

Flow Matching Ergodic Coverage

Max Muchen Sun, Allison Pinosky, Todd Murphey

Center for Robotics and Biosystems, Northwestern University, Evanston, IL 60208

Email: msun@u.northwestern.edu

Project website: <https://murpheylab.github.io/lqr-flow-matching/>

Abstract—Ergodic coverage effectively generates exploratory behaviors for embodied agents by aligning the spatial distribution of the agent’s trajectory with a target distribution, where the difference between these two distributions is measured by the ergodic metric. However, existing ergodic coverage methods are constrained by the limited set of ergodic metrics available for control synthesis, fundamentally limiting their performance. In this work, we propose an alternative approach to ergodic coverage based on flow matching, a technique widely used in generative inference for efficient and scalable sampling. We formally derive the flow matching problem for ergodic coverage and show that it is equivalent to a linear quadratic regulator problem with a closed-form solution. Our formulation enables alternative ergodic metrics from generative inference that overcome the limitations of existing ones. These metrics were previously infeasible for control synthesis but can now be supported with no computational overhead. Specifically, flow matching with the Stein variational gradient flow enables control synthesis directly over the score function of the target distribution, improving robustness to the unnormalized distributions; on the other hand, flow matching with the Sinkhorn divergence flow enables an optimal transport-based ergodic metric, improving coverage performance on non-smooth distributions with irregular supports. We validate the improved performance and competitive computational efficiency of our method through comprehensive numerical benchmarks and across different nonlinear dynamics. We further demonstrate the practicality of our method through a series of drawing and erasing tasks on a Franka robot.

I. INTRODUCTION

Many robotic tasks, such as search and rescue [36], wildlife surveys [44], and precision agriculture [37], require robots to generate motions to explore spatial distributions like thermal signals, population distributions, or soil contaminant levels. Exploration over distributions is challenging, as robots must prioritize high-density regions while also investigating lower-density areas. To address the long-horizon and non-myopic nature of exploration tasks, ergodic coverage generates optimized trajectories whose spatial distribution aligns with the target distribution, providing a formal, information-theoretic notion of coverage that is crucial for exploration tasks. Over the past decade, ergodic coverage has been applied to a wide range of applications, including environmental monitoring [24, 43], tactile sensing [2, 3], learning from demonstrations [23, 45, 48], shared control [12, 13], and embodied active learning [40, 42].

Ergodic coverage solves an optimization problem that minimizes the *ergodic metric*, which measures the difference between the empirical spatial distribution of the trajectory and the target distribution [35], while respecting the agent’s dynamic constraints. Existing methods are often based on stan-

dard trajectory optimization techniques, such as projection-based methods [33] and augmented Lagrange multiplier methods [10]. However, to be compatible with standard trajectory optimization methods, the ergodic metric, even though it evaluates the difference between distributions, must be defined over trajectories. It must also be explicitly formulated and differentiated as runtime cost function. As a result, existing ergodic metrics, such as the most commonly used Fourier ergodic metric [31] and recently developed kernel-based metrics [17, 48], need to be derived specifically to meet this requirement. On the other hand, commonly used statistical measures in machine learning, such as Kullback-Leibler divergence and optimal transport metrics, are infeasible as the ergodic metric for control synthesis using the standard trajectory optimization methods. This limitation fundamentally restricts the performance of existing ergodic coverage methods.

From a different perspective, ergodic coverage can also be viewed as a sampling process: generating ergodic trajectories resembles sampling from the target distribution, where the samples are constrained spatially and temporally under the embodied agent’s dynamics. Recently, continuous flow-based sampling approaches, such as score function-based sampling [29, 47] and continuous normalizing flows [7], have gained significant attention due to their accuracy and computational efficiency, as well as their compatibility with a wide range of statistical discrepancy measures, such as the Kullback-Leibler (KL) divergence [29] and optimal transport metrics [14]. Furthermore, the flow matching method from Lipman et al. [27] enables scalable and efficient optimization of a parameterized flow model based on a reference flow, achieving state-of-the-art performance in generative modeling [41, 49].

In this work, we bring the advantages of flow-based sampling and flow matching to ergodic coverage to overcome the limitations of existing ergodic metrics and trajectory optimization methods. Our main contributions are twofold:

- 1) We formally derive the flow matching problem from Lipman et al. [27] for ergodic coverage and show that it is equivalent to a linear quadratic regulator (LQR) problem. This theoretical result leads to a closed-form iterative optimization algorithm for generating ergodic trajectories with a reference flow from existing flow-based sampling methods (see Fig. 1).
- 2) We derive three reference flows for ergodic coverage: the standard Fourier ergodic metric [31], the Stein variational gradient flow [29], and the Sinkhorn divergence flow [14].




Fig. 1: Similarity between flow-based generative inference and flow matching ergodic coverage.

We show that our method is compatible with the standard ergodic metric and enables alternative ergodic metrics previously infeasible for control synthesis without computational overhead. Compared to state-of-the-art methods, the alternative metrics improve the performance of ergodic coverage over unnormalized distributions and distributions with irregular supports.

Lastly, we demonstrate our method on a Franka robot through a series of drawing and erasing tasks similar to [30, 6]. The code of our implementations and videos of the experiments can be found at our project website: <https://murphey-lab.github.io/lqr-flow-matching/>.

II. PRELIMINARIES AND RELATED WORK

A. Notation

We denote an n -dimensional search space as $\mathcal{X} \subseteq \mathbb{R}^n$. The trajectory of the agent is defined as a mapping $s : [0, T] \mapsto \mathcal{X}$, where T is the time horizon of the trajectory. The trajectory is governed by the agent's dynamics, $\dot{s}(t) = f(s(t), u(t))$, with $u(t) \in \mathcal{U} \subset \mathbb{R}^m$ as the control. The probability density function of the *target distribution* is denoted as $q : \mathcal{X} \mapsto \mathbb{R}_0^+$. We denote the space of all probability distributions over the domain $\mathcal{X}$ as $\mathcal{P}$.

B. Overview of ergodic coverage

The goal of ergodic coverage is to match the spatial distribution of the robot trajectory $s(t)$ with the target distribution $q(x)$. The spatial distribution of a trajectory is defined by extending the definition of *empirical distribution*.

Definition 1 (Trajectory empirical distribution). *Given a trajectory $s(t)$, the associated empirical distribution is:*

$$p_{[s]}(x) = \frac{1}{T} \int_0^T \delta(x - s(t)) dt, \quad (1)$$

where $\delta(x)$ is a Dirac delta function with the following inner product property:

$$\int_{\mathcal{X}} \delta(x - s) f(x) dx = f(s), \quad \forall s \in \mathcal{X}, \quad \forall f \in C(\mathcal{X}), \quad (2)$$

the set $C(\mathcal{X})$ contains all continuous functions over $\mathcal{X}$.

Note that the Dirac delta function $\delta(x)$ is often *heuristically* represented as a point-wise mapping:

$$\delta(x) = \begin{cases} +\infty, & \text{if } x = 0, \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

But the Dirac delta function is formally a generalized function (also called a distribution) that is not defined as a point-wise mapping but based on the inner product property shown in (2) [26].

Definition 2 (Ergodic system [38, 52]). *A dynamic system $s(t)$ is ergodic with respect to $q(x)$ if and only if:*

$$\lim_{T \rightarrow \infty} \int_{\mathcal{X}} g(x) dp_{[s]}(x) = \int_{\mathcal{X}} g(x) dq(x), \quad \forall g \in C(\mathcal{X}). \quad (4)$$

Lemma 1. *A necessary condition for Definition 2 is:*

$$\lim_{T \rightarrow \infty} p_{[s]}(x) = q(x), \quad \forall x \in \mathcal{X}. \quad (5)$$

Definition 3 (Discrepancy measure). A discrepancy measure $D : \mathcal{P} \times \mathcal{P} \mapsto \mathbb{R}_0^+$ is a non-negative functional, such that $D(p(x), q(x)) = 0$ if and only if $p(x) = q(x), \forall x \in \mathcal{X}$.

