SMALL: Formal Methods for Safe, Efficient, and Transferable Learning-enabled Autonomy

Yiannis Kantaros, Dept. of Electrical and Systems Engineering, Washington University in St. Louis

Motivation: Learning-based techniques, such as deep reinforcement learning (DRL), neural model predictive control (NMPC), Large Language Models (LLMs) have been extensively used to synthesize high-level planners or low-level controllers for systems tasked with temporal and logical mission requirements (often specified using formal languages such as Linear Temporal Logic (LTL)).
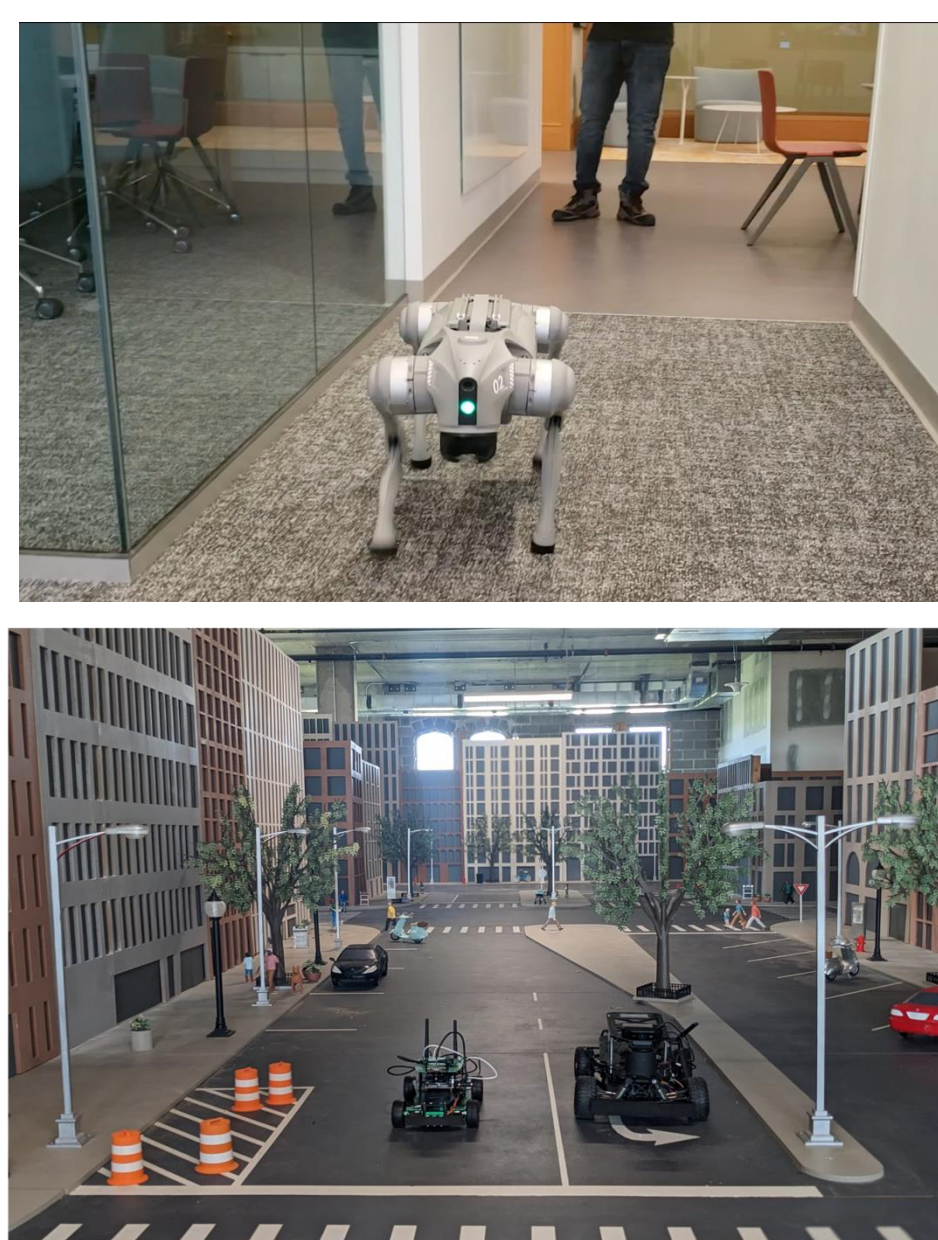
Key Challenges

