

Stein Variational Ergodic Surface Coverage with SE(3) Constraints

Jiayun Li^{1,2}, Yufeng Jin^{1,3}, Sangli Teng⁴, Dejian Gong¹, Georgia Chalvatzaki^{1,2}

Abstract—Surface manipulation tasks require robots to generate trajectories that comprehensively cover complex 3D surfaces while maintaining precise end-effector poses. Existing ergodic trajectory optimization (TO) methods demonstrate success in coverage tasks, while struggling with point-cloud targets due to the nonconvex optimization landscapes and the inadequate handling of SE(3) constraints in sampling-as-optimization (SAO) techniques. In this work, we introduce a preconditioned SE(3) Stein Variational Gradient Descent (SVGD) approach for SAO ergodic trajectory generation. Our proposed approach comprises multiple innovations. First, we reformulate point-cloud ergodic coverage as a manifold-aware sampling problem. Second, we derive SE(3)-specific SVGD particle updates, and, third, we develop a preconditioner to accelerate TO convergence. Our sampling-based framework consistently identifies superior local optima compared to strong optimization-based and SAO baselines while preserving the SE(3) geometric structure. Experiments on a 3D point-cloud surface coverage benchmark and robotic surface drawing tasks demonstrate that our method achieves superior coverage quality with tractable computation in our setting relative to existing TO and SAO approaches, and is validated in real-world robot experiments.

I. INTRODUCTION

Trajectory optimization (TO) for robots operating on complex surfaces with vision feedback represents a fundamental challenge in modern robotics, with critical applications spanning manufacturing automation [1], surgical robotics [2], and autonomous maintenance systems [3]. In such scenarios, robots must generate motion plans that simultaneously satisfy kinematic constraints, achieve comprehensive coverage of task-relevant regions, and maintain precise geometric relationships between the end-effector and surface geometry.

Ergodic TO provides a principled mathematical framework for addressing this challenge by generating trajectories that asymptotically visit regions in proportion to their information content [4]. Unlike myopic exploration strategies that focus on immediate information gain [5], ergodic methods optimize the long-term spatial distribution of robot visitation, ensuring both coverage completeness and task-oriented exploration.

However, when applied to discrete point-cloud targets with SE(3) end-effector constraints, existing ergodic TO methods face three critical limitations. First, point-cloud




Fig. 1. Robot drawing the letters 'ICRA' followed by a heart symbol on the cylindrical surface of a pot using a pen mounted on its end effector, guided by point cloud input.

ergodic coverage formulations generate highly non-convex optimization landscapes with numerous local minima that trap continuous optimization-based planners. Second, while ergodic TO has been extended to SE(3) spaces [6], existing Sampling As Optimization (SAO) methods using sampling or flow matching [7] designed to overcome local optima fail to properly handle the SE(3) manifold geometry, limiting their applicability to surface manipulation tasks which require very precise pose control. Third, flow-based and score-based SAO methods suffer from severe ill-conditioning in long-horizon TO problems, while effective preconditioning strategies for these approaches remain largely unexplored.

To address these limitations, we develop a theoretically grounded preconditioned SE(3) Stein Variational Gradient Descent (SVGD) framework for SAO ergodic TO with discrete point-cloud input. Our approach integrates ergodic optimization metrics with manifold-aware SVGD to enable parallel, geometrically consistent trajectory generation for surface manipulation tasks. We refer to this method as **Task-space Stein Variational Ergodic Coverage (TSVEC)**.

Our key contributions are as follows.

- **SVGD on the SE(3) manifold:** We derive a principled extension of SVGD to the SE(3) manifold.
- **Preconditioned ergodic TO:** We reformulate surface ergodic coverage as a nonlinear least-squares problem and introduce a Gauss-Newton (GN) preconditioner compatible within the SE(3) SVGD framework.
- **Comprehensive experimental validation:** We benchmark the proposed framework against strong baselines, showing consistent performance gains, and further validate its effectiveness in real-world robotic experiments.

II. RELATED WORK

A. Ergodic Coverage and Trajectory Optimization

Ergodic TO was introduced by Mathew and Mezić [4] as a framework for generating trajectories that visit regions pro-

Corresponding author: Jiayun Li (jiayun.li@tu-darmstadt.de).

Yufeng Jin and Sangli Teng contributed equally to this work.

This work was supported by the German Research Foundation (DFG) Emmy Noether Programme (CH 2676/1-1), the EU's Horizon Europe project "ARISE" (Grant no.: 101135959), the German Federal Ministry of Education and Research (BMBF) Project "RiG" (Grant no.: 16ME1001).

¹Department of Computer Science, TU Darmstadt, Darmstadt, Germany.

²Hessian.AI, Darmstadt, Germany.

³Honda Research Institute Europe GmbH, Offenbach, Germany.

⁴University of California, Berkeley, USA.

portional to a target information distribution. Miller et al. [8] demonstrated ergodic exploration for distributed information gathering, while subsequent work extended the approach to SE(2) spaces [9] and developed TO formulations [10].

A key challenge in ergodic TO is the initialization sensitivity problem. Current point-cloud formulations create highly nonconvex optimization landscapes where distant regions of interest cannot effectively attract trajectories from poor initializations [11]. This stems from the spectral decomposition of the ergodic metric, which creates numerous local minima that trap continuous optimization-based optimizers. Heat equation-driven approaches, like HEDAC [12], attempt to address this through diffusion-based information propagation, but perform poorly in continuous TO scenarios due to objective function imbalances that leave the optimization heavily influenced by initialization, failing to overcome the attraction basins of local minima [13].

Recent advances have extended ergodic methods to higher-dimensional spaces and SAO-based approaches. Sun et al. [6] developed fast ergodic search with kernel functions and demonstrated applications on SE(3) spaces. However, their approach focuses on single-point estimation with continuous target distributions and lacks integration with parallel sampling frameworks, necessary for exploring multiple trajectory modes. Sampling methods using SVGD [14] are computationally slow and lack awareness of SE(3) manifold constraints. Flow-based ergodic criteria using Sinkhorn divergence or Maximum Mean Discrepancy can handle multiple modes [7], [13], but are prohibitively expensive for TO with large point-cloud inputs, being limited to sparse subsampled representations. Safety-critical extensions [15] and tactile surface exploration [16] have demonstrated the versatility of ergodic methods, but computational scalability remains a fundamental limitation, restricting applications to control synthesis rather than full trajectory optimization.

