

Stein Variational Probabilistic Roadmaps

Alexander Lambert, Brian Hou, Rosario Scalise, Siddhartha S. Srinivasa, and Byron Boots

Abstract—Efficient and reliable generation of global path plans are necessary for safe execution and deployment of autonomous systems. In order to generate planning graphs which adequately resolve the topology of a given environment, many sampling-based motion planners resort to coarse, heuristically-driven strategies which often fail to generalize to new and varied surroundings. Further, many of these approaches are not designed to contend with partial-observability. We posit that such uncertainty in environment geometry can, in fact, help *drive* the sampling process in generating feasible, and probabilistically-safe planning graphs. We propose a method for Probabilistic Roadmaps which relies on particle-based Variational Inference to efficiently cover the posterior distribution over feasible regions in configuration space. Our approach, Stein Variational Probabilistic Roadmap (SV-PRM), results in sample-efficient generation of planning-graphs and large improvements over traditional sampling approaches. We demonstrate the approach on a variety of challenging planning problems, including real-world probabilistic occupancy maps and high-dof manipulation problems common in robotics. Video, additional material and results can be found here: <https://sites.google.com/view/stein-prm>.

I. INTRODUCTION AND RELATED WORK

Probabilistic roadmaps approximate a continuous configuration space with a discrete set of sampled configurations [1, 2]. Theoretical bounds have been developed for the number of uniform or low-discrepancy samples required to ensure near-optimal paths [2–4]. However, such bounds require many samples to generate a path through challenging “narrow passage” problems [5], yielding dense roadmaps that cannot be searched efficiently at query time.

Thus, roadmaps should be sparse while preserving the connectivity of the underlying free configuration space. Our key insight is that the distribution of roadmap samples should match the distribution over feasible states. We formulate this problem as particle-based variational inference [6]. Our Stein Variational Probabilistic Roadmap (SV-PRM) approach deterministically optimizes the samples of a sparse roadmap to efficiently cover the space. This method only requires that the distribution is differentiable, an assumption satisfied by

This work was (partially) funded by the National Institute of Health R01 (#R01EB019335), National Science Foundation CPS (#1544797), National Science Foundation NRI (#1637748), the Army Research Laboratory (#W911NF-16-2-0008), the Office of Naval Research, the RCTA, Amazon, and Honda Research Institute USA. Brian Hou is partially supported by a NASA Space Technology Research Fellowship.

All authors are with the Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA 98195 {[lambert6](mailto:lambert6@cs.uw.edu), [bhoul](mailto:bhou@cs.uw.edu), [rosario](mailto:rosario@cs.uw.edu), [siddh](mailto:siddh@cs.uw.edu), [bboots](mailto:bboots@cs.uw.edu)}@cs.uw.edu. Rosario Scalise is also an Oak Ridge Associated University Fellow with DEVCOM Army Research Laboratory. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government.




Fig. 1: We generate PRMs by sampling from a posterior feasibility distribution using particle-based Variational Inference. Placement of graph vertices is governed by a gradient flow, promoting uncertainty-guided exploration of the state space prior to graph construction. The approach results in sample-efficient path planning using both (a) real-world probabilistic occupancy maps and (b) cost functions for manipulation tasks. In Fig. 1(b), particle vertices are represented by orange robot configurations, with the goal configuration shown in green.

continuous Bayesian occupancy maps [7–9] from simultaneous localization and mapping (SLAM) or otherwise assumed by popular trajectory optimization algorithms [10–13]. We demonstrate the efficacy of our approach on 2D navigation and 7D manipulation problems (Fig. 1).

A complementary thread of work biases the PRM sampling distribution to focus on important regions of the configuration space. These approaches typically require additional uniform samples to provide baseline coverage of easily-sampled regions. Heuristics from collision geometry or topology are most effective on 2D problems and increase the expense of generating a single sample [14–17]. Recent work has trained neural networks to propose samples around predicted shortest paths [18, 19] or bottleneck nodes [20]. However, these approaches primarily focus on the single-query setting, i.e. with a specific start and goal configuration pair. Critical PRMs [21] are sparse roadmaps designed for the same multi-query setting targeted by SV-PRMs. This approach biases toward samples with high betweenness centrality, training a neural network to predict such nodes from local environment features. Neural network-based approaches require a dataset of similar planning environments, while SV-PRMs perform variational inference with a differentiable probabilistic feasibility function.

A key question with biased samplers is the choice of environment representation. This is often in the form of an occupancy map [17, 20], point cloud [19], or other complete representation [18]. Explicitly constructing these representations in configuration space can be challenging for many robotics settings, especially with partial-observability.

In contrast, classical geometry heuristics only require access to a binary collision checker to evaluate candidate samples [14–16], while Critical PRMs use local features to predict criticality [21]. Similarly, SV-PRMs rely only on access to a differentiable probabilistic feasibility function, and avoid explicitly evaluating the function over the entire state space.

SV-PRM relies on the SVGD method for variational inference (Section IV). Although SVGD has been used for model-predictive control [22, 23] and trajectory optimization [22, 24], we leverage its favorable properties for generating global planning graphs.

