

# Search at Scale: Improving Numerical Conditioning of Ergodic Coverage Optimization for Multi-Scale Domains

Yanis Lahrach<sup>1\*</sup>, Christian Hughes<sup>2\*</sup>, and Ian Abraham<sup>2</sup>

**Abstract**—Recent methods in ergodic coverage planning have shown promise as tools that can adapt to a wide range of geometric coverage problems with general constraints, but are highly sensitive to the numerical scaling of the problem space. The underlying challenge is that the optimization formulation becomes brittle and numerically unstable with changing scales, especially under potentially nonlinear constraints that impose dynamic restrictions, due to the kernel-based formulation. This paper proposes to address this problem via the development of a scale-agnostic and adaptive ergodic coverage optimization method based on the maximum mean discrepancy metric (MMD). Our approach allows the optimizer to solve for the scale of differential constraints while annealing the hyperparameters to best suit the problem domain and ensure physical consistency. We also derive a variation of the ergodic metric in the log space, providing additional numerical conditioning without loss of performance. We compare our approach with existing coverage planning methods and demonstrate the utility of our approach on a wide range of coverage problems.

## I. INTRODUCTION

Effective coverage planning is at the core of many robotic automation problems such as mapping, exploration, and inspection. However, current methods require hyper-specific tuning to the coverage problem, resulting in a wide range of algorithmic adaptations. For example, coverage planning at the micro-scale level, e.g., in part inspection, requires highly precise control and planning [1], while macro-level coverage planning, e.g., for weather monitoring, requires localized measurements over large spaces while satisfying differential velocity constraints [2].

Ergodic trajectory optimization has emerged as a popular general approach for coverage planning that minimizes the distance between the time-averaged statistics of a trajectory and a target spatial distribution while integrating differential dynamics constraints [3]. More recent advancements have expanded the use-case of ergodic coverage to more complex spaces via kernel-based methods [4]–[6]. However, due to their reliance on kernel functions, these approaches are heavily dependent on the appropriate selection of hyperparam-

<sup>1</sup>Yanis Lahrach is with the ICTEAM Institute and the Ecole Polytechnique at the University of Louvain (UCLouvain), Belgium. This work was carried out in part while Y. Lahrach was a visiting student at Yale University, Department of Mechanical Engineering. [yanis.lahrach@student.uclouvain.be](mailto:yanis.lahrach@student.uclouvain.be)

<sup>2</sup>Christian Hughes and Ian Abraham are with the Department of Mechanical Engineering, Yale University, New Haven, CT 06520, USA. [christian.hughes@yale.edu](mailto:christian.hughes@yale.edu), [ian.abraham@yale.edu](mailto:ian.abraham@yale.edu)

\*Equal contribution.

This material is supported by Office of Naval Research grant N000142512440. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Office of Naval Research.

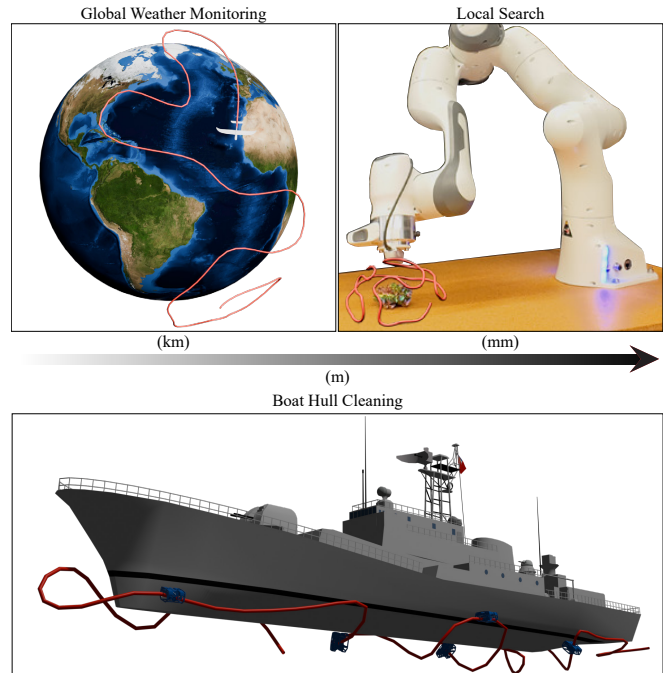


Fig. 1. **Scale-Agnostic Ergodic Search.** Our single, scale-invariant optimizer enables proportional coverage guarantees for tasks spanning vastly different spatial scales. A trajectory is generated for a Franka Emika Panda robot inspecting a small rock (right) and for a drone surveying the Atlantic Ocean for global-scale weather monitoring (left).

eters and domain scale, leading to brittleness and numerical conditioning issues at optimization time. Furthermore, tuning the hyperparameters also alters the physical interpretation of coverage, i.e., a robot’s sensor footprint, which requires careful consideration as inappropriate selection can lead to unstable optimizations and trajectories that do not prioritize coverage sufficiently.

In this paper, we present a reformulation of the ergodic coverage trajectory optimization problem that is invariant to the spatial scale of the search space, decoupling hyperparameters from the physical scale of the task, without compromising performance. Our approach uses the maximum-mean discrepancy (MMD) metric to compute ergodicity (derived in [6]) with various modifications to ensure smooth and scalable optimization. First, we propose an automatic normalization of the search domain that is decoupled from the dynamic constraints. Second, we allow the solver to adapt the time between differential constraints, e.g., between trajectory points, to best suit the coverage problem domain, enabling improved coverage in spatially sparse domains. Third, we propose

annealing techniques for automatic hyperparameter selection to adapt kernel length-scales to the given footprint of a robot without loss of physical meaning. Last, we propose a log variation of the MMD metric that is robust to the problem scale. This new approach to ergodic MMD optimization enables practical ergodic coverage for applications at any scale, from silicon coating in microelectronics to oceanic weather monitoring (as shown in Fig. 1).

*In summary, the primary contribution of this work is an ergodic trajectory optimizer that:*

- improves numerical conditioning in coverage planning across a range of geometries and spatial scales,
- maintains strict adherence to continuous dynamical constraints, and
- facilitates reliable hyperparameter selection that generalizes across diverse scales.

The paper is outlined as follows: Section II reviews related work, Section III covers preliminary background information on kernel-based methods and ergodic coverage planning, Section IV derives the modified ergodic coverage optimization, and Sections V and VI contain the results and conclusions, respectively.