B. Trajectory Optimization as Inference on SE(3)

The formulation of TO as probabilistic inference represents a fundamental paradigm shift in motion planning. Attias [17] pioneered this approach, followed by Toussaint [18] who established the foundation for continuous robot trajectory optimization. CHOMP (Covariant Hamiltonian Optimization for Motion Planning) [19] and GPMP (Gaussian Process Motion Planning) [20] further advanced this framework using functional gradients and factor graphs, respectively. While effective in Euclidean spaces, their extension to SE(3) manifolds remains challenging due to the complex geometric constraints.

SE(3) TO faces fundamental limitations in existing approaches. GPMP extensions to Lie groups [21] operate primarily in logarithm tangent space using locally linear approximations, which ignore the intrinsic manifold structure of SE(3) and can lead to accumulation of linearization errors over long horizons [22], [23]. This tangent space treatment fails to properly handle the non-commutative nature of rotations and parallel transport operations essential for SE(3) geometry. While recent work on Riemannian direct optimiza-

tion [24] for Lie group trajectories shows efficiency gains, its integration with SVGD frameworks remains unclear.

C. Equality-constrained / Manifold SVGD

Stein Variational Gradient Descent (SVGD) [25] provides a non-parametric approach for approximating complex distributions through particle-based variational inference. Although there exist manifold extensions, including Riemannian SVGD [26] for general Riemannian manifolds, projected SVGD [27] for subspace constraints, and orthogonal-gradient SVGD [28] for equality constraints used in TO applications [29], none of these are specifically designed for the SE(3) Lie group structure. These general manifold methods suffer from computational inefficiency when applied to SE(3) TO, as they fail to exploit the specific geometric properties and computational shortcuts available for matrix Lie groups.

Our proposed TSVEC addresses these limitations through a direct SE(3) manifold formulation of SVGD that exploits the manifold structure, combined with a trajectory-specific metric preconditioner, to enable robust exploration while maintaining computational efficiency on large point clouds.

III. PRELIMINARY

A. Ergodic Coverage Optimization Problem

A central component of ergodic TO is the *ergodic metric*, which measures the discrepancy between the spatial statistics of a trajectory and a target trajectory distribution. The target trajectory distribution $\varphi(\mathbf{x})$ is first encoded in the leading S Fourier basis as coefficients

$$\mathbf{c}_i = \int_{\Pi} \varphi(\mathbf{x}) F_i(\mathbf{x}) d\mathbf{x}, \quad i = 1, \dots, S, \quad (1)$$

where $F_i(\mathbf{x})$ are cosine Fourier basis functions defined over the bounded workspace $\Pi \subset \mathbb{R}^d$. Similarly, the time-averaged trajectory statistics are represented in the same basis as $c(\mathbf{X})_i$. With this representation, the ergodic objective can be written as a nonlinear least-squares problem:

$$V_e(\mathbf{X}) = \frac{1}{2} \sum_{i=1}^S \Lambda_i (c(\mathbf{X})_i - \mathbf{c}_i)^2, \quad (2)$$

where Λ_i are frequency-dependent weights. Minimizing $V_e(\mathbf{X})$ enforces that the trajectory $\mathbf{X} = \{\mathbf{x}_t \in \text{SE}(3)\}_{t=1}^{N_t}$, of length N_t , becomes ergodic with respect to the target distribution, i.e., that the time spent in each region is proportional to its utility. This formulation retains a nonlinear and nonconvex structure but is amenable to TO methods [8], [15]. It is worth noting that for discrete domains such as point clouds, one can first diffuse the point cloud to build a smooth distribution, then adopt the graph Fourier transform, i.e., the eigendecomposition of the discrete Laplace–Beltrami operator; Fig. 2 shows a diffusion example on a torus; see [16] for further details.




Fig. 2. From left to right, the figure shows a colored point cloud on a torus model, the results after graph diffusion, and the projection onto a graph Fourier basis.

B. Stein Variational Inference

SVGD updates a set of N particles $\{x_i\}$ so that their distribution progressively matches a target distribution $p(x)$, possibly unnormalized. Rather than directly optimizing the KL divergence in parameter space, SVGD finds the direction that decreases KL the fastest within a function space defined by a kernel $k(x, y)$, resulting in a kernelized update that transports particles toward the target. This yields a closed-form update rule ϕ^* which is a vector field shown in Eq. (3) that moves each particle x_i according to both the kernel-smoothed target distribution's score $\nabla \log p(x_i)$ and a repulsive kernel force $\nabla_1 k(x_i, x)$. We use the notation ∇_i to denote the gradient with respect to the i -th variable for a multi-input function. The repulsive force can be understood as the maximum entropy term in the general variational inference framework.

$$\phi^*(x) = \frac{1}{N} \sum_{i=1}^N \left[k(x_i, x) \nabla \log p(x_i) + \nabla_1 k(x_i, x) \right] \quad (3)$$

C. SE(3) Preliminary

To derive SVGD on SE(3), we need the algebraic and differential operations on this manifold. We consider perturbations of the special Euclidean group SE(3) in the form of right multiplication. For a matrix $\mathbf{P} \in \text{SE}(3)$, we define the perturbation as $\mathbf{P} \oplus \epsilon = \mathbf{P} \exp(\epsilon^\wedge)$, where $\epsilon \in \mathbb{R}^6$ is a small perturbation vector, $(\cdot)^\wedge$ denotes the hat operator mapping from \mathbb{R}^6 to the Lie algebra $\mathfrak{se}(3)$, and \exp is the matrix exponential. Similarly, we can define SE(3) subtraction as $\mathbf{P}' \ominus \mathbf{P} = \log(\mathbf{P}^{-1} \mathbf{P}')^\vee \in \mathbb{R}^6$, where the vee operator $(\cdot)^\vee$ maps from $\mathfrak{se}(3)$ to \mathbb{R}^6 . We use Ad and ad to denote the (group/algebra) adjoint operators on SE(3).