Definition 4 (Ergodic metric). Given a discrepancy measure D , the ergodic metric $E : \mathcal{S}_T \times \mathcal{P} \mapsto \mathbb{R}_0^+$ is defined as:

$$E(s(t), q(x)) = D(p_{[s]}(x), q(x)). \quad (6)$$

Lemma 2. Based on Lemma 1, a necessary condition for Definition 2 is:

$$\lim_{T \rightarrow \infty} E(s(t), q(x)) = 0. \quad (7)$$

Note that an exact ergodic system exists only at the limit of infinite time horizon. In practice, however, ergodic coverage approximately synthesizes an ergodic system with finite time horizon T through the following optimization problem.

Definition 5 (Ergodic coverage).

$$\begin{aligned} u(t)^* &= \arg \min_{u(t)} E(s(t), q(x)), \\ \text{s.t. } s(t) &= s_0 + \int_0^t f(s(t'), u(t')) dt', \quad t \in [0, T]. \end{aligned} \quad (8)$$

Solving (8) is not easy for two main reasons. First, many commonly used discrepancy measures, such as the KL-divergence, are computationally intractable to evaluate and optimize directly. This issue is further complicated by the fact that $p_{[s]}$ is an empirical distribution—where the Dirac delta function cannot be evaluated numerically—and the target distribution $q(x)$ in practice is often represented not as a standard probability density function, but instead in other forms such as unnormalized utility functions, samples, or discrete density grids. Second, optimization of the ergodic metric must respect the dynamics constraints of the system, and system dynamics can be highly nonlinear in practice, which introduces extra computational challenges.

In the next section, we will review existing ergodic metrics and the corresponding control synthesis methods, focusing on how they address these challenges and their limitations.

C. Review of existing ergodic coverage methods

In [31], the ergodic metric is specified as a Sobolev space distance using Fourier basis functions, which we name the *Fourier ergodic metric*. This formula offers a computationally tractable approximation of the ergodic metric for numerical optimization by truncating a finite number of Fourier basis functions. A closed-form model predictive control formula with an infinitesimally small planning horizon is proposed to approximate a solution to the constrained trajectory optimization problem. However, this method does not generate optimal ergodic coverage trajectories with a fixed time horizon, which is crucial for practical applications [10, 48].

Alternative trajectory optimization methods have been used to optimize the Fourier ergodic metric over long horizons. A projection-based trajectory optimization algorithm is introduced by Miller and Murphey [34], and a method based on augmented Lagrange multiplier, which can incorporate other

constraints such as time-optimality and collision avoidance, is proposed by Dong et al. [10]. In [32], a hybrid system-based model predictive control algorithm is introduced with longer planning horizons, which can be extended for multi-agent systems [1]. In [45], the evaluation of the Fourier ergodic metric is further accelerated through tensor-train decomposition, especially for applications in high-dimensional spaces. Lastly, it is worth noting that ergodic coverage does not have a unique optimal solution, as there exist multiple or even an infinite number of optimal trajectories with a given ergodic metric. To address this unique property of ergodic coverage, a variational inference-based trajectory optimization framework is proposed in Lee et al. [25] using the Stein variational gradient descent algorithm.

On the other hand, several alternative metrics other than the Fourier ergodic metric have been proposed. A kernelized ergodic metric is proposed by [48], which approximates the L^2 distance between distributions under mild regularity conditions and can be optimized through a projection-based algorithm similar to [34]. The kernel ergodic metric has a better scalability compared to the Fourier ergodic metric, and it can be extended to Lie groups. In [18], maximum mean discrepancy (MMD), a statistical measure formulated in the reproducing kernel Hilbert space for two-sample tests, is proposed as the ergodic metric. The MMD metric only requires samples from the target distribution instead of probability density functions and can also be extended to Lie groups. Trajectory optimization for the MMD-based metric can be solved using the same augmented Lagrange multiplier-based method in [10]. Lastly, an alternative metric is introduced in [20] using heat equation, which leads to a better balance between global and local exploration compared to the Fourier ergodic metric, and has been applied to multi-robot aerial survey [24] and tactile sensor-based coverage [6]. However, the trajectory optimization problem is solved through a similar model predictive control formula with an infinitesimally small time horizon as in [31]. Thus, the resulting trajectories do not generate optimal ergodic coverage performance with a fixed time horizon.

D. Flow-based sampling

[Overview] Consider the problem of sampling from a target distribution $q(x)$, while only having access to samples $\{x_i\}$ from an initial distribution $p_0(x)$. One intuition might be to find the optimal transformation $\phi^*(x) : \mathcal{X} \mapsto \mathcal{X}$, such that the transformed samples $\{\phi^*(x_i)\}$ match the statistics of the target distribution $q(x)$.

Furthermore, the transformation $\phi(x)$ can be constructed as a time-dependent transformation $\phi(\tau, x)$ through a time-dependent vector field $g(\tau, x)$ as:

$$\frac{d}{d\tau} \phi(\tau, x) = g(\tau, \phi(\tau, x)), \quad \phi_0(x) = x, \quad \tau \in [0, T], \quad (9)$$

where $\phi(\tau, x)$ is called a *flow* and $g(\tau, x)$ is the *flow vector field*. The flow creates a *probability density path* $p(\tau, x)$

governed by the Fokker–Planck equation [22]:

$$\frac{d}{d\tau}p(\tau, x) = \nabla \cdot (p(\tau, x)g(\tau, x)), \quad p(0, x) = p_0(x) \quad (10)$$

Given the set of samples $\{x'_i\}$ from the initial distribution $p_0(x)$, the flow also creates a path $x_i(\tau)$ for each sample:

$$\frac{d}{d\tau}x_i(\tau) = g(\tau, x_i(\tau)), \quad x_i(\tau) = x'_i, \quad \tau \in [0, T]. \quad (11)$$

Therefore, flow-based sampling can be formulated as finding the optimal flow vector field to transform the initial distribution p_0 toward the target distribution $q(x)$.

Definition 6 (Flow-based sampling).

$$g^*(\tau, x) = \arg \min_{g(\tau, x)} \lim_{\tau \rightarrow T} D(p(\tau), q), \quad (12)$$

where D is a discrepancy measure between distributions.

[Related work] Instead of directly solving (12), most existing flow-based sampling methods construct the flow vector field as a descent direction of the discrepancy measure $D(p, q)$, in which case the transformation of the distribution is equivalent gradient descent in the space of probabilistic measures. The Stein variational gradient descent (SVGD) algorithm [28, 29] constructs the Stein variational gradient flow as the steepest descent direction of the KL-divergence in a reproducing kernel Hilbert space. Another widely used family of flow-based sampling methods is based on Wasserstein gradient flows, which originate from the seminal work of Jordan et al. [22]. Wasserstein gradient flows construct the steepest descent direction using the Wasserstein metric, where the discrepancy measure $D(p, q)$ can be specified as optimal transport (OT) [50] metrics, maximum mean discrepancy (MMD) [5], and more recently the Sinkhorn divergence [14], which interpolates between MMD and OT metrics [11].

E. Flow matching

The flow vector field can also be modeled as a parametric function approximator (e.g., neural networks), which leads to the continuous normalizing flows (CNFs) framework [7]. However, the standard maximum likelihood training of CNFs requires numerical simulation of the flow dynamics, which can be prohibitively expensive. To address this issue, *flow matching* is proposed by Lipman et al. [27] as a computationally efficient optimization paradigm for CNFs.

Given a dataset $\{\bar{x}_i\}$ as the samples from the target distribution $q(x)$, the overall goal of CNFs is to learn to generate novel samples from $q(x)$ by transforming samples $\{x_i\}$ from an initial distribution $p_0(x)$, often specified as a Gaussian distribution. The flow matching framework optimizes a parameterized flow vector field $g(\tau, x; \theta)$ through the following optimization problem (equation 5, [27]).