## II. RELATED WORK

### A. Robotic Exploration

Approaches to autonomous exploration generally prioritize one of two main objectives: coverage over the search space and maximal information gain. Coverage-based methods aim to generate trajectories that ensure every reachable region in the search environment is visited. Traditional approaches often solve this as a Traveling Salesperson Problem (TSP) over a discretized grid, where visiting a cell signifies full coverage of that region [7], [8]. While these methods offer theoretical guarantees of complete coverage, they often lack compatibility with dynamic and kinematic constraints used in robotic systems.

Information-driven optimization approaches have emerged as a popular strategy in robotics due to their prioritization of efficient exploration over the search space [9]. These approaches optimize the robot's trajectories according to a spatial information distribution that represents the predicted information gain of visiting different regions [10]. Early work in this field involved integrating probabilistic models into control policies to minimize uncertainty across different regions of the search space [11]. However, while successful for short-term tasks with reliable information, these approaches can struggle to provide coverage of the search space over longer time-horizons and in uncertain scenarios. Recent work in ergodicity-based exploration methods have been shown to compensate for this lack of coverage in information-theoretic settings.

### B. Ergodic Search

Ergodic coverage methods seek to minimize the distance between a robot's time-averaged trajectory and a given spatial information distribution. Ergodicity-based methods have demonstrated the ability to produce trajectories with

a theoretical guarantee of complete coverage over continuous domains as time approaches infinity [3]. Traditional ergodic optimizers minimize a spectral ergodic metric based on Fourier basis functions that has been shown to generate effective coverage trajectories for a wide variety of information-gathering tasks [12]–[14]. However, these standard approaches to ergodicity-based planning are limited to bounded, well-defined Euclidean spaces, which limits their applicability in environments with complex geometries.

Recent advancements have enabled ergodic coverage in environments with complex geometries using a variety of approaches. These methods either derive basis functions from the domain [4], solve partial differential equations (PDEs) for ergodic control policies [15], or numerically approximate basis functions via the Laplace-Beltrami Operator [16]. However, each of these solutions requires specific tuning to the domain: [4] does not apply to unstructured domains (meshes), [15] requires computing a PDE to obtain a local ergodic controller and does not enable trajectory planning, and [16], [17] requires approximation of basis functions that can become computationally challenging.

More recent work has found success in generating ergodic coverage trajectories over arbitrarily-shaped domains by leveraging ideas from optimal transport [5] and two-sample maximum mean discrepancy (MMD) [6]. These methods formulate ergodic coverage as a distribution matching problem between a general form of the time-averaged statistics and domain samples rather than any structured domain definition. However, these approaches are highly sensitive to proper selection of hyperparameters (see Fig.2) as well as domain scales (see Fig.5) that lead to significant performance degradation [18]. In this paper, we propose a series of modifications to the ergodic MMD metric that decouple the dual responsibilities of kernel hyperparameters and improve the condition of the problem at arbitrary spatial scales.

## III. PRELIMINARIES

In this section, we recall (i) time-averaged visitation statistics and ergodicity, (ii) kernel mean embeddings, and (iii) a Maximum Mean Discrepancy (MMD) objective for coverage.

### A. Setup and notation

Let  $\Omega \subseteq \mathbb{R}^d$  denote a measurable search domain with Borel  $\sigma$ -algebra  $\mathcal{B}$ , and let  $\mathcal{P}(\Omega)$  be the set of Borel probability measures on  $\Omega$ . A robot evolves on a state space  $\mathcal{X} \subseteq \mathbb{R}^n$  with a differentiable map  $g : \mathcal{X} \rightarrow \Omega$  (e.g., task-space pose or surface parameterization) and returns the visited location at time  $t$ :

$$\omega_t^x := g(x_t) \in \Omega.$$

Over a finite horizon  $T \in \mathbb{N}$ , we write the discrete trajectory as  $x = \{x_0, \dots, x_{T-1}\}$ <sup>1</sup> that are constrained by the function

$$f(x_{t+1}, x_t) = 0, \quad (1)$$

<sup>1</sup>One can view the trajectory as an ordered list of sensor locations.

where  $f$  is a function that enforces physical consistency, e.g., dynamics and initial/final conditions.

### B. Time-averaged visitation and ergodicity

*Definition 1 (Time-averaged (empirical) visitation):* The visitation statistics of a trajectory  $x$  in  $\Omega$  are

$$\rho_{x,T} = \frac{1}{T} \sum_{t=0}^{T-1} \delta_{\omega_t^x} \in \mathcal{P}(\Omega), \quad \omega_t^x = g(x_t), \quad (2)$$

where  $\delta_\omega$  denotes the Dirac measure at  $\omega$ .

*Definition 2 (Ergodicity):* A (possibly infinite) trajectory  $x$  is *ergodic* with respect to a target utility  $\mu \in \mathcal{P}(\Omega)$  if, for every continuous test function  $\phi \in \mathcal{C}(\Omega)$ ,

$$\lim_{T \rightarrow \infty} \int_{\Omega} \phi(\omega) d\rho_{x,T}(\omega) = \int_{\Omega} \phi(\omega) d\mu(\omega).$$

Equivalently: time averages equal  $\mu$ -weighted space averages.

In finite time, we aim to design trajectories whose empirical statistics  $\rho_{x,T}$  match a given  $\mu$  as closely as possible while satisfying physical constraints. To do so, an ergodic metric  $\mathcal{E} : \mathcal{X}^T \times \mathcal{P}(\Omega) \rightarrow \mathbb{R}$  is optimized for trajectories  $x$  subject to physical constraints (1). More specifically, the ergodic coverage problem is given as

$$\min_x \mathcal{E}(x; \mu) \quad \text{s.t.} \quad \forall t \quad f(x_{t+1}, x_t) = 0, h(x) \leq 0 \quad (3)$$

where  $f$  enforces physical constraints and  $h$  are inequality constraints. In the following section, we derive a variation of the ergodic metric based on maximum mean discrepancy and kernel mean embeddings that enables us to specify coverage problems on a wide variety of geometries and domains.

### C. Kernel mean embeddings

We first define kernel mean embeddings of a distribution. A function  $k : \Omega \times \Omega \rightarrow \mathbb{R}$  is (symmetric) positive-definite if, for any  $m \in \mathbb{N}$ , points  $\{\omega_i\}_{i=1}^m \subset \Omega$  and coefficients  $\{\alpha_i\}_{i=1}^m \subset \mathbb{R}$ ,  $\sum_{i,j=1}^m \alpha_i \alpha_j k(\omega_i, \omega_j) \geq 0$ . To such a kernel is associated a unique reproducing kernel Hilbert space  $(\mathcal{H}_k, \langle \cdot, \cdot \rangle_{\mathcal{H}_k})$  with feature map  $\phi : \Omega \rightarrow \mathcal{H}_k$ ,  $\phi(\omega) = k(\cdot, \omega)$ , satisfying the reproducing property  $f(\omega) = \langle f, \phi(\omega) \rangle_{\mathcal{H}_k}$  for all  $f \in \mathcal{H}_k$ .