For a smooth function F that maps SE(3) to SE(3), its right-perturbed Jacobian can be written as a matrix $\mathbf{J}F(\mathbf{P}) \in \mathbb{R}^{6 \times 6}$:

$$\mathbf{J}F(\mathbf{P}) = \lim_{\epsilon \rightarrow 0} \frac{F(\mathbf{P} \oplus \epsilon) \ominus F(\mathbf{P})}{\epsilon}, \quad (4)$$

where the limit is taken component-wise, we refer to [30] for more details on SE(3), including the chain rule and basic Lie algebra operations. We consider the left-invariant inner product on the tangent space $\mathcal{T}_{\mathbf{P}}\text{SE}(3) \cong \mathfrak{se}(3)$ as $\langle \mathbf{P}\xi_1, \mathbf{P}\xi_2 \rangle_{\mathbf{P}} = \langle \xi_1, \xi_2 \rangle_{\mathbf{P}} = \frac{1}{2} \text{Tr}(\xi_1^\wedge (\xi_2^\wedge)^\top)$ for $\xi_1^\wedge, \xi_2^\wedge \in \mathfrak{se}(3)$. This right-perturbation model naturally preserves the group structure and provides a consistent framework

for optimization on SE(3). Since SE(3) is non-compact, there is, in general, no bi-invariant Riemannian metric, so the Riemannian exponential need not coincide with the Lie exponential. Nevertheless, one can perform second-order optimization by working in canonical coordinates and using the Cartan-Schouten connections to define gradients/Hessians; this yields Newton methods with local quadratic convergence using only the expansion of Lie exponentials, see [31].

IV. METHOD

A. Stein Variational Gradient Descent on SE(3)

To derive SVGD on SE(3) from first principles, we define a small perturbation (i.e., a transport map T) applied to particles on SE(3) that represent the variational distribution q . One natural choice is to directly adopt the SE(3) right addition $T(\mathbf{P}) = \mathbf{P} \oplus \tau \phi(\mathbf{P})$. Here, $\phi : \text{SE}(3) \rightarrow \mathbb{R}^6$ denotes the SVGD vector field on SE(3), and the updates are performed with a small step size τ . We follow the standard approach to derive SVGD by examining the steepest descent direction of the perturbed KL divergence:

$$\text{KL}(q_{[T]} \| p) = \mathbb{E}_{\mathbf{P} \sim q} [\log q(T(\mathbf{P})) - \log p(\mathbf{P}) + \log Z], \quad (5)$$

where $q_{[T]}$ denotes the variational distribution after applying the transport map T , and Z is the normalization constant, which is invariant under perturbations and can therefore be neglected. Starting from equation (5), we apply the change of variables formula to transform the analysis from the particle variational distribution to the unnormalized posterior distribution. Specifically, by applying T^{-1} to the particles, we can work with the inverse perturbation that directly acts on the target distribution rather than the variational distribution. This transformation allows us to derive the steepest descent direction with respect to the target distribution, which can be shown to be:

$$\nabla_\tau \text{KL}(q_{[T]} \| p) = -\mathbb{E}_{\mathbf{P} \sim q} [\nabla_\tau \log p_{[T^{-1}]}(\mathbf{P})], \quad (6)$$

where the terms inside the expectation are

$$\nabla_\tau \log p_{[T^{-1}]}(\mathbf{P}) = \nabla_\tau \log p(T(\mathbf{P})) + \nabla_\tau \log |\det(\mathbf{J}T(\mathbf{P}))|. \quad (7)$$

These derivations follow closely the supplementary material of [25].

The first term in Eq. (7) is equal to the gradient of the log-likelihood function $\text{grad} \log p(\mathbf{P})$. For an objective function $f(\mathbf{P})$ with domain on SE(3), the gradient with respect to right perturbations is computed by first obtaining the Euclidean gradient $\nabla f(\mathbf{P}) \in \mathbb{R}^{4 \times 4}$ and then projecting it onto the tangent space $\mathcal{T}_{\mathbf{P}}\text{SE}(3)$. Given the definition of the inner product, the following directional derivative holds:

$$\nabla_\tau f(\mathbf{P} \oplus \tau \epsilon) \Big|_{\tau=0} = \langle \text{grad} f(\mathbf{P}), \mathbf{P} \epsilon^\wedge \rangle_{\mathbf{P}}, \quad (8)$$

where $\epsilon \in \mathbb{R}^6$ is a perturbation vector. The gradient $\text{grad} f(\mathbf{P})$ can be computed by:

$$\text{grad} f(\mathbf{P}) = \mathbf{P} \begin{bmatrix} \boldsymbol{\Omega} & \mathbf{R}^\top \frac{\partial f}{\partial \mathbf{r}} \\ \mathbf{0}_{1 \times 3} & 0 \end{bmatrix}, \quad (9)$$

where

$$\Omega = \mathbf{R}^\top \frac{\partial f}{\partial \mathbf{R}} - \frac{\partial f}{\partial \mathbf{R}}^\top \mathbf{R} \quad (10)$$

and \mathbf{P} is parameterized as:

$$\mathbf{P} = \begin{bmatrix} \mathbf{R} & \mathbf{r} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \quad (11)$$

with $\mathbf{R} \in \text{SO}(3)$ the orientation and $\mathbf{r} \in \mathbb{R}^3$ the position. For further technical details about the gradient of $\text{SE}(3)$ functions, see [32] (Section 8.1) and [33].

The second term involving the derivative of the log-determinant can be expressed using the standard identity:

$$\nabla_\tau \log |\det(\mathbf{J}\mathbf{T}(\mathbf{P}))| = \text{Tr}[\mathbf{J}\mathbf{T}(\mathbf{P})^{-1} \nabla_\tau \mathbf{J}\mathbf{T}(\mathbf{P})] \quad (12)$$

The $\mathbf{J}\mathbf{T}(\mathbf{P})$ approaches the identity for sufficiently small τ ; however, its derivative with respect to τ is not constant. One can show this by using the definition of the transport map Jacobian in Eq. (4) with an additional perturbation ϵ for Jacobian computation. Firstly, we take the derivative with respect to τ :

$$\nabla_\tau \mathbf{J}\mathbf{T}(\mathbf{P}) \approx \phi(\mathbf{P} + \epsilon^\wedge \mathbf{P})^\wedge (I + \epsilon^\wedge) - (I + \epsilon^\wedge) \phi(\mathbf{P})^\wedge \quad (13)$$

$$\approx [\mathbf{J}\phi(\mathbf{P})\epsilon]^\wedge + \text{Lie}[\phi(\mathbf{P})^\wedge, \epsilon^\wedge], \quad (14)$$

where $\text{Lie}[\cdot]$ is the Lie bracket. Then, we extract the Jacobian by extracting the first-order term of ϵ :

$$\nabla_\tau \mathbf{J}\mathbf{T}(\mathbf{P}) = \mathbf{J}\phi(\mathbf{P}) + \text{ad}_{\phi(\mathbf{P})^\wedge}. \quad (15)$$

