Learning Implicit Priors for Motion Optimization

Julen Urain²*, An T. Le²*, Alexander Lambert¹*, Georgia Chalvatzaki², Byron Boots¹ and, Jan Peters²

Abstract-Motion optimization is an effective framework for generating smooth and safe trajectories for robotic manipulation tasks. However, it suffers from local optima that hinder its applicability, especially for multi-objective tasks. In this paper, we study this problem in light of the integration of Energy-Based Models (EBM) as guiding priors in motion optimization. EBMs are probabilistic models with unnormalized energy functions that represent expressive multimodal distributions. Due to their implicit nature, EBMs can easily be integrated as datadriven factors or initial sampling distributions in the motion optimization problem. This work presents a set of necessary modeling and algorithmic choices to effectively learn and integrate EBMs into motion optimization. We present a set of EBM architectures for learning generalizable distributions over trajectories that are important for the subsequent deployment of EBMs. Moreover, we investigate the benefit of including smoothness regularization in the learning process to improve motion optimization. In addition to gradient-based solvers, we also propose a stochastic method for trajectory optimization with learned EBMs. We provide extensive empirical results in a set of representative tasks against competitive baselines that demonstrate the superiority of EBMs as priors in motion optimization scaling up to 7-dof robot pouring that can be easily transferred to the real robotic system. Videos and additional details are available at https://sites.google. com/view/implicit-priors.

I. INTRODUCTION

Motion planning is a fundamental property for autonomous robots to achieve task-specific goals. In the context of autonomous robot manipulation, the trajectories that the robot should execute should satisfy several constraints, e.g., approaching the goal while avoiding collisions with itself and the world and joint limits. Naturally, a complex motion plan can be viewed as a multi-objective optimization problem along a specific time horizon. In this work, we study motion planning in light of motion optimization methods, when multiple objectives need to be satisfied for accomplishing an end task. In contrast with sampling-based motion planning methods [1], [2], motion optimization methods cast the motion planning problem into a trajectory optimization one [3]-[6], which optimize an initial trajectory or trajectory distribution iteratively by gradient descent or stochastic optimization until convergence.

Motion optimization is an inherently local optimization method that relies on the initialization, iteratively making



(a) EBM for trajectories

(b) Pouring Task

Fig. 1: Implicit trajectory distributions are learned from demonstrations using EBMs. These guide motion optimization to produce feasible trajectories for new problems. (a) An obstacle avoidance energy function, with generated optimal trajectories towards different goal locations. (b) A robot manipulator using learned EBMs to pour a cup within a cluttered scene.

local updates at every optimization step. Hence, due to the possible non-convexity of the cost function, these optimization methods suffer from local minima. Additionally, if the cost function is sparse, it might be hard to get the proper information for reaching low-cost regions, and therefore, the initial proposal may barely improve.

A way to avoid the local minima traps in trajectory optimization is to include a set of informative priors in the trajectory optimization. Classical motion optimization methods usually employ a set of handcrafted initialization proposals to warm-start the optimization. In Covariant Hamiltonian Optimization for Motion Planning (CHOMP) [3], the initial trajectory is a straight line that moves from the initial to the goal configuration. Stochastic Trajectory Optimization for Motion Planning (STOMP) [4], on the other hand, considers a normal initial sampling distribution with a structured covariance matrix modeled to increase entropy in the center of the state-based trajectory, with minimum variance at the start- and end-points. Alternatively, engineered costs might help guide or regularize the optimization. However, these handcrafted heuristics can be difficult to tune correctly in practice, and can do little to alleviate multimodality of complex manipulation tasks.

In a different vein, to capture the inherent multimodality of multi-objective motion planning tasks, a line of work proposes to learn trajectory distributions from datato guide the optimization process away from local minima. Explicitly, these methods fit a density model to a set of provided demonstrated trajectories [7] and then use these models as the guiding priors for motion optimization [8], [9]. However, these approaches often struggle in capturing multi-modality in high-dimensional spaces, leading to sub-optimal solutions.

In this work, we study the modeling of priors for motion optimization as *implicit models* [10]. Namely, we propose

^{*} All authors contributed equally.

¹ Paul G. Allen School of Computer Science and Engineering (CSE), University of Washington (USA)

² Intelligent Autonomous Systems Lab, Technische Universität Darmstadt (Germany)

This project has received funding from the European Unions Horizon 2020 research and innovation programme under grant agreement No # 820807 (SHAREWORK) and DFG project, PE 2315/11-1 (MET-RIC4IMITATION)

to use Energy Based Models (EBM) [11] as our density models, a class of implicit functions, that can represent unnormalized density functions that can easily capture multimodalities in data. EBMs can be represented in arbitrary latent spaces and then, transform the distribution to a higher dimensional space [12]. This property allows us to represent configuration space trajectory distributions given an EBM in a lower-dimensional task space. Additionally, due to their exponential nature, EBMs can be easily combined, allowing the composition of multiple priors, representing sub-tasks of the manipulation task, into a single structured prior.

We propose a motion optimization framework using as priors implicit functions that are modular, learnable, differentiable, and composable. Using our learned EBMs as priors, we can integrate multimodal information that can bias and guide the optimization process towards finding a feasible and smooth solution in complex tasks. Concretely, in this paper, we a) show how training with smoothness regularization and task-specific EBM design choices such as object-centric or phase conditioning can benefit motion optimization, b) explain how EBMs can naturally be integrated as structured priors in a motion optimization problem, c) propose a novel stochastic trajectory optimization method using Gaussian Process (GP) priors, and, d) demonstrate the applicability and effectiveness of our motion optimization framework using EBMs as structured priors to guide the optimization problem both in simulation and real-world robotic tasks.

II. RELATED WORK

Motion optimization. While sampling-based motion planning algorithms have gained significant traction [1], [2], they are typically computationally expensive, hindering their application in real-world problems. Moreover, these methods cannot guarantee smoothness in the trajectory execution, resulting in jerky robot motions that must be post-processed before applying them on a high degree-of-freedom (DoF) robot. Contrarily to sampling-based motion planners, trajectory optimization methods can integrate collisions and motion smoothness constraints, optimizing over an initially proposed trajectory based on a set of cost functions. CHOMP and its variants [13]-[15] optimize a cost function using covariant gradient descent over an initial suboptimal trajectory that connects the starting and goal configuration. STOMP proposed to optimize over non-differentiable constraints by drawing stochastic samples from a set of noisy trajectories, making STOMP increasingly dependent on the parameters of the noisy distribution. Due to their assumptions, these methods may not find an optimal solution or even fail to converge. TrajOpt [5] addresses the computational complexity of CHOMP and STOMP, that require a fine discretization of the trajectory for collision checking, proposing a sequential quadratic program with continuous collision-time collision checking. In Gaussian Process Motion Planning (GPMP) [6], motion optimization is cast as a probabilistic inference problem. A trajectory is parameterized as a function of continuous-time that maps to robot states, while a GP is used as a prior distribution to encourage trajectory smoothness,

while a likelihood function encodes feasibility. The trajectory is calculated via *maximum a-posteriori* (MAP) estimation from the posterior distribution of trajectories from the GP prior and the likelihood function.

Learning for motion generation. Learning-based methods for robot motion generation hold the promise of generalizing robot skills [16]. Reinforcement learning provides reactive policies that usually overfit a single task [17], while imitation learning aims to provide policies learned through demonstrations [18]. For the latter, behavioral cloning (BC) methods try to explicitly replicate the demonstrated trajectory [19], [20], but the acquired policies do not extrapolate outside the data distribution, are usually unimodal, and tend to fail when task conditions change. However, these explicit models have been integrated into motion optimization frameworks as a way to guide the optimization process, to tackle issues of computational complexity and trapping to local minima [8], [9]. Recently, various works explored the integration of learned priors for motion planning [21]–[25]. Unfortunately, the explicit models representing the data distribution are prone to collapse into single modes and do not adequately capture the multimodality of complex manipulation tasks; therefore, their contribution is limited to simple use-cases of robot motion generation. Hence, implicit models seem a more viable option [20], [26].

Highly connected to our work are the fields of Inverse Optimal Control (IOC) [27], [28] and Inverse Reinforcement Learning (IRL) [29]–[31] methods. In IOC and IRL, given a set of demonstrations, we aim to learn the cost function that these demonstrations are trying to maximize. In particular, Maximum Entropy IRL (MaxEntIRL) [31] models the demonstrations as Gibbs distribution, whose energy is given by the unknown cost function, i.e., it approximates an EBM. Although our work draws connections with MaxEntIRL, and investigates the learning of implicit models for robotic manipulations tasks, our contributions are distinctive. Crucially, we study EBMs as priors for motion optimization, proposing targeted ways for learning and integrating EBMs in robot motion optimization.

III. BACKGROUND

Energy based models are probabilistic models that represent a probability density function p_{θ} up to an unknown normalizing constant,

$$p_{\boldsymbol{\theta}}(\boldsymbol{x}) = \frac{1}{Z_{\boldsymbol{\theta}}} \exp(-E_{\boldsymbol{\theta}}(\boldsymbol{x})), \qquad (1)$$

with $x \in \mathcal{X}$ the input variable, $E_{\theta} : \mathcal{X} \to \mathbb{R}$ the energy function and $Z_{\theta} = \int_{\mathcal{X}} \exp(-E_{\theta}(x)) dx$ the normalization constant. In an EBM, the energy function E_{θ} is usually modelled by an arbitrary neural network parameterized by θ . Due to the arbitrary shape of the energy function E_{θ} , the computation of Z_{θ} is usually intractable.

Training of EBMs. Even if EBMs have been widely applied in a different set of applications, one of the most common applications is to use them as generative models. Given a dataset \mathcal{D} : $\{x_d\}_{0:N}$, we aim to fit our parameterized density model to a dataset. There exist several algorithms for learning [32]. In our work, we use Contrastive Divergence (CD) [33]. In CD, the Maximum Likelihood Estimation (MLE) objective is re-written to maximize the probability of the samples from the data compared to randomly sampled *negative* data points,

$$\mathcal{L}_{\text{CD}} = \mathbb{E}_{p_{\mathcal{D}}(\boldsymbol{x})} \left[E_{\boldsymbol{\theta}}(\boldsymbol{x}) \right] - \mathbb{E}_{q(\boldsymbol{x})} \left[E_{\boldsymbol{\theta}}(\boldsymbol{x}) \right], \quad (2)$$

with p_D being the distribution of the data and q the distribution of the negative samples. Optimally, the negative samples' distribution should match the EBM, $q(x) \propto \exp(-E_{\theta}(x))$. Nevertheless, in practice multiple different approaches are considered to sample the negative samples, such as uniform sampling, Langevin Markov Chain Monte Carlo (MCMC) [34] over the EBM [10], or learning an approximated sampler [35], [36]. It is a good practice to adapt the negative sampling distribution to the problem domain.