*Assumption 1 (Stationary, isotropic kernels):* We use stationary, isotropic kernels of the form

$$k(\omega, \omega') = \kappa \left( \frac{-d(\omega, \omega')}{h} \right), \quad h > 0, \quad (4)$$

with length (bandwidth) parameter  $h$  and distance metric  $d : \Omega \times \Omega \rightarrow \mathbb{R}$ . A canonical choice is the Gaussian/RBF kernel  $k(\omega, \omega') = \exp(-\|\omega - \omega'\|_2^2/h)$ .

*Definition 3 (Kernel mean embedding [19]):* For  $p \in \mathcal{P}(\Omega)$ , the kernel mean embedding is the Bochner integral

$$m_p(\cdot) := \mathbb{E}_{\omega \sim p} [k(\cdot, \omega)] \in \mathcal{H}_k, \quad (5)$$

which represents  $p$  as a smooth ‘‘coverage field’’ in  $\mathcal{H}_k$ .

*Remark 1 (Preview: scale sensitivity):* Because (4) depends on  $\|\omega - \omega'\|/h$ , a uniform change of physical units alters the effective resolution unless  $h$  is co-tuned. Section IV introduces a normalization that decouples similarity from units and lets us reuse a *single numerical*  $h$  across scales.

### D. Maximum Mean Discrepancy (MMD)

*Definition 4 (Maximum Mean Discrepancy):* Given  $p, q \in \mathcal{P}(\Omega)$  with embeddings  $m_p, m_q \in \mathcal{H}_k$ , the squared MMD is

$$\text{MMD}_k^2(p, q) := \|m_p - m_q\|_{\mathcal{H}_k}^2. \quad (6)$$

*Proposition 1 (Expectation form of MMD [6]):* Let  $\mathbb{E}_{p,q}[k] := \iint k(\omega, \omega') dp(\omega) dq(\omega')$ . Then

$$\text{MMD}_k^2(p, q) = \mathbb{E}_{p,p}[k] - 2\mathbb{E}_{p,q}[k] + \mathbb{E}_{q,q}[k]. \quad (7)$$

If  $k$  is characteristic, then  $\text{MMD}_k(p, q) = 0 \iff p = q$ .

*Proposition 2 (Ergodic MMD objective):* Let  $\rho_{x,T}$  be the time-averaged visitation measure from (2) and let  $\mu \in \mathcal{P}(\Omega)$  be a target utility. The *ergodic MMD objective*

$$\mathcal{E}_k(x; \mu) := \text{MMD}_k^2(\mu, \rho_{x,T}) \quad (8)$$

is minimized when the finite-time statistics of  $x$  match  $\mu$  in the Hilbert space distance.

*Proposition 3 (Finite-sample MMD estimator):* Suppose  $\mu$  is represented by samples  $\{\omega_i\}_{i=1}^M \sim \mu$  and write  $\omega_t^x = g(x_t)$ . An estimator of (8) is

$$\begin{aligned} \mathcal{E}_k(x; \mu) = & \frac{1}{T^2} \sum_{t,t'=0}^{T-1} k(\omega_t^x, \omega_{t'}^x) - \frac{2}{TM} \sum_{t=0}^{T-1} \sum_{i=1}^M k(\omega_t^x, \omega_i) \\ & + \frac{1}{M^2} \sum_{i=1}^M \sum_{j=1}^M k(\omega_i, \omega_j), \end{aligned} \quad (9)$$

where the last term is constant w.r.t.  $x$  and can be precomputed.

The finite-sample objective in (9) has three parts: (i) the self-similarity term  $\frac{1}{T^2} \sum_{t,t'} k(\omega_t^x, \omega_{t'}^x)$  penalizes trajectories that revisit the *same* places (it encourages spatial *dispersion*); (ii) the cross term  $-\frac{2}{TM} \sum_t \sum_i k(\omega_t^x, \omega_i)$  pulls the trajectory toward high-utility samples (it drives *coverage*); (iii) the last term depends only on  $\mu$  and is constant w.r.t.  $x$ . The kernel bandwidth  $h$  sets the robot’s effective visitation footprint: large  $h$  rewards broad, coarse coverage; small  $h$  rewards fine, local coverage.

We optimize (9) subject to dynamics and path constraints as described in (3). For clarity we first write the gradient entirely in the search space  $\Omega$ . If  $k \in C^1$ , the map  $x \mapsto \mathcal{E}_k(x; \mu)$  is differentiable via the chain rule. For the RBF kernel  $k(\omega, \omega') = \exp(-\|\omega - \omega'\|^2/(2h))$ , the derivative of  $\mathcal{E}_k$  with respect to  $x_t$  is given as

$$\frac{\partial \mathcal{E}_k}{\partial x_t} = \frac{\partial g}{\partial x_t}^\top \frac{\partial \mathcal{E}_k}{\partial \omega_t^x}. \quad (10)$$

Expanding out the latter term gives the following

$$\begin{aligned} \frac{\partial \mathcal{E}_k}{\partial \omega_t^x} = & \frac{2}{T^2} \sum_{t'} \frac{\omega_t^x - \omega_{t'}^x}{h} k(\omega_t^x, \omega_{t'}^x) \\ & - \frac{2}{TM} \sum_i \frac{\omega_t^x - \omega_i}{h} k(\omega_t^x, \omega_i). \end{aligned} \quad (11)$$

Because the kernel is written in physical distances, a unit scaling  $\omega \mapsto e\omega$  changes the factors in (11): the displacement  $(\omega_s - \omega_t^2)$  scales linearly by  $e$ , while  $k(eu, ev) =$

$\exp(-e^2\|u - v\|^2/h)$  changes exponentially. Consequently the gradient magnitude behaves like  $\sim (e/h) \exp(-e^2 \cdot /h)$ : for large  $e$  (fixed  $h$ ) it decays toward zero, and for small  $e$  it grows—thus the optimization varies with scale.

#### IV. ERGODIC COVERAGE AT SCALE

We improve the numerical conditioning of kernel-based ergodic optimization without altering the physical meaning of constraints. The key idea is to (i) evaluate kernel similarity in a *dimensionless* domain, (ii) adapt differential dynamics and path constraints via optimized time, and (iii) derive a log-sum-exp (LSE) surrogate that yields smooth, size-robust gradients. A practical mapping from a desired *physical* footprint to a normalized bandwidth is given later in Sec. IV-B.