- **Sample inefficiency** in learning effective controllers (design time).
- **Lack of safety/performance guarantees** of learned controllers, especially when deployed in new/unseen domains (test time).
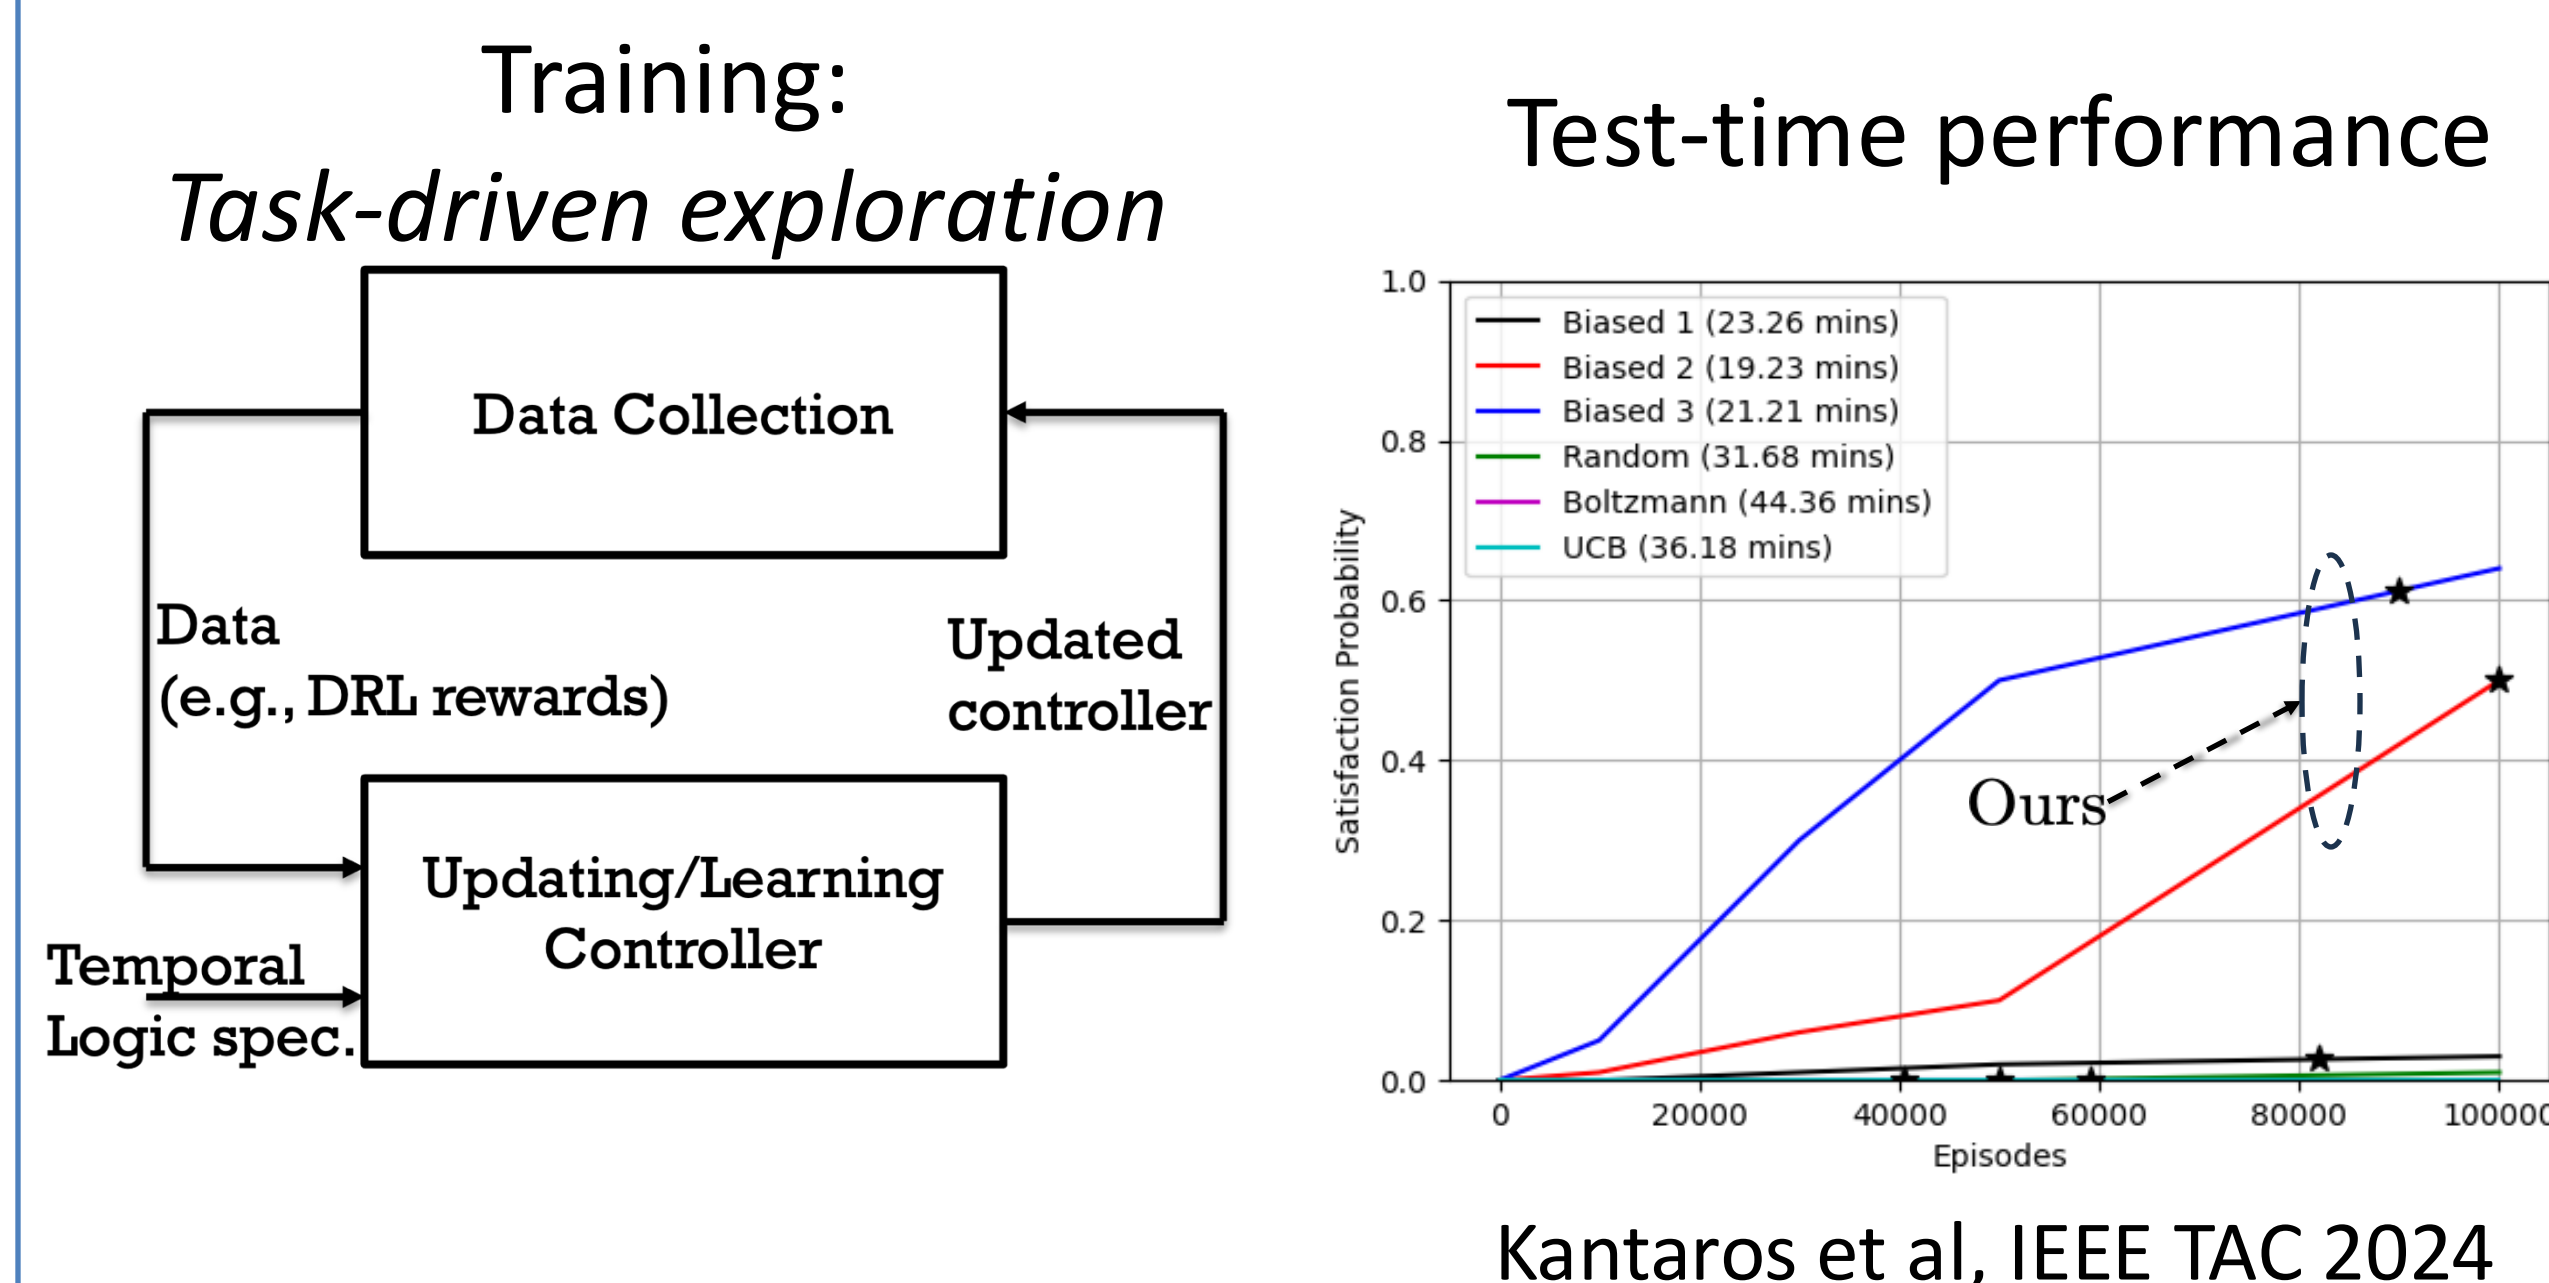
Technical approach (key ideas):