Sampling from the EBM. In contrast with most generative models [37], [38], given the implicit nature of EBMs, sampling from an EBM cannot be done explicitly. Instead, there exist multiple algorithms to generate samples from the learned energy function; from Rejection Sampling [39] to MCMC [40] with the most popular being sampling with Langevin Dynamics.

Trajectory optimization. Denoting the system state at time t to be $x_t \in \mathbb{R}^d$, we can define a discrete-time trajectory as the sequence $\tau \triangleq (x_0, x_1, ..., x_{T-1}, x_T)$ over a planning horizon T. For a given start-state x_0 , trajectory optimization aims to find the optimal trajectory τ^* which minimizes an objective function $c(\tau, x_0)$. This function might include a cost c_{g} on distance to a desired goal-state, x_g , along with a cost c_{obs} on collisions to promote obstacle avoidance. An additional term is often incorporated to penalize non-smooth trajectories, which we denote as c_{sm} . The overall objective can then be expressed as the sum of these individual terms

$$\boldsymbol{\tau}^* = \arg\min_{\boldsymbol{\tau}} c(\boldsymbol{\tau}, \boldsymbol{x}_0, \boldsymbol{x}_g) \tag{3}$$

$$= \arg\min_{\boldsymbol{\tau}} c_{obs}(\boldsymbol{\tau}) + c_g(\boldsymbol{\tau}, \boldsymbol{x}_g) + c_{sm}(\boldsymbol{\tau}, \boldsymbol{x}_0).$$
(4)

Summarizing the *context* parameters of the planning problem as $\mathcal{E} = [\mathbf{x}_0, \mathbf{x}_g, ...]^{\top}$, the objective function can also be written more generally to include any number of cost terms: $c(\boldsymbol{\tau}, \mathcal{E}) = \sum_i c_i(\boldsymbol{\tau}, \mathcal{E}).$

Gradient-based strategies for trajectory optimization typically resort to second-order iterative methods similar to Gauss-Newton [6], [41], or use pre-conditioned gradientdescent [3] to find a locally optimal solution to the objective. However, this requires that the cost function be once- or twice-differentiable, leading to carefully handcrafted cost terms such as truncated signed-distance fields (t-SDF), which must be pre-computed for a given environment.

On the other hand, sampling-based approaches resort to stochastics generation of candidate trajectories using a proposal distribution. These samples are then evaluated on the objective and weighted according to their relative performance [4], [42]. Such gradient-free optimization can operate on discontinuous costs (which may arise from contact between surfaces, for example). However, the inherent stochasticity may lead to undesirable oscillatory behavior and require additional heuristics to achieve satisfactory performance [43].

Planning as inference. The duality between probabilistic inference and optimization for planning and control has been widely explored [6], [44], [45]. With open-loop trajectory optimization, in particular, we view the trajectory τ as a random variable and first consider the target distribution

$$p(\boldsymbol{\tau}; \mathcal{E}) = \frac{1}{Z} \prod_{i} p_i(\boldsymbol{\tau}; \mathcal{E}), \qquad (5)$$

where each p_i term consists of an individual probability factor (which can also be a prior term on τ). Optimization can then be formulated as a *maximum a posteriori* (MAP) problem, where we seek to find $\tau^* = \arg \max_{\tau} p(\tau; \mathcal{E})$. This can be done by minimizing the negative-log of the distribution

$$\boldsymbol{\tau}^* = \boldsymbol{\tau}_{\text{MAP}} = \arg\min_{\boldsymbol{\tau}} - \log\prod_i p_i(\boldsymbol{\tau}; \mathcal{E}).$$
(6)

Assuming that these probability densities belong to the exponential family, we can relate them to the previous cost terms: $p_i(\tau; \mathcal{E}) \propto \exp(-c_i(\tau; \mathcal{E}))$. Substituting these into Eq. (6) recovers Eq. (3). This perspective is important, as it justifies the use of optimization methods for inference, and permits the natural integration of trajectory distributions in the optimization problem.

IV. LEARNING EBMS FOR MOTION OPTIMIZATION

Given a new context \mathcal{E} , performing inference over the posterior distribution in Eq. (5) and Eq. (6) requires that we define a prior distribution of trajectories, $p(\tau; \mathcal{E})$. Given a dataset $\mathcal{D} = \{\tau_j, \mathcal{E}_j\}_{j=1:N}$, we propose to model and *learn* such prior trajectory distributions from collected data. We define this distribution as an EBM,

$$p_{\theta}(\boldsymbol{\tau}|\mathcal{D};\mathcal{E}) = \frac{1}{Z} \exp(-E_{\theta}(\boldsymbol{\tau},\mathcal{E})).$$
(7)

with model parameters θ . In practice, \mathcal{E} represents the planning contexts, e.g., goal targets, obstacles positions, trajectory phase, etc. The dataset \mathcal{D} may consist of a collection of expert demonstrations on different environments, and we aim to fit a density function representing the data distribution with CD Eq. (2). While the learned prior distribution is based on a set of demonstrations, we desire to adapt to novel scenarios beyond the demonstrated examples. The learned distribution is expected to provide informative samples beyond the demonstrated cases. Notably, instead of learning a monolithic EBM-based prior, we can factor this prior distribution depending on various aspects of the problem

$$p_{\theta}(\boldsymbol{\tau}|\mathcal{D};\mathcal{E}) \propto \prod_{i} \exp(-E_{\theta_{i}}(\boldsymbol{\tau},\mathcal{E}_{i})).$$
 (8)

Such a factored distribution allows us to leverage composability, learn modular EBM factors independently, and combine them as needed for novel scenarios and planning



Fig. 2: Learned EBM (blue) and its gradient (orange) in a 1D dataset. (Left) EBM trained with vanilla CD. (Right) EBM trained with CD loss + denoising regularizer (9).

problems [12]. Furthermore, to properly learn EBMs for motion optimization, we need to make multiple algorithmic and modeling choices. In the following, we introduce a set of proposed choices to properly learn and represent high-dimensional, long-horizon multimodal trajectory distributions via EBMs that can be beneficial for their deployment in motion optimization.

Smoothing EBMs for gradient-based optimization. To deploy EBMs in motion optimization, we need a smooth energy landscape. However, the CD objective of Eq. (2) generates an energy landscape with multiple plateaus, with high energy values in regions where there are no data and a plateau of low-energy in the regions of the data points. While the energy landscape may capture well the distribution of the demonstrations, it might not be helpful for gradientbased sampling or optimization, with gradients close to zero in the plateau and high gradients in the cliffs (Fig. 2). We propose adding a denoising score matching loss [46], [47] as regularization to smoothen the energy landscape and improve gradient information. Denoising score matching first generates a noisy sample given a data sample $\tilde{x} \sim p(\tilde{x}|x) =$ $\mathcal{N}(\tilde{\boldsymbol{x}}|\boldsymbol{x}, \boldsymbol{\sigma}^2 \boldsymbol{I})$, as $\tilde{\boldsymbol{x}} = \boldsymbol{x} + \boldsymbol{\sigma}\boldsymbol{\epsilon}$, with $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \boldsymbol{I})$, and, then, matches the score function, $\nabla_{\boldsymbol{x}} E_{\boldsymbol{\theta}}$ to denoise the sample $\tilde{\boldsymbol{x}}$ back to \boldsymbol{x}

$$\mathcal{L}_{\text{DSM}} = \mathbb{E}_{p_{\mathcal{D}}(\boldsymbol{x},\mathcal{E})} \mathbb{E}_{p(\tilde{\boldsymbol{x}}|\boldsymbol{x})} \left[||\boldsymbol{\epsilon} - \nabla_{\boldsymbol{x}} E_{\boldsymbol{\theta}}(\tilde{\boldsymbol{x}},\mathcal{E})||_{2}^{2} \right].$$
(9)

The loss (9) encourages the gradient of the EBM to point towards the data distribution contrarily to the CD loss.

Task-specific EBMs for Motion Optimization Here, we introduce a set of model choices to represent EBMs for motion optimization, as making proper choices on the EBM architecture improves the data representation capacity and the generalization of the learned models.

Object-centric EBMs. Learning task-conditioned motion models is a vital tool for representing task-adaptive motion behaviors. In our work, we propose to learn object-centric EBM that are useful for representing desired movements in manipulation tasks that involve objects, conditioning the learned EBM on the objects' poses.

Phase-conditioned EBM. Learning density models directly on the trajectory level requires modeling in a $T \times D$ dimensional space, with T temporal horizon and D-dimensional state-space. Learning phase-conditioned density models is particularly challenging, as with long-horizon trajectories, the dimension of the input space increases, and the learning of an EBM in that space might be challenging. Additionally, T forces the temporal horizon to be discretized with a specific frequency, and it might be hard to use it for different discretization frequencies or even to represent continuoustime distributions. The usability of phase-conditioned priors for motion optimization is necessary for long-horizon tasks. Therefore, we need to learn phase-conditioned EBMs

$$p(\boldsymbol{x}|\alpha) \propto \exp(-E_{\boldsymbol{\theta}}(\boldsymbol{x},\alpha)),$$
 (10)

with x the state and α being the phase. The phase represents a continuous variable moving from 0 to 1, encoding the temporal evolution of the manipulation task. The phaseconditioned EBM represents the state-occupancy distribution for different instances of the manipulation task. Nevertheless, the phase-conditioned EBM lacks any temporal relation between temporally adjacent points, generating non-smooth trajectories. To confront this effect, we propose combining phase-conditioned EBMs with trajectory smoothing costs to represent smooth trajectory distributions

$$p(\boldsymbol{\tau}) \propto \exp\left(-\sum_{k} E_{\boldsymbol{\theta}}(\boldsymbol{x}_{k}, \alpha_{k}) + (\boldsymbol{x}_{k} - \boldsymbol{x}_{k+1})^{2}\right).$$
 (11)

V. COMBINING EBMS AND MOTION OPTIMIZATION

Although we may be able to learn data-driven prior distributions, as described in the previous section, we still require reliable inference and optimization methods to derive optimal trajectories τ^* given a new planning problem expressed by \mathcal{E} . In the following section, we focus on methods for evaluation and inference on learned EBMs for trajectory distributions. This includes techniques for sampling and optimization which account for the compose-ability of our EBM functions, as well as a stochastic trajectory optimization method suited for the planning tasks considered here.