##### A. Scale Normalization and Physical Consistency

Let  $\Omega \subset \mathbb{R}^d$  be the search domain and  $\{\omega_i\}_{i=1}^M \subset \Omega$  target samples with nonnegative weights  $\{\pi_i\}_{i=1}^M$  (normalized so  $\sum_i \pi_i = 1$ ). We define the environment extent

$$e = \max_{\ell \in \{1, \dots, d\}} \left( \max_i [\omega_i]_\ell - \min_i [\omega_i]_\ell \right) > 0,$$

and normalized coordinates  $\hat{\omega} = \omega/e \in \hat{\Omega}$ . For a trajectory  $x = \{x_t\}_{t=0}^{T-1}$  in physical units, we have  $\hat{g}(x_t) = \hat{\omega}_t^x = \omega_t^x/e$ .

We compute kernel similarity on  $\hat{\Omega}$  with a bandwidth  $h > 0$ ,

$$k_h(\hat{u}, \hat{v}) = \exp(-\|\hat{u} - \hat{v}\|_2^2/h). \quad (12)$$

In physical coordinates, this is  $k_h(u, v) = \exp(-\|u - v\|_2^2/(he^2))$ . Because distances in  $\hat{\Omega}$  are  $\mathcal{O}(1)$ , the *same numerical*  $h$  can be re-used across very different physical scales. When a *fixed physical footprint* is required, we map it to a normalized bandwidth and anneal toward the physically consistent footprint.

##### B. Bandwidth Annealing and Physical Footprint

Keeping the *normalized* bandwidth  $h$  fixed as the environment scale  $e$  changes is equivalent to changing a *physical* bandwidth to

$$h_{\text{phys}} = he^2,$$

so the effective physical footprint of visitation grows/shrinks with  $e$ . In other words, a constant  $h$  on the normalized domain does *not* preserve the robot's physical sensing/actuation footprint.

To preserve physical meaning, we choose a target physical bandwidth  $h_{\text{phys}}^*$  (fixed across scales) and set

$$h_{\text{norm}}^* = \frac{h_{\text{phys}}^*}{e^2}.$$

If we start optimization directly with a small  $h_{\text{norm}}^*$  on a large domain, then  $k_h(u, v) \approx 0$  for most pairs and the terms are dominated by self-similarity; gradients are nearly zero and the solver stalls. We therefore use a *continuation* (annealing) by beginning with a large normalized bandwidth  $h_0$  that produces wide, informative kernels, and decrease it geometrically to  $h_{\text{norm}}^*$ .

*Annealing schedule:* In practice, we define an annealing schedule that the optimization follows to facilitate continual solver cost reduction. Let  $e$  be the environment extent,  $h_{\text{norm}}^* = h_{\text{phys}}^*/e^2$  the target normalized bandwidth, and choose an large initial  $h_0$  start, e.g.,  $h_0 = 0.05$ . Over  $K$  outer iterations, we use a geometric series

$$h_k = h_0 \left( \frac{h_{\text{norm}}^*}{h_0} \right)^{\frac{k}{K-1}}, \quad k = 0, \dots, K-1, \quad (13)$$

optimizing the same loss at each  $h_k$  and warm-starting from the previous solution. The overall result is that the solver makes meaningful updates to the solution and places trajectory visitation points in key regions of interest while eventually converging into a length scale that is physically consistent to the robot. Next, we propose a variation to the differential dynamic constraints to adapt according to the solver.

##### C. Adaptive Time Differential Constraints

To ensure that differential constraints have the ability to scale to the domain size we present a time-adaptive variation of the optimization (3). Specifically, we add in a time-scaling component  $\Delta_t$  to the physical constraints  $f$  that governs the behavior between trajectory points  $x_{t+1}, x_t$ . This allows us to write the differential physics constraint as

$$f(x_{t+1}, x_t, \Delta_t) = 0 \quad (14)$$

where  $\Delta_t$  can be optimized. For example, given the state vector  $x_t = [p_t^\top, v_t^\top]^\top$  where  $p_t$  is the position and  $v_t$  is a bounded velocity, the constraint (14) is written with Euler integration as

$$f(x_{t+1}, x_t, \Delta_t) = p_{t+1} - (p_t + \Delta_t v_t) = 0.$$

Since  $\Delta_t > 0$ , we can transform the set of  $\Delta_t$  as a trajectory in the log-space  $\log(\Delta) = \{\log \Delta_t\}_{t=0}^{T-1}$  leading to the modified optimization

$$\min_{x, \log(\Delta)} \mathcal{E}(x; \mu) \quad \text{s.t.} \quad \forall t \ f(x_{t+1}, x_t, \Delta_t) = 0, h(x) \leq 0. \quad (15)$$

Optimizing over  $\log \Delta_t$  has the added advantage of being numerically smooth without the need for additional constraints and keeping differential physics constraints consistent.

In practice, an adaptive  $\Delta_t$  has several additional advantages that expands the use ergodic coverage planning more challenging problems. For instance, (i) *Resolution where it matters*: non-uniform  $\Delta_t$  increases waypoint density in high-utility/high-curvature zones and reduces aliasing of small structures. (ii) *Constraint-compatible*: trajectory velocity bounds remain enforced via  $\|x_{k+1} - x_k\| \leq v_{\text{max}} \Delta_t$ .<sup>2</sup> (iii) *Scale-robust*: combined with the normalized domain, the spatial sampling adapts to spatial sparsity of regions of exploration across orders of magnitude without retuning.

Our final modification to the ergodic coverage optimization is in the ergodic MMD objective itself which alleviates gradient scaling under optimization.

<sup>2</sup>More complex integration schemes can be used without loss of generality.

#### D. The Log-Surrogate Ergodic MMD

In this section, we define a log variation of the MMD objective that provides a smooth surrogate function.

