

Volumetric Ergodic Control

Jueun Kwon, Max M. Sun, and Todd Murphey

Abstract—Ergodic control synthesizes optimal coverage behaviors over spatial distributions for nonlinear systems. However, existing formulations model the robot as a non-volumetric point, whereas in practice a robot interacts with the environment through its body and sensors with physical volume. In this work, we introduce a new ergodic control formulation that optimizes spatial coverage using a volumetric state representation. Our method preserves the asymptotic coverage guarantees of ergodic control, adds minimal computational overhead for real-time control, and supports arbitrary sample-based volumetric models. We evaluate our method across search and manipulation tasks—with multiple robot dynamics and end-effector geometries or sensor models—and show that it improves coverage efficiency by more than a factor of two while maintaining a 100% task completion rate across all experiments, outperforming the standard ergodic control method. Finally, we demonstrate the effectiveness of our method on a robot arm performing mechanical erasing tasks. Project website: <https://murpheylab.github.io/vec/>

I. INTRODUCTION

Ergodic control generates robot trajectories such that the time-averaged spatial statistics match a target distribution, ensuring that time spent in any region is proportional to the target density. This produces non-myopic and efficient coverage that prioritizes high-density regions while guaranteeing asymptotic coverage of the search space. Ergodic control has been shown to be effective across a wide range of tasks including autonomous information gathering [1]–[3], manipulation [4]–[6], and robot learning [7]–[10].

Many robotics tasks depend not only on position but on how the robot’s volumetric body and sensors interact with the world. In manipulation, tool and end-effector geometry are critical [11]–[13]. In autonomous information gathering, the sensor’s field of view, range, and occlusions govern what can be observed and are essential for generating information-rich measurements [14], [15]. However, standard ergodic control methods model the robot as a non-volumetric point mass [1], [16], [17], ignoring the physical volume of the robot’s body or sensors that interact with the environment. While some extensions of ergodic control model the interaction beyond a point mass [5], [18], they remain limited to specific applications and do not generalize to a broader range of geometries and sensor models.

In this work, we extend the standard ergodic control formulation [16] to incorporate volumetric state representations. We also introduce a sample-based volumetric representation

The authors are with the Department of Mechanical Engineering, Northwestern University, Evanston, IL 60208, USA. Email: jueun@northwestern.edu. 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.

that accommodates arbitrary robot bodies or sensor models, provided these models can be expressed as samples (see Fig. 1 for examples). We call our method *volumetric ergodic control (VEC)*. Our method preserves the ability of standard ergodic control to handle nonlinear dynamics and guarantee asymptotic coverage, while adding minimal computational overhead for real-time control. We benchmark VEC through a series of tasks against standard ergodic control, including erasing tasks with varying end-effector shapes and search tasks with a differential-drive robot and a quadcopter. Experimental results show that our method improves ergodic coverage efficiency by more than a factor of two while achieving a 100% task completion rate across experiments. Lastly, we demonstrate the effectiveness of VEC on a Franka robot performing mechanical erasing tasks. In summary, the contributions of this paper are threefold:

- 1) Ergodic control formulation based on volumetric states;
- 2) Demonstrating that VEC is compatible with and benefits from sample-based representations;
- 3) Validation of VEC in simulation and hardware, demonstrating improved coverage efficiency and reliability.

II. RELATED WORK

A. Ergodic control

The standard ergodic control method is introduced in [16], which defines a Sobolev norm-based metric to evaluate the discrepancy between the spatial statistics of a trajectory and a target distribution, and proposes a receding-horizon control algorithm with an infinitesimal time horizon. This framework is extended in [1] by applying the iterative linear quadratic regulator (iLQR) for control optimization. Since then, several new ergodic metric formulations and control optimization algorithms have been proposed to address limitations of the standard ergodic control, including improved optimization efficiency [19], multi-agent coordination [20], scalability [4], [17], and coverage on curved surfaces [6], [21].

A small number of extensions attempt to incorporate robot or sensor geometry, such as via kernels with Kullback-Leibler (KL) divergence [18] or by decomposing manipulators into diffusion-driven agents [5]. However, these methods often assume simple and limited robot and sensor geometry and are restricted to narrow applications. Our method fills this gap by directly incorporating a volumetric state representation into the formulation of the standard ergodic metric [16], producing more effective coverage trajectories while preserving the benefits of ergodic control.




Fig. 1. **Erasing results using volumetric ergodic control (VEC).** The first column shows the target distributions (*airplane, Eiffel Tower, fire, heart, lock*) and remaining columns show VEC applied with different end-effector geometries (*scissors, star, sword, lightning, trophy*), which are modeled as volumetric states with the sample-based representation. VEC reasons over the volumetric states and produces efficient and thorough coverage trajectories.

B. Volumetric representations in motion planning

Beyond ergodic control, volumetric representations are essential in many other areas of motion planning. In manipulation, task performance depends on the geometry of tools and end-effectors, including their contact surfaces, orientations, and compliance [11], [12]. In autonomous information gathering, the sensor’s field of view, which varies with altitude and attitude and is often anisotropic, affects what can be observed and what information is collected [14], [15]. In structural inspection and environmental monitoring, feasible coverage is constrained by visibility and viewing geometry [22], [23]. In agriculture and surface disinfection, coverage quality depends on per-pass width and per-area intensity, linking robot speed and altitude directly to task performance [24]. These examples highlight the need for volumetric reasoning in ergodic control.