Structured Planning Priors. Since we cannot sample from EBMs directly, we need to initialize samples by first drawing from an initial distribution, and then perform sequential updates to approximate the relevant modes. Further, learning task-specific EBM model-components is useful for portions of the objective function which are hard to define. However, we may insist on biasing our sampling given the structure of the planning problem, and incorporate known, well-defined requirements such as goal-seeking behavior and smoothness. This can be addressed by incorporating relevant distributions which are known *a-priori*, to the contextual-prior in Eq. (7)

$$p_{\boldsymbol{\theta}}(\boldsymbol{\tau}|\mathcal{D};\mathcal{E}) \propto p_0(\boldsymbol{\tau};\mathcal{E}) \prod_i \exp(-E_{\boldsymbol{\theta}_i}(\boldsymbol{\tau},\mathcal{E}_i)),$$
 (12)

where $p_0(\tau; \mathcal{E})$ is a general trajectory-based prior. Similarly to [6], [48], we can directly integrate goal-reaching and smoothness into this distribution, which can then be directly sampled to initialize optimization. This distribution is composed using a state-based start prior $p_s(x) = \mathcal{N}(\mu_s, \Sigma_s)$, goal-directed prior $p_g(x) = \mathcal{N}(\mu_g, \Sigma_g)$, where μ_s, μ_g are the expected start and goal configurations, respectively. Smoothness is defined using a time-correlated Gaussian-Process prior p^{gp} derived from a linear time varying motion model (see Appendix A in [49] for details). These terms can then be composed

$$p_0(\boldsymbol{\tau}; \mathcal{E}) = p_s(\boldsymbol{x}_0) \ p_g(\boldsymbol{x}_T) \prod_{t=0}^{T-1} p_t^{gp}(\boldsymbol{x}_t, \boldsymbol{x}_{t+1})$$
 (13)

$$\propto \exp\left(-\frac{1}{2}\left|\left|\boldsymbol{\tau}-\boldsymbol{\mu}_{\boldsymbol{\tau}}\right|\right|_{\mathcal{K}}^{2}
ight)$$
 (14)

where $\mathcal{E} = \{\mu_s, \mu_g\}, \mathcal{K}$ is a block-diagonal covariance matrix, and μ_{τ} the straight-line trajectory defined in configuration space. This results in a closed-form, explicit distribution with low entropy near start and goal regions, and highentropy towards the middle of a trajectory, which is ideal for goal-centric planning tasks. In contrast to [4], the integrated dynamics includes higher-order terms, such that velocities and accelerations can also be sampled in a principled manner. In practice, we can efficiently generate large quantities of these time-correlated trajectories due to our parallelized GPU implementation. Note, however, it is intractable to incorporate explicit priors on behaviors such as obstacle avoidance, e.g. configuration space obstacle avoidance prior. Hence, in practice we must resort to a combination of implicit and explicit priors to generate feasible trajectories from Eq. (12). Stochastic Trajectory Optimization with GP-Priors. We can iteratively update the time-correlated sampling prior. described in the previous section, to fit the modes of a learned EBM and allow us to sample optimal trajectories in a new planning context. The optimizer, which we call Stochastic Gaussian Process Motion Planning (StochGPMP), is closely related to the importance sampling scheme used by CEM and MPPI but uses goal-directed GP distributions. We refer the reader to the Appendix on the project site for details on the derivation and specific update procedure. When used in our experiments, we select to update the mean of each component GP in the distribution. This can be easily performed in configuration space, although it requires an initial goal configuration to be approximated using Inverse-Kinematics.

VI. EXPERIMENTAL EVALUATION

Experiment I: Simulated Planar Navigation. We begin by testing our framework on a simple planar navigation problem, where a holonomic robot must reach a goal location while avoiding obstacles. We assume that the start, goal, and obstacle locations are known for a given planning problem, but the obstacle geometries (ex. size, shape) are unknown. We want to learn an implicit distribution that captures the collision-free trajectories which lead to a particular goal. Here, we investigate two possible sources of empirical data: (1) sparsely populated point-distributions in free-space and (2) a set of expert trajectory distributions. The former can be seen as a stand-in for free-space measurements taken from a depth sensor, e.g., a lidar. In practice, we generate these points by sampling uniformly throughout the 2-D plane and rejecting points in collision. The expert trajectories are generated by running the StochGPMP planner at a high covariance for many (e.g., 500) iterations, producing a sparse set of trajectories for different goal locations and obstacle placements. We collect data for 512 different environments and produce trajectories from 15 random goal locations per environment, with 5 trajectories per goal. Both sources are depicted in Fig. 3. Here, the learned EBM can be expressed as $E_{\theta}(x, \{x_{obs}^i\}_{i=1}^N)$, where x is a particular 2D-state, and x_{obs} the position of an obstacle (here, N = 3).



Fig. 3: Examples of positive sampling distributions for a single obstacle environment. (Left) Uniform point-based sampling with rejection. (Right) Expert trajectory distributions for randomly sampled goal locations along the top and right side, with start location in the bottom-left corner.

Examples of the learned EBM for both cases are shown in Fig. 4. The resulting energy functions, in either case, manage to effectively capture the demonstration distributions, conditioned on new obstacle locations. With Section V in mind, we can generate trajectories near the modes of the full target distribution, accounting for smoothness and distance-to-goal, by first sampling from a structured prior and optimizing on the EBM to capture key modes of the distribution. We compare this method of implicit trajectory generation to a standard Behavioral Cloning (BC) baseline, where the learned policy outputs the current velocity, $\dot{q} = f(q; \mathcal{X}_o, x_q)$ which is conditioned on the set of obstacle poses $\mathcal{X}_o = \{x_o\}$ and the goal location x_q . The trajectories generated from this baseline can be seen in Fig. 5, obstacle avoidance is minimal. We perform a quantitative analysis on the EBM methods by measuring the success rate on the validation set as a function of optimization iterations needed by the planner (with the BC baseline fixed). This is shown in Table I.

Opt. iters.	0	5	10	25	50
EBM-Free-space	0.556	0.470	0.643	0.747	0.852
EBM-Expert Traj.	0.556	0.690	0.791	0.847	0.877
Behavioral cloning	0.04	-	-	-	-

TABLE I: Planar navigation: Average success rate per environment, as a function of optimization iterations. A planning trajectory is deemed successful if it ends within a radius of 1.5 from the goal, without hitting the underlying obstacles.

Experiment II: Simulated Planar Manipulator. This experiment studies the performance improvement in robotics manipulation tasks when introducing our expressive implicit priors in trajectory optimization. First, this experiment demonstrates how well the EBMs provide an informative prior landscape that biases the trajectory optimizer towards better solution regions than standard Gaussian priors. Second, it shows how the learned implicit priors on the task space result in superior performance when warm-starting the trajectory optimization, compared against explicit priors in configuration space, commonly done in literature [8].

We create a simulated planar robot-arm with 3 dofs operated on a plane containing fixed red obstacles. The planner has to find an optimal low-jerk trajectory that guides the planar arm to grasp the cubic object and inserts the object into a 2-walled cubby without colliding with the cubby or the red obstacles



Fig. 4: Learned obstacle-EBMs conditioned on novel obstacle locations. (Left) free-space point sampling and (right) expert trajectory distributions, with multi-goal planning solutions depicted by blue trajectories. Discontinuities and implicit obstacle surfaces are well captured using sparse free-space point-samples during training, whereas distributions of trajectory-based demonstrations can be captured neatly by the EBMs. The latter provides a convenient "guiding" energy function for a new context, improving samples derived from multi-modal stochastic trajectory optimization.



Fig. 5: Examples of trajectory traces using the behavioral cloning baseline on the planar navigation problem. Although rolled-out sample trajectories are generally well-behaved, they fail to avoid obstacles.

(Fig. 6). Our learned EBMs represent the distribution of possible grasps and insertions that are used to generate trajectories that will bias the trajectory optimization problem towards the optimal solution. The learned energy distribution landscape is displayed in Fig. 7, where the centers are the projected 2D-origin of the object-centric frame. We train both EBMs using standard training methods as described in Section IV. When combining both implicit priors, the intuition is that the optimizer will plan the configuration trajectory that biases the end-effector toward the low-energy regions associated jointly with the most natural grasp and insert postures, given the current environmental setting. In all experiments, the grasp and insertion objectives are superpositioned with other standard costs such as smoothness, obstacle avoidance, and joint limits to completely define the trajectory optimization problem as in (3).

We compare our method against the two baselines, where the first employs Gaussian distributions as priors to define the object's grasp and insertion potentials in the task space (see Table II) as commonly done in [3], [4]. For the second baseline, we first learn a behavioral cloning model $\dot{q} = f(q; x_o, x_g)$ where x_o, x_g represents the object's grasp pose and cubby pose, while the demonstrated configuration-space trajectories are generated for simplicity from a planner. Next, we warm start the optimizer with the trajectories sampled from the behavioral cloning model, where each trajectory is rolled out by integrating the model velocity output from the initial robot configuration. We expect that the second baseline is a stronger baseline due to the integration of an explicit



Fig. 6: Simulated planar arm executing the grasp & insert tasks.

Cost terms	Mathematical expression
Joint limit	$\max(\boldsymbol{q} - \boldsymbol{q}_{max}, 0) + \max(\boldsymbol{q}_{min} - \boldsymbol{q}, 0)$
Sphere Obstacle Avoidance	$\sum_{o=1}^{K} \ \max(\boldsymbol{x}_{o} + \epsilon - \boldsymbol{x}(\boldsymbol{q}), 0)\ ^{2}, K \text{ is object number}$
Pose potential	$\frac{1}{2} \ \boldsymbol{x}(\boldsymbol{q}) - \boldsymbol{x}_n \ ^2$

TABLE II: Elementary cost terms



Fig. 7: The energy landscape of grasp & insert learned models. (Left) Learned energy for grasping points. (Right) Learned energy for placing in the cubby. Both energies are learned in the object-centric frame of the graspable object and the placing cubby. We generate the training data by motion planning. To cover all the possible placing trajectories, we run the motion planner given several initial end-effector poses.



Fig. 8: Success rate (%) of grasp & insert task over optimization steps. Surprisingly, even with few optimization steps of 50 & 100, our implicit priors' success rate is already higher comparing to the baselines. This shows that using informative and expressive EBMs can greatly assist the planner even with only few optimization steps. As optimization steps increase, our informative implicit priors still consistently outperforms previous methods.

prior. Note that for the two baselines, we randomly sample a desired grasp pose on the object's surface for setting up the grasp potential cost, and we keep fixed the target placement pose at the center of the cubby.

To show the strength of our implicit priors in trajectory optimization, we measure the success rate of successful grasping and insertions without collisions with obstacles or the target cubby. The success rates are reported over an increasing number of optimization steps in Fig. 8 to emphasize the benefit of our implicit priors in the multiple stages of the optimization, from a few steps till convergence. We run experiments with each method four times. Due to the stochasticity of the planner, we run 10 planning instances for each experiment and method, where at each instance we randomly sample a pair of a reachable cubic object pose and a target pose on the cubby. For all instances, the trajectory is fixed to 80 time-steps and the robot arm always starts from the same initial configuration.