*Proposition 4 (Log-surrogate EMMD):* Let  $\hat{\Omega}$  be the problem domain with  $\mu \in \mathcal{P}(\hat{\Omega})$  and  $\rho_{x,T} \in \mathcal{P}(\hat{\Omega})$  be the target distribution and time-average trajectory statistics respectively. Then, given a locally optimal trajectory  $x^*$ , the following holds

$$\begin{aligned} \tilde{\mathcal{E}}_h(x; \mu) &= \log \sum_{t,t'} k_h(\hat{\omega}_t^x, \hat{\omega}_{t'}^x) - 2 \log \sum_{t,i} k_h(\hat{\omega}_t^x, \hat{\omega}_i) \\ &\leq \mathcal{E}_h(x; \mu). \end{aligned} \quad (16)$$

*Proof:* Taking the log of the ergodic MMD objective in (9) and using concavity rules for log we get

$$\begin{aligned} \log(\mathcal{E}_h(x; \mu)) &\leq \mathcal{E}_h(x; \mu) \\ &\leq \log \sum_{t,t'} k_h(\hat{\omega}_t^x, \hat{\omega}_{t'}^x) - 2 \log \sum_{t,i} k_h(\hat{\omega}_t^x, \hat{\omega}_i) \\ &= \tilde{\mathcal{E}}_h(x; \mu) \end{aligned}$$

where we remove the third log term as it is a constant and the normalizing factors of  $T, M$ . ■

The surrogate MMD objective has a few properties that makes it advantageous over the MMD objectives. First, the expression (16) is numerically better behaved as it is equivalent to the *log-sum-exp* (LSE) given an RBF kernel. As a result (via logs of sums), produces globally normalized “soft-attention” weights, and yields bounded, size-robust gradients on the normalized domain  $\hat{\Omega}$ . Additionally, the log-sum-exp computation prevents over and underflows for especially large numbers that further improves the conditioning of ergodic trajectory optimization.

*Lemma 1 (Nonnegativity and shared global minima):* ] Let  $m_\rho, m_\mu \in \mathcal{H}_k$  be the kernel mean embeddings of  $\rho_{x,T}$  and  $\mu$ . Then

$$A = \|m_\rho\|_{\mathcal{H}_k}^2, \quad B = \langle m_\rho, m_\mu \rangle_{\mathcal{H}_k}, \quad C = \|m_\mu\|_{\mathcal{H}_k}^2,$$

and by Cauchy–Schwarz  $B^2 \leq AC$ , hence  $\tilde{\mathcal{E}}_h(x; \mu) = \log(AC/B^2) \geq 0$ . Moreover,  $\tilde{\mathcal{E}}_h(x; \mu) = 0$  iff  $m_\rho$  and  $m_\mu$  are collinear; for *probability* measures and a characteristic kernel, this holds iff  $\rho_{x,T} = \mu$ . Therefore  $\tilde{\mathcal{E}}_h$  and  $\text{MMD}_k^2$  share the same global minimum (value 0) and the same set of global minimizers.

*Proof:* The identities for  $A, B, C$  follow from standard algebra. Cauchy–Schwarz gives  $B^2 \leq \|m_\rho\|^2 \|m_\mu\|^2 = AC$ , hence  $\tilde{\mathcal{E}}_h = \log(AC/B^2) \geq 0$ . Equality in C–S holds iff  $m_\rho = \lambda m_\mu$  for some  $\lambda \in \mathbb{R}$ . Since the embedding is linear in the measure and both  $\rho_{x,T}$  and  $\mu$  have unit mass,  $\lambda = 1$  and  $m_\rho = m_\mu$  iff  $\rho_{x,T} = \mu$  for a characteristic kernel. ■

*Remark 2:* While  $\tilde{\mathcal{E}}_h$  and  $\text{MMD}_k^2$  have the same global minimum and zero set, they are not generally order-equivalent away from the minimum (logs act termwise). We use  $\tilde{\mathcal{E}}_h$  as a numerically well-conditioned *surrogate* with identical optimal solutions when perfect matching is feasible.

Interestingly, the log-surrogate MMD objective acts like a soft-attention gradient representation. Let us define the




Fig. 2. **Sensitivity of Kernel-Based Methods to Length Scale Selection.** Trajectories are generated for identical 2D environments at spatial scale = 1 with varying length scales ( $h$ ). Length scale selection can lead to underfitting (if too high) or overfitting (if too low) to the utility samples, resulting in degraded performance or complete optimization failure.

positive sums

$$\begin{aligned} F_{xx} &:= \sum_{r,s} k_h(\hat{\omega}_r^x, \hat{\omega}_s^x), & F_{x\mu} &:= \sum_{r,i} k_h(\hat{\omega}_r^x, \hat{\omega}_i), \\ F_{\mu\mu} &:= \sum_{i,j} k_h(\hat{\omega}_i, \hat{\omega}_j) \end{aligned}$$

and the normalized weights

$$\alpha_{ts} := \frac{k_h(\hat{\omega}_t^x, \hat{\omega}_s^x)}{F_{xx}}, \quad \beta_{ti} := \frac{k_h(\hat{\omega}_t^x, \hat{\omega}_i)}{F_{x\mu}}.$$

Note that  $\alpha_{ts}, \beta_{ti} \in (0, 1)$  and  $\sum_s \alpha_{ts} \leq 1$ ,  $\sum_t \alpha_{ts} \leq 1$ ,  $\sum_i \beta_{ti} \leq 1$  by global normalization.

*Lemma 2 (Gradient for general metric):* Assume  $d$  is differentiable and write  $\nabla_{\hat{u}} d(\hat{u}, \hat{v})$  for the gradient in the first argument. Then

$$\begin{aligned} \partial_{\hat{\omega}_s^x} \tilde{\mathcal{E}}_h &= -\frac{1}{h} \sum_t (\alpha_{st} + \alpha_{ts}) \nabla_{\hat{u}} d(\hat{\omega}_s^x, \hat{\omega}_t^x) \\ &\quad + \frac{2}{h} \sum_i \beta_{si} \nabla_{\hat{u}} d(\hat{\omega}_s^x, \hat{\omega}_i). \end{aligned} \quad (17)$$

Gradients w.r.t. the decision variables follow by the chain rule,  $\partial_{x_s} \tilde{\mathcal{E}}_h = \left(\frac{\partial \hat{g}}{\partial x}(x_s)\right)^\top \partial_{\hat{\omega}_s^x} \tilde{\mathcal{E}}_h$ .

