
A Unifying Variational Framework for Gaussian Process Motion Planning

Lucas Cosier^{*1,2} Rares Iordan^{*1} Sicelukwanda Zwane¹ Giovanni Franzese⁶
James T. Wilson³ Marc Peter Deisenroth¹ Alexander Terenin^{3,4,5} Yasemin Bekiroglu^{1,7}
¹University College London ²ETH Zürich ³Imperial College London ⁴University of Cambridge
⁵Cornell University ⁶TU Delft ⁷Chalmers University of Technology

Abstract

To control how a robot moves, motion planning algorithms must compute paths in high-dimensional state spaces while accounting for physical constraints related to motors and joints, generating smooth and stable motions, avoiding obstacles, and preventing collisions. A motion planning algorithm must therefore balance competing demands, and should ideally incorporate uncertainty to handle noise, model errors, and facilitate deployment in complex environments. To address these issues, we introduce a framework for robot motion planning based on variational Gaussian processes, which unifies and generalizes various probabilistic-inference-based motion planning algorithms, and connects them with optimization-based planners. Our framework provides a principled and flexible way to incorporate equality-based, inequality-based, and soft motion-planning constraints during end-to-end training, is straightforward to implement, and provides both interval-based and Monte-Carlo-based uncertainty estimates. We conduct experiments using different environments and robots, comparing against baseline approaches based on the feasibility of the planned paths, and obstacle avoidance quality. Results show that our proposed approach yields a good balance between success rates and path quality.

1 Introduction

Motion planning refers to the process by which a robot finds a path of motion from a start to a goal state. This path should be collision-free, satisfy constraints such as joint or torque limits, maintain a desired end-effector orientation, and fulfill other task-specific requirements. The path should also be smooth, avoiding abrupt or otherwise sharp motions. To compute the path, an algorithm needs to be able to handle potentially high-dimensional robot configuration spaces. This may involve uncertainty, due to inherent sensor noise, external disturbances, and other effects. The algorithm also needs to cope with complex environment dynamics.

Modern motion planning algorithms handle these challenges using different mathematical principles. *Sampling-based* planners, such as rapidly-exploring random trees (Kuffner and LaValle, 2000; Cheng and LaValle, 2002; Karaman and Frazzoli, 2011; Berenson et al., 2011), work by exploring the space of possible trajectories via stochastic processes, such as space-filling trees. *Optimization-based* planners (Ratliff et al., 2009; Schulman et al., 2013, 2014; Zucker et al., 2013; Kalakrishnan et al., 2011; Mukadam et al., 2016; Lirui Wang, 2020) calculate a motion plan by solving a constrained optimization problem defined on the space of possible paths. *Probabilistic-inference-based* planners (Dong et al., 2016; Attias, 2003; Toussaint and Storkey, 2006; Toussaint and Goerick, 2010; Yu and Chen, 2023; Mukadam et al., 2018) reinterpret motion planning as a Bayesian inference problem over spaces of paths, enabling motion plans to be computed using inference algorithms.

Once a path is computed, one must ensure it can be executed successfully on the physical robot. Using probabilistic representations to handle uncertainty is

becoming increasingly prominent in a number of related areas, such as imitation learning. For instance, *probabilistic motion primitives* (Paraschos et al., 2013) parameterize Gaussian trajectory distributions to learn from expert demonstrations, and provide related functionality such as blending trajectories together. Ewerton et al. (2020) extend this by adding a mixture model to handle distinct candidate trajectories, with uncertainty around each trajectory. Urain et al. (2022) integrate expert demonstrations to create stochastic motion planning primitives using energy-based-models. These developments highlight the usefulness of probabilistic representations to address uncertainty in related robotics problems.

In this work, we study probabilistic-inference-based motion planning using Gaussian processes. We (i) introduce a unifying *variational Gaussian process motion planning (vGPMP)* framework, which represents the motion planning problem as a *variational inference problem*—an optimization problem on a suitable space of probability distributions. This problem is parameterized by a set of *waypoints*, which describe locations through which the motion plans should go through, on average: these waypoints can be optimized to find motion plans for a given task. Furthermore, we (ii) show, using this variational-inference-based viewpoint, that Gaussian process motion planners are stochastic extensions of optimization-based planners. As a result, Gaussian process motion planners implemented in this manner share the same capabilities as optimization-based planners, and additionally provide uncertainty estimates due to their stochastic construction, which can be propagated downstream.

Our framework makes it possible to directly incorporate both hard and soft constraints, including joint limits and collision constraints, into the motion planning objective as well as task-independent properties, such as smoothness. We enable both interval-based and Monte-Carlo-sampling-based quantification of uncertainty, each in a scalable manner. The resulting framework augments and simplifies preceding Gaussian-process-based motion planners (Dong et al., 2016; Huang et al., 2017), giving practitioners a straightforward-to-implement way to construct motion plans with uncertainty.

2 Motion Planning using Gaussian Processes

Gaussian processes (GPs) are random functions $f : \mathcal{X} \rightarrow \mathbb{R}$, where the output value of the function at any finite set of input variables $\mathbf{x} \in \mathcal{X}^n$ follows a Gaussian distribution. A Gaussian process is characterized by a mean function $\mu : \mathcal{X} \rightarrow \mathbb{R}$ and a positive semi-definite

kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, such that, if $f \sim \mathcal{GP}(\mu, k)$, then $\mathbf{f} = f(\mathbf{x}) \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{K}_{\mathbf{x}\mathbf{x}})$ with mean $\boldsymbol{\mu} = \mu(\mathbf{x})$ and covariance $\mathbf{K}_{\mathbf{x}\mathbf{x}} = k(\mathbf{x}, \mathbf{x})$. Let (\mathbf{x}, \mathbf{y}) be the data, and to ease notation assume the prior mean is $\mu = 0$.

If we consider a Gaussian likelihood $p(\mathbf{y} | \mathbf{x}) = \mathcal{N}(f(\mathbf{x}), \sigma^2 \mathbf{I})$ where $\sigma^2 > 0$ is the noise variance, then the conditional distribution of the random function f given the training dataset (\mathbf{x}, \mathbf{y}) is also a Gaussian process. Letting $\mathbf{K}_{(\cdot)\mathbf{x}} = k(\cdot, \mathbf{x})$, we have $f | \mathbf{y} \sim \mathcal{GP}(\mu_{f|\mathbf{y}}, k_{f|\mathbf{y}})$ where

$$\mu_{f|\mathbf{y}} = \mathbf{K}_{(\cdot)\mathbf{x}}(\mathbf{K}_{\mathbf{x}\mathbf{x}} + \sigma^2 \mathbf{I}_n)^{-1} \mathbf{y} \quad (1)$$

$$k_{f|\mathbf{y}} = k(\cdot, \cdot') - \mathbf{K}_{(\cdot)\mathbf{x}}(\mathbf{K}_{\mathbf{x}\mathbf{x}} + \sigma^2 \mathbf{I})^{-1} \mathbf{K}_{\mathbf{x}(\cdot')}. \quad (2)$$

To train the GP, the hyperparameters of the kernel need to be learned. This is typically done by maximizing the marginal likelihood using standard gradient-based optimization methods (Rasmussen and Williams, 2006).

In the context of motion planning, the input space will always represent time, namely $\mathcal{X} = \mathcal{T} = \mathbb{R}$, and the output space will represent unconstrained joint values and possibly their derivatives, which we write as $\Theta = \mathbb{R}^d$. The GPs we consider will therefore be random functions $f : \mathcal{T} \rightarrow \Theta$, and the data will represent joint values and derivatives at a particular time.

