, Hard Brake, Lane Change Left, and Lane Change Right.*

Kinematics of robots

- **Physics:** Implements a Kinematic Bicycle Model with rear-axle integration to match simulator dynamics.
- **Precision:** Achieves **high-fidelity synchronization** between the planner and simulator by utilizing identical governing equations to eliminate drift.
- **Collision Checking:** High-resolution (15Hz) validation using the Separating Axis Theorem (SAT) to prevent box overlapping and "tunneling" between frames.

Motion-LM as a predictor

Motion-LM Architecture

- **Autoregressive Transformer Decoder** treating motion as a language modeling task.
- **13x13 Verlet-based discretization grid (169 tokens)** for kinematic consistency.
 - a. **Acceleration-based tokenization** utilizing a greedy search algorithm over the trajectory dataset.
- Agent encoder processing **2-second historical** windows for temporal context.
- Utilizes a **ladder mask** to enforce **temporal causality**, ensuring future tokens do not influence past predictions.

Training and Data

- Dataset consisting of **1000 episodes of human interactions (IDM+MOBIL)** in ramp-merging scenarios.
- Captured **diverse behaviors**, including aggressive merging and chaotic longitudinal maneuvers.
- Implements a **smooth Gaussian bias in the loss function** to penalize near-target actions less severely than divergent ones.


Interactive Inference and Planned Clamping

- **Functions as a world model** to evaluate the safety of proposed CAV maneuvers.
- **Supports multiple rollouts** via parallel input sampling to handle uncertainty and diverse agent reactions.
- **Causal reasoning:** Predicts specific HDV responses (e.g., yielding vs. ignoring) to robot actions.

Implementation

Environment Infrastructure

- **Simulation core:** Customized `highway-env` library with a specialized `MergeInteractionEnv` subclass.
- **Verification:** 15Hz sub-step resolution within each planning step to eliminate **ghost collisions**.
- **Physics alignment:** Mathematical parity between the planner's Bicycle Model and the simulator's internal integrator, reducing state divergence to **36cm**.
- **Safety buffer:** Inclusion of a **15cm safety envelope**. Crashed vehicles are removed.




Baseline (IDM/MOBIL)

Preliminary results


Comparative Analysis: 10 Rollouts experiments with same seed for comparison

Metric	Baseline (IDM/MOBIL)	PBS-static HDV	MotionToken-MAPF
Success Rate	<u>0.58</u>	0.47	0.58
Vehicle Collisions	0.22	0.46	<u>0.32</u>
Time to Finish Task	7.69 sec	4.49 sec	<u>5.03 sec</u>
Avg. Cycle Latency	-	0.64 sec	<u>1.15 sec</u>

Demos



PBS-static HDV



MotionToken-MAPF

Conclusions and future work

Conclusions

- Successfully implemented a **token-based Transformer** (Motion-LM) for autoregressive trajectory forecasting.
- Developed a **hierarchical PBS framework with multi-phase A*** search, establishing a robust baseline for multi-agent coordination along with forecaster
- Created a **modular architecture that adapts to diverse driver behaviors** and scales effectively to varied traffic scenarios.
- Full results are pending (detailed in the project report), but preliminary forecasting performance is promising.

Future Work

- **Environmental generalization:** Transition to diverse environment requiring explicit map metadata integration, moving beyond fixed-coordinate awareness.
- **Uncertainty:** Incorporate multimodal rollouts to tune PBS priority constraints based on predictive confidence and probabilistic occupancy.
- **Robustness:** Further robustify the planning-to-execution pipeline to ensure maneuvers remain resilient against residual kinematic drift in real-world conditions.