*Proof:* Since  $k_h(\hat{u}, \hat{v}) = \exp(-d(\hat{u}, \hat{v})/h)$ ,  $\nabla_{\hat{u}} k_h(\hat{u}, \hat{v}) = -(1/h) \nabla_{\hat{u}} d(\hat{u}, \hat{v}) k_h(\hat{u}, \hat{v})$ . Differentiating  $F_{xx}$  w.r.t.  $\hat{\omega}_s^x$  and accounting for occurrences as first and second argument while dividing by  $F_{xx}$  to obtain  $\partial \log F_{xx}$ . Similarly  $\partial \log F_{x\mu}$  follows by the chain rule. Combining all terms yields  $\partial \tilde{\mathcal{E}}_h = \partial \log F_{xx} - 2 \partial \log F_{x\mu}$ . ■

*Corollary 1 (Euclidean specialization):* If  $d(\hat{u}, \hat{v}) = \|\hat{u} - \hat{v}\|_2^2$ , then  $\nabla_{\hat{u}} d(\hat{u}, \hat{v}) = 2(\hat{u} - \hat{v})$  and

$$\partial_{\hat{\omega}_s^x} \tilde{\mathcal{E}}_h = -\frac{2}{h} \sum_t (\alpha_{st} + \alpha_{ts}) (\hat{\omega}_s^x - \hat{\omega}_t^x) + \frac{4}{h} \sum_i \beta_{si} (\hat{\omega}_s^x - \hat{\omega}_i). \quad (18)$$

We can show that with the log-surrogate MMD function, we can obtain bounded gradients that are independent of the number of trajectory knot points  $T$  and the number domain samples  $M$ .

*Lemma 3 (Bounded, size-robust gradients):* On a bounded normalized domain  $\hat{\Omega}$  with diameter  $D = \text{diam}(\hat{\Omega})$ ,

$$\|\partial_{\hat{\omega}_s} \log F_{x\mu}\| \leq \frac{2D}{h}, \quad \|\partial_{\hat{\omega}_s} \log F_{xx}\| \leq \frac{4D}{h},$$

hence  $\|\partial_{\hat{\omega}_s} \tilde{\mathcal{E}}_h\| \leq \frac{6D}{h}$ . These bounds are independent of the number of knots  $T$  and samples  $M$ .

*Proof:* Use (17) (or (18)) together with  $\sum_t (\alpha_{st} + \alpha_{ts}) \leq 2$ ,  $\sum_i \beta_{si} \leq 1$ , and  $\|\hat{\omega} - \hat{\omega}'\| \leq D$ . ■

Last, we show that the log-surrogate ergodic MMD objective is scale-invariant.

*Proposition 5 (Scale invariance on  $\hat{\Omega}$ ):* Let  $\omega = e\hat{\omega}$  be a uniform rescaling of the physical domain and  $\hat{g}(x) = g(x)/e$ . Then  $\tilde{\mathcal{E}}_h$  and its gradients are invariant under rescaling.

*Proof:* On  $\hat{\Omega}$ , pairwise distances in  $k_h$  depend only on  $\hat{\omega} = \omega/e$ , so a uniform scaling  $\omega \mapsto s\omega$  leaves  $\hat{\omega}$  unchanged. The sums  $F_{xx}$  and  $F_{x\mu}$ , hence  $\tilde{\mathcal{E}}_h = \log F_{xx} - 2\log F_{x\mu}$  and the softmax weights  $\alpha, \beta$ , are unchanged. By the chain rule,  $\partial_{x_t} \tilde{\mathcal{E}}_h = (\partial \hat{g} / \partial x_t)^\top \partial_{\hat{\omega}_t} \tilde{\mathcal{E}}_h$ , and  $\partial \hat{g} / \partial x_t = (1/e) \partial g / \partial x_t$  cancels the physical scaling inside  $\partial_{\hat{\omega}_t}$ , yielding identical gradients (cf. Lemma 3). ■

With this last proof, we formulate the final ergodic trajectory optimization as

$$\min_{x, \log(\Delta)} \tilde{\mathcal{E}}_h(x; \mu) \quad \text{s.t. } \forall t f(x_{t+1}, x_t, \Delta_t) = 0, h(x) \leq 0. \quad (19)$$

where  $h$  is annealed given (13) per iteration. In implementation, we solve the above optimization using an Augmented Lagrangian method [20], [21].

## V. RESULTS

### A. Environmental Scale in Ergodic Kernel-Based Methods

We first evaluate the dependence of existing kernel-based methods on the length parameter by conducting an ablation study (shown in Fig.2) on the variation of length parameter selection for the traditional ergodic MMD solver and the fast ergodic kernel solver provided in [4]. Here, we find that variations in the kernel length parameter can cause drastic changes in how the optimization criteria are defined, such as the visitation footprint in MMD. Given that the kernel length parameter's primary function is to denote the distance between points, different scales will require different selections of  $h$  for proper representation of the spatial data, but modification of the length scale to adequately represent the spatial scale of the environment comes at the cost of altering the ergodic trajectory's behavior as well. These results show that existing methods' reliance on spatial scale leads to a tedious trial-and-error process for appropriate hyperparameter selection for each new environment explored, which nullifies the geometrical generalization that these methods offer.

Likewise, we show that our method is able to adapt 1) its footprint on a large-scale search problem in Fig. 3 while preserving a more restrictive sensor footprint determined by the length-scale  $h$ . In addition, having  $\Delta_t$  be a free variable allows us to plan coverage trajectories over highly sparse domains like a collection of islands in Fig. 4. Here, we see that with a fixed  $\Delta_t$ , trajectory placement can become




Fig. 3. **Bandwidth annealing preserves a physical footprint and stabilizes optimization at scale.** Left: starting with a large normalized bandwidth  $h_0 = 0.05$  yields broad similarity and well-conditioned gradients on a large domain. Right: after continuation to the *target physical* footprint of 1 m (which corresponds to  $h^* = 1.5 \times 10^{-6} = h_{\text{phys}}^*/e^2$ ), the optimized trajectory (red,  $T = 500$  knots) achieves fine, proportional coverage of the bunny while constraints are enforced in physical units; domain extent  $e = 1000$  m.

challenging over the spatially sparse map which an adaptive  $\Delta_t$  allows differential constraints to adapt while satisfying coverage performance.

### B. Quantitative Analysis and Spatial Scalability

Next, we compare the performance of our scale-agnostic ergodic trajectory optimization approach against several state-of-the-art coverage planners over a 3D bunny mesh, as shown in Fig.5. Coverage trajectories were generated using the traditional ergodic MMD, traveling salesperson, flow-matching ergodic [5], and scale-agnostic EMMD planners across three different spatial scales (1, 100, and 10,000), each with identical maximum trajectory length and initial position constraints for the given scale.

Our scale-agnostic formulation consistently achieved high coverage and generates smooth trajectories across all scales. In addition, the performance of our approach at each scale was nearly identical to that of the canonical ergodic MMD solver at scale = 1, which verifies the scale-invariance property. In contrast, as the spatial scale increased, the canonical solver's performance degraded significantly at scale = 100 and failed to converge at scale = 10,000.