We notice that most fail cases of the baselines are usually due to the collision with the cubby wall. It is because radial priors such as Gaussians do not encode the insertion dynamics, and the planner has to balance between the collision avoidance and the target pose objectives in a narrow placement target. We also observe that while warm starting with explicit priors does help the planner find a better solution than the Gaussian prior, sometimes the rollout trajectory of the behavioral cloning model suffers from covariate shift, that worsens the performance of the planner leading to fail cases.

Experiment III: Robot pouring amid obstacles. This experiment studies the performance gains of introducing learned implicit priors for a more complex manipulation task in a higher-dimensional space. Namely, we investigate the integration of EBMs in trajectory optimization for a pouring task in the presence of obstacles with a 7dof LWR-Kuka robot arm. This experiment investigates (i) the benefit of including smoothness regularization in the EBM training, (ii) the advantages of phase-conditioned EBM w.r.t. learning the EBM in trajectory space, and (iii) the generalization of our EBMs in the context of the pouring task regarding arbitrary pouring places, and in the presence of obstacles. Instances of our method's performance are available in Fig. 9 for the simulated task, and Fig. 1 for our zero-transfer to the actual robotic setup. We add additional details on the problem setup in the Appendix on the project site.

To learn the pouring EBM, we recorded 500 trajectory demonstrations of the pouring task. The demonstrations were generated using a set of handcrafted policies. The demonstrations were initialized in arbitrary initial configurations and performed the pouring from different positions, generating a multimodal distribution of demonstrated trajectories. The demonstrations have been directly recorded for the glass pose and centered in the frame of the pouring pot. To properly encode the temporal information in the data, we learn a time conditioned EBM, $E_{\theta}(\boldsymbol{x}|\alpha)$. In our problem, \boldsymbol{x} is a 6dimensional state, representing the 3D position of the bottom and tip of the glass w.r.t. the pouring pot frame. Centering the EBM to the pouring pot's frame allows us to generalize the EBM to arbitrary pot poses. Additionally, we include the denoising regularization and compare its performance and compare to a baseline without the proposed regularization. We compare against three baselines. First, a solver without any prior, to appropriately evaluate the benefits of adding guiding priors. Second, a phase-conditioned EBM without smoothness regularization combined with the optimizer. Third, an EBM that is directly learned in the trajectory space to investigate the benefit of phase-conditioned EBMs in trajectory optimization. Given the learned pouring prior, we optimize the task trajectory by gradient descent. The objective function is defined by the composition of a set of cost functions-fixed initial configuration, fixed target configuration, trajectory smoothness, obstacle avoidance, keep the glass pointing up to avoid spilling and pour inside the pot (Table II). The learned EBM is added as an additional



Fig. 9: A visualization of the pouring in cluttered task. The robot should avoid the obstacles, pour in the pot over the table and come back to the initial position.



Fig. 10: Success rate (%) for the pouring task. (Left): Experiment without collision obstacles. The success rate of the motion optimization without any learned EBM is almost the maximum. We observe that the phase EBM with regularization can improve the performance slightly. (Right): Experiment with obstacles in the environment. We observe a clear benefit of using phase-based EBM in contrast with trajectory-based EBM. Training the EBM in high dimensional space requires too many samples, and it is difficult to get smooth EBMs representing the demonstrations. The EBM trained with regularization improves the obtained results with respect to non-regularized EBM. For any initial position of the trajectory particles, the regularized EBM provides an informative gradient towards the demonstrations, while in the case of non-regularized EBM, we found the gradient was informative only in the close vicinity of the demonstrations.

factor in the optimization problem. We optimize using a tempering scheme, giving more importance to the prior at the beginning and reducing its influence at the end of the optimization process.

To evaluate the pouring EBM, we measure the success rate of the motion optimization to perform the pouring task successfully. We consider an episode to be successful if it satisfies three conditions: (i) there are no collisions, (ii) the robot does not spill outside the pot, and (iii) the robot pours inside the pot. We report performance both in obstructed and obstacle-free environments. We run 50 episodes for each case, we randomize the position of the pouring pot and the obstacles on each episode. The obtained results are presented in Fig. 10. We can clearly observe the benefit of both integrating smoothing regularizers in the training of the EBM and modelling the EBMs as phase-conditioned models. The presented results are complemented with the attached video for the simulated tasks, which additionally showcases our successful transfer to the real robotic system. Note that our learned priors do not suffer for sim2real mismatch, as the only requirement is the knowledge of the object's pose, which in our case, we resolve through motion tracking.

VII. DISCUSSION AND CONCLUSIONS

Energy-based models (EBM) have multiple desirable properties that make them obvious candidates to be integrated as data-driven priors in robotics. EBMs can represent highly multimodal distributions, and due to their implicit nature, they can naturally be composed and integrated into trajectory optimization problems. Nevertheless, the high dimensionality and variability of the tasks in robotics hinder the learning of generalizable EBMs that could apply to a wide range of complex manipulation tasks. In this paper, we demonstrate that learning EBMs as implicit priors can guide motion optimization for common robotics problems. We introduce a set of modeling choices to represent generalizable EBMs that could be adapted to an arbitrary set of manipulation tasks. We propose alternative training losses to learn smoother EBMs that could provide well-conditioned gradients into optimization problems. Moreover, we introduce novel structured sampling procedures to combine data-driven models with explicit priors. Our experimental results validated the efficacy and applicability of our proposed algorithmic choices for motion optimization. In future works, we aim to explore the problem of learning conditioned EBMs given higher dimensional contexts, such as images or text.

REFERENCES

- J. J. Kuffner and S. M. LaValle, "Rrt-connect: An efficient approach to single-query path planning," in *IEEE ICRA*, 2000.
- [2] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics and Automation*, 1996.
- [3] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *IEEE ICRA*, 2009.
- [4] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "Stomp: Stochastic trajectory optimization for motion planning," in *IEEE ICRA*, 2011.
- [5] J. Schulman *et al.*, "Motion planning with sequential convex optimization and convex collision checking," *IJRR*, 2014.
- [6] M. Mukadam, J. Dong, X. Yan, F. Dellaert, and B. Boots, "Continuoustime gaussian process motion planning via probabilistic inference," *IJRR*, 2018.
- [7] J. Urain, M. Ginesi, D. Tateo, and J. Peters, "Imitationflow: Learning deep stable stochastic dynamic systems by normalizing flows," in 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 5231–5237, IEEE, 2020.
- [8] D. Koert, G. Maeda, R. Lioutikov, G. Neumann, and J. Peters, "Demonstration based trajectory optimization for generalizable robot motions," in *IEEE-RAS 16th Humanoids*, 2016.
- [9] M. A. Rana *et al.*, "Towards robust skill generalization: Unifying learning from demonstration and motion planning," in *CoRL*, PMLR, 2017.
- [10] Y. Du and I. Mordatch, "Implicit generation and generalization in energy-based models," *arXiv preprint arXiv:1903.08689*, 2019.
- [11] Y. LeCun, S. Chopra, R. Hadsell, M. Ranzato, and F. Huang, "A tutorial on energy-based learning," *Predicting structured data*, 2006.
- [12] J. Urain, A. Li, P. Liu, C. D'Eramo, and J. Peters, "Composable energy policies for reactive motion generation and reinforcement learning," in *R:SS*, 2021.
- [13] K. He, E. Martin, and M. Zucker, "Multigrid chomp with local smoothing," in *IEEE-RAS 13th Humanoids*, 2013.
- [14] A. Byravan et al., "Space-time functional gradient optimization for motion planning," in *IEEE ICRA*, 2014.
- [15] Z. Marinho, A. Dragan, A. Byravan, B. Boots, S. Srinivasa, and G. Gordon, "Functional gradient motion planning in reproducing kernel hilbert spaces," arXiv preprint arXiv:1601.03648, 2016.
- [16] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *IJRR*, 2013.
- [17] J. Ibarz, J. Tan, C. Finn, M. Kalakrishnan, P. Pastor, and S. Levine, "How to train your robot with deep reinforcement learning: lessons we have learned," *IJRR*, 2021.
- [18] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, J. Peters, et al., "An algorithmic perspective on imitation learning," *Foundations* and *Trends*® in *Robotics*, 2018.
- [19] F. Torabi, G. Warnell, and P. Stone, "Behavioral cloning from observation," in *Proceedings of the 27th IJCAI*, 2018.
- [20] T. Zhang, Z. McCarthy, O. Jow, D. Lee, X. Chen, K. Goldberg, and P. Abbeel, "Deep imitation learning for complex manipulation tasks from virtual reality teleoperation," in *IEEE ICRA*, 2018.

- [21] B. Ichter, J. Harrison, and M. Pavone, "Learning sampling distributions for robot motion planning," in *IEEE ICRA*, 2018.
- [22] J. Ortiz-Haro, J.-S. Ha, D. Driess, and M. Toussaint, "Structured deep generative models for sampling on constraint manifolds in sequential manipulation," in *CoRL*, PMLR, 2022.
- [23] C. Chamzas *et al.*, "Learning sampling distributions using local 3d workspace decompositions for motion planning in high dimensions," in *IEEE ICRA*, 2021.
- [24] A. H. Qureshi, Y. Miao, A. Simeonov, and M. C. Yip, "Motion planning networks: Bridging the gap between learning-based and classical motion planners," *IEEE Transactions on Robotics*, 2020.
- [25] R. Cheng, K. Shankar, and J. W. Burdick, "Learning an optimal sampling distribution for efficient motion planning," in *IEEE/RSJ IROS*, 2020.
- [26] Z. Kingston, M. Moll, and L. E. Kavraki, "Exploring implicit spaces for constrained sampling-based planning," *IJRR*, 2019.
- [27] S. Levine and V. Koltun, "Continuous inverse optimal control with locally optimal examples," arXiv preprint arXiv:1206.4617, 2012.
- [28] C. Finn, S. Levine, and P. Abbeel, "Guided cost learning: Deep inverse optimal control via policy optimization," in *International conference* on machine learning, PMLR, 2016.
- [29] A. Y. Ng, S. J. Russell, et al., "Algorithms for inverse reinforcement learning," in *ICML*, 2000.
- [30] P. Abbeel and A. Y. Ng, "Apprenticeship learning via inverse reinforcement learning," in *ICML*, 2004.
- [31] B. D. Ziebart, A. L. Maas, J. A. Bagnell, A. K. Dey, *et al.*, "Maximum entropy inverse reinforcement learning," in AAAI, 2008.
- [32] Y. Song and D. P. Kingma, "How to train your energy-based models," arXiv preprint arXiv:2101.03288, 2021.
- [33] G. E. Hinton, "Training products of experts by minimizing contrastive divergence," *Neural computation*, 2002.
- [34] G. Parisi, "Correlation functions and computer simulations," Nuclear Physics B, vol. 180, no. 3, pp. 378–384, 1981.
- [35] R. Kumar, S. Ozair, A. Goyal, A. Courville, and Y. Bengio, "Maximum entropy generators for energy-based models," *arXiv preprint* arXiv:1901.08508, 2019.
- [36] W. Grathowhl, J. Kelly, M. Hashemi, M. Norouzi, K. Swersky, and D. Duvenaud, "No mcmc for me: Amortized sampling for fast and stable training of energy-based models," *arXiv preprint arXiv:2010.04230*, 2020.
- [37] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *NeurIPS*, 2014.
- [38] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *ICLR* (Y. Bengio and Y. LeCun, eds.), 2014.
- [39] W. R. Gilks and P. Wild, "Adaptive rejection sampling for gibbs sampling," *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 1992.
- [40] C. Andrieu, N. De Freitas, A. Doucet, and M. I. Jordan, "An introduction to mcmc for machine learning," *Machine learning*, 2003.
- [41] M. Toussaint, "Newton methods for k-order markov constrained motion problems," arXiv preprint arXiv:1407.0414, 2014.
- [42] Z. I. Botev et al., "The cross-entropy method for optimization," in Handbook of statistics, Elsevier, 2013.
- [43] M. Bhardwaj, B. Sundaralingam, A. Mousavian, N. D. Ratliff, D. Fox, F. Ramos, and B. Boots, "Storm: An integrated framework for fast joint-space model-predictive control for reactive manipulation," in *CoRL*, PMLR, 2022.
- [44] M. Toussaint, "Robot trajectory optimization using approximate inference," in *Proceedings of the 26th annual international conference on machine learning*, pp. 1049–1056, 2009.
- [45] S. Levine, "Reinforcement learning and control as probabilistic inference: Tutorial and review," arXiv preprint arXiv:1805.00909, 2018.
- [46] P. Vincent, "A connection between score matching and denoising autoencoders," *Neural computation*, 2011.
- [47] Y. Song *et al.*, "Score-based generative modeling through stochastic differential equations," in *ICLR*, 2021.
- [48] A. Lambert and B. Boots, "Entropy regularized motion planning via stein variational inference," *RSS Workshop on Integrating Planning* and Learning, 2021.
- [49] K. V. Alwala and M. Mukadam, "Joint sampling and trajectory optimization over graphs for online motion planning," in *IEEE/RSJ IROS*, 2020.

