ETH zürich | UCL | Imperial College London | University of Cambridge | Cornell University | TU Delft | Chalmers University of Technology

PAPER

A Unifying Variational Framework for Gaussian Process Motion Planning

Lucas Cosier Rares Iordan* Sicelukwanda Zwane Giovanni Franzese*
James T. Wilson Marc Peter Deisenroth Alexander Terenin Yasemin Bekiroglu

WEBSITE

Introduction

We introduce a framework based on **variational inference for motion planning in robotics**, built using sparse Gaussian processes and pathwise conditioning. We generate a distribution of viable motion paths for the robot, taking into account self and object collisions. This is done by incorporating these constraints into the evidence lower bound used by the **variational approximation**. We then use **pathwise conditioning** to sample trajectories that can be evaluated and differentiated at arbitrary time points without further interpolation. **The framework unifies optimization-based and probabilistic-inference-based planners.**

Framework

- Unifies optimization-based and probabilistic-inference-based motion planning via variational inference
- Inducing points act as waypoints which control locations through which robot moves

Features

- Supports equality-based constraints (such as goal states), inequality-based constraints (such as joint limits), and soft constraints (such as collision avoidance)
- Simple and straightforward implementation
- Uncertainty: trajectory samples, error bars

ELBO Optimization

$$\frac{1}{2}\mathbb{E}_{q(\mathbf{u})}\mathbb{E}_{p(\mathbf{f}|\mathbf{u})}\left(\underbrace{\|\mathbf{h}_\epsilon(\mathrm{sdf}_s(\mathbf{k}_{\mathrm{fwd}}(\sigma(\mathbf{f}))))\|^2_{\mathbf{\Sigma}_{\mathrm{obs}}}}_{\text{collision term}} + \underbrace{\|\mathbf{c}(\sigma(\mathbf{f}))\|^2_{\mathbf{\Sigma}_{\mathbf{c}}}}_{\text{soft constraint terms}}\right) + D_{\mathrm{KL}}(q(\mathbf{u}') \,\|\, p(\mathbf{u}' \mid \mathbf{u}_{\mathrm{c}}))$$

- Pathwise conditioning for (negative) ELBO minimization
- Sample **u**, conditionally sample **f**, then evaluation collision likelihood of motion paths
- Customizable likelihood: **accommodates most constraints supported by other planners**
- Works with general kernels: **does not rely on specialized computational techniques**, like those for kernels arising from stochastic differential equations used in prior work

Pathwise Conditioning for Motion Planning

- Posterior consists of actual random trajectories: **these can be sampled first, then evaluated at arbitrary time points later,** without additional stochasticity or interpolation
- Motion plans **can be conditioned to go through arbitrary locations at arbitrary time points**