Definition 7 (Flow matching).

$$\theta^* = \arg \min_{\theta} \mathbb{E}_{\tau, p(\tau, x)} \|h(\tau, x) - g(\tau, x; \theta)\|^2, \quad (13)$$

where $p(\tau, x)$ is a probability density path toward the target distribution $q(x)$ under the reference flow vector field $h(\tau, x)$.

The reference flow can be constructed using methods from flow-based sampling, with examples including the probability flow ODE based on the score function [47] or optimal transport metrics [49].

The flexibility of flow-based sample generation combined with the computation efficiency of flow matching has emerged as a powerful framework for generative inference, with applications in image generation [41, 49], protein structure generation [19, 21], computer vision [53], trajectory prediction [54], and robot policy learning [8].

III. FLOW MATCHING ERGODIC COVERAGE

A. Problem formulation

Given a dynamic system $f(s(t), u(t))$ with initial condition s_0 , we define a *control sequence path* $u(\tau, t)$, in which there are two temporal variables: the *system time* $t \in [0, T]$ is the associated with dynamic system (e.g., the time variable for simulating the system trajectory $s(t)$), and the *flow time* $\tau \in [0, T]$ is the time associated with the evolution of the entire control sequence (e.g., the time variable in flow-based sampling). The flow dynamics of the control sequence path $u(\tau, t)$ at any system time t is defined as:

$$\frac{d}{d\tau}u(\tau, t) = v(\tau, t). \quad (14)$$

We name $v(\tau, t)$ the *control sequence flow* as it describes the evolution of the control at any system time t across the flow time τ . The control sequence path $u(\tau, t)$ generates a *system trajectory path* $s(\tau, t)$, which is defined at any system time t as:

$$s(\tau, t) = s_0 + \int_0^t f(s(\tau, t'), u(\tau, t')) dt', \quad t \in [0, T], \tau \in [0, T]. \quad (15)$$

We further define the *empirical distribution path* $p_{[s]}(\tau, x)$ based on the system trajectory path (15):

$$p_{[s]}(\tau, x) = \frac{1}{T} \int_0^T \delta(x - s(\tau, t)) dt. \quad (16)$$

Similar to (10), the flow dynamics of $p_{[s]}(\tau, x)$ is governed by:

$$\frac{d}{d\tau}p_{[s]}(\tau, x) = \nabla \cdot (p_{[s]}(\tau, x)z(\tau, x)), \quad (17)$$

where $z(\tau, x)$ is the *empirical distribution flow*. It is the flow vector field of the empirical distribution path $p_{[s]}(\tau, x)$, induced by the control sequence flow $v(\tau, t)$ under the dynamics constraints of (15). We now define *flow-based ergodic coverage* based on flow-based sampling (12).

Definition 8 (Flow-based ergodic coverage).

$$v^*(\tau, t) = \arg \min_{v(\tau, t)} \lim_{\tau \rightarrow T} D(p_{[s]}(\tau, x), q(x)), \quad (18)$$

Algorithm 1 Flow matching ergodic coverage

```

1: procedure FLOWMATCHING( $s_0, u_0(t), \eta, \Delta\tau$ )
2:    $\tau \leftarrow 0$ 
3:   while not converged do
4:     Simulate  $s(\tau, t)$  from  $s_0$  and  $u(\tau, t)$  with  $\tau$  fixed
5:     Evaluate the reference flow  $h(\tau, s(\tau, t))$  at  $\tau$ 
6:     Solve (24) for  $v(\tau, t)$ 
7:      $u_{\tau+\Delta\tau}(t) \leftarrow u(\tau, t) + \eta \cdot v(\tau, t)$ 
8:      $\tau \leftarrow \tau + \Delta\tau$ 
9:   end while
10:  return  $s(\tau, t)$  and  $u(\tau, t)$ 
11: end procedure

```

$$s.t. \ s(\tau, t) = s_0 + \int_0^t f(s(\tau, t'), u(\tau, t')) dt', \quad \frac{d}{d\tau} u(\tau, t) = v(\tau, t). \quad (19)$$

Lemma 3. *The solution of flow-based ergodic coverage (18) at $\tau = T$ is the solution of standard ergodic coverage (6).*

However, because of the dynamics constraints (19) between the system trajectory path $s(\tau, t)$ and the control sequence path $u(\tau, t)$, the above optimization problem cannot be solved directly. Instead, we formulate a *flow matching* problem for ergodic coverage based on (13).

Definition 9 (Flow matching ergodic coverage).

$$v^*(\tau, t) = \arg \min_{v(\tau, t)} \mathbb{E}_{p_{[s]}(\tau, x)} \|h(\tau, x) - z(\tau, x)\|^2, \quad \forall \tau \in [0, T], \quad (20)$$

where $h(\tau, x)$ is a reference flow vector field that generates a probability density path toward $q(x)$ from $p_{[s]}(\tau, x)$.

The reference flow can be constructed the same way as in existing flow-based sampling methods and we will specify three reference flows in Section III-C.

The intuition behind the flow matching ergodic coverage formulation (20) is to find the optimal *control sequence flow* $v(\tau, t)$ such that the induced empirical distribution flow $z(\tau, x)$ closely match the reference flow $h(\tau, x)$ at any flow time τ , thus generating a path of the trajectory empirical distribution toward the target distribution. However, to solve the flow matching problem, we still need to derive the relationship between the control sequence flow $v(\tau, t)$ and the induced empirical distribution flow $z(\tau, x)$ under the dynamics constraints (19). In the next section, we will show that $z(\tau, x)$ and $v(\tau, x)$ are governed by a linear dynamic system. Therefore, (20) is equivalent to a linear quadratic regulator (LQR) problem—which we name *linear quadratic flow matching*—that can be solved in closed-form at any flow time τ .

B. Linear quadratic flow matching

We start the derivation of the linear quadratic flow matching problem by substituting the definition of $p_{[s]}(\tau, x)$ into the flow matching objective (20):

$$\mathbb{E}_{p_{[s]}(\tau, x)} \|h(\tau, x) - z(\tau, x)\|^2$$

$$\begin{aligned}
&= \int_{\mathcal{X}} \left(\frac{1}{T} \int_0^T \delta(x - s(\tau, t)) dt \right) \|h(\tau, x) - z(\tau, x)\|^2 dx \\
&= \frac{1}{T} \int_0^T \left(\int_{\mathcal{X}} \delta(x - s(\tau, t)) \|h(\tau, x) - z(\tau, x)\|^2 dx \right) dt \\
&= \frac{1}{T} \int_0^T \|h(\tau, s(t)) - z(\tau, s(t))\|^2 dt \\
&= \frac{1}{T} \int_0^T \|\tilde{h}(\tau, t) - \tilde{z}(\tau, t)\|^2 dt. \quad (21)
\end{aligned}$$

We now derive the system time dynamics of $\tilde{z}(\tau, t)$ under the control sequence flow $v(\tau, t)$ at any flow time τ . As defined in (17), $z(\tau, x)$ is the flow vector field of $p_{[s]}(\tau, x)$, since the probability density path $p_{[s]}(\tau, x)$ is an empirical distribution path (16) over the system trajectory path $s(\tau, t)$ (15), we have the following based on (11):

$$\frac{d}{d\tau} s(\tau, t) = z(\tau, s(\tau, t)) = \tilde{z}(\tau, t). \quad (22)$$

