Enhancing Lattice-based Motion Planning with Introspective Learning and Reasoning

Mattias Tiger[a], David Bergström[a], Andreas Norrstig[a], Fredrik Heintz[a]

(a) Reasoning and Learning Lab, Department of Computer and Information Science, Linköping University, Linköping, Sweden

Abstract

Robust and safe robots operating outside of controlled environments require not only awareness of the robot's own actions, but also its current ability to execute those actions and how the execution is expected to vary. In motion planning, this translates to well-calibrated safety-margins for collision checking. Safety is expressed as probabilistic safety in the context of motion under uncertainty.

In this work we are concerned with introspective learning and reasoning about controller performance over time. Normal controller execution of the different actions is learned using machine learning techniques with explicit uncertainty quantification, for safe usage in safety-critical applications. Monitoring for anomalies verifies, at runtime, that the learned models remain valid for collision checking use in the motion planner.

While a motion plan is being executed, the presented approach allows for explicit awareness of controller performance under normal circumstances, and detection of unexpected and potentially dangerous behavior in abnormal circumstances. Having learned a more accurate model of execution variability, the safety margins can be reduced while maintaining the same safety as before. This, in turn, also leads to shorter and more efficient plans. Evaluation is made on the nonlinear dynamics of a quadcopter in accurate 3D simulation, using system identified parameters.

Introduction

Uncontrolled Real-world Environments

One of the challenges in such environments is **varying uncertainty over time**, which can be decomposed into three factors:

- **Dynamic** perception **uncertainty** (e.g. occlusion, GPS shadowing) **Estimation/prediction** uncertainty may vary
- **Dynamic** action **stochasticity** (e.g. wind or actuator degradation) Action **execution** uncertainty may vary
- **Dynamic** environment **unknowns** (the environment evolves) (e.g. new personal transports) **Rules** and **distributions** may change

Challenges in Robotics / Our Approach

Challenges in Robotics	Our Approach
What is **normal** behavior?	**Learn models** of execution variations from action sequences
Is the robot behaving **normally**?	**Monitor executions** with respect to models during deployment
Safe, but not task effective?	Use **models** with **explicit uncertainty** quantification for tight safety-bounds
Are learned models **safe** to use?	**Monitor models** with respect to executions during deployment

Our Contributions

We present a general approach for enhancing lattice-based motion planning methods with (1) learning models of normal motion primitive execution, (2) using the learned models to improve collision checking effectiveness during planning and to (3) efficiently monitor the motion primitive executions.

Since both collision checking and abnormality detection are safety-critical the learning is performed using machine learning techniques with explicit uncertainty quantification from **probabilistic machine learning**. The monitoring for abnormalities also verifies at runtime that the learned models remain valid for collision checking use in the motion planner.

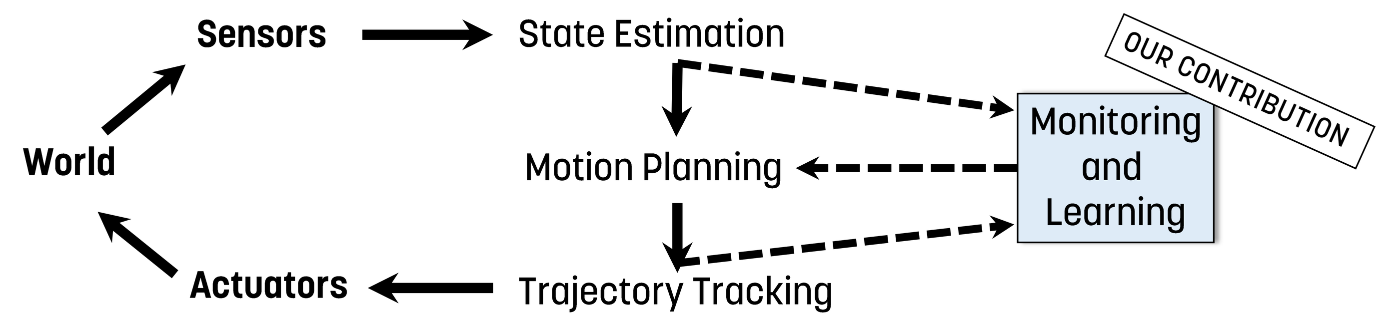

Method

The motion planner in [2] plans safe (collision free) trajectories with respect to both static and dynamic obstacles. That the plans are indeed safe to execute depends on how well-calibrated the safety margins are. A safety margin can be divided into three main parts, reflecting the uncertainties from the robot's perspective of:

- the state of the world
- the control execution
- the behavior of others

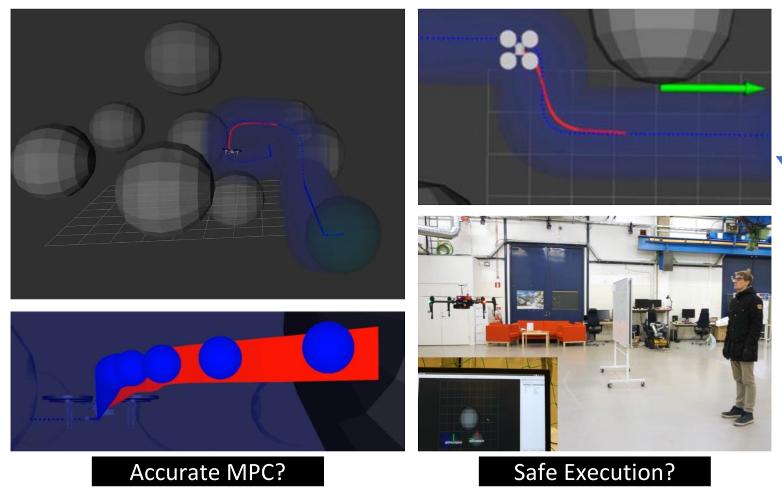

$$\mathcal{O}_{x_0(t)} = \mathcal{O}^*_{x_0(t)} * \mathcal{O}_{\text{safety}}(t)$$

Robot geometry

Safety margin

$$\mathcal{O}_A * \mathcal{O}_B = \{p_A + p_B \,|\, p_A \in \mathcal{O}_A,\, p_B \in \mathcal{O}_B\}$$

$$\mathcal{O}_{\text{safety}}(t) = \mathcal{O}_{\text{perception}}(t) * \mathcal{O}_{\text{control}}(t) * \mathcal{O}_{\text{others}}(t)$$

Lattice-based motion planning use graph search to find full trajectories made up of local trajectories (motion primitives). Motion primitives are optimized offline using numerical optimal control. We consider 104 different motion primitives in total. For every motion primitive a_i, we collect all trajectories consisting of three motion primitives, where a_i is the middle one. Using methodology from probabilistic machine learning, we estimate a unimodal model that spans the variability of a_i.