The trace of the $\mathfrak{se}(3)$ adjoint is 0; therefore $\nabla_\tau \log |\det(\mathbf{J}\mathbf{T}(\mathbf{P}))| = \text{Tr}[\mathbf{J}\phi(\mathbf{P})]$, which is the maximum entropy term of SVGD on $\text{SE}(3)$. By using the notation of gradient to represent the trace of the Jacobian element-wise, the steepest descent direction is

$$\nabla_\tau \text{KL}(q_{[T]}||p) = -\mathbb{E}_{\mathbf{P} \sim q} \left[\sum_{i=1}^6 \langle \Phi_i(\mathbf{P}), \mathbf{P} I_i^\wedge \rangle_{\mathbf{P}} \right], \quad (16)$$

where $\Phi_i(\mathbf{P}) = \phi_i(\mathbf{P}) \text{grad} \log p(\mathbf{P}) + \text{grad} \phi_i(\mathbf{P})$ represents the gradient contribution from the i -th component of the vector field ϕ , with $\{I_i^\wedge\}_{i=1}^6$ forming the standard basis for $\mathfrak{se}(3)$. To obtain a closed-form solution for discrete particles, we apply the kernel trick (reproducing property) introduced in SVGD [25] to represent the vector field ϕ using a kernel function k , and we constrain the gradient norm to be unity with respect to the kernel-induced norm. We obtain the contribution of the gradient from \mathbf{P}_i to \mathbf{P}' , as given in Eq. (19), when evaluating the vector field at \mathbf{P}' . It should be noted that the gradient is still represented at point \mathbf{P} , therefore, a parallel transport scheme is needed.

In the context of SVGD on $\text{SE}(3)$ manifolds, parallel transport of tangent vectors between different tangent spaces is essential for maintaining geometric consistency when comparing or aggregating gradient information across particles. For the group $\text{SE}(3)$, we must transport tangent vectors from one tangent space $\mathcal{T}_{\mathbf{P}}\text{SE}(3)$ to another $\mathcal{T}_{\mathbf{P}'}\text{SE}(3)$ to ensure




Fig. 3. The $\text{SE}(3)$ SVGD mechanism for sharing gradient information on the manifold: computing the natural gradient (black arrow) and applying parallel transport to obtain the final update direction (red arrow).

that the kernel-induced interactions between particles respect the underlying manifold structure.

The parallel transport operation can be efficiently computed using the adjoint representation of the Lie group. Specifically, for a tangent vector $\xi \in \mathcal{T}_{\mathbf{P}}\text{SE}(3) \cong \mathfrak{se}(3)$, its parallel transport to $\mathcal{T}_{\mathbf{P}'}\text{SE}(3)$ is given by:

$$\mathcal{P}_{\mathbf{P} \rightarrow \mathbf{P}'}(\xi) = \text{Ad}_{\mathbf{P}'^{-1}\mathbf{P}}\xi \quad (17)$$

where $\text{Ad}_{(\cdot)}$ denotes the adjoint operator to change the basis for vectors in different tangent spaces. This formulation ensures that the transported vectors maintain their geometric meaning while accounting for the relative transformation between the two poses. Therefore, the target vector field is:

$$\phi^*(\mathbf{P}') = \frac{1}{N} \sum_{i=1}^N \mathcal{P}_{\mathbf{P}_i \rightarrow \mathbf{P}'}(\phi(\mathbf{P}_i, \mathbf{P}')), \quad (18)$$

$$\phi(\mathbf{P}_i, \mathbf{P}') = k(\mathbf{P}_i, \mathbf{P}') \text{grad} \log p(\mathbf{P}_i) + \text{grad}_1 k(\mathbf{P}_i, \mathbf{P}'). \quad (19)$$

The idea of $\text{SE}(3)$ SVGD is illustrated in Fig. 3, highlighting the role of parallel transport in aggregating gradients consistently on the manifold.

B. Ergodic TO as Stein Variational Inference on $\text{SE}(3)$

In TO tasks with long horizons, particularly in the ergodic TO setting, trajectories may span several hundred steps. This leads to severe ill-conditioning of the optimization problem, thereby necessitating the development of additional preconditioners to accelerate convergence [19], [20].

To perform surface ergodic trajectory optimization with $\text{SE}(3)$ SVGD, we formulate our inference problem using a Boltzmann likelihood $p(\mathbf{X}) \propto \exp(-V(\mathbf{X}))$, where the energy $V(\mathbf{X})$ is the sum of four terms: V_s (smoothness), V_a (surface normal alignment), V_f (surface attachment), and V_e (ergodic metric) along the trajectory $\mathbf{X} = \{\mathbf{x}_t \in \text{SE}(3)\}_{t=1}^{N_t}$. We assume that all Lie operations (e.g., gradient, parallel transport, wedge, and vee operators) are computed element-wise for each time step. We also assume access to a continuously differentiable signed distance function (SDF) $f_{\text{sdf}} \in C^2$ whose gradient is normalized, i.e., $\|\nabla f_{\text{sdf}}\|_2 = 1$. The weighted energy terms are

$$V_s(\mathbf{X}) = \frac{w_s}{2} \sum_{t=2}^{N_t-1} \|\mathbf{x}_{t+1} \ominus \mathbf{x}_t - \mathbf{x}_t \ominus \mathbf{x}_{t-1}\|_2^2, \quad (20a)$$

$$V_a(\mathbf{X}) = \frac{w_a}{2} \sum_{t=1}^{N_t} \|\mathbf{x}_t(z)^\top \nabla f_{sdf}(\mathbf{x}_t) - 1\|_2^2, \quad (20b)$$

$$V_f(\mathbf{X}) = \frac{w_f}{2} \sum_{t=1}^{N_t} \|f_{sdf}(\mathbf{x}_t)\|_2^2, \quad (20c)$$

$$V_e(\mathbf{X}) = \frac{w_e}{2} \sum_{i=1}^S \Lambda_i (c(\mathbf{X})_i - \mathbf{c}_i)^2. \quad (20d)$$