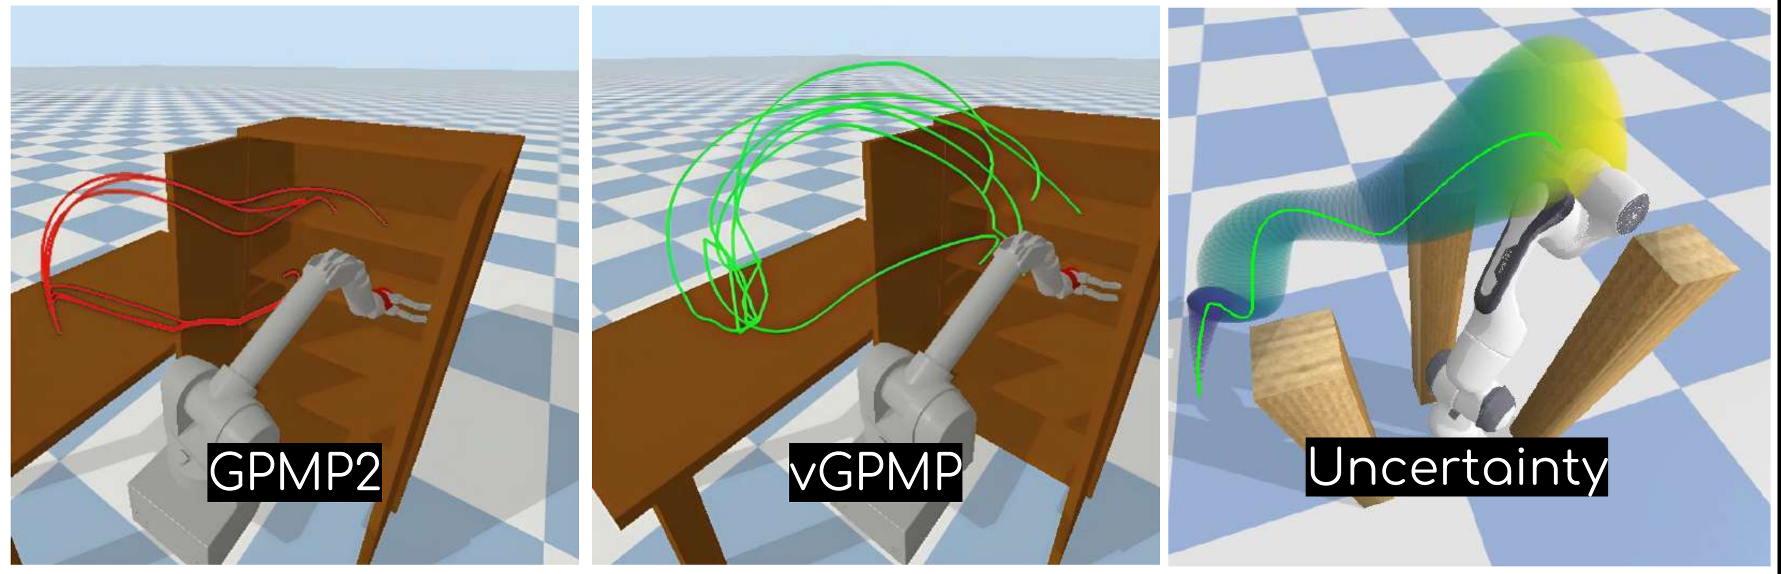

Figure 1: vGPMP (ours, left) and GPMP2 motion plans (middle), illustration of uncertainty (right)

- Path smoothness properties controlled by kernel

Grasping Alignment

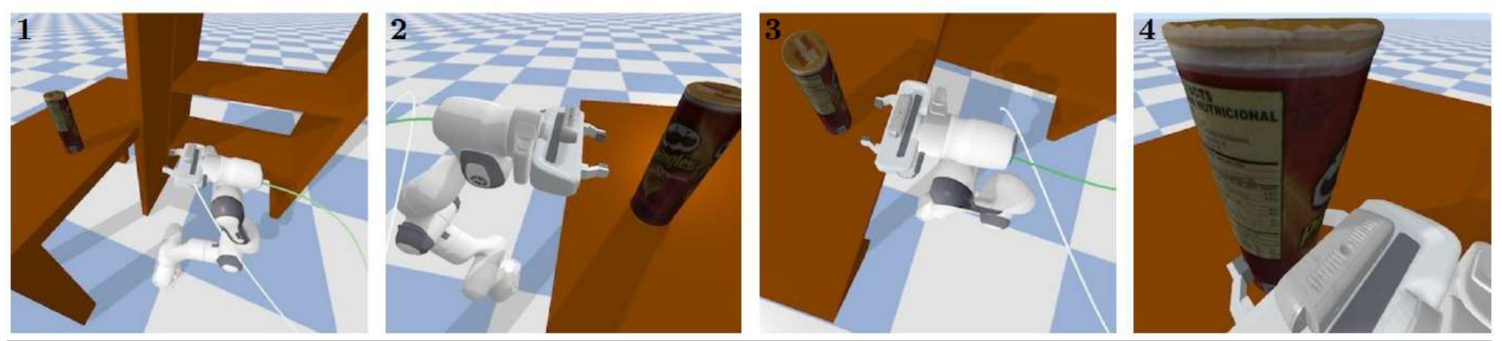

Figure 2: Randomly sampled motion plan trajectory for grasping a can

Sim-to-Real Transfer

Figure 3: Randomly sampled motion plan trajectory executed on a real robot

Main References

1. M. Mukadam, X. Yan, and B. Boots. Gaussian process motion planning. *In International Conference on Robotics and Automation, 2016.*
2. J. T. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. P. Deisenroth. Efficiently Sampling Functions from Gaussian Process Posteriors. *In International Conference on Machine Learning, 2020.*
3. J. T. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. P. Deisenroth. Pathwise Conditioning of Gaussian Processes. *Journal of Machine Learning Research, 2021.*

Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain.

Correspondence to: lcosier@ethz.ch