- Data-collection methods actively improving control performance
- Uncertainty quantification methods to monitor (calibrated) confidence in unseen domains → probabilistic assurance guarantees
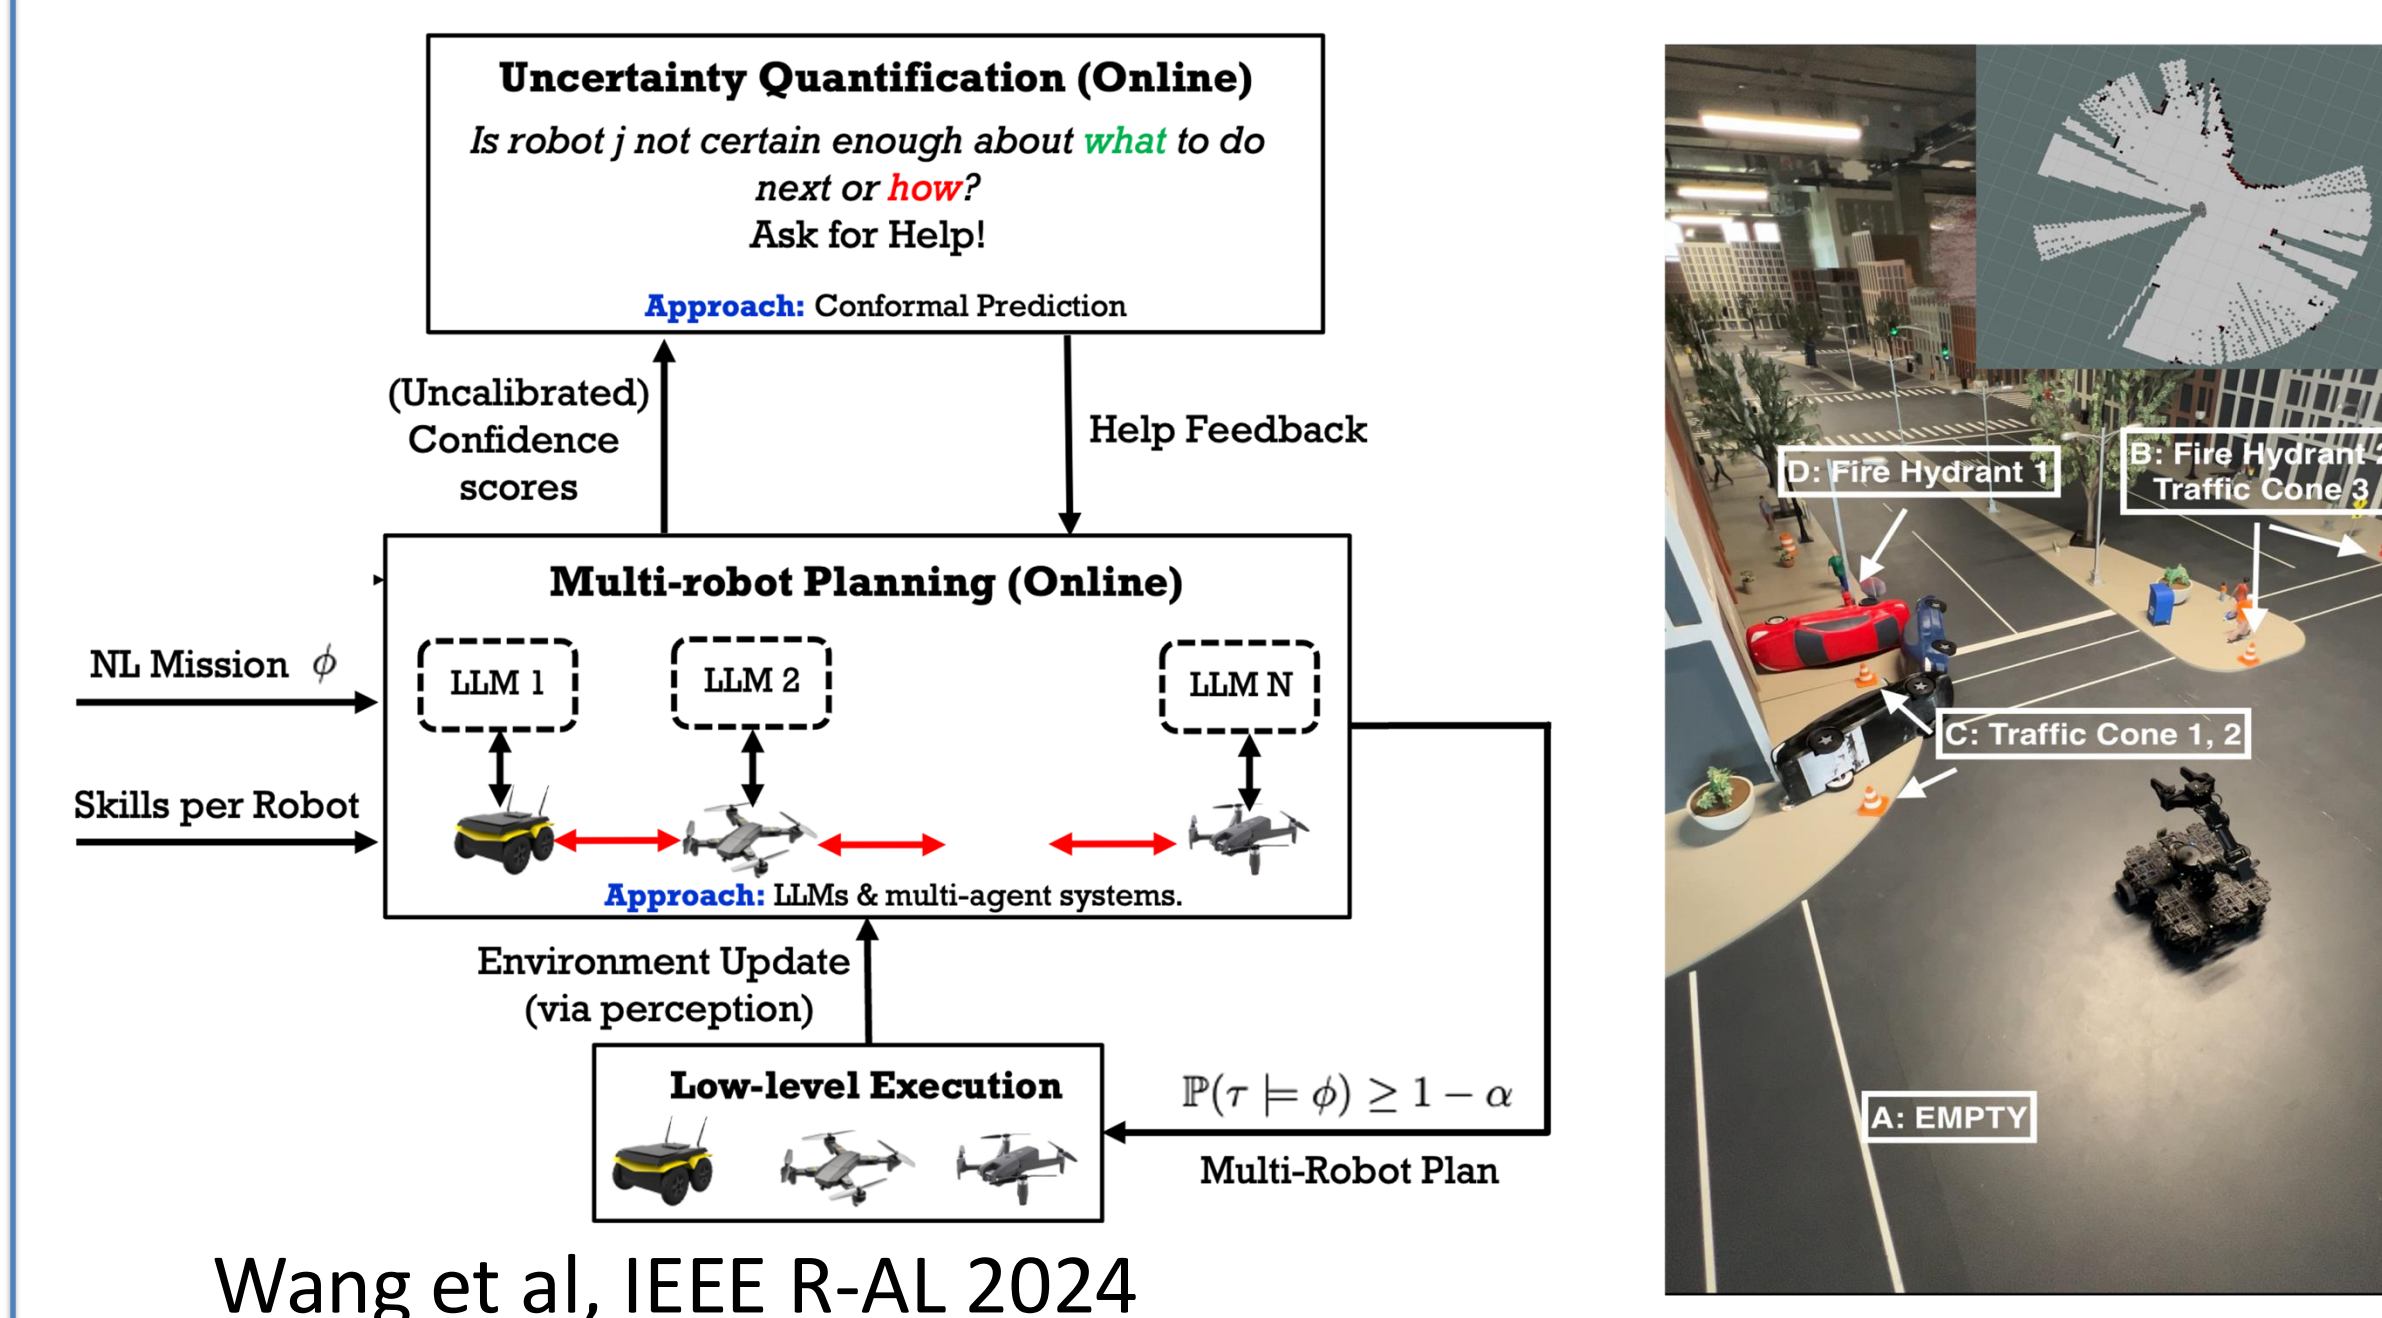
Scientific Impact