The smoothness term $V_s(\mathbf{X})$ is designed to minimize pose acceleration; The surface normal alignment energy $V_a(\mathbf{X})$ encourages the z -axis of the via-points to align with the normals of the SDF. In contrast, the surface attachment energy $V_f(\mathbf{X})$ ensures that the trajectory remains close to the SDF. Finally, the ergodic metric serves as the objective function introduced in Section III-A. Note that all energy terms are constructed in a sum-of-squares form and can be written compactly as $V(\mathbf{X}) = \frac{1}{2} r(\mathbf{X})^\top r(\mathbf{X})$, which enables Gauss–Newton (GN)-style preconditioning with $\mathbf{H} = \mathbf{J}_r^\top \mathbf{J}_r$, where $\mathbf{J}_r = \mathbf{J}r(\mathbf{X})$ can be derived in the same manner as in Eq. (4). The structure of the Hessian for smoothness energy resembles that of the CHOMP planner [19]. When formulated on the Lie group, the induced smoothness metric remains coordinate-invariant, as in CHOMP; however, it is no longer constant. By directly applying the preconditioner with the gradient and retraction, one can obtain an efficient manifold-aware TO algorithm. We will evaluate it in section V (GN method). Therefore, the preconditioned SVGD component $\hat{\psi}$ can be implemented following the methods introduced in [34], [35]. For a set of N_p trajectory particles with preconditioning matrices $\{\mathbf{H}_i \in \mathbb{R}^{6N_t \times 6N_t}\}_{i=1}^{N_p}$, we adopt the block-diagonal preconditioning method proposed in [34], the update direction is $\alpha \in \mathbb{R}^{6N_t}$ by solving the following linear system

$$\left(\sum_{i=1}^{N_p} \mathbf{H}_i k(\mathbf{X}_i, \mathbf{X}')^2 + \text{grad}_1 k(\mathbf{X}_i, \mathbf{X}')^{\otimes 2} \right) \alpha = \phi^*(\mathbf{X}')^\vee, \quad (21)$$

where the term $\text{grad}_1 k(\mathbf{X}_i, \mathbf{X}')^{\otimes 2}$ denotes the Euclidean outer product of the parallel transported kernel gradient $\mathcal{P}_{\mathbf{X}_i \rightarrow \mathbf{X}'}(\text{grad} k(\mathbf{X}_i, \mathbf{X}')^\vee) \in \mathbb{R}^{6N_t}$. This expression represents a direct extension of the approach proposed in [34]. The outer product of the kernel gradient can be interpreted as additional metric terms induced by the kernel structure. The kernel-weighted averaging of Gauss–Newton metrics provides a computational compromise necessary for algorithmic feasibility. Moreover, the computation can be further simplified through kernel design.

We adopt a sum of SE(3) kernels that compares the trajectory at each individual time step, which both simplifies the computation and mitigates potential performance degradation

caused by high-dimensional particles.

$$k(\mathbf{X}, \mathbf{X}') = \sum_{t=1}^{N_t} \exp\left(-\frac{1}{l} \|\mathbf{x}_t \ominus \mathbf{x}'_t\|_2^2\right), \quad (22)$$

where l is the kernel length scale hyperparameter. All gradients can be efficiently computed by first performing automatic differentiation and then projecting onto the tangent space using the gradient in Eq. (9). The Jacobian of the residual can be computed analytically using the fundamental Jacobian of the Lie group and the chain rule; to avoid a cluttered explanation, we refer the reader to [30] for further technical details. To ensure strict constraint compliance, the current framework can be readily extended to an Augmented Lagrangian constrained optimization approach; however, for simplicity, we leave this for future work.

V. EXPERIMENT

A. Benchmark Test

To evaluate the performance of our proposed ergodic TO algorithm, we designed a comprehensive benchmark using point cloud problems with colored point clouds as regions of interest (ROI), as shown in Fig. 4. Each mesh was sourced from standard mesh datasets and uniformly sampled according to surface area to generate evenly distributed point clouds. The ROIs were intentionally positioned at distant locations to rigorously test the TO algorithms' effectiveness.

We conducted two groups of comparative experiments: single-point estimation methods (i.e., deterministic optimization that outputs a single trajectory/mode rather than a particle set) and particle-based methods. Within each group, all TO algorithms employed identical hyperparameters, initialization strategies, and shared random seeds to ensure fair comparison. To mitigate pathological optimization on a discrete point cloud that requires normal computation using K-nearest-neighbor, all TO methods utilized a neural SDF to represent object geometry. The neural SDF was implemented using an efficient JAX framework, enabling object training completion within 10 seconds. Ergodic metrics were computed directly on point clouds using a CUDA-accelerated batched nearest neighbor package (JAXkd). All experiments were conducted on a single RTX 4080 GPU using single-precision arithmetic for compatibility with the CUDA-accelerated point cloud KNN implementation (K=60).

For the single point estimation group: We selected LBFGS from the Optax library with backtracking line search as the off-the-shelf optimizer for unconstrained TO. Gradients were computed using JAX autodiff, with quaternion components projected to the unit sphere after each update, serving as a vanilla baseline unaware of the manifold structure. We also employed the IPOPT optimizer via its official Python wrapper, coupled with the autodiff pipeline to represent second-order TO in linearized coordinates. SE(3) trajectories were parameterized in Lie algebra space and mapped back to the manifold via the exponential map for objective function evaluation. The CHOMP-like TO algorithm proposed in Section IV-B was implemented in JAX and denoted as GN,