- [50] T. D. Barfoot, C. H. Tong, and S. Särkkä, "Batch continuous-time trajectory estimation as exactly sparse gaussian process regression.," in *Robotics: Science and Systems*, vol. 10, pp. 1–10, Citeseer, 2014.
- [51] S. Särkkä, A. Solin, and J. Hartikainen, "Spatiotemporal learning via infinite-dimensional bayesian filtering and smoothing: A look at gaussian process regression through kalman filtering," *IEEE Signal Processing Magazine*, vol. 30, no. 4, pp. 51–61, 2013.
- [52] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, "Information theoretic mpc for model-based reinforcement learning," in 2017 IEEE International Conference on Robotics and Automation (ICRA), pp. 1714–1721, IEEE, 2017.
- [53] N. Wagener, C.-A. Cheng, J. Sacks, and B. Boots, "An online learning approach to model predictive control," *arXiv preprint* arXiv:1902.08967, 2019.
- [54] 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*, vol. 12, 2016.
- [55] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer school on machine learning*, pp. 63–71, Springer, 2003.
- [56] K. Kolur, S. Chintalapudi, B. Boots, and M. Mukadam, "Online motion planning over multiple homotopy classes with gaussian process inference," in 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 2358–2364, IEEE, 2019.
- [57] A. Lambert, A. Fishman, D. Fox, B. Boots, and F. Ramos, "Stein variational model predictive control," *arXiv preprint arXiv:2011.07641*, 2020.
- [58] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, highperformance deep learning library," 2019.
- [59] A. Paraschos, C. Daniel, J. R. Peters, and G. Neumann, "Probabilistic movement primitives," *Advances in neural information processing* systems, vol. 26, 2013.
- [60] R. A. Shyam, P. Lightbody, G. Das, P. Liu, S. Gomez-Gonzalez, and G. Neumann, "Improving local trajectory optimisation using probabilistic movement primitives," in 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 2666–2671, IEEE, 2019.
- [61] T. Osa, "Multimodal trajectory optimization for motion planning," *The International Journal of Robotics Research*, vol. 39, no. 8, pp. 983–1001, 2020.
- [62] T. Kim, G. Park, J. Bae, and W. Lee, "Derivative action control: Smooth model predictive path integral control without smoothing," arXiv preprint arXiv:2112.09988, 2021.
- [63] L. Petrović, J. Peršić, M. Seder, and I. Marković, "Cross-entropy based stochastic optimization of robot trajectories using heteroscedastic continuous-time gaussian processes," *Robotics and Autonomous Systems*, vol. 133, p. 103618, 2020.
- [64] L. Petrović, I. Marković, and I. Petrović, "Mixtures of gaussian processes for robot motion planning using stochastic trajectory optimization," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2022.

APPENDIX I Stochastic Gaussian Process Motion Planning

A. Gaussian Process Trajectory Prior

We can define a prior on the space of smooth, continuous-time trajectories $\boldsymbol{x}(t)$ using a Gaussian Process [6], [50]: $\boldsymbol{x}(t) \sim \mathcal{GP}(\boldsymbol{\mu}(t), \mathcal{K}(t, t'))$, with mean function $\boldsymbol{\mu}$ and covariance function \mathcal{K} . This distribution is parameterized by a discrete set of support states $\boldsymbol{\tau} = [\boldsymbol{x}_0, \boldsymbol{x}_1, ..., \boldsymbol{x}_N]^\top$, which constitutes a discrete-time trajectory defined along times $\boldsymbol{t} = [t_0, t_1, ..., t_N]$. We can then consider the distribution over $\boldsymbol{\tau}$ as $p(\boldsymbol{\tau}) = \mathcal{N}(\boldsymbol{\mu}, \mathcal{K})$, with mean $\boldsymbol{\mu} = [\boldsymbol{\mu}(t_0), ..., \boldsymbol{\mu}(t_N)]^\top$ and covariance matrix $\mathcal{K} = [\mathcal{K}(t_n, t_m)]|_{nm, 0 \le n, m \le N}$.

As described in [50], [51], a GP prior can be constructed in a principled manner by considering the structured covariance function created by a linear time-varying stochastic differential equation (LTV-SDE)

$$\dot{\boldsymbol{x}} = \mathbf{A}(t)\boldsymbol{x}(t) + \mathbf{u}(t) + \mathbf{F}(t)\mathbf{w}(t)$$
(15)

where $\mathbf{u}(t)$ is the control input, $\mathbf{A}(t)$ and $\mathbf{F}(t)$ the time-varying system matrices, and $\mathbf{w}(t)$ is a disturbance following the white-noise process $\mathbf{w}(t) \sim \mathcal{GP}(0, \mathbf{Q}_c \delta(t - t'))$, where \mathbf{Q}_c is the power-spectral density matrix.

Given a prior on start state $p_s(\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s)$ and end-goal position $p_g(\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g)$, the trajectory prior can be defined as the product of unary and binary factors

$$p_{\text{GP}}(\boldsymbol{\tau}) \propto \exp\left(-\frac{1}{2}||\boldsymbol{\tau}-\boldsymbol{\mu}||_{\boldsymbol{\mathcal{K}}}^{2}
ight)$$
 (16)

$$\propto p_s(\boldsymbol{x}_0) p_g(\boldsymbol{x}_N) \prod_{i=0}^{N-1} p_i^{gp}(\boldsymbol{x}_i, \boldsymbol{x}_{i+1})$$
(17)

Following [50], each factor is then defined as

$$p^{gp}(\boldsymbol{x}_{i}, \boldsymbol{x}_{i+1}) \propto \exp\left(-\frac{1}{2} \|\boldsymbol{\Phi}_{i,i+1}(\boldsymbol{x}_{i} - \boldsymbol{\mu}_{i}) - (\boldsymbol{x}_{i+1} - \boldsymbol{\mu}_{i+1})\|_{\mathbf{Q}_{i,i+1}}^{2}\right),$$

$$p_{s}(\boldsymbol{x}_{0}) \propto \exp\left(-\frac{1}{2} \|\boldsymbol{x}_{0} - \boldsymbol{\mu}_{s}\|_{\boldsymbol{\Sigma}_{s}}^{2}\right),$$

$$p_{g}(\boldsymbol{x}_{N}) \propto \exp\left(-\frac{1}{2} \|\boldsymbol{x}_{N} - \boldsymbol{\mu}_{g}\|_{\boldsymbol{\Sigma}_{g}}^{2}\right)$$
(18)

where $\Phi_{i,i+1}$ is the state transition matrix and $\mathbf{Q}_{i,i+1}$ the variance between times t_i and t_{i+1} . For a first-order LTV-SDE, where the state consists of position and velocity terms $\boldsymbol{x} = [\mathbf{x}, \dot{\mathbf{x}}]^{\top}$, $\mathbf{x}, \dot{\mathbf{x}} \in \mathbb{R}^d$, the state transition and covariance matrices equate to

$$\mathbf{\Phi}_{i,i+1} = \begin{bmatrix} \mathbf{I} & \Delta t_i \mathbf{I} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad \mathbf{Q}_{i,i+1} = \begin{bmatrix} \frac{1}{3} \Delta t_i^3 \mathbf{Q}_c & \frac{1}{2} \Delta t_i^2 \mathbf{Q}_c \\ \frac{1}{2} \Delta t_i^2 \mathbf{Q}_c & \Delta t_i \mathbf{Q}_c \end{bmatrix}$$

where $\Delta t_i = t_{i+1} - t_i$. Sampling from this distribution using the time-correlated GP covariance results in trajectories which have low variance near the start and goal positions (depending on the values of Σ_s and Σ_g), and higher variance near the middle of the horizon length. An example of smooth, goal-directed trajectories sampled from a one-dimensional GP-prior is depicted in Fig. 11.

B. Stochastic Optimization for Motion Planning

As described in the main text, we can express the motion-planning problem as finding the trajectory which minimizes the negative log-likelihood of a factored probability distribution

$$\boldsymbol{\tau}^* = \arg\min_{\boldsymbol{\tau}} -\log p_{\rm GP}(\boldsymbol{\tau}) \prod_k p_k(\boldsymbol{\tau})$$
(19)

with non-negative functions $p_k(\tau) \propto \exp(-E_k(\tau))$ encoding additional constraints, such as obstacle collision costs and joint limit violations.