Furthermore, the system trajectory path $s(\tau, t)$ is governed by the dynamics shown in (15) and the control sequence path $u(\tau, t)$ is governed by the ODE shown in (14). Therefore, with a sufficiently small perturbation $\epsilon > 0$ at a fixed τ , we have:

$$\begin{aligned}
&s(\tau, t) + \epsilon \cdot \tilde{z}(\tau, t) \\
&= s_0 + \int_0^t f\left(s(\tau, t') + \epsilon \cdot \tilde{z}(\tau, t'), u(\tau, t') + \epsilon \cdot v(\tau, t')\right) dt',
\end{aligned}$$

from which¹ we have the dynamics of $\tilde{z}(\tau, t)$ as:

$$\begin{aligned}
\tilde{z}(\tau, t) &= \frac{d}{d\epsilon} [s(\tau, t) + \epsilon \cdot \tilde{z}(\tau, t)]_{\epsilon=0} \\
&= \frac{d}{d\epsilon} \left[\int_0^t f\left(s(\tau, t') + \epsilon \cdot \tilde{z}(\tau, t'), u(\tau, t') + \epsilon \cdot v(\tau, t')\right) dt' \right]_{\epsilon=0} \\
&= \int_0^t \frac{d}{d\epsilon} \left[f\left(s(\tau, t') + \epsilon \cdot \tilde{z}(\tau, t'), u(\tau, t') + \epsilon \cdot v(\tau, t')\right) \right]_{\epsilon=0} dt' \\
&= \int_0^t A(\tau, t') \cdot \tilde{z}(\tau, t') + B(\tau, t') \cdot v(\tau, t') dt', \quad (23)
\end{aligned}$$

where

$$\begin{aligned}
A(\tau, t') &= \frac{\partial}{\partial s} f(s(\tau, t'), u(\tau, t')), \\
B(\tau, t') &= \frac{\partial}{\partial u} f(s(\tau, t'), u(\tau, t')).
\end{aligned}$$

From (23), we can see that the system time dynamics of $\tilde{z}(\tau, t)$ and $v(\tau, t)$ are governed by a *time-varying linear system* at a fixed flow time τ , which is the linearization of the system dynamics $f(s(t), u(t))$. We can now define the *linear quadratic flow (LQ) matching* problem based on this result.

Definition 10 (Linear quadratic flow matching). *At any flow time τ , the problem is defined as*

$$v^*(\tau, t) = \arg \min_{v(\tau, t)} \int_0^T \|\tilde{h}(\tau, t) - \tilde{z}(\tau, t)\|_Q^2 + \|v(\tau, t)\|_R^2 dt \quad (24)$$

¹Note that there is no perturbation on the initial state s_0 since it is fixed.




Fig. 2: The bandwidth of the RBF kernel in Stein variational gradient flow affects the resulting ergodic trajectories. Overly small bandwidth could cause the mode collapse issue while overly large bandwidth could also lead to under-exploration.

$$\text{s.t. } \tilde{z}(\tau, t) = \int_0^t A(\tau, t') \cdot \tilde{z}(\tau, t') + B(\tau, t') \cdot v(\tau, t') dt', \quad (25) \quad \text{where}$$

where Q and R are regularization matrices.

The LQ flow matching problem (24) follows the standard linear quadratic regulator formula; thus, it is convex and can be solved in closed-form by solving the continuous-time Riccati equation [16]. To generate an ergodic coverage trajectory, the linear quadratic flow matching problem can be solved at a given flow time τ (starting from $\tau=0$) to generate the optimal control sequence flow $v^*(\tau, t)$, from which we can forward simulate the control sequence path $u(\tau, t)$ using Euler's method:

$$u(\tau + \Delta\tau, t) = u(\tau, t) + \eta \cdot v^*(\tau, t), \quad (26)$$

where $\Delta\tau$ and η are the integration step size for the flow time and the control sequence path, respectively. We summarize the overall process in Algorithm 1. Note that the projection-based trajectory optimization method from [34] also solves an LQR problem in each iteration, but the LQR formulation in [34] is different from ours (24) and must be analytically derived for a specific ergodic metric, while our linear quadratic flow matching formula is compatible with any reference flow without adaptation. Lastly, some flow specifications, such as in continuous normalizing flows, require a finite terminal flow time $\mathcal{T}$, while others, such as Stein variational gradient flows, converge as $\mathcal{T} \rightarrow \infty$. The latter case reduces the dependency on explicit integration of the flow time τ , which improves the numerical stability. In the next section, we will provide three specifications for the reference flow, none of which requires a finite flow time range $\mathcal{T}$.

C. Specifications of reference flows

We first derive the reference flow for the widely used Fourier ergodic metric from Mathew and Mezić [31]. We then specify the reference flow using the Stein variational gradient flow [29] and Sinkhorn divergence flow [11].

[Fourier ergodic metric flow] Without loss of generality, we define the normalized Fourier basis function over a n -dimensional rectangular space $\mathcal{S} = [0, L_1] \times \dots \times [0, L_n]$:

$$f_{\mathbf{k}}(x) = \frac{1}{h_{\mathbf{k}}} \prod_{i=1}^n \cos\left(\frac{k_i \pi}{L_i} x_i\right) \quad (27)$$

$$\begin{aligned} x &= [x_1, x_2, \dots, x_n] \in \mathcal{S} \\ \mathbf{k} &= [k_1, \dots, k_n] \in \mathcal{K} \subset \mathbb{N}^n \\ \mathcal{K} &= [0, 1, \dots, K_1] \times \dots \times [0, 1, \dots, K_n], \end{aligned}$$

and $h_{\mathbf{k}}$ is the normalization term such that the norm of each basis function is 1. The Fourier ergodic metric is defined based on the Sobolev distance between two distributions:

$$D(p, q) = \sum_{\mathbf{k} \in \mathcal{K}} \lambda_{\mathbf{k}} (p_{\mathbf{k}} - q_{\mathbf{k}})^2, \quad (28)$$

$$p_{\mathbf{k}} = \langle p(x), f_{\mathbf{k}}(x) \rangle, \quad q_{\mathbf{k}} = \langle q(x), f_{\mathbf{k}}(x) \rangle, \quad \lambda_{\mathbf{k}} = (1 + \|\mathbf{k}\|)^{-\frac{n+1}{2}}.$$

We can substitute $p(x)$ as the trajectory empirical distribution $p_{[s]}(\tau, x)$, which leads to the following based on the inner product property of the Dirac delta function:

$$p_{\mathbf{k}}(\tau) = \langle p_{[s]}(\tau, x), f_{\mathbf{k}}(x) \rangle = \frac{1}{T} \int_0^T f_{\mathbf{k}}(s(\tau, t)) dt. \quad (29)$$

Following [4], we can derive the reference flow for the Fourier ergodic metric as a Wasserstein gradient flow:

$$\begin{aligned} h(\tau, x) &= \frac{d}{dx} \left[\frac{\partial}{\partial p} D(p_{[s]}(\tau, x), q(x)) \right] \\ &= \frac{d}{dx} \left[\sum_{\mathbf{k} \in \mathcal{K}} 2\lambda_{\mathbf{k}} (p_{\mathbf{k}}(\tau) - q_{\mathbf{k}}) f_{\mathbf{k}}(x) \right] \\ &= \sum_{\mathbf{k} \in \mathcal{K}} 2\lambda_{\mathbf{k}} (p_{\mathbf{k}}(\tau) - q_{\mathbf{k}}) \cdot \frac{d}{dx} f_{\mathbf{k}}(x). \end{aligned} \quad (30)$$

Note that, the reference flow for the Fourier ergodic metric can be evaluated in closed-form.

[Stein variational gradient flow] The Stein variational gradient flows optimizes the KL-divergence between $p_{[s]}(\tau, x)$ and the target distribution $q(x)$:

$$D_{KL}(p_{[s]}(\tau, x), q(x)) = \mathbb{E}_{p_{[s]}(\tau, x)} \left[\log \left(\frac{p_{[s]}(\tau, x)}{q(x)} \right) \right]. \quad (31)$$

It is defined as:

$$h(\tau, x) = \mathbb{E}_{p_{[s]}(\tau, x')} [\mathcal{A}_q k(x, x')], \quad (32)$$

where $k(x, x')$ is a kernel function² and $\mathcal{A}$ is the Stein

²It is often specified as a radial basis function (RBF) in practice.