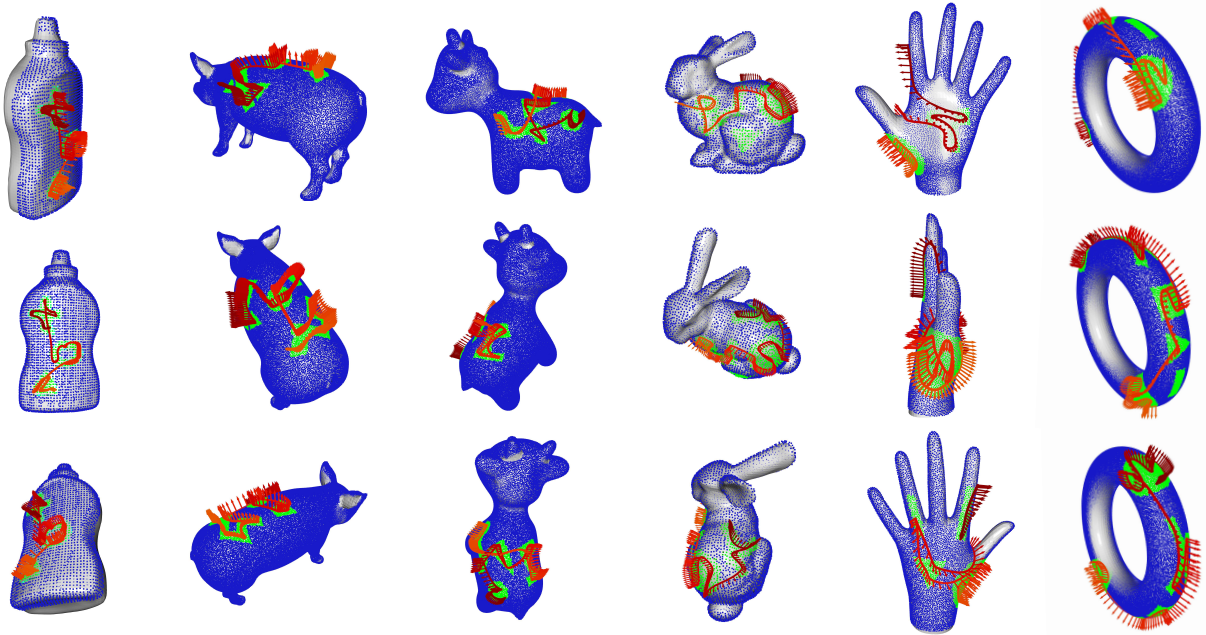


Fig. 4. Six benchmark problems. From left to right: Mustard Bottle (simplified as Must), Pig, Spot, Bunny, Hand, and Torus. The background point cloud is shown in blue, ROIs in green, and the SDF in silver to enhance visualization. The TSVEC-planned trajectory is depicted as a solid line transitioning from light red to dark red, with normal vectors indicating time steps 1–200. Only the Z-axis is shown to avoid visual clutter. The objects are arranged according to ROI curvature variation, from flat to highly curved.

which uses backtracking line search and adaptive damping as globalization.

For the particle-based group: We implemented our method, TSVEC, as well as the vanilla Stein ergodic TO introduced in [14], denoted as Stein Ergodic (SE), using JAX. To make SE work on SE(3), we adopted the SE(3) gradient, retraction, and parallel transport mechanisms. To showcase that the performance gain is not due to accidentally better initialization, we implemented Batch GN, which runs a batch of parallel GN as in single-point estimation without the Stein variational mechanism.

The trajectory length N_t is 200 steps, generating an optimization problem with 1400 decision variables. The weights assigned in V_s, V_a, V_f , and V_e are 5.0, 3.0, 3.0, and 0.1, respectively. To distribute the effect of the trajectory realistically, we apply a Gaussian smoothing kernel, simulating the agents’ radius. The stopping criterion was $1e^{-4}$ for single precision compatibility. We used straight line initialization for all methods connecting the furthest two colored point clouds, as this might be the only prior knowledge available about the scenario. The start poses were initialized and aligned with the world coordinate frame. For particle-based methods, we added small white noise with 0.005 variance to each via point, as SVGD-based optimizers generate infinitely large kernel gradients if two particles overlap. We use point cloud diffusion to generate a globally smooth vector field, convexly combined with local sharp features to guide the optimization. The number of principal components was set to 300 to capture complex surface patterns.

The theoretically sound merit function for SVGD line search or stopping criteria is the kernelized Stein discrepancy

(KSD) discussed in [36], which provides a principled way to quantify how far the particle distribution deviates from the target in kernel space. However, KSD evaluation is expensive and numerically unstable due to its dependence on higher-order derivatives. Therefore, no reasonable stopping criterion exists for TSVEC and SE. Thus, we ran them for maximum iterations: 1000 for TSVEC and 9999 for SE. Since vanilla SE suffers from convergence issues, we allowed it ten times more iterations. The step size for TSVEC is 0.1, and 0.02 for SE. All particle methods used 100 particles. We tested this within our framework and found that using 500 particles does not incur memory issues on a single GPU. Kernel lengths were set to 0.05, except for the torus (0.1), as it needs more exploration. We found that the mean adaptive kernel-length heuristic was less effective than expected, and therefore removed it from the current framework. Results for methods tested on the benchmark were run 10 times, and we reported their mean values in TABLE I. We do not emphasize computation time as a criterion, since IPOPT requires dense Hessian computation via automatic differentiation and is therefore extremely slow. Instead, we use the iteration count and the final objective value as the key evaluation criteria. A dash (‘-’) indicates method failure.

The table is organized by the curvature variation of the scenario surface. As shown, L-BFGS rarely achieves satisfactory performance, with direct quaternion clipping causing premature termination. While IPOPT performs considerably better, it outperforms the GN planner only once in the low surface curvature case. GN consistently surpasses both baselines in convergence iterations and final objective values. Moreover, Batch GN consistently outperforms single GN,

which aligns with expectations. Most notably, TSVEC consistently achieves superior performance across all scenarios, obtaining significantly better objective values than all competing approaches. In contrast, SE suffers from convergence difficulties, as the 200-step trajectory results in a severely ill-conditioned optimization problem. During experimentation, we observed that while the ergodic objective’s mathematical form remains unchanged, it exhibits slight non-smoothness due to the nearest-neighbor search required for its evaluation.

TABLE I
PERFORMANCE COMPARISON ACROSS BENCHMARK SCENARIOS.