II. PROBABILISTIC ROADMAPS AND SAFETY

In the conventional PRM algorithm [1], a graph $G = \{V, E\}$ with vertices V and edges E is generated by (1) randomly sampling a vertex candidate from the continuous configuration space $\mathbf{x}^i \sim \mathcal{X}$, (2) performing a collision check on the sample, (3) adding $\mathbf{x}^i$ to V if $\mathbf{x}^i$ is collision-free, and (4) connecting $\mathbf{x}^i$ to existing neighboring vertices in the graph $\mathbf{x}^j \in N(\mathbf{x}^i)$ if the corresponding edge $e_{ij} : \mathbf{x}^i \mathbf{x}^j$ is also collision-free. A path on G connecting start and goal vertices $(\mathbf{x}_s, \mathbf{x}_g) \in V$ is defined as a sequence of connected vertices $\xi = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_f)$ where $\mathbf{x}_1 = \mathbf{x}_s$, $\mathbf{x}_f = \mathbf{x}_g$. The edge cost is a function $c : E \rightarrow \mathbb{R}^+$, and the path cost is the sum of edge costs over the path: $c(\xi) = \sum_{e \in \xi} c(e)$. Denoting the set of paths on G as Ξ , planning on G is then the problem of finding the lowest-cost path $\xi^* = \arg \min_{\xi \in \Xi} c(\xi)$.

A vertex $\mathbf{x}^i \in \mathcal{X}$ is deemed *feasible* (i.e. “collision-free”) if it satisfies some constraint $h(\mathbf{x}^i) \leq 0$. The safe set is defined by all such points, $\mathcal{S} \subseteq \mathcal{X}$, and a safe path denoted by the continuous function $\xi_S = \xi_S(t) \in \mathcal{S}$ where $t \in [0, T]$, $\xi_S(0) = \mathbf{x}_s$ and $\xi_S(T) = \mathbf{x}_g$. The path cost is then the integral of point-wise costs along the path: $c(\xi_S) = \int_{t=0}^T c(\xi_S(t)) dt$. Any start-goal pair $(\mathbf{x}_s, \mathbf{x}_g) \in \mathcal{S}$ is deemed reachable if a corresponding path ξ_S exists. The graph G approximates the topology of $\mathcal{S}$, since it is constructed such that $V, E \in \mathcal{S}$. A “good” graph G is then one which satisfies the following properties: (1) Any point $\mathbf{x}^i \in \mathcal{S}$ can be connected to G with a set of collision-free edges $E^i = \{e_{ij} : \mathbf{x}^i \mathbf{x}^j \mid \mathbf{x}^j \in N(\mathbf{x}^i)\}$. (2) For any start and goal $\mathbf{x}_s, \mathbf{x}_g \in \mathcal{S}$, the augmented graph $G' = (V \cup \{\mathbf{x}_s, \mathbf{x}_g\}, E \cup \{E^s, E^g\})$ contains an optimal path ξ^* with a cost bounding the optimal path cost in the safe set: $c(\xi^*) - c(\xi_S^*) \leq \delta$, where $\delta \geq 0$. For query-time efficiency, the graph G should have minimal size while satisfying (1) and (2). The sampling strategy for generating V is then vital in achieving these properties.

III. POSTERIOR-GUIDED ROADMAP GENERATION

In order to generate candidate vertices in an informed and sample-efficient manner, we desire a sampling distribution having high probability in the safe set $\mathcal{S}$, and low probability elsewhere. Specifically, we can represent the probability of a given point $\mathbf{x} \in \mathbb{R}^d$ being collision-free by the *feasibility* likelihood¹ $p(z = 1 \mid \mathbf{x}; \Theta)$ with parameter $\Theta \in \mathbb{R}^m$. Here,

¹We use this notation for convenience, maintaining the following equivalence: $p(z = 1 \mid \mathbf{x}; \Theta) = p(h(\mathbf{x}) \leq 0 \mid \mathbf{x}; \Theta) = p(\mathbf{x} \in \mathcal{S} \mid \mathbf{x}; \Theta)$.

the occupancy indicator variable $z \in \{0, 1\}$ labels a given location as being in-collision ($z = 0$) or collision-free ($z = 1$). Using Bayes’ Rule, we can obtain a posterior probability over collision-free space:

$$p(\mathbf{x} \mid z = 1; \Theta) = \eta p(z = 1 \mid \mathbf{x}; \Theta) p(\mathbf{x}) \quad (1)$$

where $p(\mathbf{x})$ is a prior probability, and η a normalizing factor. This formulation bears particular significance in the partially-observable setting, where we do not have access to ground truth occupancy labels z , but can make measurements using a probabilistic model accounting for sensor noise. In this case, the likelihood distribution $p(z = 1 \mid \mathbf{x}; \Theta)$ can be interpreted as a probabilistic map of unoccupied regions [25]. However, as we shall see, such feasibility likelihoods may also correspond to negatively-exponentiated costs found within the context of trajectory-optimization. In both cases, we can derive a target posterior distribution to inform a sampling scheme for PRMs.