2.1 Gaussian Process Motion Planning

By modeling the distribution of feasible trajectories as functions that map time to robot states, GPs provide a principled approach for formulating the motion planning problem through the lens of probabilistic inference. Gaussian processes were introduced to motion planning by Dong et al. (2016), who propose to choose a GP prior which is given by a *linear time-varying stochastic differential equation*, namely

$$d\mathbf{f}(t) = (\mathbf{A}(t)\mathbf{f}(t) + \mathbf{u}(t)) dt + \mathbf{F}(t) d\mathbf{w}(t) \quad (3)$$

where $\mathbf{A}(t)$ and $\mathbf{F}(t)$ are time-varying matrices, $\mathbf{u}(t)$ is the control input to the system, and $\mathbf{w}(t)$ is a Wiener process. The solution of this equation defines the GP prior. The primary motivations for this choice are that it (i) encodes spatial smoothness, and (ii) gives rise to structured covariance matrices, which are block-tridiagonal because the stochastic differential equation (SDE) possesses a Markov property. Dong et al. (2016) show that this accelerates computation and enables fast maximum a posteriori inference.

To complete the model, Dong et al. (2016) introduce a likelihood which encourages the robot to avoid collisions and constraint violations. Let $\mathbf{k}_{\text{fwd}} : \Theta \rightarrow \Psi$ be the forward kinematics function, where Ψ represents joint locations and orientations. Given the forward kinematics output, following Ratliff et al. (2009), they

approximate the robot’s body by a set of s spheres, and calculate a signed distance field for each sphere: let $\text{sdf}_s : \Psi \rightarrow \mathbb{R}^s$ be the function representing this calculation. Let $h_\varepsilon(x) = \max(-x + \varepsilon, 0)$ be the hinge loss, where $\varepsilon > 0$ is called the *safety distance* parameter. Let c be a linear penalty function for values exceeding a given joint limit threshold. By convention, h_ε and c act component-wise on vectors. Let $\|\mathbf{x}\|_{\Sigma}^2 = \mathbf{x}^T \Sigma^{-1} \mathbf{x}$ be the Mahalanobis norm, and let Σ_{obs} and Σ_c be diagonal matrices. With this notation, Dong et al. (2016) propose the likelihood $p(\mathbf{e} \mid \mathbf{f})$ proportional to

$$\exp\left(-\frac{1}{2}\left(\underbrace{\|\mathbf{h}_\varepsilon(\text{sdf}_s(\mathbf{k}_{\text{fwd}}(\mathbf{f})))\|_{\Sigma_{\text{obs}}}}_{\text{collision term}}^2 + \underbrace{\|c(\mathbf{f})\|_{\Sigma_c}}_{\text{soft constraint term}}^2\right)\right) \quad (4)$$

where \mathbf{e} represents the probability of *collision-free events with motion constraints*. From here, Dong et al. (2016) approximate the posterior distribution as follows: (i) apply maximum a posteriori inference, deriving efficient algorithms which use the SDE’s Markov property to recast the resulting optimization objective as inference in a sparse factor graph. Once such a trajectory is obtained, (ii) apply Laplace approximations to the posterior distribution, and (iii) interpolate joint configuration in-between the states at times $t_1, \dots, t_T \in \mathcal{T}$. At run-time, the output values of \mathbf{f} are clamped to enforce constraints. In total, the computations performed produce an optimal trajectory \mathbf{f}^* , and a Gaussian distribution representing trajectories around it that quantifies uncertainty.

3 Variational Gaussian Process Motion Planning

Computing motion plans using GPs involves balancing multiple competing factors, including obstacle avoidance, smoothness, and conformity to joint limits. These factors depend simultaneously on the motion planner’s hyperparameters, which must be tuned, and on the formulation of the objective itself. To provide practitioners with precise control over what the planner produces and enable them to modify it as needed to introduce other objectives for specific tasks, we develop a unifying mathematical and algorithmic formalism for Gaussian process motion planning.

3.1 Sparse Gaussian Processes as a Unifying Framework

From a probabilistic inference perspective, one of the challenges in Gaussian Process motion planning is the need to accommodate non-Gaussian likelihoods. This means that the motion planning posterior, unlike the prior, is not a GP, which means one cannot perform exact probabilistic inference. One must therefore use

numerical methods to extract means, covariances, and Monte Carlo samples from the posterior distribution.

Sparse Gaussian Processes (Hensman et al., 2013, 2015; Titsias, 2009) are a state-of-the-art method for applying GPs in settings with non-Gaussian likelihoods. They approximate the exact non-Gaussian posterior with the best Gaussian approximation, in the sense of Kullback–Leibler divergence. Originally proposed to tackle scalability issues which arise when working with large datasets, they provide a broad and flexible framework for designing approximate probabilistic inference schemes using GPs. We now apply these notions to design a unifying framework for motion planning.

We start with an *arbitrary* Gaussian process prior $f \sim \mathcal{GP}(\mu, k)$, generalizing the framework of Dong et al. (2016) beyond stochastic differential equations. Thus, for any collection of *times* $\mathbf{T} = [t_0, t_1, \dots, t_n]$, the random vector evaluated on this subset $\mathbf{f} = [\mathbf{f}(t_0), \mathbf{f}(t_1), \dots, \mathbf{f}(t_n)]$ follows a joint Gaussian distribution whose mean and covariance are determined by μ and k . Simplifying, suppose that we are also given a likelihood $p(\mathbf{e} \mid \mathbf{f})$, which incorporates constraints and collision avoidance, such as the one presented in Section 2.1—we will discuss likelihoods later in Section 3.2. These ingredients specify the exact posterior distribution, which is non-Gaussian.

To handle non-Gaussianity, following the sparse GP framework (Snelson and Ghahramani, 2006; Quiñero-Candela and Rasmussen, 2005; Titsias, 2009), we introduce the *variational approximation* $f \mid \mathbf{u}$, which is a GP conditioned on the event $f(\mathbf{z}) = \mathbf{u}$, where $\mathbf{z} \in \mathcal{T}^m$ are a set of locations called the *inducing points*, and \mathbf{u} are called the *inducing variables*. Let $p(\mathbf{u}) = \mathcal{N}(\mathbf{0}, \mathbf{K}_{\mathbf{z}\mathbf{z}})$ be the prior at the \mathbf{z} -locations. Further, let $q(\mathbf{u}) \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{u}}, \boldsymbol{\Sigma}_{\mathbf{u}})$ be the distribution of the \mathbf{u} -values, which we assume to be a free-form multivariate Gaussian whose mean and covariance $\boldsymbol{\mu}_{\mathbf{u}}$ and $\boldsymbol{\Sigma}_{\mathbf{u}}$ will be learned using optimization. The gives $f \mid \mathbf{u} \sim \mathcal{GP}(\mu_{f|\mathbf{u}}, k_{f|\mathbf{u}})$ with


$$\mu_{f|\mathbf{u}} = \mu(\cdot) + \mathbf{K}_{(\cdot)\mathbf{z}} \mathbf{K}_{\mathbf{z}\mathbf{z}}^{-1} (\mathbf{u} - \boldsymbol{\mu}_{\mathbf{u}}) \quad (5)$$

$$k_{f|\mathbf{u}} = k(\cdot, \cdot) - \mathbf{K}_{(\cdot)\mathbf{z}} \mathbf{K}_{\mathbf{z}\mathbf{z}}^{-1} \mathbf{K}_{\mathbf{z}(\cdot)} \quad (6)$$


To learn the parameters of $q(\mathbf{u})$, we minimize the Kullback–Leibler divergence between $f \mid \mathbf{u}$ and the true posterior. One can show (Matthews et al., 2016) that this is equivalent to maximizing the *evidence lower bound* (ELBO)

$$\mathbb{E}_{q(\mathbf{u})} \mathbb{E}_{p(f|\mathbf{u})} \log p(\mathbf{e} \mid \mathbf{f}) - D_{\text{KL}}(q(\mathbf{u}) \parallel p(\mathbf{u})) \quad (7)$$

which defines our general framework for probabilistic-inference-based motion planning. This objective includes an expectation: since sampling from the involved distributions is tractable, it can be optimized



(a) Sparse Gaussian process



(b) Motion plans via GP samples

Figure 1: Intuitive illustration of representing motion paths with GPs. (a) Sparse GP posterior for 1 joint with inducing variables $\mathbf{u} = [\mathbf{u}_c, \mathbf{u}']$ denoted by ‘o’, where $\mathbf{u}_c = [\boldsymbol{\theta}_0, \boldsymbol{\theta}_1]$ and \mathbf{u}' represent the $M = 4$ inducing locations $\mathbf{z} \in \mathcal{T}$ subject to the learning process, shown in-between $t = 0$ and $t = 1$. We interpret the inducing points as *waypoints* through which the random motion plans travel. (b) Candidate end-effector trajectories for the Franka Panda robot and associated uncertainty, which is illustrated as the orange shaded region.

using stochastic optimization algorithms. We use standard gradient-based techniques because they are simple, widely-used throughout machine learning, and work in practice: algorithmic details are given in Appendix A.