III. PRELIMINARIES ON ERGODIC CONTROL

A. Notation

Without loss of generality, we assume the search space, denoted as \mathcal{X} , is an N -dimensional rectangular space $\mathcal{X} = [0, L_1] \times \dots \times [0, L_d] \subset \mathbb{R}^d$, where L_i is the length of the i -th dimension. We denote the continuous-time robot trajectory as $s : [0, T] \mapsto \mathcal{S}$, and the control trajectory as $u : [0, T] \mapsto \mathcal{U}$, where $\mathcal{S} \subset \mathbb{R}^n$ is the state space of the robot, and $\mathcal{U} \subset \mathbb{R}^m$ is the space of feasible control signals. We denote the continuous-time dynamics of the robot as $\dot{s}(t) = f(s(t), u(t))$.

Note that the search space \mathcal{X} and the robot state space \mathcal{S} are not necessarily the same. For example, for an aerial robot conducting information gathering over a planar surface, the search space is 2-dimensional, while the state space of the robot is at least 6-dimensional. In such cases, there exists

a differentiable mapping that maps the state of the robot to a point in the search space, such as extracting the x and y position of the aerial robot during control optimization. For simplicity, we assume the state space coincides with the search space in the standard ergodic control problem.

B. Ergodic control

The ergodic metric quantifies the discrepancy between the time-averaged spatial statistics of a trajectory—formulated as an empirical distribution—and the target distribution, denoted as $q: \mathcal{X} \mapsto \mathbb{R}_0^+$.

Definition 1 (Dirac delta function): A Dirac delta function is a generalized function [25]—not characterized by point-wise mapping but by its inner product property—defined as follows:

$$\int f(x)\delta(x-s)dx = f(s), \quad \forall s \in \mathcal{X}, \forall f \in \mathcal{C}(\mathcal{X}), \quad (1)$$

where $\mathcal{C}(\mathcal{X})$ denotes the space of continuous functions over the domain \mathcal{X} .

Definition 2 (Trajectory empirical distribution): The trajectory empirical distribution is defined as:

$$p_s(x) = \frac{1}{T} \int_0^T \delta(x-s(t))dt, \quad (2)$$

where δ is the Dirac delta function.

Direct evaluation of the discrepancy between the trajectory empirical distribution (2) and the target distribution using standard methods such as the KL divergence is challenging due to the presence of the Dirac delta function in (2). To address this challenge, the standard ergodic metric introduced in [16] is based on a Sobolev norm-based discrepancy measure which leverages the inner product property of the Dirac delta function. The Sobolev norm is formulated using the normalized Fourier basis functions as a set of orthonormal bases in the function space, which are defined below.

Definition 3 (Normalized Fourier basis functions): A normalized Fourier basis function, characterized by a vector index of natural numbers $\mathbf{k} = [k_1, \dots, k_d] \in \mathbf{K} \subset \mathbb{C}^d$, is defined as:

$$f_{\mathbf{k}}(x) = \frac{1}{h_{\mathbf{k}}} \prod_{i=1}^d \cos\left(\frac{k_i \pi}{L_i} x_i\right), \quad (3)$$

where $h_{\mathbf{k}}$ ensures that $\int f_{\mathbf{k}}(x)^2 dx = 1$.

Since the normalized Fourier basis functions form a set of orthonormal bases, we can decompose the probability density function of the target distribution $q(x)$ as follows:

$$q(x) = \sum_{\mathbf{k} \in \mathbf{K}} \phi_{\mathbf{k}} f_{\mathbf{k}}(x), \quad \phi_{\mathbf{k}} = \int q(x) f_{\mathbf{k}}(x) dx, \quad (4)$$

where $\phi_{\mathbf{k}}$ is the Fourier coefficient of $q(x)$. While the exact orthonormal decomposition requires an infinite number of Fourier basis functions, in practice, we can approximate the decomposition with a finite number of Fourier coefficients.

More importantly, we can decompose the trajectory empirical distribution in the same way:

$$p_s(x) = \sum_{\mathbf{k} \in \mathbf{K}} c_{\mathbf{k}} f_{\mathbf{k}}(x), \quad c_{\mathbf{k}} = \int p_s(x) f_{\mathbf{k}}(x) dx, \quad (5)$$

where the inner product property of the Dirac delta function leads to closed-form evaluation of the Fourier coefficients:

$$\begin{aligned} c_{\mathbf{k}} &= \int \left[\frac{1}{T} \int_0^T \delta(x-s(t)) dt \right] f_{\mathbf{k}}(x) dx \\ &= \frac{1}{T} \int_0^T \left[\int f_{\mathbf{k}}(x) \delta(x-s(t)) dx \right] dt \\ &= \frac{1}{T} \int_0^T f_{\mathbf{k}}(s(t)) dt. \end{aligned} \quad (6)$$

Based on the Fourier coefficients, the standard ergodic metric is defined as follows.