Scen.	Method	V_s	V_a	V_f	V_e	V	Iter.	Time(s)
Must	LBFGS	2.26e-2	7.02e-4	2.31e-4	1.87e-1	2.10e-1	153	3.16
	IPOPT	1.67e-2	9.48e-3	1.06e-3	1.07e-1	1.34e-1	1000	600
	GN	7.84e-3	1.67e-3	9.82e-4	1.41e-1	1.51e-1	58	0.392
	SE	1.37e-1	2.23e+0	2.98e-2	3.86e-1	2.79e+0	9999	220
Must	Batch GN	6.32e-3	1.81e-3	8.84e-4	1.24e-1	1.33e-1	116	12.1
	TSVEC	1.78e-3	1.98e-4	2.12e-4	8.88e-2	9.10e-2	1000	138
	SE	1.37e-1	2.23e+0	2.98e-2	3.86e-1	2.79e+0	9999	220
Pig	LBFGS	-	-	-	-	-	-	-
	IPOPT	3.37e-3	4.87e-3	1.62e-4	3.95e-1	4.03e-1	499	128
	GN	3.29e-3	6.28e-4	4.57e-5	2.97e-2	3.37e-2	52	0.441
	SE	1.47e-2	1.63e-1	6.35e-3	6.18e-1	8.02e-1	9999	249
Pig	Batch GN	2.69e-3	3.97e-4	4.51e-5	2.68e-2	2.99e-2	123	14.0
	TSVEC	1.16e-3	1.33e-4	7.68e-6	9.42e-3	1.07e-2	1000	145
	SE	1.47e-2	1.63e-1	6.35e-3	6.18e-1	8.02e-1	9999	249
Spot	LBFGS	1.97e-2	3.32e-2	5.44e-3	5.11e-1	5.69e-1	137	9.58
	IPOPT	5.20e-3	6.94e-3	4.28e-4	3.16e-1	3.28e-1	1000	290
	GN	2.04e-3	7.91e-4	4.02e-5	1.87e-1	1.90e-1	40	0.359
	SE	1.22e-1	8.45e-1	9.57e-3	1.28e+0	2.26e+0	9999	244
Spot	Batch GN	1.58e-3	3.84e-4	2.29e-5	1.79e-2	1.99e-2	119	13.6
	TSVEC	1.17e-3	2.54e-4	9.04e-6	6.86e-3	8.29e-3	1000	144
	SE	1.22e-1	8.45e-1	9.57e-3	1.28e+0	2.26e+0	9999	244
Bunny	LBFGS	7.21e-3	6.24e-3	7.15e-4	1.34e-1	1.48e-1	262	3.88
	IPOPT	5.42e-3	4.50e-3	5.15e-4	1.16e-1	1.26e-1	199	38.8
	GN	5.56e-3	2.69e-3	3.73e-4	7.74e-2	8.60e-2	36	0.317
	SE	2.06e-2	8.98e-2	4.44e-3	2.44e-1	3.59e-1	9999	224
Bunny	Batch GN	4.44e-3	2.00e-3	2.61e-4	3.65e-2	4.32e-2	107	11.5
	TSVEC	3.81e-3	1.59e-3	2.19e-4	1.78e-2	2.34e-2	1000	138
	SE	2.06e-2	8.98e-2	4.44e-3	2.44e-1	3.59e-1	9999	224
Hand	LBFGS	7.34e-3	5.10e-3	1.38e-3	5.22e-1	5.36e-1	70	3.04
	IPOPT	3.61e-3	2.16e-3	2.88e-4	3.07e-1	3.13e-1	256	61.3
	GN	1.01e-2	2.29e-3	3.12e-3	1.50e-1	1.66e-1	49	0.412
	SE	2.39e-2	1.86e-2	2.04e-3	6.79e-1	7.24e-1	9999	230
Hand	Batch GN	6.76e-3	1.51e-3	5.08e-4	9.97e-2	1.08e-1	150	16.3
	TSVEC	4.49e-3	6.81e-4	2.83e-4	6.83e-2	7.38e-2	1000	139
	SE	2.39e-2	1.86e-2	2.04e-3	6.79e-1	7.24e-1	9999	230
Torus	LBFGS	3.71e-2	1.95e-2	1.98e-3	7.41e-1	8.00e-1	114	9.47
	IPOPT	9.31e-3	1.26e-2	1.03e-3	6.39e-1	6.62e-1	1000	274
	GN	3.66e-3	5.56e-4	9.50e-5	1.78e-1	1.83e-1	60	0.525
	SE	5.87e-2	4.44e-1	1.72e-2	9.16e-1	1.44e+0	9999	254
Torus	Batch GN	3.09e-3	4.58e-4	5.51e-5	8.74e-2	9.10e-2	294	32.6
	TSVEC	2.99e-3	6.88e-4	2.63e-4	5.68e-2	6.08e-2	1000	140
	SE	5.87e-2	4.44e-1	1.72e-2	9.16e-1	1.44e+0	9999	254

B. Real World Task

To evaluate TSVEC’s performance in real-world surface manipulation, we deployed a Franka Emika Panda robot to track trajectories corresponding to SE(3) motions planned by our framework. The target point cloud represents hand-written random geometries and characters, obtained by sampling from the mesh and colored using a rendering tool. A cooking pot with a 22 cm diameter and 16 cm height was used to simulate a cylindrical curved surface. To execute the task, we first converted TSVEC task-space trajectories into joint-space commands using a sequential MuJoCo inverse-kinematics (IK) planner with damped least squares, warm-started from the previous time step’s solution. The resulting joint trajectories were then executed on the robot via joint-




Fig. 5. Left: SE(3) planning result on the hand-colored point cloud surface. Right: the corresponding sequential IK solution. The viewing angle matches that of Fig. 1, allowing a direct comparison between the planned and real-world results.

space position control. The real-world tracking results are presented in Fig. 1.

When applied to the drawing task, the GN planner is largely unable to produce recognizable characters, as it frequently becomes trapped in local minima and lacks sufficient exploration. This limitation is particularly critical in ergodic trajectory optimization, where broader exploration is necessary to discover better suboptimal solutions; instead, GN trajectories tend to converge directly to poor local minima. In contrast, TSVEC generates trajectories that are mostly recognizable and substantially reduces the number of trials required. An example of the planning result for the letter “A” is shown in Fig. 5.

VI. CONCLUSIONS

This work introduced TSVEC, a preconditioned SE(3) SVGD framework for ergodic trajectory optimization on discrete point cloud surfaces. By casting the problem as manifold-aware sampling-as-optimization, our method addresses a key difficulty of local gradient-based ergodic trajectory optimization on point cloud targets, namely its tendency to become trapped in poor local minima under highly non-convex surface distributions.

The proposed approach combines SE(3)-aware Stein variational updates, which maintain geometric consistency in particle interactions, with Gauss–Newton-style preconditioning that improves the conditioning of long-horizon optimization problems. Together, these design choices make the resulting SAO framework substantially more robust to the trivial local optima that often affect optimization-based alternatives. Across the benchmark scenarios considered in this paper, TSVEC consistently achieved better coverage quality and lower final objective values than strong baselines, including IPOPT, L-BFGS, and SE(3)-aware Gauss–Newton methods.