Alternatively, we can view the problem of *stochastic* motion planning in terms of distribution matching, where we seek to match a proposal trajectory distribution $q(\tau)$ to some target distribution $q^*(\tau)$. This is perhaps most clearly described by the information theoretic derivation of model-predictive path integral (MPPI) control in [52]. Here, we consider a system with general discrete-time dynamics defined by $X = F(\mathbf{x}_0, V)$, for a state-sequence $X = (\mathbf{x}_0, \mathbf{x}_1, ..., \mathbf{x}_N)$, control-trajectories



Fig. 11: An example GP prior for 1D trajectories. The dashed line is the mean constant trajectory. The 32 solid lines are sampled trajectories from the GP prior.

 $V = (\mathbf{v}_0, \mathbf{v}_1, ..., \mathbf{v}_N)$, and start-state \mathbf{x}_0 . The target distribution over control-trajectories $V = (\mathbf{v}_0, \mathbf{v}_1, ..., \mathbf{v}_N)$ is then maximumentropy Gibbs distribution

$$q^*(V) = \frac{1}{\eta} \exp\left(-\frac{1}{\lambda}S(V)\right) p(V)$$
(20)

where $p(V) = \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$, a zero-mean centered normal distribution, $S(\cdot)$ the cost function, λ a scaling term, and η a normalizing factor. Note that $q^*(V)$ can be interpreted as a *posterior* distribution, with a likelihood function provided by $\exp\left(-\frac{1}{\lambda}S(V)\right)$ and prior p(V). Given the open-loop characterization of the system, we can consider distributions directly over V, leaving the state-dynamics $F(\cdot, \cdot)$ as part of the cost-evaluation function $S = E \circ F(\mathbf{x}_0, V)$, for some function E.

For a kinematic system, the control sequence can be defined as a time-derivative of state, ex. $V = (\dot{\mathbf{x}}_0, \dot{\mathbf{x}}_1, ..., \dot{\mathbf{x}}_N)$, and we can instead consider the joint distribution over *both* state-trajectories and their derivatives: $\boldsymbol{\tau} = (X, V) = (\mathbf{x}_0, \dot{\mathbf{x}}_0, \mathbf{x}_1, \dot{\mathbf{x}}_1, ..., \mathbf{x}_N, \dot{\mathbf{x}}_N)$. As such, we can similarly define the optimal distribution to be

$$q^*(\boldsymbol{\tau}) = \frac{1}{\eta} \exp\left(-\frac{1}{\lambda} E(\boldsymbol{\tau})\right) p_F(\boldsymbol{\tau})$$
(21)

where the dynamics F are now integrated into the prior distribution $p_F(\tau)$, describing the un-controlled stochastic dynamics. Fortunately, we have already derived an example of such a distribution in Eq. (16), using the dynamics in Eq. (15). Leveraging the time-correlated prior covariance \mathcal{K} , we can define the nominal dynamics distribution as $p_F(\tau) = \mathcal{N}(\mathbf{0}, \mathcal{K})$.

Using the proposal distribution $q(\tau; \theta) = \mathcal{N}(\mu, \mathcal{K})$ with parameters $\theta = (\mu, \mathcal{K})$, we can then minimize the KL-divergence with the target distribution

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \operatorname{KL}\left(q^*(\boldsymbol{\tau}) \parallel q(\boldsymbol{\tau}; \boldsymbol{\theta})\right) \tag{22}$$

Introducing the nominal trajectory distribution $p_F(\tau)$ into the equation

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \int q^*(\boldsymbol{\tau}) \log \frac{q^*(\boldsymbol{\tau})}{q(\boldsymbol{\tau};\boldsymbol{\theta})} \frac{p_F(\boldsymbol{\tau})}{p_F(\boldsymbol{\tau})} d\boldsymbol{\tau}$$
(23)

$$= \arg\min_{\boldsymbol{\theta}} - \int q^*(\boldsymbol{\tau}) \log \frac{q(\boldsymbol{\tau}; \boldsymbol{\theta})}{p_F(\boldsymbol{\tau})} d\boldsymbol{\tau}$$
(24)

where the cross-entropy objective in the last line results from independence of the entropy $\int q^*(\tau) \log \frac{q^*(\tau)}{p_F(\tau)} d\tau$ from θ .

In order to derive a practical algorithm for solving this objective, we can simplify the problem by optimizing for only the mean trajectory $\theta = \mu$, keeping \mathcal{K} as a fixed covariance. Sampling trajectories from $q^*(\tau)$ directly is typically intractable. Following the derivation in [52], we can instead apply importance sampling such that the proposal distribution $q(\tau; \theta)$ can

be used for evaluating the expectation in Eq. (23). We then arrive at the update rule to the proposal distribution parameter

$$\boldsymbol{\mu}_{new} = \boldsymbol{\mu} + \frac{\mathbb{E}_{\boldsymbol{\tau} \sim q(\boldsymbol{\tau};\boldsymbol{\mu})} \left[\exp\left(-\frac{1}{\lambda} E(\boldsymbol{\tau}) + \alpha(\boldsymbol{\tau})\right) (\boldsymbol{\tau} - \boldsymbol{\mu}) \right]}{\mathbb{E}_{\boldsymbol{\tau} \sim q(\boldsymbol{\tau};\boldsymbol{\mu})} \left[\exp\left(-\frac{1}{\lambda} E(\boldsymbol{\tau}) + \alpha(\boldsymbol{\tau})\right) \right]}$$
(25)

where the term $\alpha(\tau) = -\frac{1}{2}\mu \mathcal{K}^{-1}\tau$ results from the importance-sampling factor.

The expectations are approximated via Monte Carlo sampling, where K trajectories are drawn directly from the proposal distribution: $\{\boldsymbol{\tau}^k\}_{k=0}^K \sim q(\boldsymbol{\tau}; \boldsymbol{\mu})$. Using the notation $\tilde{E}(\boldsymbol{\tau}) = E(\boldsymbol{\tau}) - \lambda \alpha(\boldsymbol{\tau})$, we have

$$\boldsymbol{\mu}_{new} = \boldsymbol{\mu} + \frac{\sum_{k=0}^{K} \exp\left(-\frac{1}{\lambda} \tilde{E}(\boldsymbol{\tau}^{k})\right) (\boldsymbol{\tau}^{k} - \boldsymbol{\mu})}{\sum_{k=0}^{K} \exp\left(-\frac{1}{\lambda} \tilde{E}(\boldsymbol{\tau}^{k})\right)}$$
(26)

$$=\boldsymbol{\mu} + \sum_{k=0}^{K} w^{k} (\boldsymbol{\tau}^{k} - \boldsymbol{\mu})$$
(27)

where w^k is the softmax-weight for τ with respect to the other samples. Similarly to [53], we can also incorporate a step-size, or smoothing parameter $\gamma \in [0, 1]$

$$\boldsymbol{\mu}_{new} = \boldsymbol{\mu} + \gamma \sum_{k=0}^{K} w^k (\boldsymbol{\tau}^k - \boldsymbol{\mu})$$
(28)

$$= (1 - \gamma)\boldsymbol{\mu} + \gamma \sum_{k=0}^{K} w^k \boldsymbol{\tau}^k.$$
⁽²⁹⁾

This update can be then applied in an iterative fashion until a converged trajectory distribution is reached. The resulting *posterior* distribution approximation, $q_{final}(\tau)$, will be parameterized by a mean trajectory $\boldsymbol{\mu}_{final} = [\boldsymbol{x}_0, \boldsymbol{x}_1, ..., \boldsymbol{x}_N]^{\top}$ (where a state \boldsymbol{x}_i includes position and velocity elements, $\boldsymbol{x}_i = (\mathbf{x}_i, \dot{\mathbf{x}}_i)$, for example). This trajectory provides the motion plan in configuration-space to be used for execution on the system.

Note that, unlike MPPI, there is no need to perform *roll-outs* of a dynamics model to generate trajectory samples. Since the goal-directed stochastic dynamics are integrated into the proposal distribution via the time-correlated covariance matrix \mathcal{K} , we can sample trajectories directly from the system to acquire sequences of positions, velocities, etc. This avoids the linear-with-time complexity of sequential Markov sampling using single-step dynamics. Furthermore, the sampling-based optimization scheme of StochGPMP does not require that the cost function $E(\cdot)$ to be differentiable. For motion planning problems, this permits the use of hard constraints such as discontinuous occupancy functions (to be used as obstacle costs) or joint limits, for instance.

C. GP Interpolation

A main advantage of maintaining a Gaussian Process trajectory distribution is the ability to query the posterior at *any time* within the planning horizon, *i.e.* $t_0 < t' < t_N$. As mentioned in [54], this can be used to densify the mean trajectory μ of the optimized proposal distribution, smoothly interpolating motion commands between timesteps. Specifically, given an optimized mean trajectory μ_{final} over times $t = [t_0, t_1, ..., t_N]$, we can interpolate to a dense time series $t_d = [t_0, t'_1, ..., t'_{M-1}, t'_M]$ where N < M and $t'_{j+1} - t'_j < t_{i+1} - t_i \forall i, j$. This is done by conditioning the mean of the posterior [55]

$$\boldsymbol{\mu}_d = \boldsymbol{\mathcal{K}}(\boldsymbol{t}_d, \ \boldsymbol{t}) \, \boldsymbol{\mathcal{K}}^{-1} \boldsymbol{\mu}, \tag{30}$$

where the kernel matrix $\mathcal{K}(t_d, t) = [\mathcal{K}(t'_m, t_n)]|_{0 \le m \le M, 0 \le n \le N}$ can be constructed as in Section I-A by considering the Markovian dependencies between neighboring states for transition matrices $\Phi_{i,j}$, $\Phi_{j,i+1}$ and covariances $Q_{i,j}$, $Q_{j,i+1}$, where $t_i \le t'_j \le t_{i+1}$. The interpolation step may be desirable for execution on systems requiring higher temporal frequencies. Also, the projection term $\mathcal{K}(t_d, t) \mathcal{K}^{-1}$ can be pre-computed and stored prior to planning. As with \mathcal{K} , this matrix is sparse and can benefit from efficient storage and matrix-multiplication using libraries for sparse linear algebra.

In GPMP [6], [54], interpolation is applied in-the-loop to finely resolve collision costs, before projecting the evaluations to the support states. This may warrant a similar approach in StochGPMP, where densification of the mean trajectory μ would allow sampling of more finely resolved trajectories for evaluation (note that the softmax weights w^k evaluate an entire trajectory τ^k , so projection back to support states would be unnecessary here). However, since we would still need to sample dense trajectories from the full distribution for evaluation, this approach would not provide any significant gain in space or runtime complexity. Therefore, the temporal resolution of the optimized mean trajectory μ is initially chosen to adequately resolve obstacle features.

D. Algorithm

Given a start state $x_0 \in \mathbb{R}^b$ and goal state $x_g \in \mathbb{R}^b$ defined in configuration space having b dimensions, as well as a cost-function $E(\cdot)$, we seek to find the optimal trajectory μ^* which avoids high-cost regions and defines a motion plan which reaches the goal.