Definition 4 (Standard ergodic metric): The standard ergodic metric between a trajectory $s(t)$ and a target distribution $q(x)$ is defined as:

$$\mathcal{E}(s(t), q(x)) = \sum_{\mathbf{k} \in \mathbf{K}} \lambda_{\mathbf{k}} (c_{\mathbf{k}} - \phi_{\mathbf{k}})^2, \quad (7)$$

where $\{\lambda_{\mathbf{k}}\}$ is a convergent sequence:

$$\lambda_{\mathbf{k}} = (1 + \|\mathbf{k}\|^2)^{-\frac{d+1}{2}}, \quad (8)$$

which makes the formulation (7) equivalent to a Sobolev norm with a negative index of $\frac{1+d}{2}$.

Ergodic control addresses the control optimization problem of minimizing the ergodic metric.

Definition 5 (Standard ergodic control):

$$u(t)^* = \arg \min_{u(t)} \mathcal{E}(s(t), q(x)), \quad (9)$$

$$\text{s.t. } s(t) = s_0 + \int_0^t f(s(t), u(t)) dt, \quad \forall t \in [0, T]. \quad (10)$$

The optimization problem (9) can be solved through standard optimal control methods, such as nonlinear programming, iLQR, and augmented Lagrange multiplier.

IV. VOLUMETRIC ERGODIC CONTROL

A. Volumetric state representation

Standard ergodic control formulations model the robot as a non-volumetric point, neglecting the fact that in practice, the robot interacts with the environment through its body or a sensor, both of which occupy physical volume.

Definition 6 (Volumetric state representation): A volumetric state representation, defined as a state-dependent probability density function $g: \mathcal{X} \times \mathcal{S} \mapsto \mathbb{R}_0^+$:

$$\int g(x, s) dx = 1, \quad \forall s \in \mathcal{S}, \quad (11)$$

captures how the robot's body or sensors, occupying physical volume, interact with the environment.

Definition 7 (Volumetric empirical distribution): Given a robot trajectory $s(t)$ with time horizon $[0, T]$ and the volumetric state representation $g(x, s)$, the trajectory empirical distribution is defined as:

$$p_s^v(x) = \frac{1}{T} \int_0^T g(x, s(t)) dt, \quad (12)$$

which is the spatial distribution associated with a robot trajectory under the influence of the volumetric state representation.

B. Volumetric ergodic control

In addition, the Fourier coefficients used in the ergodic metric must be redefined to apply to this volumetric empirical distribution. We derive the Fourier coefficients for the volumetric empirical distribution below:

$$c_{\mathbf{k}}^v = \int_{\mathcal{X}} f_{\mathbf{k}}(x) p_s^v(x) dx \quad (13)$$

$$= \int_{\mathcal{X}} f_{\mathbf{k}}(x) \left[\frac{1}{T} \int_0^T g(x, s(t)) dt \right] dx \quad (14)$$

$$= \frac{1}{T} \int_0^T \left[\int_{\mathcal{X}} f_{\mathbf{k}}(x) g(x, s(t)) dx \right] dt \quad (15)$$

$$= \frac{1}{T} \int_0^T f_{\mathbf{k}}^v(s(t)) dt. \quad (16)$$

We call $f_{\mathbf{k}}^v(x)$ the *volumetric Fourier basis function*.

Definition 8 (Volumetric Fourier basis): Given a volumetric state representation $g(s, x)$, the corresponding volumetric Fourier basis function is defined as:

$$f_{\mathbf{k}}^v(s(t)) = \int f_{\mathbf{k}}(x) g(x, s(t)) dx = \mathbb{E}_{x \sim g(\cdot, s(t))} [f_{\mathbf{k}}(x)], \quad (17)$$

which is the expected Fourier basis function with respect to the spatial distribution of the volumetric state representation.

Comparing the Fourier coefficients with volumetric representation (16) with the standard Fourier coefficients in (6), we can see that the Fourier coefficients in both cases are the time averages of the corresponding Fourier basis function over the state trajectory of the robot. We can now define the volumetric ergodic metric and the corresponding ergodic control problem.

Definition 9 (Volumetric ergodic metric): Given a volumetric state representation, the corresponding ergodic metric between a trajectory $s(t)$ and a target distribution $q(x)$ is defined as:

$$\mathcal{E}^v(s(t), q(x)) = \sum_{\mathbf{k} \in \mathbf{K}} \lambda_{\mathbf{k}} (c_{\mathbf{k}}^v - \phi_{\mathbf{k}})^2, \quad (18)$$

Definition 10 (Volumetric ergodic control): Given a volumetric state representation, the corresponding ergodic control problem is defined as:

$$u(t)^* = \arg \min_{u(t)} \mathcal{E}^v(s(t), q(x)), \quad (19)$$

$$\text{s.t. } s(t) = s_0 + \int_0^t f(s(t), u(t)) dt, \quad \forall t \in [0, T]. \quad (20)$$

As shown in (18), the volumetric ergodic metric, and therefore the corresponding control optimization problem, shares the same structure as the standard ergodic control problem (7). Thus, any control optimization algorithm for the standard ergodic control can be directly applied to solve the volumetric ergodic control problem.

C. Sample-based volumetric representation

While the volumetric formulation above is defined in continuous space, in practice it is often more convenient to represent the robot's volumetric state representation through a finite set of samples. Such sample-based volumetric representations naturally arise in applications such as ray-traced camera fields of view, end-effector contact points, or discretized body volumes. We approximate the state-dependent spatial density $g(\cdot, s)$ with an empirical measure induced by a finite set of samples that depend on the robot state s :

$$g(x, s) = \frac{1}{N} \sum_{i=1}^N \delta(x - h_i(s)). \quad (21)$$

Each sample is mapped from the robot state to the search space through a differentiable mapping $h_i : \mathcal{S} \mapsto \mathcal{X}$.