In the real-world surface drawing experiments, TSVEC further produced recognizable character trajectories with little hand tuning, whereas the optimization-based planners we tested often converged prematurely to trajectories with poor visual structure. Taken together, these results show that TSVEC is an effective and robust method for ergodic trajectory optimization on discrete point cloud surfaces, particularly in settings where non-convexity and SE(3) geometric constraints make purely local optimization difficult.

REFERENCES

- [1] B. Maric, A. Mutka, and M. Orsag, "Collaborative human-robot framework for delicate sanding of complex shape surfaces," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2848–2855, 2020.
- [2] J. Fu, I. Burzo, E. Iovene, J. Zhao, G. Ferrigno, and E. De Momi, "Optimization-based variable impedance control of robotic manipulator for medical contact tasks," *IEEE Transactions on Instrumentation and Measurement*, vol. 73, pp. 1–8, 2024.
- [3] P. Pfändler, K. Bodie, G. Crotta, M. Pantic, R. Siegart, and U. Angst, "Non-destructive corrosion inspection of reinforced concrete structures using an autonomous flying robot," *Automation in Construction*, vol. 158, p. 105241, 2024.
- [4] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D: Nonlinear Phenomena*, vol. 240, no. 4–5, pp. 432–442, 2011.
- [5] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell, "Curiosity-driven exploration by self-supervised prediction," in *International conference on machine learning*. PMLR, 2017, pp. 2778–2787.
- [6] M. M. Sun, A. Gaggari, P. Trautman, and T. Murphey, "Fast ergodic search with kernel functions," *IEEE Transactions on Robotics*, 2025.
- [7] M. M. Sun, A. Pinosky, and T. Murphey, "Flow matching ergodic coverage," in *Proceedings of Robotics: Science and Systems*, 2025.
- [8] L. M. Miller, Y. Silverman, M. A. MacIver, and T. D. Murphey, "Ergodic exploration of distributed information," *IEEE Transactions on Robotics*, vol. 32, no. 1, pp. 36–52, 2015.
- [9] L. M. Miller and T. D. Murphey, "Trajectory optimization for continuous ergodic exploration on the motion group $se(2)$," in *52nd IEEE Conference on Decision and Control*. IEEE, 2013, pp. 4517–4522.
- [10] —, "Trajectory optimization for continuous ergodic exploration," in *2013 American Control Conference*. IEEE, 2013, pp. 4196–4201.
- [11] E. Ayvali, H. Salman, and H. Choset, "Ergodic coverage in constrained environments using stochastic trajectory optimization," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 5204–5210.
- [12] S. Ivić, B. Crnković, and I. Mezić, "Ergodicity-based cooperative multiagent area coverage via a potential field," *IEEE transactions on cybernetics*, vol. 47, no. 8, pp. 1983–1993, 2016.
- [13] C. Hughes, H. Warren, D. Lee, F. Ramos, and I. Abraham, "Ergodic trajectory optimization on generalized domains using maximum mean discrepancy," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 01–07.
- [14] D. Lee, C. Lerch, F. Ramos, and I. Abraham, "Stein variational ergodic search," in *Proceedings of Robotics: Science and Systems*, 2024.
- [15] C. Lerch, D. Dong, and I. Abraham, "Safety-critical ergodic exploration in cluttered environments via control barrier functions," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 10205–10211.
- [16] C. Bilaloglu, T. Löw, and S. Calinon, "Tactile ergodic coverage on curved surfaces," *IEEE Transactions on Robotics*, 2025.
- [17] H. Attias, "Planning by probabilistic inference," in *International workshop on artificial intelligence and statistics*. PMLR, 2003, pp. 9–16.
- [18] M. Toussaint, "Robot trajectory optimization using approximate inference," in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 1049–1056.
- [19] 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.
- [20] 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.
- [21] J. Dong, M. Mukadam, B. Boots, and F. Dellaert, "Sparse gaussian processes on matrix lie groups: A unified framework for optimizing continuous-time trajectories," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 6497–6504.
- [22] J. L. Blanco-Claraco, "A tutorial on $se(3)$ transformation parameterizations and on-manifold optimization," *arXiv preprint arXiv:2103.15980*, 2021.
- [23] R. Wagner, O. Birbach, and U. Frese, "Rapid development of manifold-based graph optimization systems for multi-sensor calibration and slam," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 3305–3312.
- [24] S. Teng, T.-Y. Lin, W. A. Clark, R. Vasudevan, and M. Ghaffari, "Riemannian direct trajectory optimization of rigid bodies on matrix lie groups," in *Proceedings of Robotics: Science and Systems*, 2025.
- [25] Q. Liu and D. Wang, "Stein variational gradient descent: A general purpose bayesian inference algorithm," *Advances in neural information processing systems*, vol. 29, 2016.
- [26] C. Liu and J. Zhu, "Riemannian stein variational gradient descent for bayesian inference," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [27] P. Chen and O. Ghaffari, "Projected stein variational gradient descent," *Advances in Neural Information Processing Systems*, vol. 33, pp. 1947–1958, 2020.
- [28] R. Zhang, Q. Liu, and X. Tong, "Sampling in constrained domains with orthogonal-space variational gradient descent," *Advances in Neural Information Processing Systems*, vol. 35, pp. 37108–37120, 2022.
- [29] T. Power and D. Berenson, "Constrained stein variational trajectory optimization," *IEEE Transactions on Robotics*, 2024.
- [30] J. Sola, J. Deray, and D. Atchuthan, "A micro lie theory for state estimation in robotics," *arXiv preprint arXiv:1812.01537*, 2018.
- [31] R. Mahony and J. H. Manton, "The geometry of the newton method on non-compact lie groups," *Journal of Global Optimization*, vol. 23, no. 3, pp. 309–327, 2002.
- [32] T. D. Barfoot, *State estimation for robotics*. Cambridge University Press, 2024.
- [33] N. Boumal, *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023.
- [34] G. Detommaso, T. Cui, Y. Marzouk, A. Spantini, and R. Scheichl, "A stein variational newton method," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [35] D. Wang, Z. Tang, C. Bajaj, and Q. Liu, "Stein variational gradient descent with matrix-valued kernels," *Advances in neural information processing systems*, vol. 32, 2019.
- [36] Q. Liu, J. Lee, and M. Jordan, "A kernelized stein discrepancy for goodness-of-fit tests," in *International conference on machine learning*. PMLR, 2016, pp. 276–284.