As described in the main body of the paper, the function $E(\cdot)$ can be composed as the sum of different functions, defining separate constraints and penalties *i.e.* $E(\cdot) = \sum_{p} E_{p}(\cdot)$. For encoding task-based constraints, these cost components typically consist of a function $G : \mathbb{R}^{\ell \times N} \to \mathbb{R}$, where ℓ is the dimension of the task space and N the trajectory length, as well as the forward-kinematic mapping $FK : \mathbb{R}^{b \times N} \to \mathbb{R}^{\ell \times N}$, such that the cost component function is the composition $E_{p}(\cdot) = G(FK(\cdot))$. It should be emphasized that neither the function G or FK need to be differentiable for samplingbased optimization. It is often desirable to define a goal position within task space, such as the desired end-effector pose $\mathbf{p}_{g} \in SE(3)$ and velocity $\dot{\mathbf{p}}_{g} \in SE(3)$ at the goal for a reaching task on a robot manipulator. In this, case one may simply use inverse kinematics to initialize the goal state in configuration space: $\mathbf{x}_{g} = IK(\mathbf{p}_{g}, \dot{\mathbf{p}}_{g})$. We additionally need to define covariance matrices $\boldsymbol{\Sigma}_{s} \in \mathbb{R}^{b \times b}$ and $\boldsymbol{\Sigma}_{g} \in \mathbb{R}^{b \times b}$ for the start and goal states, respectively.

We additionally need to define covariance matrices $\Sigma_s \in \mathbb{R}^{b \times b}$ and $\Sigma_g \in \mathbb{R}^{b \times b}$ for the start and goal states, respectively. These may explicitly represent the actual uncertainty in these positions, or they can be viewed as inverted-weights which scale the importance of generating samples which start at x_0 and end at x_g . To initialize the sampling distribution $q(\tau; \mu, \mathcal{K})$, we first set the mean μ to be the straightline trajectory between x_0 and x_g , equidistant along the time sequence $t = [t_0, t_1, ..., t_N]$ and with constant velocity. The GP covariance matrix is initialized at the start of the optimization using the procedure in Section I-A. Although \mathcal{K} remains fixed in the current implementation, it is conceivable to gradually decrease the GP variance by scaling the \mathbf{Q}_c matrix according to some schedule, thereby decreasing the entropy of the distribution for denser sampling and improved estimation of q^* . In Algorithm 1, we describe Stochastic Gaussian-Process Motion Planning (StochGPMP), an algorithm for single-trajectory/uni-modal trajectory optimization.

Algorithm 1: StochGPMP

```
// Initialize mean trajectory
\boldsymbol{\mu} \leftarrow InitMean(\boldsymbol{x}_0, \boldsymbol{x}_a, \boldsymbol{t})
// Construct GP covariance
\mathcal{K} \leftarrow InitCov(\Sigma_0, \Sigma_a, \mathbf{Q}_c, t)
while Not Converged do
Sample \{\boldsymbol{\tau}^k\}_{k=1}^K \sim q(\boldsymbol{\tau}; \boldsymbol{\mu}, \boldsymbol{\mathcal{K}})
       // Batched computation
       for i = 1, 2, ..., K do in parallel
              // Evaluate trajectory cost
              E^k = E(\boldsymbol{\tau}^k)
       end
       for i = 1, 2, ..., K do in parallel
             // Compute weights including important-sampling factor w^k \leftarrow \frac{\exp\left(-\frac{1}{\lambda}\tilde{E}(\tau^k)\right)}{\sum_{j=0}^{K}\exp\left(-\frac{1}{\lambda}\tilde{E}(\tau^j)\right)}
       end
      // Update proposal distribution \pmb{\mu} \leftarrow \pmb{\mu} + \gamma \sum_{k=0}^{K} w^k (\pmb{\tau}^k - \pmb{\mu})
end
// Optional: GP-interpolation for trajectory densification
\boldsymbol{\mu}_d \leftarrow \boldsymbol{\mathcal{K}}(\boldsymbol{t}_d, \boldsymbol{t}) \boldsymbol{\mathcal{K}}^{-1} \boldsymbol{\mu}
```

E. Multi-goal, Multi-modal Motion Planning

Motion planning is generally a non-convex optimization problem, where different homotopy classes can be induced by the presence of obstacles [56], [57]. Furthermore, in the case of multi-goal planning, we may have a distribution of possible goals for a given environment or set of obstacles. Thus, most motion planning approaches will only return an approximate

Input: Start state x_0 , goal state x_g , cost function $E(\cdot)$, temperature λ , start-factor cov. Σ_0 , goal-factor cov. Σ_g , \mathbf{Q}_c , step-size γ , time sequence t, dense time sequence t_d (optional)

Algorithm 2: Multi-StochGPMP

Input: Start state x_0 , goal states \mathbf{X}_g , cost function $E(\cdot)$, temperature λ , start cov. matrix Σ_0 , goal cov. matrices Ξ_g , planning GP variance matrix \mathbf{Q}_c , prior GP variance matrices $\mathbf{Q}_{c,0}$, step-size γ , time sequence t, dense time sequence t_d (optional)

```
// Prior const-vel. mean trajectories
M_0 \leftarrow InitMeans(x_0, \mathbf{X}_q, t)
// Prior GP cov. matrices
K_0 \leftarrow InitCovs(\Sigma_0, \Xi_a, \mathbf{Q}_{c,0}, t)
// Initialize mean trajectories
\boldsymbol{M} \sim \mathcal{N}(\boldsymbol{M}_0, \boldsymbol{K}_0)
// Construct planning cov. matrices
K \leftarrow InitCovs(\Sigma_0, \Xi_a, \mathbf{Q}_c, t)
for iter = 1 : iter_{max} do
     // Sample trajectories
     \mathbf{T} \sim q(\mathbf{T}; \boldsymbol{M}, \boldsymbol{K})
     // Evaluate trajectory costs
     E_{\mathbf{T}} \leftarrow E(\mathbf{T})
     // Compute weights across sampling dimension
     \mathbf{W} \leftarrow softmax(-\frac{1}{\lambda}E_{\mathbf{T}})
     // Update proposal distribution
     \boldsymbol{M} \leftarrow \boldsymbol{M} + \gamma \mathbf{W} \otimes (\mathbf{T} - \boldsymbol{M})
end
// Evaluate solutions
E_{\boldsymbol{M}} \leftarrow E(\boldsymbol{M})
// Pick best trajectory
i^* \leftarrow \arg\min_i [E_i \in E_M]
\mu^* \leftarrow \mu_{i^*} \in M
// Optional: GP-interpolation for trajectory densification
\boldsymbol{\mu}_d \leftarrow \mathcal{K}(\boldsymbol{t}_d, \boldsymbol{t}) \, \mathcal{K}_*^{-1} \boldsymbol{\mu}^*
```

locally optimal solution. We would ideally like to leverage parallel computation to efficiently provide many feasible solutions for the multi-goal and multi-modal setting.

We can implement a vectorized version of Algorithm 1 for the multi-goal, multi-modal setting using PyTorch [58]. This permits GPU parallelization across different goals, motion plans, trajectory sample generation and evaluation. For a given environment, we may have N_g number of goal states: $\mathbf{X}_g = [\mathbf{x}_g^1, \mathbf{x}_g^2, ..., \mathbf{x}_g^{N_g}]^\top \in \mathbb{R}^{N_g \times b}$, with goal covariance matrices $\mathbf{\Xi}_g = [\mathbf{\Sigma}_g^1, \mathbf{\Sigma}_g^2, ..., \mathbf{\Sigma}_g^{N_g}]^\top \in \mathbb{R}^{N_g \times b \times b}$. For each goal state, we may wish to independently optimize for multiple motion plans of the same length, which can be represented by the the mean trajectory tensor $M^{n_g} = [\boldsymbol{\mu}_{n_g}^1, \boldsymbol{\mu}_{n_g}^2, ..., \boldsymbol{\mu}_{n_g}^{N_g}]^\top \in \mathbb{R}^{N_p \times N \times b}$, where $n_g \in [1, 2, ..., N_g]$ is the goal index, and N_p the number of plans-per-goal. This can yet be further expanded to include plans across all goals: $M = [M^1, M^2, ..., M^{N_g}]^\top \in \mathbb{R}^{N_g \times N_p \times N \times b}$. In this case, the collection of planning trajectories M is the parameter to be optimized. We can also construct a global covariance matrix, which might contain independent GP-covariances for each goal location (with possibly different noise levels set by \mathbf{Q}_c): $\mathbf{K} = [\mathbf{K}^1, \mathbf{K}^2, ..., \mathbf{K}^{N_g}]^\top$. We can then define a vectorized proposal distribution q from which we can draw trajectory samples across goals and motion plans: $\mathbf{T} \sim q(\mathbf{T}; \mathbf{M}, \mathbf{K})$ where $\mathbf{T} \in \mathbb{R}^{N_g \times N_p \times K \times N \times b}$ for K number of samples per motion plan. These samples can then be evaluated by the cost-function $E(\cdot)$, provided that this is also suited to vectorized processing (which is the case for a neural network model in PyTorch, for instance).

To initialize the mean trajectories M, we randomly sample from a starting prior distribution: $M \sim \mathcal{N}(M_0, K_0)$, where $M_0 = [\mu_0^1, \mu_0^2, ..., \mu_0^{N_g}]^\top \in \mathbb{R}^{N_g \times N \times b}$ are the constant-velocity, straight-line trajectories from the start state to each goal state, and $K_0 = [\mathcal{K}_0^1, \mathcal{K}_0^2, ..., \mathcal{K}_0^{N_g}]^\top \in \mathbb{R}^{N_g \times (N \times b) \times (N \times b)}$ the independent covariance matrices. These in turn can be





(a) 0 Iterations

(b) 100 Iterations

Fig. 12: Example of multi-StochGPMP applied to the multi-goal setting for a planar navigation problem. The GP covariance matrices are constructed using a constant-velocity stochastic transition model [50]. Three goal locations (green dots) are used to initialize the mean sampling trajectories (blue), with three Gaussian process distributions per goal. Trajectory samples are generated in batch (red), resolving local collisions induced by the presence of obstacles (black). Smooth mean-trajectory solutions are generated after a few planning iterations, covering different paths to the goals.

constructed from goal-specific covariances Ξ_g and state-transition variance matrices $\mathbf{Q}_{c,0} = [\mathbf{Q}_{c,0}^1, \mathbf{Q}_{c,0}^2, ..., \mathbf{Q}_{c,0}^{N_g}]^\top$. In practice, each initial GP covariance \mathcal{K}_0^i is set to be higher than the respective planning covariance \mathcal{K}^i , therefore increasing likelihood of covering different homotopy classes (if they exist) while maintaining dense, local sampling for planning updates (similar to having initial exploratory behavior).