The inducing points \mathbf{u} can be understood as *waypoints* through which the robot should move: this is illustrated graphically in Figure 1. By optimizing the ELBO, the motion planning algorithm must therefore find waypoints which ensure the stochastic trajectories comply with the necessary requirements. To understand how to assure this, we study (i) how to specify the likelihood $p(\mathbf{e} | \mathbf{f})$ to incorporate these and other factors. We then study (ii) how to extract uncertainty for use by upstream systems.

3.2 Incorporating Constraints, Smoothness, and Collision Avoidance

A motion planner must generate plans adhering to various constraints, such as desired start-and-end-states and collision-avoidance. In the variational Gaussian Process motion planning framework, this involves integrating information into the posterior distribution using conditional probability, typically through the likelihood $p(\mathbf{e} | \mathbf{f})$. We now describe how this works for a number of different forms of information.

Joint Limits and Inequality-based Constraints

To respect the robot’s joint limits, we need to incorporate an inequality constraint into the GP model. Previous, non-variational Gaussian Process motion planning approaches (Dong et al., 2016) generally enforce such constraints softly by adding an extra term to the optimization objective. This does not guarantee the constraints are satisfied, which necessitates post-processing. The most natural alternative, which guarantees con-

straints to be satisfied, would be to modify the prior using conditional probability, but this unfortunately breaks Gaussianity of the GP. We therefore instead propose to leverage the variational formulation and enforce constraints using a bijective non-linear transformation. For box constraints of the form

$$\theta_i^{(\min)} \leq \theta_i \leq \theta_i^{(\max)} \quad (8)$$

this can be done by passing the GP through a scaled and shifted sigmoid function $\sigma : \Theta \rightarrow \Theta_c$, where Θ_c is constrained joint space. This amounts to simply replacing \mathbf{f} with $\sigma(\mathbf{f})$ in the likelihood, and returning motion plans based on the transformed GP sample paths, which are guaranteed to satisfy all joint constraints. One can similarly handle other constraints which map bijectively into a box, such as triangular constraints. This straightforward solution highlights the value of a variational formulation, which makes such aspects computationally simple to add.

Start States, Goal States, and Equality-based Constraints

To ensure our trajectories pass through the (unconstrained) start and goal states $\boldsymbol{\theta}_0, \boldsymbol{\theta}_1$ at times $t = 0$ and $t = 1$, we condition the variational process $f | \mathbf{u}$ on the events $f(0) = \boldsymbol{\theta}_0$ and $f(1) = \boldsymbol{\theta}_1$. This is done by partitioning the inducing variables as $\mathbf{u} = [\mathbf{u}_c, \mathbf{u}']$, where \mathbf{u}_c are the inducing variables representing start and goal states, and \mathbf{u}' are the remaining inducing variables. We then simply set \mathbf{u}_c to the values needed, and exclude them from optimization. We illustrate this in Figure 1 where \mathbf{u}_c are set to $[0, 0]$ at $t = 0, t = 1$ in the 1D case. This approach enables one to handle general equality-based constraints.

Collision Avoidance, Velocity Limits, and Additional Soft Constraints To incorporate constraints that do not have a convenient mathematical form, we




Figure 2: A sampled trajectory for grasping a can. Frames 1-2, 2-3, 3-4 depict respectively the end-effector alignment with the final pose, the approach stage, and the grasping stage.

include likelihood-based soft constraints in a manner that mirrors non-variational Gaussian Process motion planning algorithms, such as the collision-avoidance term proposed by Dong et al. (2016). In such cases, we replace \mathbf{f} with $\sigma(\mathbf{f})$ to account for inequality constraints as described previously. This enables us to handle collision avoidance, velocity constraints and other soft constraints, which are implicitly enforced by penalizing violations via the optimization objective, providing a flexible modeling tool that can be adapted to the task at hand. For example, to enforce grasping, we can include the terms

$$\|g_s - k_{\text{fwd}}(\sigma(\mathbf{f}))\|_{\Sigma_{\text{grasp}}}^2 \quad (9)$$

inside the exponential function in the log likelihood (10). Here, g_s is a *good graspable pose*, which consists of a vector and matrix that, together, represent the position and orientation of the end effector: we subtract element-wise and calculate the norm in the respective product space. The matrix Σ_{grasp} plays a role similar to Σ_{obs} , and quantifies the relative importance of this term in the likelihood. In total, this term effectively allows us to condition the Gaussian process on values in kinematic space, whereas the GP itself lives in joint space. During the training phase, we draw a batch of samples, calculate the joint’s position and orientation for each sample using forward kinematics, and compute the norm. This likelihood term therefore guides the samples towards the desired end-effector position and rotation, all as part of the variational optimization problem that produces the motion plan.

As a proof-of-concept, we include a grasping experiment. Alignment for viable grasping is achieved by adding an appropriate soft constraint term (9), which encourages the arm to orient itself accordingly. One can take this further, conditioning the GP at appropriate time steps to open and close the gripper of the robot to move objects from one place to another. This reinforces the value of the proposed framework, which allows us to treat different kinds of constraints—here, collision-avoidance and grasping alignment—in essentially-the-

same way mathematically. We demonstrate a motion plan incorporating these elements in Figure 2, and include a video in the supplementary material.

Smoothness Smoothness can be controlled via kernel choice and its hyperparameters. In particular, stationary kernels on \mathbb{R}^1 can be formulated to include a *length scale* parameter $\kappa > 0$, and written $k(x, x') = k\left(\frac{|x-x'|}{\kappa}\right)$. This parameter determines the characteristic distances over which the function varies (Rasmussen and Williams, 2006). Kernels on \mathbb{R}^d for $d > 1$ can be defined similarly, with one length scale per dimension. Larger length scales yield smoother trajectories, while smaller length scales yield rougher, more jagged trajectories. Through kernel choice, for instance by choosing a Matérn kernel with given smoothness (Rasmussen and Williams, 2006), one can also control properties such as the number of times a motion path is differentiable. We illustrate this in Figure 3.