Substituting (21) into the definition of volumetric Fourier basis function (17), the function can be expressed as:

$$\begin{aligned} f_{\mathbf{k}}^v(s) &= \int f_{\mathbf{k}}(x) g(x, s(t)) dx \\ &= \int f_{\mathbf{k}}(x) \left[\frac{1}{N} \sum_{i=1}^N \delta(x - h_i(s)) \right] dx \\ &= \frac{1}{N} \sum_{i=1}^N \int f_{\mathbf{k}}(x) \delta(x - h_i(s)) dx \\ &= \frac{1}{N} \sum_{i=1}^N f_{\mathbf{k}}(h_i(s)). \end{aligned} \quad (22)$$

Sample-based representations of robot geometries and sensor models are widely used in robotics. For example, point cloud representations are common in manipulation [26]–[28], while ray-casting and ray-tracing methods are standard in perception [29], [30]. Importantly, in all such cases the resulting volumetric state representations are differentiable, which ensures compatibility with gradient-based control optimization. VEC shares the same computational complexity as standard ergodic control methods with respect to the number of Fourier coefficients, grid resolution, and planning horizon [17], [31]. In addition, VEC scales linearly with the number of samples used to represent the volume. Lastly, VEC can leverage the same computational acceleration techniques developed for standard ergodic control, including parallelized computation [31] and tensor-train decomposition [4].

D. Implementation details

We apply the standard iterative linear quadratic regulator (iLQR) algorithm [1] for control optimization, with the formulation also compatible with other control optimization methods, such as augmented Lagrange multiplier [3] and

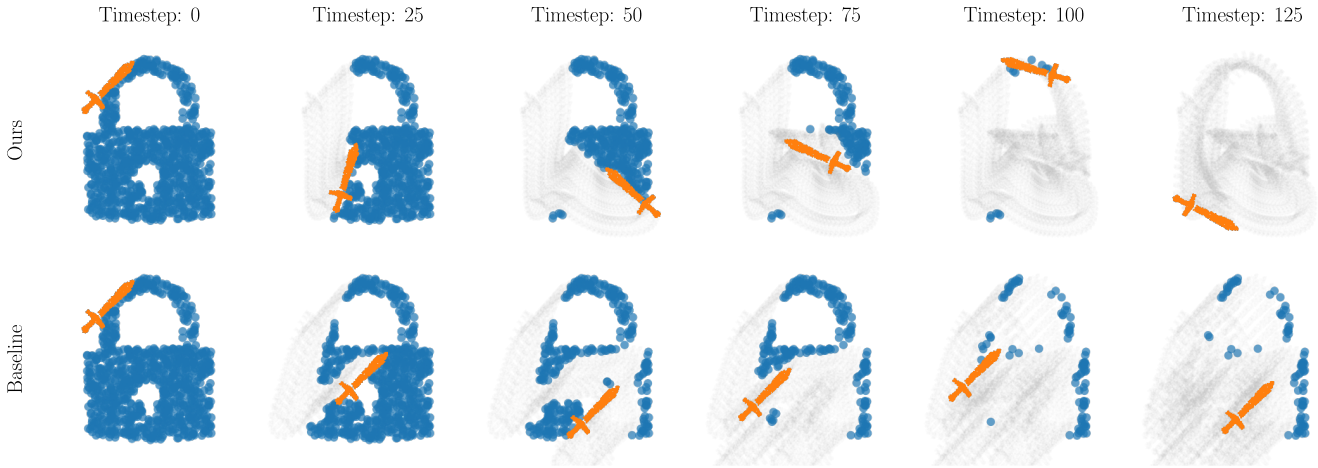


Fig. 2. **Qualitative results for erasing benchmark.** VEC effectively leverages the geometry of the end-effector through the volumetric state representation, generating both translational and rotational movements to completely erase the target in less time compared to the baseline ergodic control method.

sequential action control [19]. We implement the controller as a closed-loop receding-horizon controller based on state feedback, similar to the implementations in [19], [20], where the robot re-plans the trajectory at each timestep based on all past states. Although the Fourier coefficients of the target distribution can be updated at each timestep in practice, we fix the target distribution in our benchmark evaluations to better assess the effectiveness of the volumetric state representation relative to the baseline. We implement our method in Python using the JAX library for numerical computation.

V. EXPERIMENTAL RESULTS

A. Overview

Our benchmarks address the following three questions:

- Q1: Do standard control optimization methods, such as receding-horizon iLQR, consistently reduce the proposed volumetric ergodic metric?
- Q2: Does volumetric state representation improve the performance of ergodic control with different end-effector geometries?
- Q3: Do volumetric state representations improve the performance of ergodic control in search tasks across different sensors and robot dynamics?

B. Experimental results for Q1

[Benchmark Q1 design] We validate the effectiveness of VEC across three nonlinear dynamics and corresponding volumetric state representations:

- 1) Double-integrator: The volumetric state is a 2-D shape controlled from a randomly chosen pivot point.
- 2) Differential-drive robot: The volumetric state is modeled as a forward-facing solid-state LiDAR.
- 3) 12-DOF quadcopter: The volumetric state representation is obtained by ray-casting onto a plane.