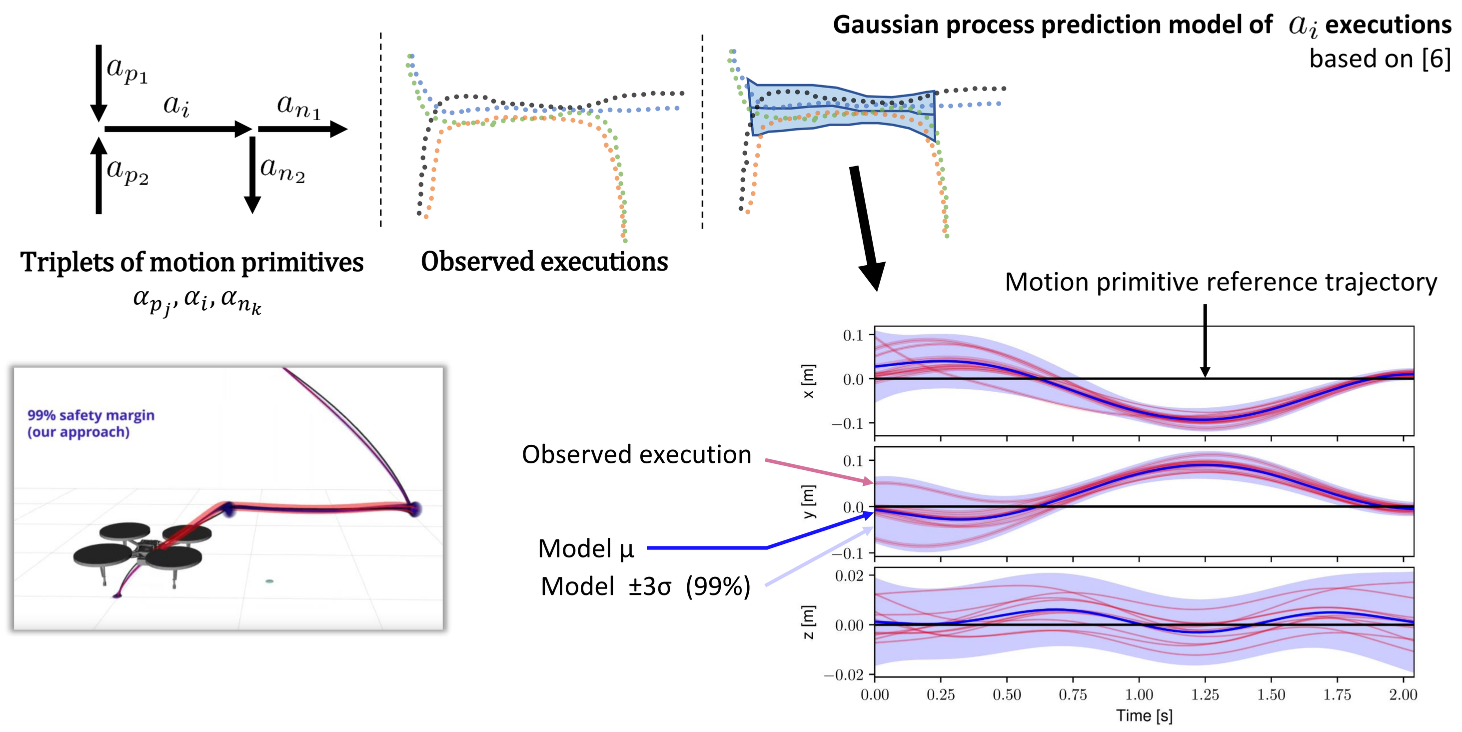

Results

The standard approach in the literature is to use a sphere for $\mathcal{O}_{\text{safety}}(t)$, capturing the aggregation of $\mathcal{O}_{\text{perception}}(t)$, $\mathcal{O}_{\text{control}}(t)$, $\mathcal{O}_{\text{others}}(t)$, with a single radius parameter invariant to time. We use this baseline, R, and construct two more with tighter bounds, R_i (different for each primitive) and $R_i(\tau)$ (also varying over time). For a meaningful comparison with our proposed variability model, we probabilistically ground the baselines by assuming a Gaussian likelihood over the observations centered on the primitive reference trajectory (i.e. relative translation is zero).

Safety margin components are defined as centered p-Probability Regions (p-PR).
Abnormality detection: Model the rate of leaving the p-PR as a random variable. Monitor that failure rate < p is likely.

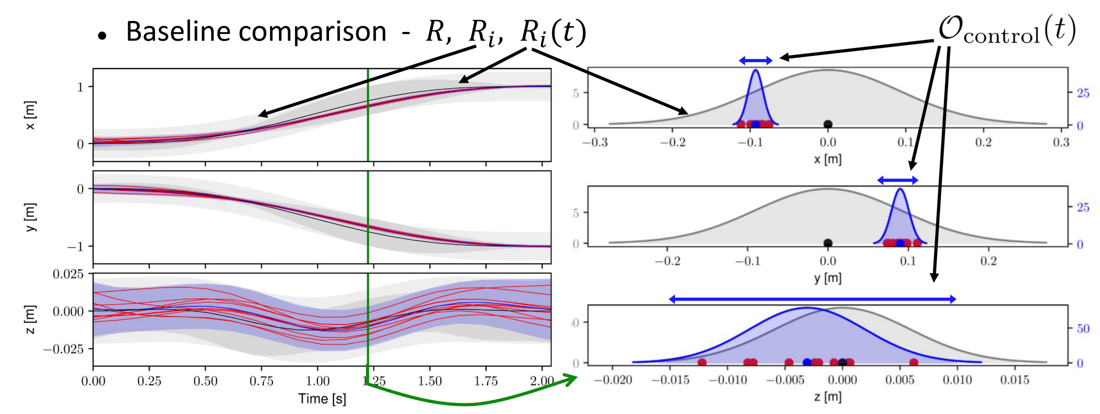

	RMSE (*meters*)
Observed Executions vs Motion Primitive	0.149 ± 0.053
Observed Executions vs $\mu_i(t)$	**0.055** ± 0.043
Motion Primitive vs $\mu_i(t)$	0.143 ± 0.052