The flow-matching ergodic solver achieved high coverage within the normalized domain, but similar to the traditional MMD approach, its performance degraded with increasing scale. This is because the flow-matching ergodic solver is based on Sinkhorn divergence, which characteristically interpolates between MMD and optimal transport metrics [5]. Similar to the traditional MMD metric, the Sinkhorn divergence metric naturally struggles with optimization at larger spatial scales. In addition, the Sinkhorn divergence at vastly different spatial scales incurs a computational cost that scales linearly with the scale of the search space. In particular, the Sinkhorn divergence relies more on its leverage of the Wasserstein distance, which assists with coverage performance at small variations in scale at the cost of slower computational speeds.

The numerical instability of Sinkhorn divergence at vastly larger scales, such as 10,000, yielded few tractable gradi-




Fig. 4. **Adaptive vs. fixed time steps on a sparse, multi-scale target.** Philippines silhouette (black) with identical trajectory length,  $T = 50$ . *Left* (Optimized  $\Delta t$ ): time is concentrated over dense archipelagos; the trajectory refines coverage on many small islands and traverses open water quickly. *Right* (Fixed  $\Delta t$ ): uniform spacing oversamples empty water and undersamples intricate regions, yielding poorer effective resolution on the islands.

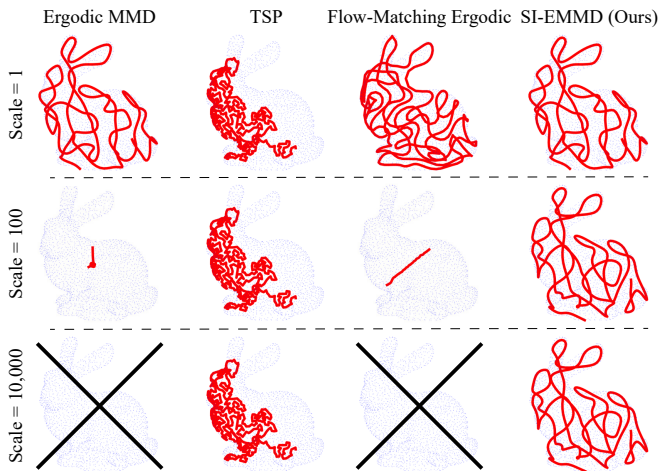


Fig. 5. **Comparison of Scale-Invariant EMD to State-of-the-Art Methods.** Quantitative comparison of our scale-agnostic method against state-of-the-art coverage planners on the Stanford bunny mesh (2,503 pts). Trajectories are constrained to an identical maximal length that scales proportionally with environment size. Where existing ergodic exploration methods fail to adapt to vastly different spatial scales, and TSP solvers struggle to generate trajectories that prioritize coverage in length-constrained settings, our method consistently achieves high coverage and maintains numerical stability across scales, while maintaining similar optimization times and coverage quality to ergodic MMD trajectories at scale = 1.

ents with which to follow, which decreased computational requirements at the cost of a viable coverage trajectory. However, at scales where the Sinkhorn divergence metric did not fail, its computational variance to spatial scale and the number of computed samples resulted in infeasibly-long computation times which we show in Table I.

Despite consistent performance across varying scales, the traveling salesperson solver demonstrated suboptimal coverage over the mesh due to its inability to globally select which point to visit. Although these solvers can guarantee full coverage of all reachable space given a sufficiently long trajectory, they do not intrinsically prioritize broad visitation within a constrained trajectory length.

For a comparative analysis of unconstrained coverage, we compare TSP’s unconstrained coverage to our proposed

TABLE I  
QUANTITATIVE ANALYSIS: COMPARISON AGAINST EXISTING METHODS IN 3D

Metric	Scale 1	Scale 100	Scale 10,000
<b>EMMD</b>			
Coverage (%)	91.69	0.52	0.00
Comp. Time (sec)	$39.66 \pm 0.32$	$39.75 \pm 0.24$	$31.84 \pm 0.24$
<b>TSP</b>			
Coverage (%)	42.31	42.39	42.39
Comp. Time (sec)	$0.015 \pm 0.003$	$0.014 \pm 0.001$	$0.013 \pm 0.001$
<b>Flow-Matching</b>			
Coverage (%)	98.60	0.88	0.00
Comp. Time (sec)	$1249.45 \pm 228.5$	$60.60 \pm 6.57$	$28.95 \pm 0.30$
<b>SI-EMMD (Ours)</b>			
Coverage (%)	91.69	91.09	91.09
Comp. Time (sec)	$40.49 \pm 0.31$	$40.43 \pm 0.25$	$39.47 \pm 0.41$

method by uniformly sampling the bunny mesh to the desired length of trajectory control knots (500). As shown in Fig.6, the unconstrained TSP path achieved near-perfect coverage, but at the expense of a trajectory more than twice as long as any length-constrained version. In addition, this trajectory proved dynamically infeasible due to its jagged and abrupt movements between adjacent points. In contrast, our scale-invariant EMD solver achieves comparable coverage while strictly adhering to smoothness and velocity constraints, thus simultaneously ensuring trajectory feasibility and preserving the proportional coverage guarantees of ergodic methods.

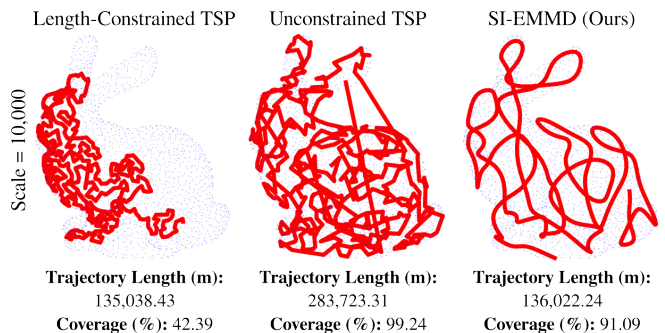


Fig. 6. **Constraint Violations for Similar Coverage with TSP.** Trajectory optimization of three different methods on the Stanford bunny model. The standard Traveling Salesperson solver (TSP) with an imposed trajectory distance constraint achieves only 42.39% coverage due to TSP’s general incompatibility with imposed constraints. When the length constraint is removed and the bunny mesh is uniformly sampled to the number of control knots (trajectory points), TSP achieves 99.24% coverage at the cost of more than double the trajectory length. The proposed scale-invariant metric achieves 91.09% coverage at significantly less trajectory length while accommodating physical constraints.