Across all three platforms, the task is to cover a randomized Gaussian mixture distribution. We test each platform for 25 trials with randomized initial states.

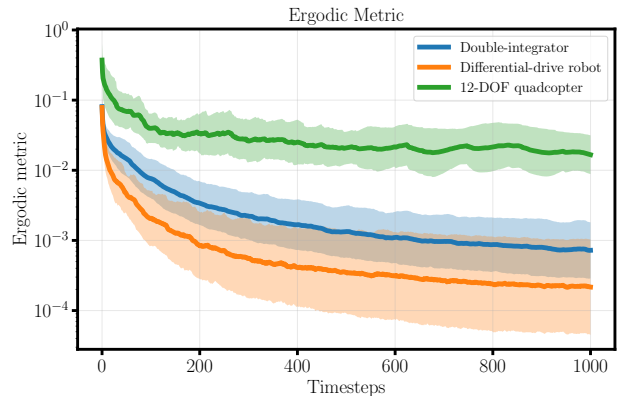


Fig. 3. **Validation of control optimization.** VEC is compatible with standard control optimization methods, such as iLQR. iLQR consistently minimizes the ergodic metric over time until convergence, across different robot dynamics, volumetric state representations, and randomized initial states. *Note that the values of the converged ergodic metric depend on the specific robot dynamics and volumetric state representation, thus should not be directly compared across platforms.*

[Benchmark Q1 results] As shown in Fig. 3, the iLQR consistently minimizes the volumetric ergodic metric across the three different platforms. These results demonstrate that the proposed volumetric ergodic control formulation is compatible with standard control optimization methods such as iLQR. Note that the converged ergodic metric value depends on the specific robot dynamics and the volumetric state representation, therefore should not be compared across robot platforms. Lastly, the average control frequencies are 8.3 Hz for the double-integrator, 9.5 Hz for the differential-drive robot, and 5.0 Hz for the 12-DOF quadcopter.

C. Experimental results for Q2

[Benchmark Q2 design] We evaluate VEC on a 2-D erasing task, where the objective is to erase all points in the planar workspace. We use the five shapes—*airplane*, *Eiffel Tower*, *fire*, *heart*, *lock*—as the target distributions and five tool geometries—*scissors*, *star*, *sword*, *lightning*, *trophy*—as the volumetric state representation, all from the open-source benchmark in [10]. In each trial, we randomize the pivot

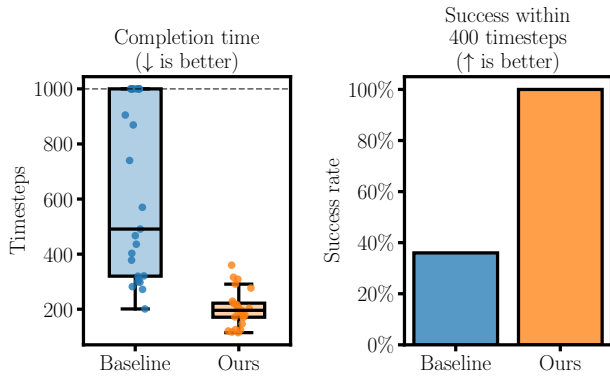


Fig. 4. **Erasing benchmark.** (Left) VEC achieves 100% task completion in less than half the time required by the baseline, which only completes 17 out of the 25 trials. (Right) VEC succeeds in all 25 trials under 400 timesteps, while the baseline only completes 9 trials within the same amount of time.

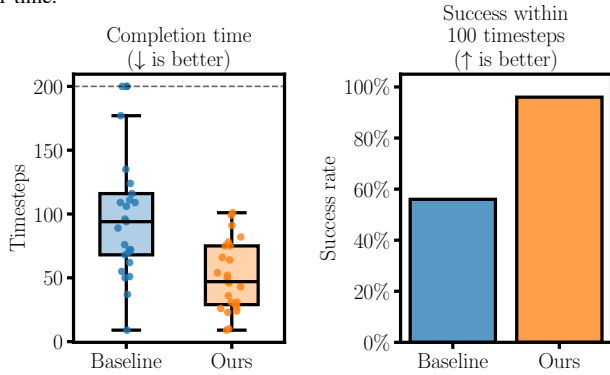


Fig. 5. **Ground search benchmark.** (Left) VEC achieves 100% task completion with on average half as many timesteps as the baseline, which only completes 22 out of the 25 trials. (Right) VEC succeeds in 24/25 trials under 100 timesteps, while the baseline only completes 14 trials within the same amount of time.

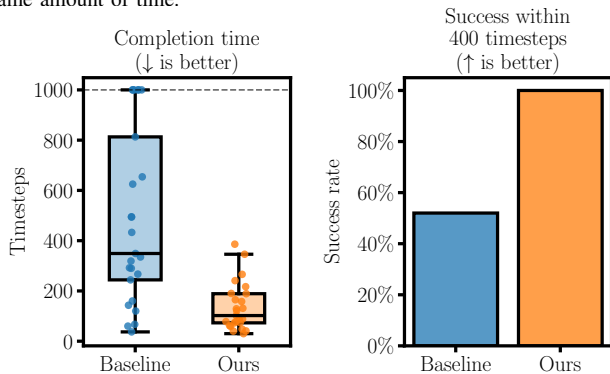


Fig. 6. **Aerial search benchmark.** (Left) VEC achieves 100% task completion in less than half the time required by the baseline, which only completes 19 out of the 25 trials. (Right) VEC succeeds in all 25 trials under 400 timesteps, while the baseline only completes 13 trials within the same amount of time.

point for controlling the tool, and the robot is modeled as a double integrator in its position and orientation. The baseline ergodic controller is implemented with the same dynamics and control optimization method, but is based on a point-based state representation at the pivot point. This ablation study focuses on the effect of volumetric state representation by keeping dynamics and control optimization identical across methods.