- Designing data-efficient learning-enabled controllers that can be safely transferred to new/unseen task and environmental domains.
- Evaluation using ground and aerial robot platforms on delivery and transportation applications. Other CPS applications include search-and-rescue, disaster relief, or exploration.
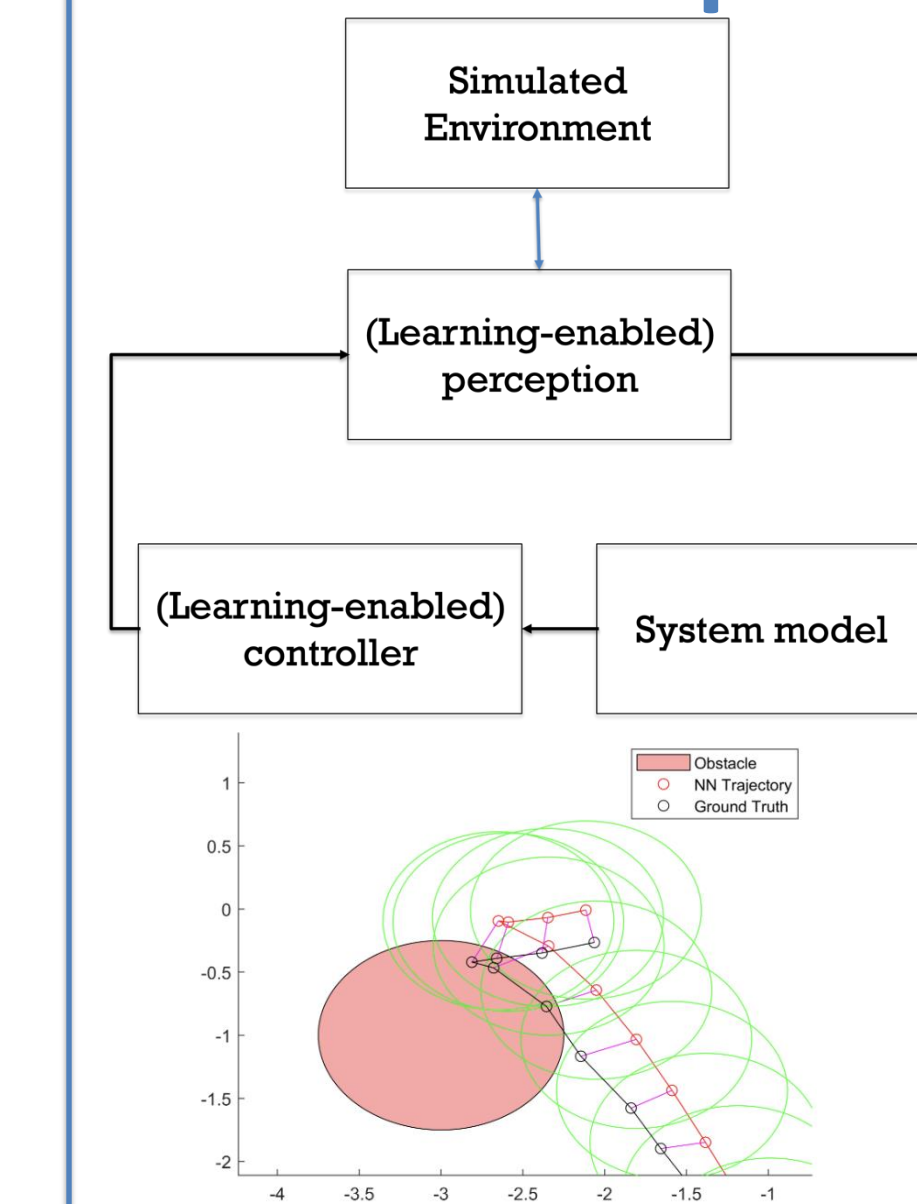
Contribution A: Sample-efficient DRL-based Control