3.3 Computation via Variational Inference and Connections with Other Planners

By combining the previous likelihood terms, we obtain the negative ELBO

$$\frac{1}{2} \mathbb{E}_{q(\mathbf{u})} \mathbb{E}_{p(\mathbf{f}|\mathbf{u})} \left(\underbrace{\|h_\varepsilon(\text{sdf}_s(k_{\text{fwd}}(\sigma(\mathbf{f}))))\|_{\Sigma_{\text{obs}}}}_{\text{collision term}}^2 + \underbrace{\|c(\sigma(\mathbf{f}))\|_{\Sigma_c}}_{\text{soft constraint terms}}^2 \right) + D_{\text{KL}}(q(\mathbf{u}') \parallel p(\mathbf{u}' | \mathbf{u}_c)) \quad (10)$$

which we minimize to obtain the optimal parameters of $q(\mathbf{u})$. This enforces equality constraints through the \mathbf{u}_c -values, inequality constraints through σ , and soft constraints through the collision and other soft constraint terms. Compared to previous GP-based approaches such as for instance Dong et al. (2016), we do not need to introduce sparse factor graphs, Laplace approximations, interpolation between time steps, or other scalability tricks. Instead, we simply carry minimization out using stochastic optimization, by sampling




Figure 3: Illustration of smoothness properties of different kernels, including the nowhere-differentiable Matérn- $1/2$, once-differentiable Matérn- $3/2$, and twice-differentiable Matérn- $5/2$ kernel.

\mathbf{u} -values, conditionally sampling \mathbf{f} -values, computing the objective, and taking gradient steps.

From (10), we deduce that probabilistic-inference-based motion planners based on variational inference are *stochastic extensions of optimization-based motion planning algorithms*, such as STOMP, CHOMP (Kalakrishnan et al., 2011; Zucker et al., 2013), and their variants. The main differences in our formulation are:

- (i) instead of a single trajectory, we now return a distribution over trajectories,
- (ii) the loss contains an extra expectation term, and
- (iii) we include an additional KL-based regularizer to ensure we do not stray too far from the prior, allowing us to control smoothness and related task-independent properties in a distributional manner.

The effect of these differences is that they introduce uncertainty, at cost of a slightly more complex formulation compared to optimization-based planners. Our optimization objective is a sum of multiple terms that reflect complementary aspects of motion planning, which need to be weighted appropriately. Empirically, we found it to be important that the entries of Σ_{obs} in (10) are large enough, to ensure the motion planner does not ignore the collision part of the objective. We present an ablation study in Appendix E on how each likelihood term affects the motion planning behavior. The formulation presented uses a unimodal variational approximation: this choice favors simplicity and ease of optimization. One could instead make a choice in favor of the opposite tradeoffs, following for instance Ewerton et al. (2020), by using a mixture model for the variational approximation.

Compared to other probabilistic-inference-based planners, such as GPMP2 (Dong et al., 2016), our formulation ensures computational tractability through sparse GPs, waypoints, and variational inference, rather than through a factor graph arising from the SDE prior. Compared to GVI (Yu and Chen, 2023), we work

with posterior function samples rather than purely in distribution-space, making our implementation much closer to that of an optimization-based planner. Our algorithm’s complexity is $\mathcal{O}(TM^2 + dM^3 + T(d + s))$ for obtaining the variational parameters which define a motion plan, $\mathcal{O}(dM^3)$ for computing error bars, and $\mathcal{O}(TM^2 + M^3)$ for both computing the mean motion plan and drawing one Monte Carlo sample when using efficient sampling via pathwise conditioning (Wilson et al., 2020, 2021), and $\mathcal{O}(T(d + s))$ to compute the likelihood. Here, T is the number of time points, M is the number of inducing points, which for our tasks is usually around a dozen, d is the dimensionality of joint space, and s is the number of spheres. Calculations are given in Appendix C.

Replanning In situations where we need to replan, we do so by starting with the previous motion plan and moving inducing points and inducing values around appropriately in order to compute a new motion plan. We also condition on the current value as initial state.

3.4 Uncertainty: Intervals and Monte Carlo Function Samples

Once a Gaussian process motion planning algorithm has completed the optimization process, the next step is to extract an actual motion plan from it. Since, in the variational framework, motion plans map bijectively to Gaussian processes, the *most likely* motion plan under the variational posterior can be obtained via the variational posterior’s mean. We can incorporate uncertainty in two ways: (a) through posterior intervals or (b) by drawing Monte Carlo random function samples. These work in the following manner:

- (a) Posterior intervals: for a given scaling level $\alpha > 0$ and independent GPs per output dimension, we can represent uncertainty through posterior intervals, which are $\mu_{f|\mathbf{u}}(x) \pm \alpha \sqrt{k_{f|\mathbf{u}}(x,x)}$. For large-enough α , this gives a high-probability estimate on the region through which motion plans will travel.

Robot	Env	M	Planners			
			vGPMP (ours)	RRT-connect	LBKPIECE	CHOMP
Franka	I	Acc	100%	100%	88.89%	18.06%
		MPL	2.75 ± 2.14	1.16 ± 1.67	1.82 ± 2.48	0.96 ± 1.95
		MC	-1.17 ± 1.73	-11.81 ± 3.23	-14.03 ± 4.55	-9.25 ± 6.94
	B	Acc	81.81%	100%	100%	14.54%
		MPL	2.60 ± 1.86	2.10 ± 2.09	2.21 ± 2.19	1.27 ± 2.01
		MC	-6.76 ± 4.69	-84.42 ± 18.63	-116.58 ± 20.05	-37.47 ± 12.77
UR10	I	Acc	100%	97.22%	86.11%	11.11%
		MPL	1.63 ± 1.09	2.42 ± 2.55	4.03 ± 4.37	0.20 ± 0.47
		MC	-56.03 ± 11.86	-281.84 ± 22.91	-479.73 ± 35.52	-5940.77 ± 108.31
	B	Acc	89.82%	85.45%	16.36%	18.18%
		MPL	3.95 ± 1.91	1.82 ± 2.22	0.93 ± 1.56	0.55 ± 0.90
		MC	-7.10 ± 5.55	-42.56 ± 13.03	-26.48 ± 9.97	-58.53 ± 6.58

Table 1: Experiments in different environments using two robots. The environments (Env) used are named Industrial (I) and Bookshelves (B), and the Measures (M) are accuracy (Acc), Mean Path Length (MPL), and Mean Clearance (MC). Mean and standard deviation are computed over the entire benchmark suite, and are averaged over five total runs.

- (b) Posterior function samples: we can also draw random function $\phi \sim \text{GP}(\mu_{f|u}, k_{f|u})$. For stationary kernels, we do so up to a small error tolerance using the *efficient sampling* technique of Wilson et al. (2020, 2021), which works by drawing approximate random functions from the prior, and transforming them into the posterior using path-wise conditioning. Once drawn, we can then plug these randomly-sampled functions into upstream systems as needed for the given task.

We emphasize that (b) is computationally tractable in part because the variational formulation can handle general kernels. These technical primitives can be used to propagate uncertainty into any upstream decision-making algorithm built atop of GP. Gaussian processes are widely used to represent uncertainty in a number of decision-making systems, including in Bayesian optimization, active learning, and reinforcement learning. Our formalism can therefore in principle facilitate the development of similar systems in the context of motion planning with further work.

4 Experiments

We now evaluate vGPMP and compare its performance and capabilities with other motion planners, including both those based on probabilistic inference and those based on other formalisms. To illustrate the unifying nature of our techniques, we consider standard collision-avoidance benchmarks. In all cases, full experiment details can be found in Appendix D.


To understand how vGPMP and other motion planners handle collision avoidance, we ran them on two kinds

of motion planning benchmarks. For the first kind, we compare with a state-of-the-art probabilistic-inference-based planner *GPMP2*, using similar hyperparameters as detailed in Appendix D, on the Barrett *WAM* robot (7 DOF) in the *Lab* environment (Mukadam et al., 2016). For the second kind, we compare with sampling-based and optimization-based planners, using benchmarks originally introduced by Schulman et al. (2013), which are based on the *Franka Panda* (7 DOF) and *UR10* (6 DOF) robots in the *Industrial* and *Bookshelves* environments, and have 36 and 55 unique motion planning problems, respectively. A visualization of these environments can be seen in Appendix D.