Entropy Regularization: 0.001
Elapsed Time/Iter: 0.0074 s

Entropy Regularization: 0.01
Elapsed Time/Iter: 0.0063 s

Entropy Regularization: 0.1
Elapsed Time/Iter: 0.0051 s

Entropy Regularization: 1.0
Elapsed Time/Iter: 0.0042 s




Fig. 3: The entropy regularization weight ϵ in Sinkhorn divergence (35) affects the resulting ergodic trajectories. Smaller values make the divergence closer to the Wasserstein distance, improving coverage performance at the cost of slower computation [11].

operator [29] defined as:

$$\mathcal{A}_q k(x, x') = k(x, x') \frac{d}{dx} \log q(x) + \frac{d}{dx} k(x, x'). \quad (33)$$

The Stein reference flow only requires access to the derivative of the log-likelihood function of the target distribution—also called the score function of the target distribution. Furthermore, the expectation term in (32) can be calculated as:

$$\mathbb{E}_{p_{[s]}(\tau, x')} [\mathcal{A}_q k(x, x')] = \frac{1}{T} \int_0^T \mathcal{A}_q k(x, s(\tau, t)) dt. \quad (34)$$

In Liu and Wang [29], it is shown that the Stein variational gradient flow is the steepest descent direction for the KL-divergence between $p_{[s]}(\tau, x)$ and $q(x)$ in a reproducing kernel Hilbert space.

The key advantage of the Stein variational gradient flow is that, it requires evaluation of the score function—the derivative of the log-likelihood of the target distribution. The evaluation of score function does not require the target distribution to have a normalized or even non-negative density function. This property has been shown to significantly improve the numerical stability of generative models in Song et al. [47]. Similarly, this property is also crucial for applications of ergodic coverage as in practice, where the target distributions are often unnormalized utility functions, such as from Gaussian processes regression [3, 25]. On the other hand, similar to Stein variational gradient descent SVGD, the choice of kernel function parameters affects the resulting ergodic trajectories. As shown in Fig. 2, in the case of a radial basis kernel function, an overly small bandwidth could lead to the mode collapse issue commonly seen in SVGD [55, 25] and an overly large bandwidth could result in trajectories that under explore regions with high density. Instead of hand-tuning the kernel parameters, a parameter selection formula is given in Liu and Wang [29] for radial basis kernel functions (See Section 5 of [29]).

Lastly, the use of Stein variational gradient flow in our method is to replace the standard Fourier ergodic metric to optimize a single trajectory. At the same time, Stein variational inference is also applied to ergodic coverage to generate

multiple optimal trajectories under the Fourier ergodic metric in Lee et al. [25].

[Sinkhorn divergence flow] The Sinkhorn divergence is a discrepancy measure based on the entropic regularized optimal transport distance $OT_\omega(p, q)$ between two distributions:

$$OT_\omega(p, q) = \min_{\pi(x, x')} \mathbb{E}_\pi [c(x, x')] + \omega KL(\pi(x, x') \| p(x)q(x')),$$

$$\text{s.t. } \int \pi(x, x') dx' = p(x) \text{ and } \int \pi(x, x') dx = q(x'),$$

where the transport plan $\pi(x, x')$ is a joint probability distribution, ω is a regularization weight term, and $c(x, x')$ is a cost function³. The entropic regularized OT distance suffers from entropic bias—the minima of $OT_\omega(p, q)$ is not 0 with $\omega > 0$. The Sinkhorn divergence [14] removes the entropic bias:

$$D_\omega(p, q) = OT_\omega(p, q) - \frac{1}{2} OT_\omega(p, p) - \frac{1}{2} OT_\omega(q, q). \quad (35)$$

While evaluating the Sinkhorn divergence directly over continuous probability density functions is computationally intractable, it can be evaluated efficiently over samples from the distributions using the Sinkhorn algorithm [9], which is an iterative algorithm based on matrix scaling. Furthermore, since all the calculations of the Sinkhorn algorithm are differentiable, the gradient of the Sinkhorn divergence can be evaluated over a batch of samples using auto-differentiation [14]. For flow matching ergodic coverage, we can discretize the system trajectory path $s(\tau, t)$ at discrete system time steps $\{t_i\}$. At a given flow time τ and a set of samples $\{x'_i\}$ from the target distribution $q(x)$, we can directly evaluate the reference flow $h(\tau, x)$ at all the discrete system time steps in a batch:

$$\{\tilde{h}(\tau, t_i)\}_i = \text{AutoDiff}_{\{s(\tau, t_i)\}_i} \left(D_\omega \left(\{s(\tau, t_i)\}_i, \{x'_j\}_j \right) \right).$$

We refer the readers to [14] for more details on the calculation of the Sinkhorn divergence and the Sinkhorn algorithm. In our experiments, we use the Geomloss package from [11] to efficiently evaluate the Sinkhorn divergence and the associated flows, which supports accelerated evaluations on GPUs.

³It is often specified as the L^p distance between x and x' in practice.




Fig. 4: **Results for Benchmark Q1.** Quantitative (left) and qualitative (right) results for the comparison between our method with existing trajectory optimization methods for the Fourier ergodic metric. Our method consistently reaches the same desired level of ergodicity with less time. The white line in the violin plot is the median of the results, and the black dot in the trajectory plot is the initial position.

As an optimal transport metric, the Wasserstein distance is particularly sensitive to the deformation of the distributions' supports [11], which makes Sinkhorn divergence flow particularly advantageous with non-smooth target distributions with local geometric features, such as distributions with irregular supports. On the other hand, the entropy regularization term ϵ in (35) affects the resulting ergodic trajectories. As shown in Fig. 3, smaller values of entropy regularization make Sinkhorn divergence closer to the Wasserstein distance, improving the coverage performance at the cost of slower computation [11].

IV. NUMERICAL BENCHMARK

A. Overview

We aim to answer the following four questions in our numerical benchmarks:

- Q1: For the Fourier ergodic metric, how does our flow matching framework perform compared to existing trajectory optimization methods in terms of ergodic coverage results and numerical efficiency?
- Q2: By integrating Stein variational gradient flows, does our flow matching framework improve the robustness of ergodic coverage when facing inaccurately normalized target distributions?
- Q3: By integrating Sinkhorn divergence flows, does our flow matching framework improve the ergodic coverage results over non-smooth target distributions?
- Q4: How does our flow matching framework handle different linear or nonlinear dynamics?
- Q5: How does the trajectory horizon affect the computation of each reference flow?

B. Baseline selection and implementation details

For Q1, we compare our method to two Fourier-based ergodic trajectory optimization methods: the projection-based trajectory optimization method from Miller and Murphey [34] and the augmented Lagrange multiplier-based method from [10]. For Q2, we use the Fourier ergodic metric and the kernelized ergodic metric from Sun et al. [48], since both metrics and the Stein variational gradient flow require access to the probability density function of the target distribution. For

the non-smooth distributions in Q3, we use the Fourier ergodic metric and the maximum mean discrepancy (MMD)-based ergodic metric from Hughes et al. [18]; the Fourier ergodic metric can be calculated over discretized target distribution representations (e.g., images), and the MMD-based metric is designed for sample-based distribution representations.

The computation related to the Fourier ergodic metric is implemented in Python using the JAX package for acceleration. We implement the continuous-time Riccati equation solver for the linear quadratic flow matching problem (24) in C++ with Python interface, same as the projection-based trajectory optimization algorithm. We use the JAX-based implementation from Dong et al. [10] for the augmented Lagrange multiplier-based trajectory optimization method. The computation related to Sinkhorn divergence is implemented in Python using the Geomloss package from Genevay et al. [14]. The code of our implementation will be released on the project website: <https://murpheylib.github.io/lqr-flow-matching/>.

C. Experiment results for Q1

[Benchmark Q1 design] We randomly generate a trimodal Gaussian mixture distribution and the initial position in each trial, and benchmark our flow matching framework (FM) alongside projection-based trajectory optimization (Projection) and augmented Lagrange multiplier-based trajectory optimization (Lagrange) over 100 trials. We set a desired value of the Fourier ergodic metric 0.005, and measure the wall clock time required for each method to reach the desired level of ergodicity (in seconds). We use second-order point mass dynamics for all the tests.