[Benchmark Q2 results] Fig. 2 shows the qualitative results of a representative trial from the benchmark. For quantitative

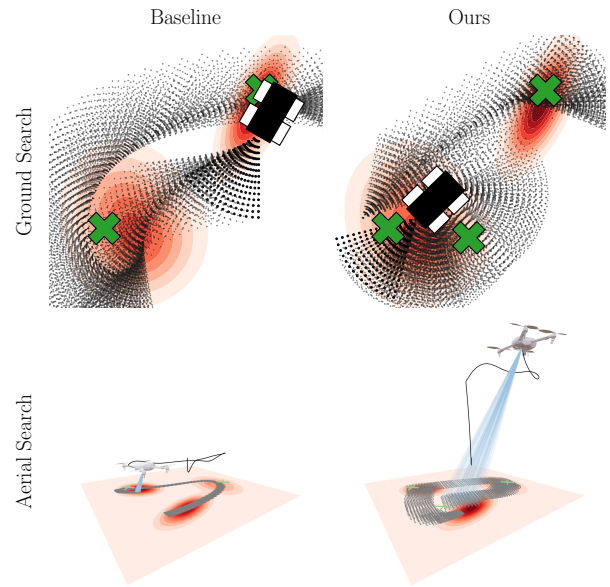


Fig. 7. **(Top) Qualitative results for ground search benchmark.** VEC leverages the volumetric LiDAR model to adapt the coverage trajectory, enabling faster and more reliable target discovery than the baseline ergodic control. **(Bottom) Qualitative results for aerial search benchmark.** VEC leverages the volumetric ray-casting camera model, enabling the quadcopter to ascend and widen its field of view, thereby completing the search more efficiently than baseline ergodic control, which remains at a constant height and performs sub-optimally.

results, Fig. 4 (Left) shows the timesteps required for task completion. VEC preserves the asymptotic coverage guarantee of ergodic control, achieving 100% task completion (25/25 trials), whereas the baseline succeeds in only 17/25 trials. Furthermore, the volumetric state representation significantly reduces the time required for task completion compared to the baseline ergodic control method. Fig. 4 (Right) shows the task completion rate under a fixed budget of 400 timesteps. VEC succeeds in all 25/25 trials within the time limit, while the baseline completes only 9/25. This highlights the importance of volumetric state representation, which enables not only convergence but also efficient execution, ensuring reliable task completion in practice. Lastly, compared to the baseline, VEC adds only minimal computational overhead: the standard ergodic control method averages 87 ms per timestep, while VEC averages 132 ms despite using 1000 points for the volumetric state representation—both sufficient for real-time closed-loop control.

D. Experimental results for Q3

[Benchmark Q3 design] We evaluate VEC on two search tasks over a planar target distribution, where the objective is to find three unknown targets. Targets are sampled from a Gaussian mixture with fixed weights and randomized means and covariances. A target is considered found once any sensor measurement enters a fixed detection radius. The Gaussian mixture serves as a prior “signs of life” distribution indicating regions where targets are more likely to appear. The ground search task uses a second-order differential-drive robot with a forward-facing solid-state LiDAR. Uniform radial sampling yields a higher concentration of points near




Fig. 8. **Franka erasing demonstration with multiple eraser shapes.** VEC accounts for the physical eraser size, producing smooth motions that efficiently erase the full target shape without redundancy. Snapshots show successful completion of the erasing task with multiple eraser geometries, including square (top), horizontal rectangle (middle), and vertical rectangle (bottom). More details are provided in the multimedia attachment.

the robot, emphasizing local coverage. The aerial search task uses a 12-DOF quadcopter equipped with a ray-casting pinhole camera grid mounted at a 20° forward tilt. Each pixel corresponds to a ray projected onto the ground plane, so the volumetric representation naturally expands, contracts, and skews with altitude and attitude. For both robots, the baseline ergodic controllers share the same dynamics and control optimization method, differing only in that they use point-based state representations. This design isolates the impact of volumetric state representation by holding dynamics and optimization identical across methods.

[Benchmark Q3 results] For ground search, qualitative results are shown in Fig. 7 (Top). Fig. 5 (Left) shows the timesteps required for task completion. VEC preserves the asymptotic coverage guarantee of ergodic control, completing all 25/25 trials compared to 22/25 for the baseline, while also requiring on average only half as many timesteps. Fig. 5 (Right) shows that under a 100-timestep budget, VEC succeeds in 24/25 trials, whereas the baseline completes only 14/25. Lastly, VEC adds minimal computational overhead: baseline averages 158 ms per timestep, while VEC averages 240 ms with 1000 points used as the volumetric state—both sufficient for real-time closed-loop control.