We evaluate behavior with respect to (i) *accuracy*, which is the percentage of tasks within the benchmark suite correctly solved by the planner, (ii) *clearance*, which quantifies collision-avoidance by measuring the distance to objects, and (iii) *path length*, which is the total distance moved by the end effector. Specifically, clearance, which we use it as a metric to assess safety, is defined as the negation of the collision term $-\frac{1}{2} \|\mathbf{h}_\varepsilon(\text{sdf}_s(\mathbf{k}_{\text{fwd}}(\sigma(\mathbf{f}))))\|_{\Sigma_{\text{obs}}}^2$ used in the likelihood. Note that clearance and path length are related: higher clearance pushes paths further away from objects, hence increasing path length. If required, an extra objective can be added to bring paths to desired lengths whilst maintaining collision avoidance.

All measurements are repeated using 5 different random seeds to assess variability. Further details on measurements are given in Appendix D.

Table 2 shows that vGPMP and GPMP2 both succeed in finding obstacle-free paths reaching the targets in all cases. Compared to GPMP2, vGPMP leads to better



(a) Start query state

(b) Goal query state

Figure 4: Example distribution of collision-free paths generated by vGPMP in the boxes environment, with a random sample shown in green.

clearance values and consequently longer paths moving away from the objects ensuring safety. Both methods require careful tuning of Σ_{obs} to ensure that this term is a sufficiently-important part of the loss to achieve collision-avoidance. Overall, in comparison to GPMP2, vGPMP performs comparably and provides a simpler computational pipeline and more general formulation.

Next, we examine how uncertainty translates to real-world execution. For vGPMP, we use Monte Carlo sampling to obtain 150 paths, from which we calculate the resulting envelope around the mean trajectory. The intervals in Figure 4 show the successful incorporation of the problem constraints as expressed in the low uncertainty at the beginning and end of the trajectory. This uncertainty grows as one moves further away from start and goal states, as well as points of collision.

Table 1 shows that vGPMP consistently leads to high success rates in reaching the goal states, achieving 91.42% across all experiments, which is significantly higher than the performance of LBKPIECE (69.78%) and CHOMP (15.69%). RRT-connect achieves a

Metric	vGPMP (ours)	GPMP2
Acc.	100%	100%
MC	-0.90 ± 1.50	-2.58 ± 2.54
MPL	2.17 ± 1.00	1.89 ± 0.98

Table 2: Comparison using WAM robot in the lab environment. Metrics shown are *accuracy*, *mean clearance*, and *mean path length*, defined in Section 4.

slightly higher success rate, 95.05% on average—however, the generated paths demonstrate worse clearance values. In most cases, vGPMP produces longer paths, since uncertainty drives paths away from obstacles for better collision avoidance.

vGPMP consistently yields better clearance values in all environments using the two robots, except with *Franka Panda* in the *bookshelves* environment, where the performance of vGPMP drops due to collisions with the floor: this could be alleviated by including an extra term in the loss, which we omitted to ensure a fair comparison with other planners that also do not consider the floor. In total, without post-processing, the uncertainty produced by vGPMP results in safer trajectories compared to those of other planners.

Finally, we apply our approach directly to a real robot. We demonstrate the feasibility of the approach by executing a trajectory generated from vGPMP given three obstacles placed around a Franka Panda robot. To execute the generated motion plans on the physical robot, we used a low-level joint impedance position controller (Franzese et al., 2021). We show the resulting trajectory in Figure 5, and provide a video in the paper’s supplementary materials.




Figure 5: Sampled trajectory executions on the real (bottom) and the simulated (top) robot showing successful transfer using a low-level controller.

5 Conclusion

We present a unifying framework for Gaussian-process-based motion planning with obstacle avoidance, using a formulation based on variational inference. The framework supports equality constraints, inequality constraints, and general soft constraints. It also allows for fine-grained control over smoothness through the prior. Our approach generalizes previous Gaussian-process-based motion planners to arbitrary kernels in a manner that simplifies computation. Additionally, by viewing variational inference as an optimization problem over a space of random motion paths, our framework also generalizes optimization-based planners to incorporate uncertainty. We evaluate the approach using different robots, environments, and tasks—including obstacle avoidance and grasping—showing that it has high accuracy in reaching target positions, while avoiding obstacles and using uncertainty to providing better clearance compared to baselines. We demonstrate the feasibility of the approach on a real robot, executing a sampled trajectory while avoiding obstacles.

References

- H. Attias. Planning by probabilistic inference. In *Artificial Intelligence and Statistics*, 2003. Cited on page 1.
- D. Berenson, S. Srinivasa, and J. Kuffner. Task space regions: A framework for pose-constrained manipulation planning. *International Journal of Robotics Research*, 2011. Cited on page 1.
- P. Cheng and S. M. LaValle. Resolution complete rapidly-exploring random trees. In *International Conference on Robotics and Automation*, 2002. Cited on page 1.
- S. Chitta, I. Sucas, and S. Cousins. Moveit! *IEEE Robotics and Automation Magazine*, 2012. Cited on page 14.
- E. Coumans and Y. Bai. PyBullet, a python module for physics simulation in robotics, games and machine learning, 2017. Cited on page 13.
- J. Dong, M. Mukadam, F. Dellaert, and B. Boots. Motion Planning as Probabilistic Inference using Gaussian Processes and Factor Graphs. In *Robotics: Science and Systems*, 2016. Cited on pages 1–6, 12, 14, 15.
- M. Ewerton, O. Arenz, and J. Peters. Assisted Teleoperation in Changing Environments with a Mixture of Virtual Guides. *Advanced Robotics*, 2020. Cited on pages 2, 6.
- G. Franzese, A. Mészáros, L. Peternel, and J. Kober. ILoSA: Interactive learning of stiffness and attractors. In *Intelligent Robots and Systems*, 2021. Cited on page 8.
- J. Hensman, N. Fusi, and N. D. Lawrence. Gaussian Processes for Big Data. In *Uncertainty in Artificial Intelligence*, 2013. Cited on page 3.
- J. Hensman, A. Matthews, and Z. Ghahramani. Scalable Variational Gaussian Process Classification. In *Artificial Intelligence and Statistics*, 2015. Cited on page 3.
- E. Huang, M. Mukadam, Z. Liu, and B. Boots. Motion planning with graph-based trajectories and Gaussian process inference. In *International Conference on Robotics and Automation*, 2017. Cited on page 2.
- M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal. STOMP: Stochastic trajectory optimization for motion planning. In *International Conference on Robotics and Automation*, 2011. Cited on pages 1, 6.
- S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *International Journal of Robotics Research*, 2011. Cited on page 1.
- M. E. Khan and D. Nielsen. Fast yet Simple Natural-Gradient Descent for Variational Inference in Complex Models. In *International Symposium on Information Theory and its Applications*, 2018. Cited on page 13.
- J. Kuffner and S. M. LaValle. RRT-connect: An efficient approach to single-query path planning. In *International Conference on Robotics and Automation*, 2000. Cited on page 1.
- D. F. Lirui Wang Yu Xiang. Manipulation trajectory optimization with online grasp synthesis and selection. In *Robotics: Science and Systems*, 2020. Cited on page 1.
- A. G. d. G. Matthews, J. Hensman, R. Turner, and Z. Ghahramani. On Sparse Variational Methods and the Kullback-Leibler Divergence between Stochastic Processes. In *Artificial Intelligence and Statistics*, 2016. Cited on page 3.
- A. G. d. G. Matthews, M. van der Wilk, T. Nickson, K. Fujii, A. Boukouvalas, P. León-Villagrà, Z. Ghahramani, and J. Hensman. GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research*, 2017. Cited on page 13.
- M. Mukadam, J. Dong, X. Yan, F. Dellaert, and B. Boots. Continuous-time Gaussian Process Motion

