2024 NSF CYBER-PHYSICAL SYSTEMS PRINCIPAL INVESTIGATORS' MEETING

CPS: Medium: Robust Learning for Perception-Based Autonomous Systems

PI: Nikolai Matni, Co-PIs: George Pappas & Kostas Daniilidis, University of Pennsylvania

https://perception-based-autonomy.github.io

Penn Engineering | GRASP Laboratory
General Robotics, Automation, Sensing & Perception Lab

How should we design perception-based control pipelines for these very different systems?

Bomb defusing robot
safety >> performance

Drone racing
performance >> safety

Where on the system architecture continuum should we be?

modular
+ "easily" verified/validated, easier to integrate priors, "better engineering"
- rigid interfaces, - limited end-to-end optimization

co-design
+ rich interfaces
+ mid-level reps give robustness
+ end-to-end opt via module co-design

end-to-end
high performance w/ end-to-end opt, adaptable +
lack safety guarantees, harder to integrate priors – opaque/hard to troubleshoot,

Challenge

- What does shift from simple to rich perceptual sensors mean for design of safety-critical CPS?

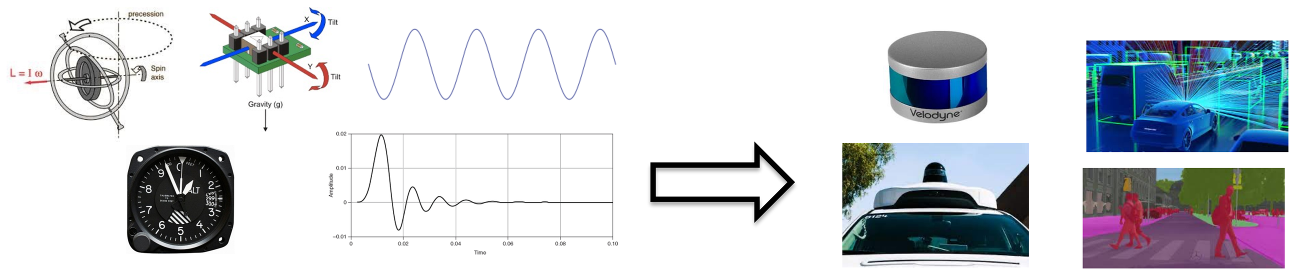

- Machine and deep learning are key, but are current methods suitable for safety-critical control loops?

Identifying & navigating tradeoffs

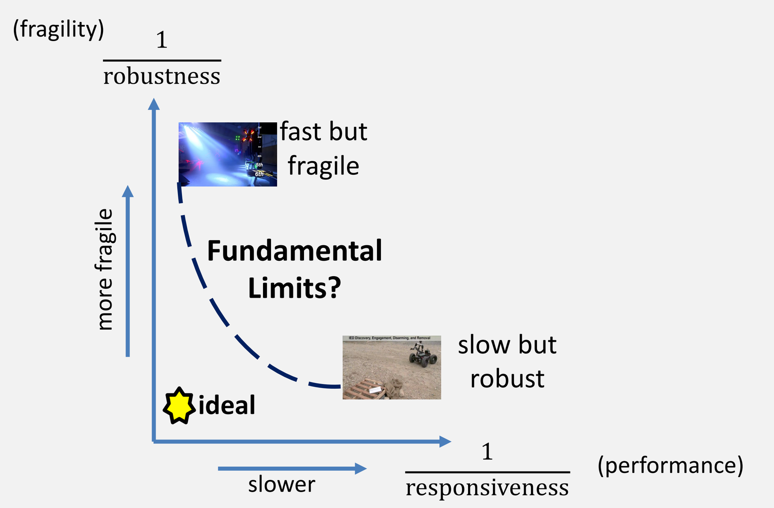

Scientific Impact

- Principled integration of learning-based perception, prediction, planning, & control
- Modular design pipeline with meaningful interfaces that allow for end-to-end metrics to be optimized
- Fundamental limits of robustness/performance in learning and perception-based cyber-physical systems

Solution

Goal: a modular approach with meaningful interfaces that allow for **end-to-end** metrics to be optimized and/or

Hypothesis: interfaces defined in terms of **uncertainty representation & propagation** between perception, prediction, planning, & control based on **mid-level representations**

Uncertainty and risk-aware perception-based control

Learning Robust Output-Feedback Control Barrier Functions from Safe Expert Demonstrations