Training:
Task-driven exploration

Test-time performance

Biased 1 (23.26 mins)
Biased 2 (19.23 mins)
Biased 3 (21.21 mins)
Random (31.68 mins)
Boltzmann (44.36 mins)
UCB (36.18 mins)

Ours

Kantaros et al, IEEE TAC 2024

Contribution B: Safe LLM-based Multi-Robot Planning

Uncertainty Quantification (Online)
Is robot j not certain enough about what to do next or how?
Ask for Help!

Approach: Conformal Prediction

(Uncalibrated) Confidence scores

Help Feedback

Multi-robot Planning (Online)

NL Mission ϕ

Skills per Robot

LLM 1 LLM 2 LLM N

Approach: LLMs & multi-agent systems.

Environment Update (via perception)

Low-level Execution

$\mathbb{P}(\tau \models \phi) \geq 1 - \alpha$

Multi-Robot Plan

Wang et al, IEEE R-AL 2024

Work-in-Progress: Safety Monitors Against Sim2Real Gap

Given: a learning-enabled controller, designed using simulation data and a system model.

How to monitor its safety when deployed in the real world on a physical robot platform?

Simulated Environment

(Learning-enabled) perception

(Learning-enabled) controller

System model

Broader Impacts:

- Enable safe deployment of CPS with learning-enabled decision-making algorithms in unseen domains with user-specified assurance guarantees.
- Applications: transportation, manufacturing, delivery.
- Design a new graduate course: *Learning and Planning in Robotics*
- Research opportunities to UG, MS, PhD students and the WashU Robotics Club.
- Summer research internships to K-12 students and teachers (Summer 2023, 2025)