- Planning via Probabilistic Inference. *International Journal of Robotics Research*, 2018. Cited on page 1.
- M. Mukadam, X. Yan, and B. Boots. Gaussian process motion planning. In *International Conference on Robotics and Automation*, 2016. Cited on pages 1, 7, 14.
- A. Paraschos, C. Daniel, J. R. Peters, and G. Neumann. Probabilistic Movement Primitives. In *Advances in Neural Information Processing Systems*, 2013. Cited on page 2.
- J. Quiñonero-Candela and C. Rasmussen. A Unifying View of Sparse Approximate Gaussian Process Regression. *Journal of Machine Learning Research*, 2005. Cited on page 3.
- C. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006. Cited on pages 2, 5, 15.
- N. Ratliff, D. Silver, and J. Bagnell. Learning to Search: Functional Gradient Techniques for Imitation Learning. *Autonomous Robots*, 2009. Cited on pages 1, 2.
- H. Salimbeni, S. Eleftheriadis, and J. Hensman. Natural Gradients in Practice: Non-Conjugate Variational Inference in Gaussian Process Models. In *Artificial Intelligence and Statistics*, 2018. Cited on page 13.
- J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel. Motion planning with sequential convex optimization and convex collision checking. *International Journal of Robotics Research*, 2014. Cited on page 1.
- J. Schulman, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, and P. Abbeel. Finding locally optimal, collision-free trajectories with sequential convex optimization. In *Robotics: Science and Systems*, 2013. Cited on pages 1, 7.
- E. Snelson and Z. Ghahramani. Sparse Gaussian Processes using Pseudo-inputs. In *Advances in Neural Information Processing Systems*, 2006. Cited on page 3.
- M. Titsias. Variational Learning of Inducing Variables in Sparse Gaussian Processes. In *Artificial Intelligence and Statistics*, 2009. Cited on page 3.
- M. Toussaint and C. Goerick. A Bayesian View on Motor Control and Planning. In *From Motor Learning to Interaction Learning in Robotics*. 2010. Cited on page 1.
- M. Toussaint and A. Storkey. Probabilistic inference for solving discrete and continuous state Markov decision processes. In *International Conference on Machine Learning*, 2006. Cited on page 1.
- J. Urain, A. T. Le, A. Lambert, G. Chalvatzaki, B. Boots, and J. Peters. Learning Implicit Priors for Motion Optimization. In *International Conference on Intelligent Robots and Systems*, 2022. Cited on pages 2, 15.
- 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. Cited on pages 6, 7, 13.
- J. T. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. P. Deisenroth. Pathwise Conditioning of Gaussian Processes. *Journal of Machine Learning Research*, 2021. Cited on pages 6, 7, 13.
- H. Yu and Y. Chen. A Gaussian Variational Inference Approach to Motion Planning. *IEEE Robotics and Automation Letters*, 2023. Cited on pages 1, 6, 15.
- M. Zucker, N. Ratliff, A. Dragan, M. Pivtoraiko, M. Klingensmith, C. Dellin, J. A. Bagnell, and S. Srinivasa. CHOMP: Covariant Hamiltonian Optimization for Motion Planning. *International Journal of Robotics Research*, 2013. Cited on pages 1, 6.

Supplementary Materials: A Unifying Variational Framework for Gaussian Process Motion Planning

A Algorithm

A full algorithmic description of vGPMP is given below.

Algorithm 1 vGPMP

- 1: **Inputs:** number of total optimization steps L , the variational distribution $q(\mathbf{u})$, the prior $p(\mathbf{u})$, input timesteps $\mathbf{X} = [t_0, t_1, \dots, t_N]$, the process $f \sim \mathcal{GP}(\mu, k)$, step size ρ , total number of samples K drawn at each ELBO estimation step, samples drawn from posterior P , optional: query timesteps with finer discretization e.g. $\mathbf{X}_{\text{new}} = [t_0, t_{0.5}, \dots, t_{N_{\text{new}}}]$, $N_{\text{new}} > N$.
 - 2: **for** $\ell = 1 : L$ **do**
 - 3: $\text{Draw } \{\mathbf{u}^k\}_{k=1}^K \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{u}_\ell}, \boldsymbol{\Sigma}_{\mathbf{u}_\ell})$
 - 4: $\mathbf{K}_{\mathbf{z}\mathbf{z}} \leftarrow k(\mathbf{u}, \mathbf{u})$
 - 5: $\mathbf{L} \leftarrow \text{CholeskyDecomposition}(\mathbf{K}_{\mathbf{z}\mathbf{z}})$
 - $\backslash\backslash$ Matrix inversion using Cholesky decomposition
 - 6: $\mathbf{K}_{\mathbf{z}\mathbf{z}}^{-1} \leftarrow \text{CholeskySolve}(\mathbf{L}, \mathbf{I})$
 - $\backslash\backslash$ Draw pathwise samples
 - 7: $\{\mathbf{f}^k\}_{k=1}^K \sim f(\mathbf{X}) + \mathbf{K}_{\mathbf{x}\mathbf{z}}\mathbf{K}_{\mathbf{z}\mathbf{z}}^{-1}(\mathbf{u} - \mathbf{f}_m)$
 - $\backslash\backslash$ Evaluate the likelihood and average over samples
 - 8: $\mathbf{E} \leftarrow \sum_{i=1}^K \underbrace{\|\mathbf{h}_\varepsilon(\text{sdf}_s(\text{k}_{\text{fwd}}(\sigma(\mathbf{f}^i))))\|_{\boldsymbol{\Sigma}_{\text{obs}}}^2}_{\text{collision term}} + \underbrace{\|\mathbf{c}(\sigma(\mathbf{f}^i))\|_{\boldsymbol{\Sigma}_{\mathbf{c}}}^2}_{\text{soft constraint terms}}$
 - $\backslash\backslash$ Compute the mean and covariances for $p(\mathbf{u})$ and $q(\mathbf{u})$
 - 9: $\boldsymbol{\mu}_{\mathbf{p}} \leftarrow \mathbf{K}_{\mathbf{z}\mathbf{z}}^{-1}\mathbf{u}_{\mathbf{c}}$
 - 10: $\boldsymbol{\mu}_{\mathbf{q}} \leftarrow [\mathbf{u}_{\mathbf{c}}, \boldsymbol{\mu}_{\mathbf{u}_s}]$
 - $\backslash\backslash$ Whiten the difference between $q(\mathbf{u})$ and $p(\mathbf{u})$
 - 11: $\mathbf{W} \leftarrow \text{TriangularSolve}(\mathbf{L}, \boldsymbol{\mu}_{\mathbf{q}} - \boldsymbol{\mu}_{\mathbf{p}})$
 - $\backslash\backslash$ Compute the KL divergence and ELBO
 - 12: $\text{D}_{\text{KL}} \leftarrow 0.5 (\log |\boldsymbol{\Sigma}_{\mathbf{u}_\ell}| - \text{tr}(\boldsymbol{\Sigma}_{\mathbf{u}_\ell}) - \dim(\mathbf{W}) + \mathbf{W}^\top \mathbf{W})$
 - 13: $\mathcal{L} \leftarrow \alpha \frac{1}{N} \sum_{i=1}^N \mathbf{E}^i + \text{D}_{\text{KL}}$
 - $\backslash\backslash$ Update the variational distribution
 - 14: $\boldsymbol{\mu}_{\mathbf{u}_{\ell+1}} = \boldsymbol{\mu}_{\mathbf{u}_\ell} + \rho [\nabla_{\mathbf{q}} \mathcal{L}]$
 - 15: $\boldsymbol{\Sigma}_{\mathbf{u}_{\ell+1}} = \boldsymbol{\Sigma}_{\mathbf{u}_\ell} + \rho [\nabla_{\mathbf{q}} \mathcal{L}]$
 - 16: **end for**
 - 17: $\text{Draw } \{\mathbf{u}^i\}_{i=1}^P \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{u}_L}, \boldsymbol{\Sigma}_{\mathbf{u}_L})$
 - $\backslash\backslash$ Draw pathwise samples from posterior or return the mean
 - 18: $\{\mathbf{f}^i\}_{i=1}^P \sim f(\mathbf{X}_{\text{new}}) + \mathbf{K}_{\mathbf{x}'\mathbf{z}}\mathbf{K}_{\mathbf{z}\mathbf{z}}^{-1}(\mathbf{u} - \mathbf{f}_m)$
 - $\backslash\backslash$ Optional: pick a trajectory of choice - here we pick the lowest collision cost trajectory
 - 19: $\boldsymbol{\theta} = \arg \min_{\{f^i\}_{i=1}^P} \left[\|\mathbf{h}_\varepsilon(\text{sdf}_s(\text{k}_{\text{fwd}}(\sigma(\mathbf{f}^i))))\|_{\boldsymbol{\Sigma}_{\text{obs}}}^2 \right]$
 - 20: **return:** $\boldsymbol{\theta}$
-