[Q1 results] The statistics of the elapsed times for all the methods across the 100 randomized trials are shown in Fig. 4 (left). The results show that our flow matching framework consistently reaches the desired level of ergodicity in less time. In Fig. 4 (right), we also show qualitative comparisons of the ergodic trajectories generated by each method over the same target distribution and under the same initial condition.

D. Experiment results for Q2

We answer Q2 with two numerical benchmarks (Q2.A and Q2.B). Details of benchmark design and results are as follow.




Fig. 5: **Benchmark Q2.A.** Quantitative (left) and qualitative (right) results show that our method is more accurate and consistent under normalization errors compared to SOTA methods. The white line in the violin plot is the median of the results, and the black dot in the trajectory plot is the initial position.




Fig. 6: **Benchmark Q2.B.** Quantitative (left) and qualitative (right) results show that our method is consistently collects more positive life signals compared to SOTA methods with unnormalized target distributions. The white line in the violin plot is the median of the results, and the black dot in the trajectory plot is the initial position.

[Benchmark Q2.A design] We follow the same randomization process in Q1. In addition, we randomly scale the target distribution between 0.1 and 10.0 to introduce normalization errors. We use our method with the Stein variational gradient flow (FM-Stein) and compare it with the Fourier ergodic metric optimized (Fourier) and the kernelized ergodic metric optimized (Kernel), all three of which require access to the probability density function of the target distribution. We optimize the Fourier ergodic metric using the augmented Lagrange multiplier-based method from Dong et al. [10] and the kernelized ergodic metric using the projection-based trajectory optimization method from Sun et al. [48], as these implementations yield the best empirical results and thus better demonstrate the properties of each ergodic metric. Each method is allowed to optimize for up to 0.5s wall clock time in each trial, and we report the Fourier ergodic metric with respect to the correctly normalized target distribution for the converged trajectory. We use second-order point mass dynamics for all the tests.

[Q2.A results] The statistics of the ergodic metric at convergence for each method are shown in Fig. 5 (left). The results show that the ergodic metric value that our method converges to is, on average, one order of magnitude lower than the two baselines, despite the inaccurately normalized target distributions and all within the same amount of wall clock time. Compared to the Fourier metric, our method is also more consistent across the 100 trials. A qualitative comparison

between the trajectories is also shown in Fig. 5 (right). The results show that both baselines fail to sufficiently explore the target distribution under normalization error, while our method successfully generates an accurate coverage trajectory.

[Benchmark Q2.B design] We test the three methods in Benchmark Q2.A in a simulated search and rescue (SAS) benchmark. Given a search space, we assume there are unknown regions within the space that emit positive life signals. A set of prior negative and positive life signals are sparsely sampled across the search space, the target distribution is the likelihood of measuring positive life signal across the search space computed using Gaussian processes (GPs) from the prior measurements, which is an unnormalized density function. The task of the benchmark is to explore the space to collect positive life signals given the GP estimation. We model the robot dynamics as a second-order differential-drive system. We conduct 100 test trials with randomized prior measurement and robot initial state.

[Q2.B results] The statistics of the number of positive life signals collected by each method are shown in Fig. 6 (left), where our method consistently collect more positive life signals compared to the two baseline SOTA methods. Furthermore, the coverage trajectories from the tested methods in a representative trial are shown in Fig. 6 (right). These results show the advantage of the Stein variational gradient flow in practical applications, where unnormalized target distributions negatively affect the performance of existing ergodic coverage



Fig. 7: **Qualitative results for Q3.A.** Our method better captures the non-smooth geometry of the target distributions by leveraging the Sinkhorn divergence flow.




Fig. 8: **Qualitative results for Q3.B.** The target distribution is shown in white color. Compared to MMD and Fourier ergodic metric, Sinkhorn divergence leads to better coverage trajectories on target distributions with irregular, discontinuous supports, which are commonly seen in real-world exploration problems.




Fig. 9: **Quantitative results for Q3.A.** Our method consistently reaches lower coverage errors compared to the baselines. The white line in the violin plot is the median of the results.




Fig. 10: **Quantitative results for Q3.B.** Our method consistently reaches lower coverage errors compared to the baselines. The white line in the violin plot is the median of the results.

methods but does not affect our flow matching method.

E. Experiment results for Q3

We answer Q3 with two numerical benchmarks (Q3.A and Q3.B). Details of benchmark design and results are as follow. **[Benchmark Q3.A design]** We choose 10 open source stock icons as the non-smooth target distributions (shown in Fig. 7) and conduct 10 trials with randomized initial positions on each icon for all the methods tested (100 trials in total).

We use our method with the Sinkhorn divergence flow (FM-Sinkhorn) and compare it to the same Fourier ergodic metric baseline from Benchmark Q3 (Fourier) and the maximum mean discrepancy metric optimized using the same augmented Lagrange multiplier-based method [17] (MMD). The target distribution is represented as discrete grids for Fourier and discrete samples (drawn using rejection sampling) for FM-Sinkhorn and MMD. Since the target distributions are non-smooth uniform distributions, we measure the coverage error




Fig. 11: **Results for Q4.** Our method generates comparable ergodic coverage trajectories across six different dynamics.

Ergodic Coverage with 3D Aircraft Dynamics




Fig. 12: **Results for Q4.** Qualitative results with 3D aircraft dynamics using both the Stein variational gradient flow and Sinkhorn divergence flow.

using the trajectory uniformity metric in Mathew and Mezić [31] (equation 4). We use second-order point mass dynamics for all the tests.

[Q3.A results] In Fig. 7, we qualitatively show trajectories for each method across the 10 tested target distributions. Our method better captures the non-smooth geometry of the target distributions by leveraging the Sinkhorn divergence flow, which is a known advantage of optimal transport metrics [39]. The quantitative results are shown in Fig. 9. On average, our method has a coverage error of 0.001, which is two times lower than MMD (0.0028) and Fourier (0.0023).

[Benchmark Q3.B design] We further test the methods from Benchmark Q3.A using the NASA Mars Water Resource Maps (MWR)⁴ as the target distribution. We model the robot dynamics as a second-order differential-drive system. We conduct 100 test trials with randomized robot initial state and evaluate the coverage error using the same metric from Benchmark Q3.A.

[Q3.B results] In Fig. 8, we qualitatively show the trajectories from a representative trial. Quantitative results are shown in Fig. 10. Different from the target distributions in Benchmark Q3.A, the target distributions from the Mars Water Resource Maps has irregular and discontinuous supports, which are common in real-world exploration tasks and create extra challenges




Fig. 13: **Results for Q5.** Elapsed time of solving the linear quadratic flow matching problem across different trajectory horizons with all three reference flows.

for ergodic coverage. From the results, we can see that ergodic coverage based on MMD is particularly affected by such target distributions, while our method using the Sinkhorn divergence flow consistently produces the best coverage trajectories across the benchmark trials, demonstrating the advantage of optimal transport-based ergodic metric in practical applications.

F. Experiment results for Q4

We specify the target distribution as a trimodal Gaussian mixture model and use our method with the Stein variational gradient flow. We fix the initial position of the robot and test our method with six different dynamics: first and second-order point mass dynamics, first and second-order differential drive dynamics, and first and second-order Dubins car dynamics. The results are qualitatively shown in Fig. 11, which show that our method generates comparable coverage trajectories across different linear and nonlinear dynamics.

Furthermore, we qualitatively test our method for 3D coverage tasks with the 3D aircraft dynamics from Lee et al. [25]. Results from both the Stein variational gradient flow and Sinkhorn divergence flow are shown in Fig. 12 and are included in our project website.

⁴<https://ammos.nasa.gov/marswatermaps/?mission=MWR>




Fig. 14: **Drawing task.** Planned object drawing trajectories (left), an example time sequence (middle), and final robot drawings (right). See videos on project website: <https://murphey-lab.github.io/lqr-flow-matching/>.