For system execution, the converged plans M are evaluated, and the final trajectory can be selected as the lowest-cost trajectory $\mu \in M$. In the case where the method is used to gather demonstrations (as described in the main text), this step can be ignored, and all solutions M stored along with contextual data \mathbf{X}_g and $E(\cdot)$. The full algorithm for Multi-StochGPMP is shown in Algorithm 2.

An example of the algorithm being applied for a 2D trajectory optimization is shown in Fig. 12. Here, a simple point-particle linear system is defined in cartesian space, with $\mathbf{Q}_c = \sigma^{-2} \mathbf{I} \in \mathbb{R}^{2 \times 2}$. The optimization is performed over position-velocity trajectories of length N = 64, with step size $\gamma = 0.5$.

F. Related Work

The use of structured, time-correlated covariance matrices for sample-based trajectory optimization has been examined in prior work [4], [59]–[62]. In the STOMP algorithm [4], a precision matrix with a band-diagonal structure is proposed for generating discrete-time trajectories. In comparison, the GP covariance in Section I-A allows for sampling of system velocities (or potentially higher derivatives) which are time-correlated with state according to a principled, system-based derivation of the prior distribution. Given the continuous-time representation, this also provides an option to densify trajectories through GP-interpolation. The factored form of the prior distribution (Eq. (16)) results in fine-grained, yet intuitive, control over the construction of the \mathcal{K} matrix, *i.e.* the start, goal, and transition covariances can be varied and set independently according to the desired coverage of the sampling distribution.

GP-based sampling priors have been examined for uni-modal motion planning in [63]. Although a similar approach is taken to Algorithm 1, we take an information-theoretic perspective in deriving the update rule in Section I-B. Here, we draw connections to existing approaches for stochastic optimal control, which only consider proposal distributions over the control space. Further, we motivate our approach towards the batch setting, where sampling and optimization can be performed in parallel. Concurrently to this work, a multi-modal planning algorithm was proposed in [64] which uses a mixture of Gaussian process distributions. Here, sampling of trajectories is performed sequentially by first drawing from the latent variable distribution, then sampling from the resulting GP component. Given that our multi-modal, multi-goal planning method in Algorithm 2 is intended for batch computation using vectorized representations, sampling and distribution updates are performed simultaneously across all modes of the proposal distribution.

APPENDIX II

EXTENDED EXPERIMENTAL EVALUATION

In the following, we provide an extended description of the experiments we performed. All the three experiments are divided in two parts: the learning of the EBM components and the trajectory optimization problem using the learned EBM

as additional implicit priors.

In the three experiments, the trajectory optimization objective function was defined by a weighted sum of a set of heuristic costs, introduced in Table II and the learned EBM components. To find the optimal trajectory, we considered both sampling based methods (Exp I and Exp II), presented in Section I and gradient based methods (Exp III). The experiments were performed to validate the possibility of using EBM as costs in trajectory optimization problems and to measure the benefit of the proposed algorithmic decisions in the performance of the obtained trajectories.

A. Experiment I: 2D point Navigation

As described in the main text, datasets for the planar navigation problem were generated by either random uniform sampling or stochastic trajectory generation with dimensions x : [-10, 10], y : [-10, 10] sized region. For random sampling, 1024 example points were generated in the free-space regions of obstacles for each environment (i.e. obstacle configuration). In the case of trajectory demonstrations, for each environment, 15 random goal locations were sampled along the axis defined by x : [-9, 9], y = 9 and x = 9, y : [-9, 9], with start location fixed at x = -9, y = -9. Trajectories were generated using the StochGPMP optimizer, with 5 trajectories of length 64 generated for each goal location. Each trajectory is selected to be the mean of each optimized StochGPMP distribution. The planning parameters were set to a timestep of dt = 0.02, temperature $\lambda = 1$, goal-factor stand-dev. $\sigma_{s} = 1 \times 10^{-2}$, start-factor stand-dev. $\sigma_g = 1 \times 10^{-2}$, GP-factor stand-dev. $\sigma_{gp} = 0.1$, and obstacle-factor stand-dev. $\sigma_{obs} = 1 \times 10^{-5}$. For sampling-based trajectory optimization, a sample size of 32 was used. A total of 512 environments were generated for each experiment. The EBM model consisted of a simple 2-hidden layer MLP (hidden width=512), with 2-D obstacle locations concatenated to the 2-dimensional inputs.

B. Experiment II: Planar Manipulator

This experiment evaluates how object-centric EBM helps solving the task of grasp and insert object into a walled-cubby. As mentioned, we consider a simulated planar manipulator where the task space is 2D. The objective is to find a smooth trajectory in the joint space from initial joint configuration q_0 to grasp the white cube and insert the white cube into the cubby while avoiding collisions. Note that the white cube in our case is a visual mesh having no collision model.

Dataset generation

In this experiment, we consider a simple grasp mechanism, where the planar robot arm can grasp the object just by touching (e.g. suction, magnetic pull mechanisms). For generating grasp points, we define four 2D uniform distributions with small margin along four edges of the cube (observing from the top view). For each side edges, we samples uniformly 10000 points and hence in total the number of positive samples are 40000. For generating insert trajectories, we sample uniformly 512 initial points in the planar task space. Then for each initial point, we optimize for GP trajectory distribution having 32 time steps using our StochGPMP planner and sample 128 collision-free trajectories from the solution distribution. In total, we have $512 \times 128 \times 32$ positive samples for training insert EBM. While the effect number of positive samples for training planar EBMs priors does not need to be large, we observe that the positive samples should have a suitable population to cover relevant task spaces to give informative guidance for motion optimization.

EBM training

The learned EBMs represents a state distribution

$$p(\boldsymbol{x}) \propto \exp\left(-E_{\boldsymbol{\theta}}(\boldsymbol{x} - \boldsymbol{x}_{\text{frame}})\right).$$
 (31)

with regard to their respective object-centric frame. For both grasp and insert case, the object-centric EBMs are trained by first transforming the dataset points into the object frame of reference in consideration, then we performance training as described in Section V. The network architecture used for both cases are a simple fully-connected network with two hidden layers having the width of 512.

Trajectory Optimization with learned EBM

We define the manipulation planning objective with multiple cost terms: (i) GP prior factor cost encouraging trajectory smoothness and connecting starting and final configuration, which can be computed by taking negative logarithm of Eq. (16), (ii) obstacle avoidance cost and (iii) joint limit cost as specified in Table II. And finally, (iv) learned grasp EBM and (v) learned insert EBM cost as:

$$c(\boldsymbol{q}) = E_{\text{grasp}}(FK(\boldsymbol{q}) - \boldsymbol{x}_{\text{cube}}) + E_{\text{insert}}(FK(\boldsymbol{q}) - \boldsymbol{x}_{\text{cubby}})$$
(32)

where $FK(\cdot)$ is forward kinematic function.

C. Experiment III: Robot Pouring amid obstacles

In this experiment, we evaluate the integration of learned EBM components for solving a pouring task amid obstacles. We consider a 7 DoF Kuka-LWR robot manipulator. Given the robot's initial joint configuration q_0 , we aim to find a trajectory that moves the robot to pour in an arbitrarily positioned pot and recover back to the initial joint configuration.

EBM training

The learned EBM represents a distribution $p(\tau)$ that provides high probability to the trajectories performing the pouring task and low probability to the rest. For this experiment, we consider modelling this distribution with a phase-conditioned EBM, introduced in Eq. (11). Representing the whole trajectory distribution with a single EBM might be hard due to the high dimensionality. A trajectory τ is represented in a $N \times T$ dimensional space, with N the state dimension and T the temporal dimension. Additionally, different trajectories might vary in temporal length, changing the temporal dimension T of the input. To deal with these limitations, the phase-conditioned EBM, $E_{\theta}(x, \alpha)$ represents the distribution of the state for a particular instant of the trajectory. The temporal instant is set by an α conditioning variable that goes from 0 representing the start of the trajectory to 1 the end of the trajectory. Then, the whole trajectory distribution is represented by

$$p(\boldsymbol{\tau}) \propto \exp\left(-\sum_{k} E_{\boldsymbol{\theta}}(\boldsymbol{x}_{k}, \alpha_{k}) + (\boldsymbol{x}_{k} - \boldsymbol{x}_{k+1})^{2}\right).$$
(33)

with k representing different instants of the trajectory and $(x_k - x_{k+1})^2$ a smoothness cost to guarantee coherence in the trajectory. For our problem, we represent the spatial dimension N in a 6 dimensional Euclidean space, $x \in \mathbb{R}^6$ consisting on the 3D position of the cup's bottom and the 3D position of the top of the cup. The position of this two points is represented in the pouring pot's reference frame. This way, the learned EBM can directly adapt to arbitrary poses in the pouring pot. Additionally, we consider as baseline an EBM trained in the whole trajectory $p(\tau) \propto \exp(-E_{\theta}(\tau))$. We chose this baseline to evaluate the benefit of a structured distribution representation w.r.t. modelling the distribution with a single EBM.

We train the model with CD loss with a uniform distribution as negative sample generation. Additionally, we add a denoising score matching loss as regularizer. As previously introduced, we expect the denoising regularizer enhance the landscape of the EBM for motion planning. We also consider as baseline an EBM trained without the regularizer to evaluate the benefits of adding it.

Trajectory Optimization with learned EBM

Given the learned pouring EBM, we evaluate the performance benefit from adding the EBM as additional cost function in the objective function. To solve the pouring task, we define an objective function with multiple cost functions. A configuration space potential cost in the initial and final configuration to encourage the initial and final configuration to be a desired pose

$$c(\boldsymbol{q}) = \frac{1}{2} \left| \left| \boldsymbol{q} - \boldsymbol{q}^* \right| \right|^2, \tag{34}$$

with q^* the target configuration. A table and obstacle avoidance cost, similar to the one presented in Table II. A trajectory smoothness cost

$$c(\boldsymbol{q}_{0:T}) = \frac{1}{T} \sum_{k=0}^{T-1} ||\boldsymbol{q}_k - \boldsymbol{q}_{k+1}||^2,$$
(35)

that encourages neighbours states to be close to each other. A pouring avoidance cost, that encourages the orientation of the cup to be pointing up without spilling in almost all the trajectory points. To compute the cost, we compute the angle between the z axis vector in the cup v_{z_glass} and the z axis on the world frame v_z , $\theta = \cos^{-1}(v_z \cdot v_{z_glass})$. The cost aims to minimize the angle

$$c(\theta) = \left| \left| \theta \right| \right|^2. \tag{36}$$

Finally, a pouring cost that encourages the glass to be close to the pot and with a certain angle to pour properly. Additionally, we add a learned EBM trained on trajectories performing the pouring task. In contrast with the previous experiments that apply sampling based optimization methods, we find the optimal trajectory by gradient descent. To reduce the possibility of getting a locally optimal solution, we initialize multiple particles and evolve all in batch. The evaluation is presented in Section VI.