B Incorporating Additional Constraints under the Variational Representation

For learning the variational distribution, we choose $q(\mathbf{u}) = \mathcal{N}(\boldsymbol{\mu}_{\mathbf{u}}, \boldsymbol{\Sigma}_{\mathbf{u}})$ with a particular representation we now describe. Letting $\mathbf{u} = [\mathbf{u}', \mathbf{u}_c]$ be the partition of \mathbf{u} into the learned and constrained parts, the matrix $\boldsymbol{\Sigma}_{\mathbf{u}}$ contains one part for \mathbf{u}' , which is padded with zeros to encode the uncertainty in the given motion planning constraints \mathbf{u}_c . One can additionally introduce noise in the latent variables, but we omit this here to ease notation.

Computational Considerations To improve computational efficiency, we use a whitened representation that only considers the learned portion of the covariance $\boldsymbol{\Sigma}_{\mathbf{u}}$. Let \mathbf{z}' and \mathbf{z}_c be the inducing points corresponding to \mathbf{u}' and \mathbf{u}_c , respectively. The part of the distribution representing \mathbf{u}_c is a degenerate Gaussian, so we omit it from the whitening process. This representation is achieved using the block partitioned representation of $\mathbf{K}_{\mathbf{z}\mathbf{z}}$, because the Cholesky decomposition of the covariance of $p(\mathbf{u}'|\mathbf{u}_c)$ can be written as $\text{Chol}(\mathbf{K}_{\mathbf{z}'\mathbf{z}'} - \mathbf{K}_{\mathbf{z}'\mathbf{z}_c}\mathbf{K}_{\mathbf{z}_c\mathbf{z}_c}^{-1}\mathbf{K}_{\mathbf{z}_c\mathbf{z}'})$: this can be seen by considering the block LU decomposition. An added benefit of this approach is that the KL term in (10) can be computed efficiently, as the Cholesky factor is already cached and used in the pathwise updates.

Velocity Modeling and Higher Order Derivatives Including constraints on derivatives, such as maintaining constant velocity, is crucial for real-world execution. By incorporating these constraints, the resulting models can accurately account for physical system limitations, ensuring the feasibility and reliability of planned trajectories during execution. In our framework, one can encode velocities as prior information by considering inducing variables along with their time derivatives, namely $\mathbf{u} = [\mathbf{u}', \dot{\mathbf{u}}', \mathbf{u}_c, \dot{\mathbf{u}}_c]$ which implies that the respective kernel matrix can be assembled by differentiating the kernel function. In many situations, we set the start and end velocities to zero. Our framework also allows angular velocities of the samples to be computed by differentiating the sample paths. This means they can be included in the likelihood via a soft constraint term which brings the velocity of sampled trajectories to desired values. This lead to, for instance, the likelihood term

$$\|\dot{\sigma}(\mathbf{f}) - \mu_{\mathbf{v}}\|_{\boldsymbol{\Sigma}_{\text{velocity}}} \quad (11)$$

Acceleration and other higher-order derivatives can also be considered following a similar argument, assuming appropriate kernel choice to ensure sufficient differentiability. For the Matérn- $5/2$ kernel, this benefit is two-fold as the GP is then Markovian: if advantageous in the setting at hand, this can be used to re-formulate the necessary computations in terms of sparse linear algebra, which makes it possible to generalize the stochastic-differential-equation-based techniques of Dong et al. (2016) to GP priors with more available derivatives.

C Time Complexity

In this section, we examine the scalability of vGPMP. Let s be the number of spheres that model the robot body. Let T be the total number of discrete time points. Let M be the number of inducing points, and d be the dimensionality of joint space. The time complexity of *one iteration* can be split up into

$$\text{Cost}_{\text{iteration}} = C_{\text{sampling}} + C_{\text{likelihood}} + C_{\text{KL}} \quad (12)$$

We begin by deriving the complexity cost of computing the likelihood for a single sample, for one single joint configuration at some timestep t .

The time complexity of the forward kinematics computation using the Denavit-Hartenberg (DH) convention and homogeneous transformation matrices is $\mathcal{O}(d)$, where d is the number of degrees of freedom of the robot arm. To compute the forward kinematics, we need to multiply the homogeneous transformation matrices for each joint, as follows:

$$\mathbf{F}_d = \prod_{i=0}^{d-1} \mathbf{F}_{i,i} \mathbf{F}_{i,i+1} \quad (13)$$

Each \mathbf{F}_i is 4×4 matrix computed using the DH parameters, and hence requires constant amount of computation. Then, an additional transform is needed to translate the sphere offsets into world coordinates. This can be easily achieved by considering some other homogeneous transform $\mathbf{H}_{i,j}$ for some sphere $c_j, j = \{1, \dots, s\}$ and some frame of reference (in world coordinates) i such that we can get the sphere placement in world coordinates, namely

$$\mathbf{H}_j = \mathbf{F}_i \mathbf{H}_{i,j} \quad (14)$$

Robot	Env	Parameters				
		$\mathbb{E}[q(\mathbf{u})]$	length scale	kernel variance	σ_{obs}	Ind. var.
Franka	I	Zeros Trained	[2, 2, 2, 2, 2, 2] Trained	0.2 Trained	0.005 Not trained	10 Trained
	B	Interp. Trained	[4, 4, 4, 3, 3, 3] Trained	0.5 Trained	0.0005 Not trained	24 Trained
UR10	I	Interp. Trained	[6, 6, 6, 6, 6, 6] Trained	0.1 Trained	0.0001 Not trained	18 Trained
	B	Upright Trained	[4, 4, 4, 4, 4, 4] Trained	0.25 Trained	0.0005 Not trained	12 Not trained
WAM	Lab	Interp. Trained	[5, 5, 5, 4, 5, 5] Trained	0.05 Trained	0.005 Not trained	14 Not trained

Table 3: Experiment configurations for various robots in different environments. The environments (Env) used are named *industrial* (I) and *bookshelves* (B). Interp. stands for *interpolated*. In all cases, set $\varepsilon = 0.05$, aside from UR10 on the *industrial* environment where it was $\varepsilon = 0.08$ and WAM on the *lab* environment where it was $\varepsilon = 0.03$. We also list which parameters are optimized during training, which are denominated as *trained*, and which are fixed during training, denominated as *not trained*.

To do this for all joints and all spheres, the time complexity is $\mathcal{O}(s)$, which brings the overall cost to $\mathcal{O}(d + s)$. After the joint configuration has given the robot pose and the sphere position in world coordinates, the likelihood can be computed. The cost can be broken up into applying the hinge loss to the signed distances, and the cost of computing the quadratic form. For the first half, the cost of computing the signed distance involves a lookup in a 3D grid, incurring constant cost per sphere. Since the hinge loss’ complexity is linear in the number of spheres, the overall time complexity of this procedure is $\mathcal{O}(s)$. Additionally, since the covariance-matrix-term in the likelihood is assumed diagonal, for a given vector of spheres of size s , the quadratic form can be computed in linear time $\mathcal{O}(s)$. At each iteration of the optimization the likelihood is evaluated at T timesteps, bringing the overall cost to $C_{\text{likelihood}} = T(d + s)$.