A. Bayesian Occupancy Maps

Bayesian inference provides a useful framework for integrating observations and incorporating uncertainty to construct occupancy maps. In this case, the target space corresponds to 2- or 3-dimensional Cartesian space, i.e. $\mathbf{x} \in \mathbb{R}^2$ or $\mathbf{x} \in \mathbb{R}^3$. The map parameter Θ can be obtained by performing Bayesian inference using collected occupancy data. Given a set of state-measurement pairs $D = \{\mathbf{x}_m, z_m\}_{m=1}^M = (X, \mathbf{z})$, a likelihood model $p(z = 1 \mid \mathbf{x}, \mathbf{w})$ and prior $p(\mathbf{w})$, the posterior over model parameters can be expressed as

$$p(\mathbf{w} \mid \mathbf{z}, X) = \eta \prod_{m=1}^M p(z_m \mid \mathbf{x}_m, \mathbf{w}) p(\mathbf{w}) \quad (2)$$

where η is a normalizing constant. The feasibility likelihood $p(z = 1 \mid \mathbf{x}; \Theta)$ can then be modeled for any point in the domain $\mathbf{x} \in \mathcal{X}$, where the parameter Θ can be chosen as a sufficient statistic from the fitted posterior in Eq. (2), e.g. $\Theta \equiv \mathbb{E}_{p(\mathbf{w} \mid \mathbf{z}, X)} [\mathbf{w}] = \mu_{\mathbf{w}}$. Approximating the posterior is non-trivial, however a variety of scalable approaches exist. Bayesian Hilbert Maps [9], for instance, model the likelihood of free-space using the sigmoid:

$$p(z = 1 \mid \mathbf{x}, \mathbf{w}) = (1 + \exp(\mathbf{w}^\top \Phi(\mathbf{x})))^{-1} \quad (3)$$

with a feature vector $\Phi(\mathbf{x})$ of radial-basis functions.

B. Feasibility Distributions in Motion Planning

Defining a state trajectory as the continuous-time function $\tau \triangleq \mathbf{x}(t) : t \rightarrow \mathbb{R}^d$, and a start state $\mathbf{x}_0$, trajectory optimization aims to find the optimal trajectory τ^* which minimizes an objective functional $\mathcal{F}(\tau; \mathbf{x}_0)$. The solution must be feasible and avoid collisions with obstacles (i.e. belong to the safe set). This condition can be imposed by including an additional obstacle penalty $\mathcal{F}_{obs}$ in the objective [10, 11]:

$$\mathcal{F}(\tau; \mathbf{x}_0) = \mathcal{F}_{task}(\tau; \mathbf{x}_0) + \alpha \mathcal{F}_{obs}(\tau) \quad (4)$$

where $\mathcal{F}_{task}(\tau; \mathbf{x}_0)$ includes all other task objectives and

constraints, such as distance-to-goal and penalty on joint-limit violations. The objective is minimized via a Gauss-Newton form of iterative gradient descent, leveraging both gradient and local curvature of the cost function for rapid convergence [10, 13, 26]. The obstacle cost at a given trajectory state $\mathbf{x} \in \tau$ can be modelled using a truncated signed-distance field (t-SDF), which penalizes a state if a Cartesian point on the robot is within an ϵ -distance from the surface of the nearest obstacle. Approximating the robot surface using a collection of body spheres $s_j, j = 1 : K$, and letting $d(\mathbf{x}, s_j)$ to be the distance from the surface of s_j to the nearest obstacle, a t-SDF value is defined by the hinge loss:

$$c(\mathbf{x}, s_j) = \begin{cases} -d(\mathbf{x}, s_j) + \epsilon & d(\mathbf{x}, s_j) \leq \epsilon \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

These costs are combined across spheres to construct an obstacle-cost vector: $\mathbf{h}(\mathbf{x}) = [c(\mathbf{x}, s_j)]_{j=1:K}$, and define the total obstacle cost as the scaled inner-product: $\mathcal{F}_{obs}(\mathbf{x}) = \alpha \|\mathbf{h}(\mathbf{x})\|^2$, where $\alpha > 0$. As in [12, 13], the obstacle cost at a given state can be interpreted as the negative log-likelihood function: $-\log p(z = 1 | \mathbf{x}; \Theta) = \alpha \|\mathbf{h}(\mathbf{x})\|^2 + \text{const}$, with parameter Θ including obstacle locations and geometry. Here, the likelihood probability of a state being collision-free is modeled as $p(z = 1 | \mathbf{x}; \Theta) \propto \exp(-\alpha \|\mathbf{h}(\mathbf{x})\|^2)$. Motion planning can then be framed as an inference problem, where the objective is to find the *maximum a-posteriori* solution:

$$\tau^* = \arg \max_{\tau} p(\tau | z = 1; \Theta) \quad (6)$$

$$= \arg \max_{\tau} \prod_{t=0}^T p(z = 1 | \mathbf{x}_t; \Theta) p(\tau) \quad (7)$$

where trajectories are discretized into a set of states $\tau = \{\mathbf{x}_t\}_{t=0:T}$ over a horizon T .