		Average normalized area (m^2)
R	(same for all primitives)	1.000
R_i	(individual for every primitive)	0.646
$R_i(t)$	(individual and time dependent)	0.329
$p(x\|t, \boldsymbol{X}^i),\ \forall i$		**0.109**

Conclusions / Summary

With increased autonomy of cyber-physical systems, the need for integrated introspection capabilities is of growing importance. Such capabilities allow a robot to self-monitor and to react to unexpected changes to circumstances in the environment. This is paramount if robots are supposed to operate safely in unstructured, dynamic and complex environments.

We present an integrated approach for learning and monitoring the execution of motion actions, motion primitives, within the lattice-based motion planning paradigm. Interesting future works is robust online learning of motion primitive execution models, as well as learning such models for different situations e.g. windy and non-windy conditions for multi-modal operations.

System Perspective

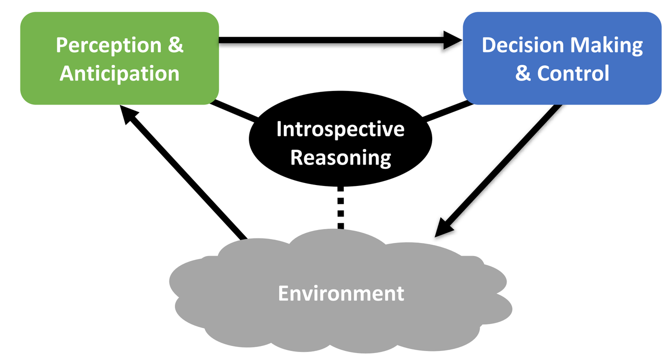

Efficiency, effectiveness and safety are competing qualities, and in safety-critical applications the required degree of safety makes it very challenging to reach useful levels of efficiency and effectiveness. To this end we investigate a holistic perspective on agent motion in complex and dynamic environments. The perception-action loop of an intelligent agent can be deconstructed into different components in several ways. In our long-term endeavour we take a holistic perspective of motion and motion patterns - mostly in terms of trajectories - and have made contributions within Perception & Anticipation, Decision Making & Control and Runtime Verification. An important contribution is an integrated whole, made possible by compatible mathematical languages and models from logic, probability theory and machine learning. We use extensive simulation as a valuable - and necessary - complement to field robotics experimentation. Below is an overview of our AI-Robotics stack in ROS (Robot Operating System).