### C. Real-World Validation and Applications

In order to validate the practical utility of our method, we deployed the scale-invariant ergodic MMD optimizer on a real-world robotic platform. A Crazyflie 2.1 drone is placed in a millimeter-scale ‘block city’ within a 3m x 3m x 1 m flight cage (as shown in Fig.7) and tasked with uniformly locating a stationary sphero hidden between the buildings. To validate our proposed method’s effectiveness at large numerical scales, the ergodic planner is provided

with the  $9\text{ m}^3$  domain in millimeter scale ( $9,000,000,000\text{ mm}^3$  search volume) and instructed to optimize a uniform-utility coverage trajectory given control barrier functions to avoid collisions with the city's buildings. Despite the significant spatial scale, our optimization method generated an effective coverage trajectory while respecting the drone's strict dynamical constraints. The generated ergodic trajectory successfully guided the drone in between the buildings according to a set of control barrier functions placed about the buildings and allowed the drone to successfully locate the target [22]. This experiment confirms the feasibility of our coverage planner to integrate with physical systems and execute complex exploration tasks with high precision.




Fig. 7. **Real-World Autonomous Object Search Using Scale-Agnostic Optimization.** Our proposed optimizer is validated on a physical Crazyflie drone to perform a real-world coverage task on a millimeter-scale 'block city'. Snapshots from the drone's downward-facing view confirm the successful execution of the coverage task and discovery of a stationary target that is hidden between buildings.

To further display the physical applications of our method to large-scale industrial tasks, we generated a coverage trajectory over the hull of a naval ship with a length of 190 m (see Fig.1). At ports, trajectories of this type can be tracked by underwater robots in order to clean the underside of ships without the need for costly and time-consuming approaches to ship maintenance, such as lifting the boats out of the water and manually washing and scrubbing the hull. The enablement of autonomous systems to perform basic coverage tasks over objects of arbitrary geometric structure and physical size yields the potential to create substantial economic and time savings in fields of maritime maintenance, materials processing, search-and-rescue, and even global-scale monitoring tasks.

## VI. CONCLUSION

We introduce a scale-invariant ergodic MMD trajectory optimizer that alleviates numerical challenges of kernel-based ergodic methods to spatial scale. Our approach leverages a novel MMD formulation that guarantees ergodic coverage over the reachable search space within a normalized domain while preserving physical constraints at their true scale. Furthermore, our method's performance is evaluated in a variety of physical implementations and simulations which showcase its light computational requirements, coverage quality, and resilience across diverse environmental scales.

## REFERENCES

- [1] A. Wahab, A. Khalid, and R. Nawaz, "Non-contact metrology inspection system for precision micro products," in *2014 International Conference on Robotics and Emerging Allied Technologies in Engineering (iCREATE)*, 2014, pp. 151–156.
- [2] A. Luca, G. Oriolo, M. Vendittelli, and S. Iannitti, *Planning Motions for Robotic Systems Subject to Differential Constraints*, 03 2005, vol. 10, pp. 145–145.
- [3] I. Abraham, A. Prabhakar, and T. D. Murphey, "Active area coverage from equilibrium," in *Algorithmic Foundations of Robotics XIII*, M. Morales, L. Tapia, G. Sánchez-Ante, and S. Hutchinson, Eds. Cham: Springer International Publishing, 2020, pp. 284–300.
- [4] M. M. Sun, A. Gaggar, P. Trautman, and T. Murphey, "Fast ergodic search with kernel functions," *IEEE Transactions on Robotics*, vol. 41, pp. 1841–1860, 2025.
- [5] M. M. Sun, A. Pinosky, and T. Murphey, "Flow Matching Ergodic Coverage," in *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, June 2025.
- [6] 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)*, 2025, pp. 01–07.
- [7] G. Laporte, H. Mercure, and Y. Nobert, "Generalized travelling salesman problem through n sets of nodes: the asymmetrical case," *Discrete Applied Mathematics*, vol. 18, no. 2, pp. 185–197, 1987.
- [8] D. L. Applegate, R. E. Bixby, V. Chvátal, and W. J. Cook, *The Traveling Salesman Problem: A Computational Study*. Princeton University Press, 2006.
- [9] F. Bourgault, A. Makarenko, S. Williams, B. Grocholsky, and H. Durrant-Whyte, "Information based adaptive robotic exploration," vol. 1, 02 2002, pp. 540 – 545 vol.1.
- [10] F. Amigoni and V. Caglioti, "An information-based exploration strategy for environment mapping with mobile robots," *Robotics and Autonomous Systems*, vol. 58, no. 5, pp. 684–699, 2010.
- [11] S. Thrun, W. Burgard, and D. Fox, "A probabilistic approach for concurrent map acquisition and localization for mobile robots," 12 1997.
- [12] A. Mavrommati, E. Tzorakoleftherakis, I. Abraham, and T. D. Murphey, "Real-time area coverage and target localization using receding-horizon ergodic exploration," *IEEE Transactions on Robotics*, vol. 34, no. 1, pp. 62–80, Feb. 2018.
- [13] G. Sartoretti, A. Rao, and H. Choset, "Spectral-based distributed ergodic coverage for heterogeneous multi-agent search," in *International Symposium Distributed Autonomous Robotic Systems*. Springer, 2021, pp. 227–241.
- [14] D. Lee, C. Lerch, F. Ramos, and I. Abraham, "Stein variational ergodic search," in *Robotics: Science and Systems*, 2024.
- [15] S. Ivić, A. Sikirica, and B. Crnković, "Constrained multi-agent ergodic area surveying control based on finite element approximation of the potential field," *Engineering Applications of Artificial Intelligence*, vol. 116, p. 105441, Nov. 2022.
- [16] C. Bilaloglu, T. Löw, and S. Calinon, "Tactile ergodic coverage on curved surfaces," *IEEE Transactions on Robotics*, vol. 41, pp. 1421–1435, 2025.
- [17] "Towards a theoretical foundation for laplacian-based manifold methods," *Journal of Computer and System Sciences*, vol. 74, no. 8, pp. 1289–1308, Dec. 2008.
- [18] S. Sirigiri, C. Hughes, I. Abraham, and F. Ramos, "Diversifying parallel ergodic search: A signature kernel evolution strategy," in *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. [Online]. Available: <https://openreview.net/forum?id=3XuUnUEI7e>
- [19] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *Journal of Machine Learning Research*, vol. 13, no. 25, pp. 723–773, 2012. [Online]. Available: <http://jmlr.org/papers/v13/gretton12a.html>
- [20] E. G. Birgin and J. M. Martínez, *Practical augmented Lagrangian methods for constrained optimization*. SIAM, 2014.
- [21] R. T. Rockafellar, "A dual approach to solving nonlinear programming problems," *Mathematical Programming*, 1973.
- [22] C. Lerch, D. Dong, and I. Abraham, "Safety-critical ergodic exploration in cluttered environments via control barrier functions," 2023. [Online]. Available: <https://arxiv.org/abs/2211.04310>