Although gradient-based controllers (ex. iLQR) and motion planners (ex. GPMP) can quickly return a feasible trajectory, the solution may only be *locally* optimal, and requires an initial trajectory or set of waypoints to prevent falling into poor optima. Such a path can be produced by sampling PRM vertices from the feasibility posterior over state-space, $\mathbf{x}^i \sim p(\mathbf{x} | z = 1; \Theta) = \eta p(z = 1 | \mathbf{x}; \Theta) p(\mathbf{x})$, and generating the shortest-path on the resulting graph.

For both cases in Section III-A and Section III-B, we would like to directly sample from the posterior distributions in order to generate viable points in regions of the state space that are likely to satisfy state-based constraints: $\mathbf{x}^i \sim p(\mathbf{x} | z = 1; \Theta)$. Unfortunately, obtaining a closed-form expression for this target distribution is intractable in general. The posterior-sampling procedure must then be approximated using methods such as Markov-Chain Monte Carlo (MCMC) or Variational Inference (VI).

IV. VARIATIONAL INFERENCE AND SVGD

Variational inference (VI) is a powerful tool for approximating challenging probability densities in Bayesian statistics. As opposed to MCMC methods, VI formulates inference as an optimization problem. A proposal distribution

$q(\mathbf{x})$, belonging to a family $\mathcal{Q}$, is chosen to minimize the KL-divergence with the target posterior distribution $p(\mathbf{x} | z)$ over latent variable $\mathbf{x}$:

$$q^*(\mathbf{x}) = \min_{q \in \mathcal{Q}} D_{KL}(q(\mathbf{x}) \| p(\mathbf{x} | z)) \quad (8)$$

Traditional VI methods typically require careful selection of the distribution class $\mathcal{Q}$, which is often chosen to have a tractable parametric form at the expense of introducing bias. A recently developed method, *Stein Variational Gradient Descent* [6, 27], avoids the challenge of determining an appropriate $\mathcal{Q}$ by leveraging a non-parametric, particle based representation of the posterior distribution. This approach approximates a posterior $p(\mathbf{x} | z)$ with a set of particles $\{\mathbf{x}^i\}_{i=1}^N, \mathbf{x}^i \in \mathbb{R}^d$. The particles are iteratively updated according to $\mathbf{x}^i \leftarrow \mathbf{x}^i + \epsilon \phi^*(\mathbf{x}^i)$, given a step-size ϵ . The function $\phi^*(\cdot)$ lies in the unit-ball of a reproducing kernel Hilbert space (RKHS). This RKHS is characterized by a positive-definite kernel $k(\cdot, \cdot)$. The term $\phi^*(\cdot)$ represents the optimal perturbation or velocity field (i.e. gradient direction) which maximally decreases the KL-divergence:

$$\phi^* = \arg \max_{\phi \in \mathcal{H}} \left\{ -\nabla_{\epsilon} D_{KL}(q_{[\epsilon\phi]} \| p(\mathbf{x} | z)) \text{ s.t. } \|\phi\|_{\mathcal{H}} \leq 1 \right\}, \quad (9)$$

where $q_{[\epsilon\phi]}$ indicates the particle distribution resulting from taking an update step. This has been shown to yield a closed-form solution [6] which can be interpreted as a functional gradient in RKHS, and can be approximated with the set of particles:

$$\hat{\phi}^*(\mathbf{x}) = \frac{1}{N} \sum_{j=1}^N \left[k(\mathbf{x}^j, \mathbf{x}) \nabla_{\mathbf{x}^j} \log p(\mathbf{x}^j | z) + \nabla_{\mathbf{x}^j} k(\mathbf{x}^j, \mathbf{x}) \right]. \quad (10)$$

Eq. (10) has two terms that control different aspects of the algorithm. The first term is essentially a scaled gradient of the log-likelihood over the posterior's particle approximation. The second term is known as the *repulsive force*. Intuitively, it pushes particles apart when they get too close to each other and prevents them from collapsing into a single mode. This allows the method to approximate complex, possibly multi-modal posteriors. For the case of a single particle, the method reduces to a standard optimization of the log-likelihood or a MAP estimate of the posterior as the repulsive force term vanishes, i.e. $\nabla_{\mathbf{x}} k(\mathbf{x}, \mathbf{x}) = 0$. SVGD's optimization structure empirically provides better particle efficiency than other popular sampling procedures, such as Markov Chain Monte Carlo [28]. The deterministic, gradient-based updates result in smooth transformations of the proposal distribution, a property which makes SVGD particularly attractive for trajectory optimization and inference.

A. Hessian-Scaled Kernels

As discussed in [27, 29], the convergence and accuracy of the SVGD algorithm can be largely improved by incorporating curvature information into the kernel. For instance, a positive-definite matrix M can be used as a metric to scale inter-particle distances inside of an *anisotropic* RBF kernel:




Fig. 2: (a)-(b) Partially and (c)-(d) fully-explored INTEL Bayesian Hilbert Map environment. (a)-(b) compares the resulting SV-PRM with conservative $\beta = 0.54$ and optimistic $\beta = 0.45$ graph generation. (c) demonstrates the improved connectivity of SV-PRM(Uniform) relative to (d) the PRM(Uniform) particles (which were used to initialize the SV-PRM in (c)).