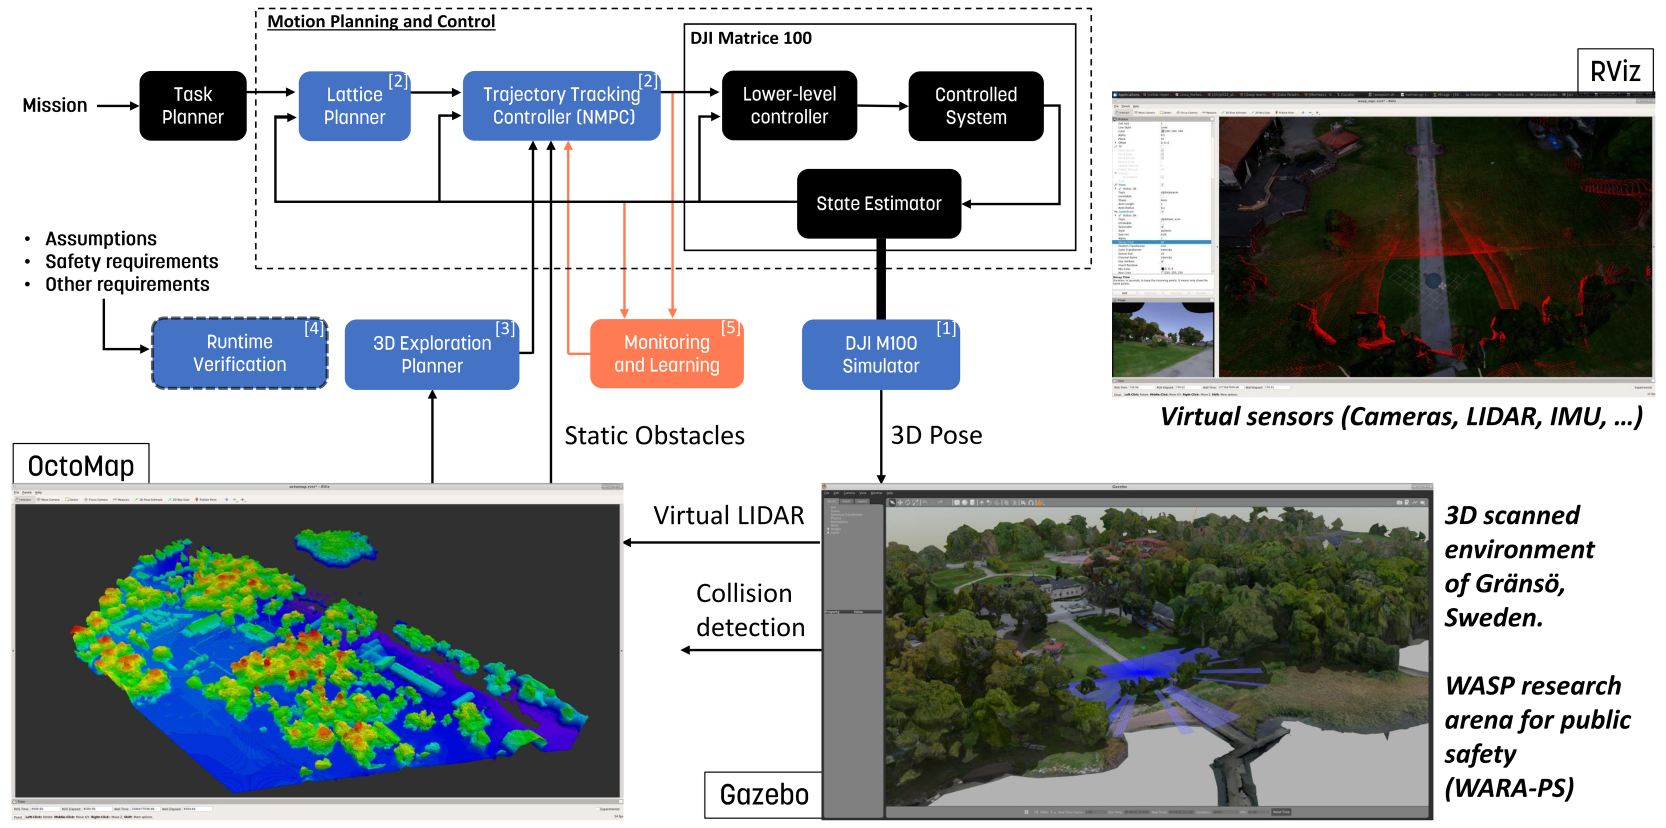

Video Paper More

References

[1] O. Andersson, O. Ljungqvist, M. Tiger. *Towards Real-Time Integrated Motion Planning and Collision Avoidance for the Matrice 100 Quadcopter*. Technical report, Division of Artificial Intelligence and Integrated Computer Systems, and Division of Automatic Control, Linköping University, Oct. 2018.

[2] O. Andersson, O. Ljungqvist, M. Tiger, D. Axehill, F. Heintz. *Receding-Horizon Lattice-Based Motion Planning with Dynamic Obstacle Avoidance*, IEEE Conference on Decision and Control, 2018.

[3] M. Selin, M. Tiger, D. Duberg, F. Heintz, P. Jensfelt. *Efficient autonomous exploration planning of large-scale 3D environments*. IEEE Robotics and Automation Letters, 2019.

[4] M. Tiger, and F. Heintz. *Incremental reasoning in probabilistic signal temporal logic*. International Journal of Approximate Reasoning 119 (2020): 325-352.

[5] M. Tiger, D. Bergström, A. Norrstig, F. Heintz. *Enhancing Lattice-Based Motion Planning With Introspective Learning and Reasoning*. IEEE Robotics and Automation Letters, 2021.

[6] M. Tiger, F. Heintz. *Online sparse Gaussian process regression for trajectory modeling*. IEEE International Conference on Information Fusion, 2015.