The performance improvement with VEC is even more drastic in the aerial search benchmark. Qualitative results are shown in Fig. 7 (Bottom). Fig. 6 (Left) shows the timesteps required for task completion. VEC completes all 25/25 trials compared to only 19/25 for the baseline, while also requiring fewer than half as many timesteps. Fig. 6 (Right) shows that under a 400-timestep budget, VEC succeeds in all 25/25 trials, whereas the baseline completes only 13/25. Similarly, VEC adds minimal computational overhead: baseline averages 213 ms per timestep, while VEC averages 276 ms per timestep with 1000 points used as the volumetric state—both remaining within real-time closed-loop requirements.

E. Erasing Demonstration with a Franka Robot

We further demonstrate the effectiveness of VEC on a Franka robot for a series of erasing tasks with multiple end-effector geometries. The target distribution is taken from the numerical erasing benchmark and hand drawn on a whiteboard. VEC is then executed using the robot’s physical eraser geometry as the volumetric state. Fig. 8 shows snapshots from three representative experiments, with videos of the full experiments provided in the multimedia attachment. VEC achieves complete erasing in all experiments. Notably, by incorporating volumetric state representation, the controller successfully alters the spatial patterns of the generated motion to accommodate the end-effector geometry, demonstrating the effectiveness of VEC in real-world conditions.

VI. CONCLUSIONS AND DISCUSSIONS

In this work, we introduce volumetric ergodic control (VEC), an extension of ergodic control based on a volumetric state representation. By introducing the volumetric Fourier basis functions, VEC replaces the standard ergodic control formulation, preserves the asymptotic coverage guarantee of ergodic control, and is solvable using standard control optimization methods such as iLQR. We further show that VEC is compatible with sample-based volumetric representations, which enables arbitrary robot and sensor geometries to be represented as samples of volumetric states. Comprehensive benchmark evaluations show that VEC consistently outperforms the standard ergodic control baselines across multiple robot types and their associated geometries or sensor models, achieving a 100% task success rate across all benchmarks while introducing minimal computational overhead for real-time closed-loop control. Hardware demonstrations on a Franka robot further showcase the efficacy of VEC in real-world conditions. Looking forward, we will extend VEC to more complex manipulation tasks—such as insertion,

grasping, and active tactile perception—where reasoning over volumetric representations is necessary.