$k(\mathbf{x}^j, \mathbf{x}^i) = \exp\left(-\frac{1}{2h}(\mathbf{x}^j - \mathbf{x}^i)^\top M(\mathbf{x}^j - \mathbf{x}^i)\right)$, where h is the bandwidth parameter. Curvature information can then be shared across particles by averaging their local Hessian evaluations. Specifically, denoting the negative Hessian matrix to be $H(\mathbf{x}) = -\nabla_{\mathbf{x}}^2 \log p(\mathbf{x} | \mathbf{z})$, we can define the metric $M = \frac{1}{N} \sum_{j=1}^N H(\mathbf{x}^j)$, which is computed using $\mathbf{x}^j$ -values from the previous iteration.

V. STEIN VARIATIONAL PRMS

As an alternative to random or purely heuristic-driven sampling for PRMs, we propose to approximate the posterior distribution of feasible space $p(\mathbf{x} | \mathbf{z} = 1; \Theta)$ using a set of SVGD particles $\{\mathbf{x}^i\}_{i=1}^N$.

In the *Inference* phase, particles are initialized by drawing samples from the prior distribution over configuration space: $\{\mathbf{x}^i\}_{i=1}^N \sim p(\mathbf{x})$. The particle distribution is then updated using successive SVGD iterations, where at each iteration, the log-likelihood and log-prior gradients are computed for each particle *in parallel*:

$$\nabla_{\mathbf{x}} \log p(\mathbf{x}^i | \mathbf{z} = 1; \Theta) = \quad (11)$$

$$\nabla_{\mathbf{x}} \log p(\mathbf{z} = 1 | \mathbf{x}^i; \Theta) + \nabla_{\mathbf{x}} \log p(\mathbf{x}^i), \quad (12)$$

Given a choice of kernel $k(\cdot, \cdot)$, the RKHS gradient in Eq. (10) is computed for each particle. The particles are then updated, and the process is repeated until convergence.

During the *Construction* phase, the feasibility of a given point $\mathbf{x}$ is imposed by a chance constraint on the feasibility-likelihood: $p(\mathbf{z} = 1 | \mathbf{x}; \Theta) \geq \beta$, where $\beta \in [0, 1]$. Vertices and edges are only accepted into the graph G if this constraint is satisfied, ensuring that they lie within the safe set $\mathcal{S}$ with a probability of at least β . Particles which violate the chance-constraint are removed from the vertex-candidate set. Edges are connected between two vertices if the edge-length is less than the connectivity radius ρ , which defines the set of nearest neighbors for each vertex. In practice, this can drastically reduce the number of edges in the graph, therefore decreasing the problem size in the planning phase.

In the *Planning* phase, a feasible path is generated by first performing a “collision-check” on a candidate edge to ensure that the constraint $p(\mathbf{z} = 1 | \mathbf{x}; \Theta) \geq \beta$ holds. Equidistant query points are generated along the edge at a fixed resolution, and edge-feasibility is approximated by ensuring

that all query points satisfy the chance constraint. If the edge is accepted into the graph G , its length is evaluated and stored. A user-specified search algorithm is then run to completion in order to generate a shortest-path solution ξ^* .

A subtle but key aspect of this VI-based PRM approach is that a candidate set of vertices are *not* generated by first fitting a parameterized distribution $q(\mathbf{x})$ to the target posterior $p(\mathbf{x} | \mathbf{z} = 1; \Theta)$, then sampling points from $q(\mathbf{x})$. Rather, the particles themselves both govern the implicit empirical distribution q and comprise the set of candidate vertices. As these vertices are *deterministically* optimized according to the SVGD algorithm, they will approach regions of the space which are more likely to be collision-free, however the repulsive-force term will ensure particle diversity and adequate coverage of the target distribution. The non-parametric nature of the algorithm allows this approach to be flexible and model highly multi-modal distributions.

This uncertainty-guided procedure for PRM generation is quite general, and does not necessitate sampling heuristics. However, the prior distribution $p(\mathbf{x})$ can be chosen to bias portions of the configuration space if required. Furthermore, this approach can exploit gradient-based information of the target distribution, and can be implemented in batch for efficient GPU parallelization.

With the β -parameter, we can control optimism in graph construction. Low values can lead to graphs which populate edges in uncertain or unexplored areas, producing optimistic path solutions with shorter paths. Conversely, higher values restrict vertices and edges to high-probability regions, resulting in longer paths. This behavior is depicted in Fig. 2 with a partial occupancy map generated from the Intel dataset.

Although SV-PRM requires a differentiable probability distribution to evaluate the gradient in Eq. (11), as discussed in Section III, this property can be satisfied with continuous Bayesian occupancy maps or is already assumed by popular trajectory optimization algorithms.