Data-driven risk verification (CDC 2021, HSCC 2021)

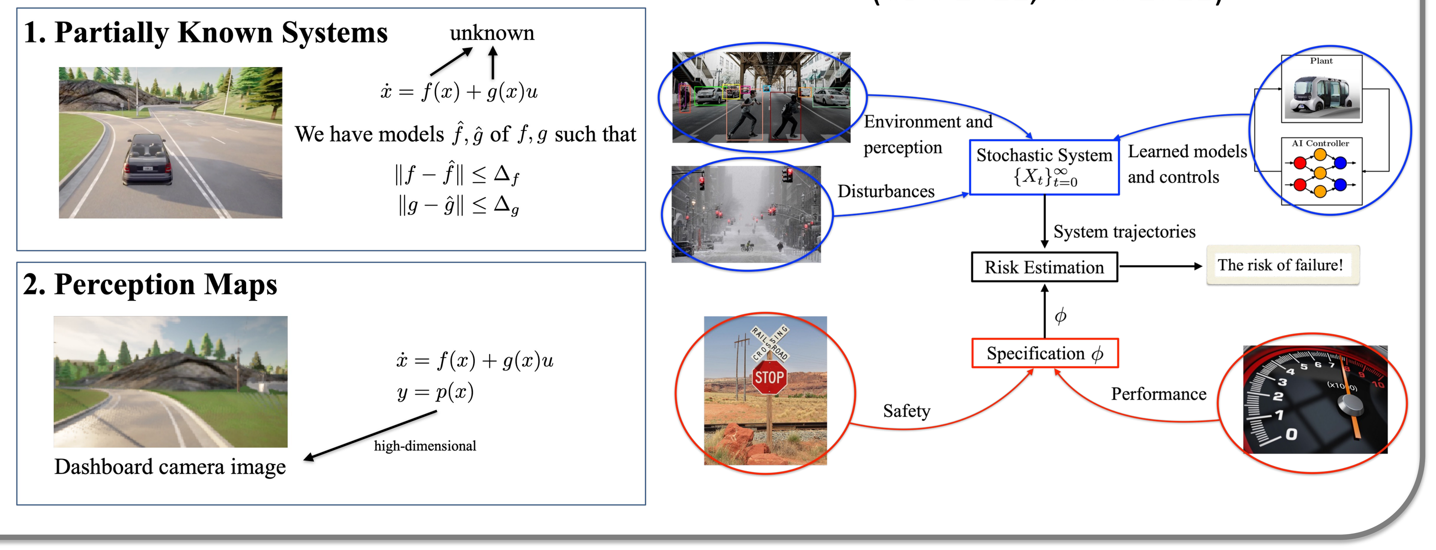

From Semantic SLAM to Semantic Navigation

Uncertainty-driven Planner for Exploration and Navigation (ICRA 2022)

Active Semantic Goal Navigation (ICLR 2022)

