

Master's Thesis Defense Announcement
Mechanical and Aerospace Engineering Department
University of Texas at Arlington

LLM-Driven Closed-Loop UAV Control with Obstacle-Aware Model Predictive Control

By: Halimcan Yasar

Thesis Advisor: Dr. Kamesh Subbarao

3:00 p.m., Thursday, August 6, 2026

WH, 413

Abstract

This thesis presents a closed-loop control architecture for uncrewed aerial vehicles (UAVs) in which a large language model (LLM) serves as a high-level decision module operating over a persistent, metric 3D world model.

Rather than generating low-level commands or open-loop plans, the LLM selects one parameterized maneuver per decision step from a small, verified library of flight primitives conditioned on a structured representation of the drone state, tracked object positions, and mission specification.

Translational motion is executed by a planar model predictive controller (MPC) with soft obstacle avoidance, using obstacle hypotheses provided by the LLM, so that safety-critical constraint handling is delegated to the optimization-based controller rather than the language model.

The architecture is evaluated in simulation using a PX4 X500 quadrotor with RGB-D sensing in ROS 2 and Gazebo across multiple LLM backends including Gemini-3.5, GPT-4o, GPT-5.2, and a locally deployed Qwen2.5-3B model.

Experiments spanning target approach, constrained navigation, yaw-aligned approach, and object search tasks indicate that constraining the LLM to symbolic maneuver selection over an explicit 3D world model yields more interpretable and more constraint-respecting behavior in the tested scenarios, while preserving flexible natural-language task specification. An MPC ablation study demonstrates that obstacle avoidance must be enforced at the planner level: with the MPC disabled, a correct obstacle declaration by the LLM is not turned into a safe trajectory and the direct-setpoint fallback contacts the obstacle, showing that a competent execution layer, and not obstacle declaration alone, is what yields safe motion.

A key emergent behavior is also observed in which the closed loop compensates for the small terminal offset introduced by the soft-constrained MPC formulation through a subsequent corrective maneuver selected by the LLM, validating the combined design without any explicit residual-compensation logic.

The findings establish that constraining the LLM to a verified action space and delegating safety enforcement to a deterministic planner are both necessary conditions for reliable LLM-driven UAV autonomy.