VI. EXPERIMENTS

We evaluate Stein Variational PRMs and relevant baselines on three planning environments: (1) a synthetic planar point-navigation problem with challenging obstacle distributions, (2) a probabilistic occupancy map generated from a real-world baseline, and (3) a high-dimensional simulated manip-

ulation problem. In all experiments, we use the anisotropic RBF kernel (Section IV) with the empirical Fisher information matrix as the metric $M = \frac{1}{N} \sum_{i=1}^N \nabla_{\mathbf{x}}^2 \log p(\mathbf{x}^i | z = 1)$ which we have observed to improve particle coverage over regular RBF kernels.

We evaluate different roadmap generation strategies, where the traditional PRM generates samples according to a (pseudo)-random number generator, potentially with cost-based rejection sampling. SV-PRM is initialized with the samples from PRM, then performs SVGD on the particle set. The connection strategy described in Section V is used for both roadmaps. The LAZYSP [30] algorithm is then used to search the roadmap, lazily evaluating edges to check whether the cost threshold would be exceeded. We parallelize SVGD evaluations using a PyTorch implementation, allowing efficient GPU-driven inference.

A. Navigation

In the synthetic CHECKERBOARD environment (Table I), we compare performance in low- and high-particle regimes across various initialization schemes. The continuous target distribution is generated by fitting a grid of RBF features to a binary occupancy map. In the real world INTEL environment (Fig. 1(a)), we fit a Bayesian Hilbert Map (BHM) [9] to LIDAR data from the Intel Research Lab in Seattle [31]. The resulting probabilistic occupancy map is the SV-PRM target distribution. We introduce a barrier function to prevent particles from exceeding the limits of configuration space.

B. 7-Dof Motion Planning

To demonstrate the scalability of SV-PRMs in higher-dimensional robotics problems, we test the approach on a simulated reaching task depicted in Fig. 1(b). Here, the start pose of the modeled Franka arm is configured to be reaching in the top-right compartment. A feasible path in joint-space $\mathbf{x} \in \mathbb{R}^7$ must be found to reach the goal configuration (green) in the lower left compartment, while avoiding the cabinet frame (grey). Similarly to Section III-B, the obstacle cost $\mathbf{h}(\mathbf{x})$ is modeled using a smoothly varying t-SDF, as used in [11], with an offset value of 0.25m, and a feasibility likelihood of $\exp(-10 \|\mathbf{h}(\mathbf{x})\|^2)$. The prior $p(\mathbf{x})$ is formulated as a uniform distribution, with exponentially decreasing probability near the system joint-limit values. As in the Navigation experiments, we compare our approach to traditional PRMs, varying the initialization type between random uniform samples and pseudo-random samples generated from a Halton sequence. Further, we add a task-based heuristic sampling distribution using a multi-modal mixture of Gaussians. Three isotropic Gaussian components are centered on the start and goal configurations, as well as on a retracted “home” away from the obstacles.

VII. RESULTS

We report *query-dependent* metrics on the CHECKERBOARD environment, which measure the benefits of using the SV-PRM for specific start-goal planning queries. We select a challenging start-goal pair that requires the roadmap




TABLE I: CHECKERBOARD environment roadmaps generated with different sampling strategies and number of samples. The planning query consists of the start (green) and goal (red), and the shortest path (blue) is visualized if it exists. Best viewed digitally.

	PRM				SV-PRM			
	U	U+Rej	H	H+Rej	U	U+Rej	H	H+Rej
N = 100	2	15	✗	✓	29	30	✓	✓
N = 250	20	30	✓	✓	30	30	✓	✓

TABLE II: Number of successful planning trials on the CHECKERBOARD environment with N=100 and 250 sampled vertices. Uniform-based samplers are evaluated across 30 random seeds. The low-discrepancy Halton-based sampler [32] is deterministic, so success is denoted with a ✓ or ✗. In the low-particle regime, uniform samples (with or without rejection) rarely capture a path between the start-goal pair. Despite these poor initializations, the SVGD optimization yields roadmaps with feasible paths in all but one trial.




Fig. 3: Maximum mean discrepancy (MMD) of SV-PRM(Uniform) with respect to the uniform feasible distribution of the ground truth fully-mapped Intel-BHM model (Fig. 2(c)). This differs from the partial map that SV-PRM optimizes the particles against, but demonstrates how particles have been distributed throughout the unmapped regions. Results are averaged across 100 trials. MMD decreases with SVGD iterations, and with increasing particle number N .

to pass through several narrow passages (Table I). Table II reports the number of successful planning trials for low- and high-particle regimes, showing that SV-PRMs better capture the connectivity of the space with the same number of samples. Compared to the larger PRMs that are needed to match the smaller SV-PRMs connectivity, this improved sample efficiency and sparsity yields faster planning times and statistically-equivalent path lengths.

On the INTEL environment, we additionally consider *query-independent* metrics that evaluate the approximation quality of the resulting roadmap without a specific planning query context. These metrics help understand the multi-query performance of SV-PRM.





Fig. 4: Roadmap coverage as a function of roadmap size, estimated by sampling 1000 collision-free states in the partial Intel-BHM map and attempting to connect them to roadmap vertices via collision-free edges. SV-PRM(Uniform) and PRM(Uniform+Rejection) achieve comparable coverage, while SV-PRM(Uniform+Rejection) achieves slightly higher coverage with greater improvement in the low-particle regime.