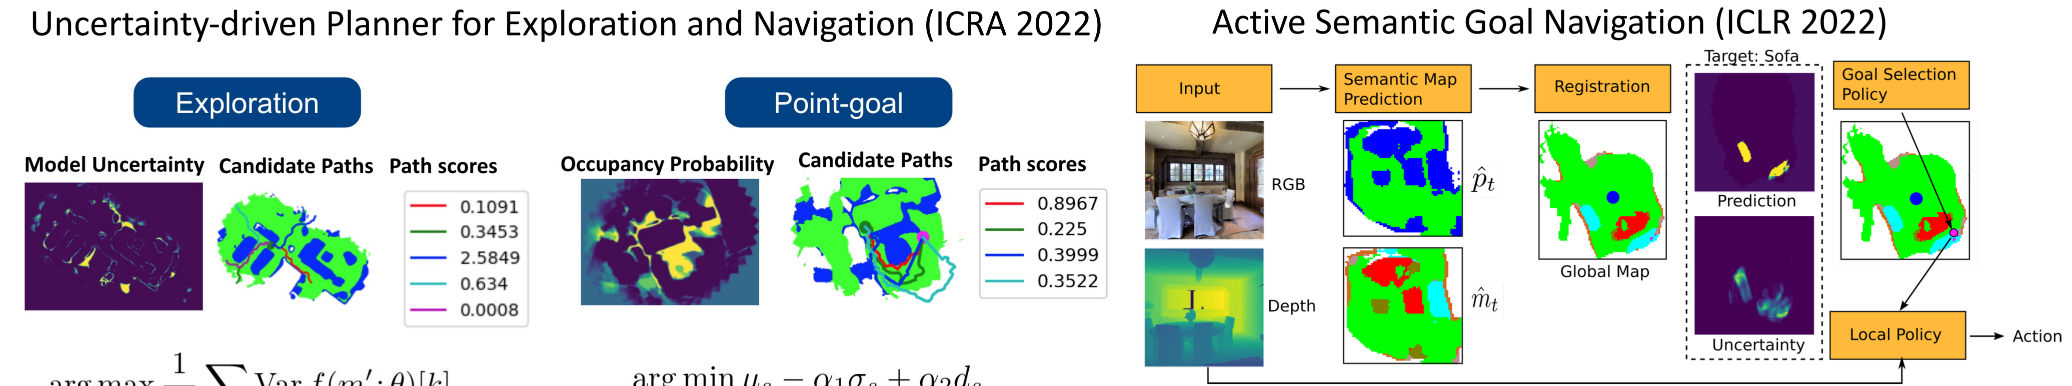

$$\arg\max_{s \in S} \frac{1}{|s|} \sum_{k \in s} \mathrm{Var}\, f(m_t^i; \theta)[k]$$

Score each path s by average model uncertainty at path locations k

$$\arg\min_{s} \mu_s - \alpha_1 \sigma_s + \alpha_2 d_s$$

Select paths using an Upper Confidence Bound on path traversability weighted by path length

1. Predicts labels in unseen semantic map regions
2. Actively learns semantic map approach
3. Uses confidence bound objectives for active long-term goal selection

Societal Impact

autonomous vehicles

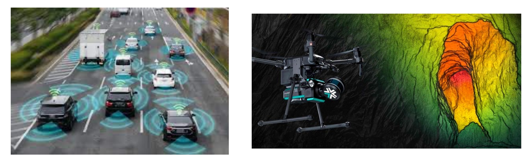

search & rescue

disaster & humanitarian relief

Fukushima "robot in hell"

Zipline drone

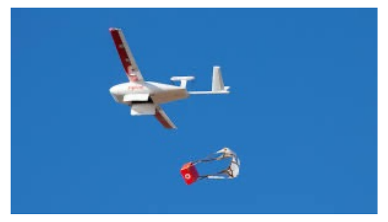

Education and Outreach

- Developed new publicly available courses on Learning for Dynamics and Control at Penn
- Ran IEEE CDC 2021 workshop on robust deep learning-based control
- Industry collaborations with Google Robotics

Broader Impact

- **Broadening participation:** Matni & Daniilidis co-advise URM students on this project (2 woman, 1 BIPOC);
- **Self-driving vehicles:** Approximately 38,000 people die every year in crashes on U.S roadways. Studies suggest self-driving vehicles can reduce this number by up to 34%: robust perception-based autonomy is a key enabling technology for self-driving vehicles.

Select bibliography

1. G. Georgakis, B. Bucher, A. Arapin, K. Schmeckpeper, N. Matni, and K. Daniilidis, Uncertainty-driven Planner for Exploration and Navigation, International Conference on Robotics and Automation (ICRA), 2022.
2. G. Georgakis*, B. Bucher*, K. Schmeckpeper, S. Singh, K. Daniilidis , Learning to Map for Active Semantic Goal Navigation, International Conference on Learning Representations (ICLR), 2022.
3. L. Lindemann, A. Robey, L. Jiang, S. Tu, and N Matni, Learning Robust Output Control Barrier Functions from Safe Expert Demonstrations, OJ, 2021.
4. L. Lindemann, L. Jiang, N. Matni, and G. J. Pappas Risk of Stochastic Systems for Temporal Logic Specifications, Submitted, 2022.
5. L. Lindemann, A. Rodionova, and G. J. Pappas, "Temporal Robustness of Stochastic Signals", in Proceedings of the Conference on Hybrid Systems: Computation and Control, 2022.
6. Matthew Cleaveland, Lars Lindemann, Radoslav Ivanov, and George J. Pappas, "Risk Verification of Stochastic Systems with Neural Network Controllers", in Artificial Intelligence Journal, 2022.

2024 NSF CPS PI Meeting
03/20-03/21

Award ID#: 2038873