The efficient sampling technique of Wilson et al. (2020, 2021) uses pathwise conditioning, which involves solving a linear system and computing the Cholesky decomposition of $\mathbf{K}_{\mathbf{z}\mathbf{z}}$. Therefore, drawing the samples amounts to $C_{\text{sampling}} = TM^2 + M^3$ for one sample. We view the cost of Fourier features as constant.

Lastly, we consider computation of the KL term. The dominant factor here is the computation of the covariance matrix inverse and determinant. We must also whiten $\Sigma_{\mathbf{u}}$. We have that (a) the cost of computing $\mathbf{K}_{\mathbf{z}\mathbf{z}}$ is $\mathcal{O}(dM^2)$, (b) the cost of computing the Cholesky decomposition and solving the linear system takes $\mathcal{O}(dM^3)$ time for a dense matrix, (c) the cost of whitening the covariance is dominated by solving the linear system, where the Cholesky factor is computed beforehand and is upper triangular, which therefore also takes $\mathcal{O}(dM^3)$. Overall, $C_{\text{KL}} = dM^3$, which gives the overall time complexity of Algorithm 1 as $\mathcal{O}(TM^2 + dM^3 + T(d + s))$.

D Experimental Details

Experiments were performed on an AMD Ryzen 7 3700x CPU. For the proposed planner, we employ independent GPs per output dimension with Matérn- $5/2$ kernels. This GP model provides a sufficient degree of differentiability to express higher order dynamics, and is expressive enough to accommodate different ranges of motion necessary in the motion plans. Experiments were repeated five times with different seeds. We use the GPFlow library (Matthews et al., 2017), with the complementary GPFlow-based pathwise sampling module (Wilson et al., 2020, 2021) and PyBullet for simulations (Coumans and Bai, 2017). vGPMP uses Adam with hyperparameters $\beta_1 = 0.8$ and $\beta_2 = 0.95$ and learning rate $\eta = 0.09$.

Since optimization is carried out on the loss directly, any standard optimizer can be used. In practice, with standard stochastic optimization we observed convergence across all test suites in 150 optimization steps or less—however, if even faster convergence is desired, one could for instance apply natural gradients with some additional computational overhead (Khan and Nielsen, 2018; Salimbeni et al., 2018).




Figure 6: Example distributions of collision-free paths generated by vGPMP in two benchmark environments. Sampled trajectories are shown in green.

We include RRT-connect, LBKPIECE and CHOMP via the MoveIt! OMPL library (Chitta et al., 2012). RRT-connect and LBKPIECE were ran with default settings, with post-processing left on to smoothen the resulting trajectory. CHOMP was ran using the configuration detailed by Dong et al. (2016).

We compare the proposed approach, vGPMP, with GPMP2 (Dong et al., 2016) in the *lab* environment, originally introduced by Mukadam et al. (2016), using the (7 degree of freedom) Barrett WAM robot. We test both planners using 24 different cases with various start and goal states. We initialize GPMP2 with parameters as outlined in Dong et al. (2016): $\Sigma_{\text{obs}} = 0.02^2 \mathbf{I}$, $\varepsilon = 0.2$ with 101 temporally equidistant states. Results are detailed in Table 2. For vGPMP, we set $\Sigma_{\text{obs}} = 0.005^2 \mathbf{I}$, $\varepsilon = 0.03$, and only 14 inducing locations. For optimization, we use 14 samples per iteration for a total of 150 iterations.

All parameters used for baselines against RRT-connect, LBKPIECE, CHOMP and GPMP2 are outlined in Table 3. The number of iterations for *Franka* and *UR10* is 130. From Table 1, in the *industrial* environment using UR10 our approach finds shorter paths in general, which is in part due to higher length scale initialization.

We evaluated the performance of vGPMP and other baseline planners in two different environments, which can be seen in Figure 6, using various start and goal states. vGPMP generates smooth trajectories that are also further away from the obstacles in comparison to baselines as illustrated in Figure 7. If we filter samples to obtain a *most-collision-free path*, this avoids the bookshelf in a safer manner than the respective path from RRT-connect. To ensure balanced comparisons, average success rates reported are calculated as a weighted average with respect to the number of problems.

Each hyperparameter in Table 3 can be interpreted as follows. The *expectation* $\mathbb{E}[q(\mathbf{u})]$ represents the mean of the prior around which paths are sampled for optimization. *Length scale* refers to the kernel length scale

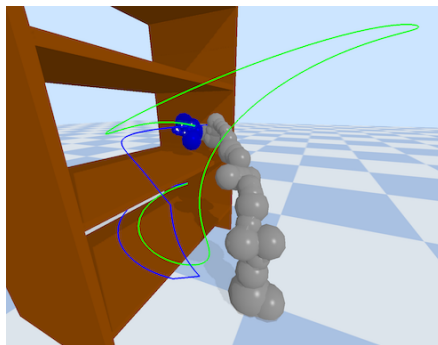





Figure 7: Comparison of trajectories from a baseline motion planner, RRT-connect (in blue) and vGPMP (in green). As in this example, in general vGPMP leads to smooth trajectories further away from the obstacles. Spheres around robot links that are used in collision detection can also be seen.




(a) Trajectory distribution for the Panda arm using all likelihood terms.



(b) Trajectory distribution for the Panda arm missing the collision likelihood.



(c) Trajectory distribution for the Panda arm missing the position likelihood.



(d) Trajectory distribution for the Panda arm missing the orientation term.

Figure 8: Visual comparison of the influence of each constraint term on optimization and the resulting trajectory distribution. In Figure 8a optimization uses all three likelihoods for optimization: collision, position and orientation. In Figure 8b the collision likelihood was omitted. In Figure 8c, Figure 8d the position and orientation part of the grasping likelihood were omitted, respectively.

hyperparameter, which controls spatial variability—see Rasmussen and Williams (2006). *Kernel variance* corresponds to the size of the outlined volume in Figure 4. The *likelihood standard deviation* term σ_{obs} controls the weight on the likelihood in the objective, with lower values representing higher weight on the likelihood. The term *Ind. Var.* represents the number of inducing variables corresponding to waypoints used during optimization.

E Ablation Study

In Section 3.2, we showed that our framework can be easily adapted to include additional constraints of interest. Here, we use the grasping setup using the additional terms introduced in (9) to perform an ablation experiment on how the contribution of each likelihood term affects the final distribution. For all experiments, the configuration of vGPMP is fixed. Figure 8a shows the final trajectory distribution following optimization using all 3 constraints: the collision, position, and orientation likelihoods, as explained earlier. Through sufficient weight on the grasping term, the resulting end effector alignment enables the robot to pick up the can. Figure 8b, Figure 8c and Figure 8d show the resulting distribution in the absence of the collision, position and orientation terms, respectively. In all 3 cases, grasping is unsuccessful. Without collision information, the grasping constraints compete against the KL divergence, and the position constraint is not enforced completely. Without position information, the trajectory is collision-free and the end effector is aligned, but is located away from the table. Finally, without orientation, the end effector is not aligned for grasping. Note that in all examples, the uncertainty contracts and then expands again. This is a consequence of predicting at timesteps not seen during training, that is, timesteps which are either in-between, before, or after the waypoints. Note that this behavior may not be present in other GP-related motion planners, particularly if they implement predictions of this form using interpolation algorithms used atop the planner’s output (Dong et al., 2016; Yu and Chen, 2023; Urain et al., 2022) rather than using the uncertainty obtained from the Gaussian process.