First, we simulate a partial map by limiting the amount of data ingested by the BHM. Fig. 2 visualizes the map when the BHM has consumed little data; the outer hallway loop has yet to be identified. The SV-PRM on this partial map reflects the resulting uncertainty in the model. Many vertices are in the likely collision-free spaces, but vertices are also distributed throughout the unmapped (and therefore uncertain) area. By choosing the cost threshold appropriately, SV-PRMs preserves more vertices and edges in the roadmap. The same SV-PRM can yield multiple candidate paths for a specific query by varying the threshold: a conservative threshold safely connects the start and goal via the known corridor, while an optimistic threshold navigates through the unknown and finds a shorter (but ultimately invalid) path. In this way, users are able to specify their degree of certainty in the map quality. SV-PRMs cull vertices accordingly to efficiently plan paths at the desired cost threshold.

To further characterize this partial map setting, we evaluate the maximum mean discrepancy (MMD) [33], which measures the difference between two distributions based on samples drawn from each of them. Fig. 3 visualizes the MMD as particles are optimized against the partial map, demonstrating that SVGD improves the MMD over time. Increasing the number of particles also lowers the MMD, although with diminishing returns.

Fig. 4 demonstrates that SV-PRMs improve the configuration space coverage relative to the corresponding random initialization. Note that this coverage metric does not characterize the connectivity of the roadmap; as demonstrated in Table I, rejection sampling can achieve good coverage without capturing the connectivity of the underlying space and yielding a feasible path. On query-dependent metrics, Fig. 5 shows that both SV-PRM initializations enjoy shorter or comparable paths for the same number of vertices as the PRM baselines, while solving a higher rate of problems.

Results for the 7-dof MANIPULATOR task are shown in Fig. 6, where we compare path statistics between SV-PRM and PRM for random and quasi-random initializations. SV-PRM manages to generate a higher degree of feasible joint-space across vertex counts for all cases, and yields lower



(a) Rate of trials with feasible solution. (b) Mean solution path lengths.

Fig. 5: INTEL environment with 100 random trials. (a) SV-PRM(Uniform+Rejection) succeeds on over 40% with 50 vertices and all 100 trials with 100 vertices. (b) SV-PRM(Uniform+Rejection) has lower solution path length than both PRM baselines, across all roadmap sizes. SV-PRM(Uniform) achieves comparable average solution path length to PRM(Uniform+Rejection), while solving a higher rate of problems.




Fig. 6: Solution statistics on the manipulation reaching experiment corresponding to Fig. 1(b). Success rate and average path length of solutions are compared between SV-PRM and the PRM baseline for 30 trials (10 per random seed) using different initialization distributions (*i.e.* priors) across varying number of vertices: $N \in \{16, 32, 64, 100, 128\}$. Particle-based posterior sampling with SV-PRM significantly improves the planning solution, particularly in the regime of few vertices.

path costs when solutions are found. As can be seen in the accompanying video, performing inference with SVGD before graph construction pushes pose configurations into collision-free regions, increasing coverage in feasible space.

VIII. DISCUSSION AND FUTURE WORK

The Stein Variational Probabilistic Roadmap proves to be a robust method for efficiently generating graphs well-suited for multi-query planning, outperforming existing biased-sampler PRM approaches. Future work will investigate how to pick the appropriate prior distributions as well as number of particles, with more emphasis on planning for manipulation as these results were very promising. Another exciting result worth investigating further is the tunable optimism parameter for partially explored environments. This is especially applicable when running the algorithm in an online setting where the particles are updated incrementally as new observations arrive, similar in spirit to the results in [9]. Finally, this approach would be an excellent complement to learned neural network samplers that leverage experience to guide sampling [18–21], where SV-PRMs can continue to optimize these biased samples or directly incorporate the scoring functions as a probabilistic cost.