G. Experiment results for Q5

We evaluate the elapsed time of solving the linear quadratic flow matching problem (24) across different trajectory time horizons (100 to 1000 time steps) for all three reference flows. From the results shown in Fig. 13, we can see the flow matching formula exhibits a linear time complexity for all three reference flows. Furthermore, flow matching with the Fourier ergodic metric flow has near identical computation time with the Stein variational gradient flow. While flow matching with Sinkhorn divergence flow takes longer time, it is still sufficient for real-time trajectory optimization.

V. HARDWARE DEMONSTRATIONS

We demonstrate the effectiveness of our method on a Franka Emika Panda robot for a series of drawing and erasing tasks. For both tasks, we use the Sinkhorn divergence flow to generate ergodic coverage trajectories over five open source stock icons as the target distributions. For the drawing task (see Fig. 14), the robot end-effector is equipped with a dry-erase marker, which is used to draw the desired trajectories on a whiteboard. For the erasing task (see Fig. 15), we manually fill in regions of the whiteboard corresponding to the desired objects and attach an eraser to the robot’s end-effector. For both hardware tasks, the trajectories are pre-planned in the x-y plane using double-integrator point mass dynamics and are executed open-loop. The robot controller operates at 3 Hz for drawing and 6 Hz for erasing.

In Fig. 14, we qualitatively show that all planned drawing trajectories were feasible for the robot to execute. The Final drawings (right) closely resemble the planned trajectories (left). In Fig. 15, we qualitatively show that the robot successfully erased all filled object drawings. We further conduct closed-loop erasing experiments with the onboard camera of the robot providing image feedback, where new drawings are added during erasing. Videos of both the open-loop and closed-loop hardware demonstrations can be found on the project

website, which illustrate the different paths executed for each object.

VI. LIMITATIONS

The proposed flow matching framework focuses on trajectory optimization, aiming to generate the most ergodic trajectory within a fixed time horizon. In contrast, feedback control-based ergodic coverage methods [20, 45, 6] offer improved computational efficiency and robustness to runtime disturbances, but typically require more time to achieve comparable coverage [48]. The Fourier ergodic metric enables online model predictive control while still guaranteeing the asymptotic convergence [31, 32], such that the robot can rapidly react to unknown perturbations during the task (e.g., noisy odometry readings) while still maintaining good coverage performance. However, it is unclear if a model predictive control variant of the flow matching formula has a similar asymptotic convergence guarantee for the Stein variational gradient flow or the Sinkhorn divergence flow. As a result, all the experiments in this work are based on long-horizon trajectory optimization, which generates a reference trajectory prior to the task, and the control actions for the robot during runtime are generated through a lower-level tracking-based controller. While the flow matching method has sufficient computational efficiency to replan during runtime, the lack of a model predictive control or feedback control formula with asymptotic convergence guarantees remains a limitation of this work.

In addition, there are limitations for the two alternative ergodic metrics introduced in this work. Despite the ability to generate ergodic trajectories over unnormalized distributions using the Stein variational gradient flow, the score function in (32) cannot be computed in closed-form if the target distribution is represented as samples. While there exist techniques such as score matching [46, 51] to estimate the score function from data, the optimization cannot be conducted in real-time,




Fig. 15: **Erasing task.** Initial filled object drawings to be erased (left), an example time sequence (middle), and final result for each erased object (right). See videos on project website: <https://murphey-lab.github.io/lqr-flow-matching/>.

and the estimated score function can be inaccurate in regions with low sample density [46]. Furthermore, while the kernel function improves computational efficiency and flexibility, the kernel function only provides local correlation between the states; thus, it could lead to sub-optimal performance over spatially disconnected target distributions, which is a common issue in Stein variational gradient flows [29]. On the other hand, the Sinkhorn divergence flow is limited by the computational cost associated with evaluating the optimal transport distance—the computation increases linearly with the number of samples from the target distribution [15], which could have a larger impact as the size of search space increases.

Lastly, while formal convergence analyses exist for standard trajectory optimization frameworks in ergodic coverage [33, 10], such analysis is currently lacking for our flow matching-based ergodic coverage framework. As a result, the step size parameters η , $\Delta\tau$, and the number of iterations in Algorithm 1 need be manually tuned in practice.

VII. CONCLUSIONS

In this work, we introduce a novel optimization paradigm for ergodic coverage based on flow matching. Our method enables alternative ergodic metrics previously infeasible for control synthesis to overcome the limitations of existing metrics with no computational overhead. We show that using the Stein variational gradient flow improves the robustness of ergodic coverage on unnormalized target distributions, and using the Sinkhorn divergence flow improves the ergodic coverage performance over non-smooth distributions with irregular supports. We further validate the effectiveness of our method through hardware demonstrations on a Franka robot.

Our work provides a new perspective for bridging statistical inference and the decision-making of embodied agents, where the motion of the agents serves as a dynamically-constrained inferred posterior. While this work focuses on two specific methods to construct the flow, recent advances in flow-based

inference and generative models suggest that the methods in this paper can be paired with any statistical inference framework that admits flow-based inference, providing opportunities for advances in embodied intelligence. The coupling between statistical inference and motion synthesis in our work could further enable the integration of ergodic coverage as a behavior characterization framework, which extends the conventional notion of trajectory tracking to distribution tracking, into learning-based methods for motion synthesis or policy representations.

ACKNOWLEDGMENTS

This work is supported by ARO grant W911NF-19-1-0233 and W911NF-22-1-0286. The views expressed are the authors' and not necessarily those of the funders.

REFERENCES

- [1] Ian Abraham and Todd D. Murphey. Decentralized Ergodic Control: Distribution-Driven Sensing and Exploration for Multiagent Systems. *IEEE Robotics and Automation Letters*, 3(4):2987–2994, 2018.
- [2] Ian Abraham, Ahalya Prabhakar, Mitra J. Z. Hartmann, and Todd D. Murphey. Ergodic Exploration Using Binary Sensing for Nonparametric Shape Estimation. *IEEE Robotics and Automation Letters*, 2(2):827–834, 2017.
- [3] Ian Abraham, Anastasia Mavrommati, and Todd Murphey. Data-Driven Measurement Models for Active Localization in Sparse Environments. In *Robotics: Science and Systems XIV*. 2018.
- [4] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savare. *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*. Springer Science & Business Media, 2008.
- [5] Michael Arbel, Anna Korba, Adil SALIM, and Arthur Gretton. Maximum Mean Discrepancy Gradient Flow. In *Advances in Neural Information Processing Systems*, volume 32. 2019.