REFERENCES

- [1] L. M. Miller, Y. Silverman, M. A. MacIver, and T. D. Murphey, “Ergodic Exploration of Distributed Information,” *IEEE Transactions on Robotics*, vol. 32, no. 1, pp. 36–52, Feb. 2016. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/7350162>
- [2] D. Lee, C. Lerch, F. Ramos, and I. Abraham, “Stein Variational Ergodic Search,” in *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024. [Online]. Available: <https://www.roboticsproceedings.org/rss20/p001.pdf>
- [3] D. E. Dong, H. Berger, and I. Abraham, “Time-optimal ergodic search: Multiscale coverage in minimum time,” *The International Journal of Robotics Research*, vol. 44, no. 10–11, pp. 1664–1683, Sep. 2025. [Online]. Available: <https://doi.org/10.1177/02783649241273597>
- [4] S. Shetty, J. Silv erio, and S. Calinon, “Ergodic Exploration Using Tensor Train: Applications in Insertion Tasks,” *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 906–921, Apr. 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/9473030/>
- [5] C. Bilaloglu, T. L ow, and S. Calinon, “Whole-Body Ergodic Exploration With a Manipulator Using Diffusion,” *IEEE Robotics and Automation Letters*, vol. 8, no. 12, pp. 8581–8587, Dec. 2023. [Online]. Available: <https://ieeexplore.ieee.org/document/10305244>
- [6] —, “Tactile Ergodic Coverage on Curved Surfaces,” *IEEE Transactions on Robotics*, vol. 41, pp. 1421–1435, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/10848343/>
- [7] A. Prabhakar and T. Murphey, “Mechanical intelligence for learning embodied sensor-object relationships,” *Nature Communications*, vol. 13, no. 1, p. 4108, Jul. 2022. [Online]. Available: <https://www.nature.com/articles/s41467-022-31795-2>
- [8] A. Pinosky and T. D. Murphey, “Embodied Active Learning of Generative Sensor-Object Models,” Oct. 2024, arXiv:2410.11130 [cs]. [Online]. Available: <http://arxiv.org/abs/2410.11130>
- [9] T. A. Berrueta, A. Pinosky, and T. D. Murphey, “Maximum diffusion reinforcement learning,” *Nature Machine Intelligence*, vol. 6, no. 5, pp. 504–514, May 2024, publisher: Nature Publishing Group. [Online]. Available: <https://www.nature.com/articles/s42256-024-00829-3>
- [10] 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. [Online]. Available: <https://www.roboticsproceedings.org/rss21/p051.pdf>
- [11] M. Cutkosky, “On grasp choice, grasp models, and the design of hands for manufacturing tasks,” *IEEE Transactions on Robotics and Automation*, vol. 5, no. 3, pp. 269–279, Jun. 1989. [Online]. Available: <https://ieeexplore.ieee.org/document/34763/>
- [12] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. Aparicio, and K. Goldberg, “Dex-Net 2.0: Deep Learning to Plan Robust Grasps with Synthetic Point Clouds and Analytic Grasp Metrics,” in *Robotics: Science and Systems XIII*. Robotics: Science and Systems Foundation, Jul. 2017. [Online]. Available: <http://www.roboticsproceedings.org/rss13/p58.pdf>
- [13] H. Ha, S. Agrawal, and S. Song, “Fit2Form: 3D Generative Model for Robot Gripper Form Design,” in *Proceedings of the 2020 Conference on Robot Learning*. PMLR, Oct. 2021, pp. 176–187, iSSN: 2640-3498. [Online]. Available: <https://proceedings.mlr.press/v155/ha21b.html>
- [14] E. Galceran and M. Carreras, “A survey on coverage path planning for robotics,” *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1258–1276, Dec. 2013. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S092188901300167X>
- [15] A. Bouman, J. Ott, S.-K. Kim, K. Chen, M. J. Kochenderfer, B. Lopez, A.-a. Agha-mohammadi, and J. Burdick, “Adaptive Coverage Path Planning for Efficient Exploration of Unknown Environments,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2022, pp. 11 916–11 923, iSSN: 2153-0866. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/9982287>
- [16] G. Mathew and I. Mezi c, “Metrics for ergodicity and design of ergodic dynamics for multi-agent systems,” *Physica D: Nonlinear Phenomena*, vol. 240, no. 4–5, pp. 432–442, Feb. 2011, publisher: Elsevier BV. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S016727891000285X>
- [17] M. M. Sun, A. Gaggari, P. Trautman, and T. Murphey, “Fast Ergodic Search With Kernel Functions,” *IEEE Transactions on Robotics*, vol. 41, pp. 1841–1860, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/10891753>
- [18] E. Ayvali, H. Salman, and H. Choset, “Ergodic coverage in constrained environments using stochastic trajectory optimization,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2017, pp. 5204–5210, iSSN: 2153-0866. [Online]. Available: <https://ieeexplore.ieee.org/document/8206410>
- [19] 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. [Online]. Available: <https://ieeexplore.ieee.org/document/8114522/>
- [20] I. Abraham and T. D. Murphey, “Decentralized Ergodic Control: Distribution-Driven Sensing and Exploration for Multiagent Systems,” *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 2987–2994, Oct. 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/8392363/>
- [21] D. Dong, A. Xu, G. Gutow, H. Choset, and I. Abraham, “Ergodic Exploration over Meshable Surfaces,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, May 2025, pp. 13 088–13 094. [Online]. Available: <https://ieeexplore.ieee.org/document/11127626>
- [22] B. Englot and F. Hover, “Sampling-Based Coverage Path Planning for Inspection of Complex Structures,” *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 22, pp. 29–37, May 2012. [Online]. Available: <https://ojs.aaai.org/index.php/ICAPS/article/view/13529>
- [23] K. Hasan, S. Ahmad, A. F. Liaf, M. Karimi, T. Ahmed, M. A. Shawon, and S. Mekhilef, “Oceanic Challenges to Technological Solutions: A Review of Autonomous Underwater Vehicle Path Technologies in Biomimicry, Control, Navigation, and Sensing,” *IEEE Access*, vol. 12, pp. 46 202–46 231, 2024. [Online]. Available: <https://ieeexplore.ieee.org/document/10477402/>
- [24] D. E. Martin, W. E. Woldt, and M. A. Latheef, “Effect of Application Height and Ground Speed on Spray Pattern and Droplet Spectra from Remotely Piloted Aerial Application Systems,” *Drones*, vol. 3, no. 4, p. 83, Dec. 2019, publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: <https://www.mdpi.com/2504-446X/3/4/83>
- [25] S. M. J. Lighthill, *An Introduction to Fourier Analysis and Generalised Functions*. Cambridge University Press, 1958.
- [26] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, “PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space,” in *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc., 2017. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2017/hash/d8bf84be3800d12f74d8b05e9b89836f-Abstract.html
- [27] C. Wang, D. Xu, Y. Zhu, R. Mart n-Mart n, C. Lu, L. Fei-Fei, and S. Savarese, “DenseFusion: 6D Object Pose Estimation by Iterative Dense Fusion,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Long Beach, CA, USA: IEEE, Jun. 2019, pp. 3338–3347. [Online]. Available: <https://ieeexplore.ieee.org/document/8953386/>
- [28] M. Shridhar, L. Manuelli, and D. Fox, “Perceiver-Actor: A Multi-Task Transformer for Robotic Manipulation,” in *Proceedings of The 6th Conference on Robot Learning*. PMLR, Mar. 2023, pp. 785–799, iSSN: 2640-3498. [Online]. Available: <https://proceedings.mlr.press/v205/shridhar23a.html>
- [29] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “NeRF: representing scenes as neural radiance fields for view synthesis,” *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, Jan. 2022. [Online]. Available: <https://dl.acm.org/doi/10.1145/3503250>
- [30] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon, “KinectFusion: Real-time dense surface mapping and tracking,” in *2011 10th IEEE International Symposium on Mixed and Augmented Reality*, Oct. 2011, pp. 127–136. [Online]. Available: <https://ieeexplore.ieee.org/document/6162880/>
- [31] M. Sun, A. Pinosky, I. Abraham, and T. Murphey, “Scale-Invariant Fast Functional Registration,” in *Robotics Research*. Cham: Springer Nature Switzerland, 2023, pp. 153–169. [Online]. Available: <https://link.springer.com/chapter/10.1007/978-3-031-25555-7-11>