REFERENCES

- [1] L. E. Kavraki, P. Švestka, J.-C. Latombe, and M. H. Overmars, “Probabilistic roadmaps for path planning in high-dimensional configuration spaces,” *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [2] S. Karaman and E. Frazzoli, “Sampling-based algorithms for optimal motion planning,” *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [3] K. Solovey, O. Salzman, and D. Halperin, “New perspective on sampling-based motion planning via random geometric graphs,” *The International Journal of Robotics Research*, vol. 37, no. 10, pp. 1117–1133, 2018.
- [4] L. Janson, B. Ichter, and M. Pavone, “Deterministic sampling-based motion planning: Optimality, complexity, and performance,” *The International Journal of Robotics Research*, vol. 37, no. 1, pp. 46–61, 2018.
- [5] D. Hsu, J.-C. Latombe, and R. Motwani, “Path planning in expansive configuration spaces,” in *IEEE International Conference on Robotics and Automation*, 1997.
- [6] Q. Liu and D. Wang, “Stein variational gradient descent: A general purpose Bayesian inference algorithm,” in *Advances in Neural Information Processing Systems*, 2016.
- [7] S. T. O’Callaghan and F. T. Ramos, “Gaussian process occupancy maps,” *The International Journal of Robotics Research*, vol. 31, no. 1, pp. 42–62, 2012.
- [8] F. Ramos and L. Ott, “Hilbert maps: Scalable continuous occupancy mapping with stochastic gradient descent,” *The International Journal of Robotics Research*, vol. 35, no. 14, pp. 1717–1730, 2016.
- [9] R. Senanayake and F. Ramos, “Bayesian Hilbert Maps for dynamic continuous occupancy mapping,” in *Conference on Robot Learning*, 2017.
- [10] N. Ratliff, M. Zucker, J. A. Bagnell, and S. S. Srinivasa, “CHOMP: Gradient optimization techniques for efficient motion planning,” in *IEEE International Conference on Robotics and Automation*, 2009.
- [11] M. Zucker, N. Ratliff, A. D. Dragan, M. Pivtoraiko, M. Klingensmith, C. M. Dellin, J. A. Bagnell, and S. S. Srinivasa, “CHOMP: Covariant hamiltonian optimization for motion planning,” *The International Journal of Robotics Research*, vol. 32, no. 9-10, pp. 1164–1193, 2013.
- [12] M. Mukadam, X. Yan, and B. Boots, “Gaussian process motion planning,” in *IEEE International Conference on Robotics and Automation*, 2016.
- [13] M. Mukadam, J. Dong, X. Yan, F. Dellaert, and B. Boots, “Continuous-time Gaussian process motion planning via probabilistic inference,” *The International Journal of Robotics Research*, vol. 37, no. 11, pp. 1319–1340, 2018.
- [14] V. Boor, M. H. Overmars, and A. F. van der Stappen, “The Gaussian sampling strategy for probabilistic roadmap planners,” in *IEEE International Conference on Robotics and Automation*, 1999.
- [15] C. Holleman and L. E. Kavraki, “A framework for using the workspace medial axis in PRM planners,” in *IEEE International Conference on Robotics and Automation*, 2000.
- [16] D. Hsu, T. Jiang, J. Reif, and Z. Sun, “The bridge test for sampling narrow passages with probabilistic roadmap planners,” in *IEEE International Conference on Robotics and Automation*, 2003.
- [17] M. Saroya, G. Best, and G. A. Hollinger, “Roadmap learning for probabilistic occupancy maps with topology-informed growing neural gas,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4805–4812, 2021.
- [18] B. Ichter, J. Harrison, and M. Pavone, “Learning sampling distributions for robot motion planning,” in *IEEE International Conference on Robotics and Automation*, 2018.
- [19] 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*, vol. 37, no. 1, pp. 48–66, 2021.
- [20] R. Kumar, A. Mandalika, S. Choudhury, and S. S. Srinivasa, “LEGO: Leveraging experience in roadmap generation for sampling-based planning,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019.
- [21] B. Ichter, E. Schmerling, T.-W. E. Lee, and A. Faust, “Learned critical probabilistic roadmaps for robotic motion planning,” in *IEEE International Conference on Robotics and Automation*, 2020.
- [22] A. Lambert, A. Fishman, D. Fox, B. Boots, and F. Ramos, “Stein variational model predictive control,” in *Proceedings of the 4th Annual Conference on Robot Learning*.
- [23] L. Barcelos, A. Lambert, R. Oliveira, P. Borges, B. Boots, and F. Ramos, “Dual online stein variational inference for control and dynamics,” in *Robotics: Science and Systems*, 2021.
- [24] A. Lambert and B. Boots, “Entropy regularized motion planning via stein variational inference,” *RSS Workshop on Integrating Planning and Learning*, 2021.
- [25] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.
- [26] 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,” *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [27] D. Wang, Z. Tang, C. Bajaj, and Q. Liu, “Stein variational gradient descent with matrix-valued kernels,” in *Advances in Neural Information Processing Systems*, 2019.

- [28] W. Y. Chen, A. Barp, F.-X. Briol, J. Gorham, M. Girolami, L. Mackey, and C. Oates, “Stein Point Markov Chain Monte Carlo,” in *International Conference on Machine Learning*, 2019.
- [29] G. Detommaso, T. Cui, A. Spantini, Y. Marzouk, and R. Scheichl, “A Stein variational Newton method,” in *Advances in Neural Information Processing Systems*, 2018.
- [30] C. Dellin and S. S. Srinivasa, “A unifying formalism for shortest path problems with expensive edge evaluations via lazy best-first search over paths with edge selectors,” in *International Conference on Automated Planning and Scheduling*, 2016.
- [31] A. Howard and N. Roy, “The robotics data set repository (radish),” 2003. [Online]. Available: <http://radish.sourceforge.net/>
- [32] J. H. Halton, “On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals,” *Numerische Mathematik*, vol. 2, no. 1, pp. 84–90, 1960.
- [33] A. Gretton, K. Borgwardt, M. J. Rasch, B. Scholkopf, and A. J. Smola, “A kernel method for the two-sample problem,” in *Advances in Neural Information Processing Systems*, 2008.