- [6] Cem Bilaloglu, Tobias Löw, and Sylvain Calinon. Tactile Ergodic Coverage on Curved Surfaces. *IEEE Transactions on Robotics*, pages 1–15, 2025.
- [7] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural Ordinary Differential Equations. In *Advances in Neural Information Processing Systems*, volume 31. 2018.
- [8] Eugenio Chisari, Nick Heppert, Max Argus, Tim Welschhold, Thomas Brox, and Abhinav Valada. Learning Robotic Manipulation Policies from Point Clouds with Conditional Flow Matching. In *8th Annual Conference on Robot Learning*, 2024.
- [9] Marco Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Advances in Neural Information Processing Systems*, volume 26. 2013.
- [10] Dayi Ethan Dong, Henry Berger, and Ian Abraham. Time-optimal ergodic search: Multiscale coverage in minimum time. *The International Journal of Robotics Research*, 2024.
- [11] Jean Feydy, Thibault Sjourne, Franois-Xavier Vialard, Shun-ichi Amari, Alain Trounev, and Gabriel Peyre. Interpolating between Optimal Transport and MMD using Sinkhorn Divergences. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, pages 2681–2690. 2019.
- [12] Kathleen Fitzsimons and Todd D. Murphey. Ergodic Shared Control: Closing the Loop on pHRI Based on Information Encoded in Motion. *J. Hum.-Robot Interact.*, 11(4):37:1–37:20, 2022.
- [13] Kathleen Fitzsimons, Aleksandra Kalinowska, Julius P Dewald, and Todd D Murphey. Task-based hybrid shared control for training through forceful interaction. *The International Journal of Robotics Research*, 39(9):1138–1154, 2020.
- [14] Aude Genevay, Gabriel Peyre, and Marco Cuturi. Learning Generative Models with Sinkhorn Divergences. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, pages 1608–1617. 2018.
- [15] Aude Genevay, Lnac Chizat, Francis Bach, Marco Cuturi, and Gabriel Peyre. Sample Complexity of Sinkhorn Divergences. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, pages 1574–1583. 2019.
- [16] John Hauser. A Projection Operator Approach to the Optimization of Trajectory Functionals. *IFAC Proceedings Volumes*, 35(1):377–382, 2002.
- [17] Christian Hughes, Houston Warren, Darrick Lee, Fabio Ramos, and Ian Abraham. Ergodic Trajectory Optimization on Generalized Domains Using Maximum Mean Discrepancy, 2024. arXiv:2410.10599 [cs].
- [18] Christian Hughes, Houston Warren, Darrick Lee, Fabio Ramos, and Ian Abraham. Ergodic Trajectory Optimization on Generalized Domains Using Maximum Mean Discrepancy, 2024. arXiv:2410.10599.
- [19] Guillaume Hugut, James Vuckovic, Kilian FATRAS, Eric Thibodeau-Laufer, Pablo Lemos, Riashat Islam, Cheng-Hao Liu, Jarrid Rector-Brooks, Tara Akhound-Sadegh, Michael M. Bronstein, Alexander Tong, and Joey Bose. Sequence-Augmented SE(3)-Flow Matching For Conditional Protein Generation. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [20] Stefan Ivi, Bojan Crnkovi, and Igor Mezi. Ergodicity-Based Cooperative Multiagent Area Coverage via a Potential Field. *IEEE Transactions on Cybernetics*, 47(8): 1983–1993, 2017.
- [21] Bowen Jing, Bonnie Berger, and Tommi Jaakkola. AlphaFold Meets Flow Matching for Generating Protein Ensembles. In *NeurIPS 2023 Generative AI and Biology (GenBio) Workshop*, 2023.
- [22] Richard Jordan, David Kinderlehrer, and Felix Otto. The Variational Formulation of the Fokker–Planck Equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.
- [23] Aleksandra Kalinowska, Ahalya Prabhakar, Kathleen Fitzsimons, and Todd Murphey. Ergodic imitation: Learning from what to do and what not to do. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3648–3654, 2021.
- [24] Luka Lana, Karlo Jakac, and Stefan Ivi. Model predictive altitude and velocity control in ergodic potential field directed multi-UAV search, 2024. arXiv:2401.02899 [cs].
- [25] Darrick Lee, Cameron Lerch, Fabio Ramos, and Ian Abraham. Stein Variational Ergodic Search. In *Robotics: Science and Systems XX*. 2024.
- [26] Sir M. J. Lighthill. *An Introduction to Fourier Analysis and Generalised Functions*. Cambridge University Press, 1958.
- [27] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow Matching for Generative Modeling. In *The Eleventh International Conference on Learning Representations*, 2022.
- [28] Qiang Liu. Stein Variational Gradient Descent as Gradient Flow. In *Advances in Neural Information Processing Systems*, volume 30. 2017.
- [29] Qiang Liu and Dilin Wang. Stein Variational Gradient Descent: A General Purpose Bayesian Inference Algorithm. In *Advances in Neural Information Processing Systems*, volume 29. 2016.
- [30] Tobias Lw, Jrmy Maceiras, and Sylvain Calinon. drozBot: Using Ergodic Control to Draw Portraits. *IEEE Robotics and Automation Letters*, 7(4):11728–11734, 2022.
- [31] George Mathew and Igor Mezi. Metrics for ergodicity and design of ergodic dynamics for multi-agent systems. *Physica D: Nonlinear Phenomena*, 240(4-5):432–442, 2011.
- [32] Anastasia Mavrommati, Emmanouil Tzorakoleftherakis, Ian Abraham, and Todd D. Murphey. Real-Time Area Coverage and Target Localization Using Receding-

- Horizon Ergodic Exploration. *IEEE Transactions on Robotics*, 34(1):62–80, 2018.
- [33] Lauren M. Miller and Todd D. Murphey. Trajectory optimization for continuous ergodic exploration on the motion group $SE(2)$. In *52nd IEEE Conference on Decision and Control*, pages 4517–4522, 2013.
- [34] Lauren M. Miller and Todd D. Murphey. Trajectory optimization for continuous ergodic exploration. In *2013 American Control Conference*, pages 4196–4201, 2013.
- [35] Lauren M. Miller, Yonatan Silverman, Malcolm A. MacIver, and Todd D. Murphey. Ergodic Exploration of Distributed Information. *IEEE Transactions on Robotics*, 32(1):36–52, 2016.
- [36] R.R. Murphy. Human-robot interaction in rescue robotics. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 34(2): 138–153, 2004.
- [37] Joe Norby, Sean Wang, Hairong Wang, Shane Deng, Nick Jones, Akshit Mishra, Catherine Pavlov, Hannah He, Sathya Subramanian, Vivek Thangavelu, Natasha Sihota, Thomas Hoelen, Aaron M. Johnson, and Gregory V. Lowry. Path to autonomous soil sampling and analysis by ground-based robots. *Journal of Environmental Management*, 360:121130, 2024.
- [38] Karl E. Petersen. *Ergodic Theory*. Cambridge University Press, 1989.
- [39] Gabriel Peyré and Marco Cuturi. Computational Optimal Transport: With Applications to Data Science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- [40] Allison Pinosky and Todd D. Murphey. Embodied Active Learning of Generative Sensor-Object Models, 2024. arXiv:2410.11130 [cs].
- [41] Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky T. Q. Chen. Multisample Flow Matching: Straightening Flows with Minibatch Couplings. In *Proceedings of the 40th International Conference on Machine Learning*, pages 28100–28127. 2023.
- [42] Ahalya Prabhakar and Todd Murphey. Mechanical intelligence for learning embodied sensor-object relationships. *Nature Communications*, 13(1):4108, 2022.
- [43] Ahalya Prabhakar, Ian Abraham, Annalisa Taylor, Millicent Schlafly, Katarina Popovic, Giovanni Diniz, Brendan Teich, Borislava Simidchieva, Shane Clark, and Todd Murphey. Ergodic Specifications for Flexible Swarm Control: From User Commands to Persistent Adaptation. In *Robotics: Science and Systems XVI*. 2020.
- [44] Kunal Shah, Grant Ballard, Annie Schmidt, and Mac Schwager. Multidrone aerial surveys of penguin colonies in Antarctica. *Science Robotics*, 5(47):eabc3000, 2020.
- [45] Suhan Shetty, João Silvério, and Sylvain Calinon. Ergodic Exploration Using Tensor Train: Applications in Insertion Tasks. *IEEE Transactions on Robotics*, 38(2): 906–921, 2022.
- [46] Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced Score Matching: A Scalable Approach to Density and Score Estimation. In *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, pages 574–584. 2020.
- [47] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*, 2021.
- [48] Max Muchen Sun, Ayush Gaggarr, Pete Trautman, and Todd Murphey. Fast Ergodic Search With Kernel Functions. *IEEE Transactions on Robotics*, 41:1841–1860, 2025.
- [49] Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, 2023.
- [50] Cédric Villani. *Optimal Transport*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer, Berlin, Heidelberg, 2009.
- [51] Pascal Vincent. A Connection Between Score Matching and Denoising Autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.
- [52] Peter Walters. *An Introduction to Ergodic Theory*. Springer Science & Business Media, 2000.
- [53] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. PointFlow: 3D Point Cloud Generation With Continuous Normalizing Flows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [54] Sean Ye and Matthew C. Gombolay. Efficient Trajectory Forecasting and Generation with Conditional Flow Matching. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2816–2823, 2024.
- [55] Jingwei Zhuo, Chang Liu, Jiaxin Shi, Jun Zhu, Ning Chen, and Bo Zhang. Message Passing Stein Variational Gradient Descent. In *Proceedings of the 35th International Conference on Machine Learning*, pages 6018–6027. 2018.