Measure Preserving Flows for Ergodic Search in Convoluted Environments

Albert Xu[†], Bhaskar Vundurthy[†], Geordan Gutow[†], Ian Abraham[*], Jeff Schneider[†], and Howie Choset[†]

[†] Carnegie Mellon University, Pittsburgh PA 15232, USA
[*]Yale University, New Haven CT 06511, USA

Abstract. Autonomous robotic search has important applications in robotics, such as the search for signs of life after a disaster. When *a priori* information is available, for example in the form of a distribution, a planner can use that distribution to guide the search. Ergodic search is one method that uses the information distribution to generate a trajectory that minimizes the ergodic metric, in that it encourages the robot to spend more time in regions with high information and proportionally less time in the remaining regions. Unfortunately, prior works in ergodic search do not perform well in complex environments with obstacles such as a building's interior or a maze. To address this, our work presents a modified ergodic metric using the Laplace-Beltrami eigenfunctions to capture map geometry and obstacle locations within the ergodic metric. Further, we introduce an approach to generate trajectories that minimize the ergodic metric while guaranteeing obstacle avoidance using measure-preserving vector fields. Finally, we leverage the divergence-free nature of these vector fields to generate collision-free trajectories for multiple agents. We demonstrate our approach via simulations with single and multi-agent systems on maps representing interior hallways and long corridors with non-uniform information distribution. In particular, we illustrate the generation of feasible trajectories in complex environments where prior methods fail.

Keywords: Multi-agent search, Ergodic search, Search with obstacles, Measure-preserving flows

1 Introduction

Robotic exploration has multiple applications including search and rescue missions [1], autonomous data gathering/monitoring [2], and surveillance/patrolling [3]. Such exploration problems typically entail coverage path planning [4–6] where the robot determines a trajectory that visits every point in a given space.

This research was supported in part by an Intelligence Community Postdoctoral Research Fellowship at Carnegie Mellon University, administered by Oak Ridge Institute for Science and Education through an interagency agreement between the U.S. Department of Energy and the Office of the Director of National Intelligence.

When *a priori* information is available, not every point needs to be visited, and it is beneficial to focus the robot's exploration on higher information regions. In this context, ergodic search [7] is a promising effort to aid the robots to spend more time in higher information regions and proportionally less in low information regions. Unlike uniform coverage approaches exemplified by the lawn mower pattern [8], ergodic search approaches naturally circle around "hot spots" of information, for instance the survivor gathering locations in a search and rescue mission [1].

However, virtually all prior efforts with ergodic search presume no or small obstacles in the search region [7,9–11]. Efforts towards obstacle avoidance utilize stochastic trajectory optimization [12] to discard candidate trajectories that intersect with obstacles, or place repulsive vector fields around the obstacles [3,13] to push the robots away. Alternatively, the authors in [14] use control barrier functions to serve as additional hard constraints for trajectory optimization. While such algorithms do avoid obstacles, they either struggle to converge to a feasible trajectory in complex environments or are unable to navigate around larger non-convex obstacles, like in the environment shown in Fig. 1. In this work, we focus on generating feasible ergodic trajectories for uniform and non-uniform information maps in complex environments for multiple robots, as illustrated in orange, green, and purple in Fig. 1.

To this end, we first compute measure-preserving vector fields for a given environment with obstacles. Assuming point-sized robots with first-order dynamics, we trace out trajectories by having these point robots flow along the measure-preserving vector fields. Using theorems from ergodic theory, we prove that as time $t \to \infty$, these trajectories are ergodic; the time-averaged statistics of the robot trajectory match the probability distribution of information on the map. Finite-time trajectories can then be generated by seeding a randomly generated fixed-length trajectory and minimizing its ergodicity using gradient-based methods.

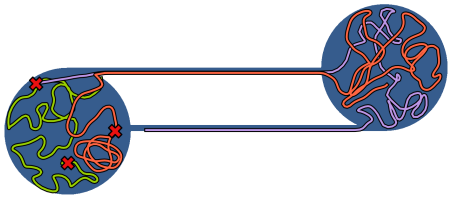


Fig. 1. An interior environment with two rooms connected by a narrow hallway. An ergodicity-minimizing trajectory generated using measure-preserving vector fields translates the three robots, starting from the red X's, to the other room through the narrow hallway for efficient information gathering.

Further, we note that the ergodic metric in prior works relies heavily on the Fourier transforms that are well-defined for Euclidean spaces. While the presence of a few scattered obstacles does not significantly impact the distance metric on Euclidean spaces, more complex environments (see Fig. 1) with narrow corridors and sharp turns significantly warp the distance metric, rendering the space non-Euclidean. Given this work primarily focuses on such complex environments, we turn to the Laplace-Beltrami eigenfunctions [15–17], which are considered to be the natural generalization of the Fourier basis functions for arbitrary manifolds.

These functions naturally encode the information of the free space without the obstacles when generating measure-preserving vector fields. Consequently, the resulting trajectories are always collision-free. Finally, an added advantage of using measure-preserving ergodic flows is that they are always divergence-free. This allows us to generalize our work to multiple agents without any additional effort and generate collision-free trajectories for all involved agents.

In summary, our contributions are as follows:

1. A formal proof that a system following the dynamics of a time-varying measure-preserving flow will be ergodic in the infinite time limit,
2. An approach to generate finite-time ergodic trajectories in the presence of obstacles using measure-preserving vector fields,
3. An amendment to the ergodic metric to handle complex environments and generate collision-free trajectories that minimize the ergodic metric, and
4. A generalization to multi-agent systems to generate joint trajectories that naturally avoid inter-agent collisions.

The rest of the paper is organized as follows. Section 2 summarizes prior ergodic search techniques that deal with obstacle and collision avoidance. Section 3 introduces the theory and proofs justifying the use of measure-preserving vector fields as well as the proposed Laplace-Beltrami eigenfunctions. Section 4 shows experimental results of simulated search trajectories on test maps, and comparisons against other methods. Lastly, Section 5 concludes the paper and poses a few future work options.

2 Literature Review

2.1 Ergodic search amidst obstacles

In the context of non-uniform information maps, ergodic search-based methods have proven to be quite effective in exploring a space [3, 4, 9]. They showcase an ability to balance exploration and exploitation, where the former attempts to visit all possible locations for new information while the latter searches high-information regions myopically [1, 8, 18]. However, historically, ergodic search in the presence of obstacles has been a challenge.

One of the first methods that tackle obstacles places repulsive vector fields around idealized circular obstacles [13]. The repulsive field strength is scaled with the distance from the robot to the obstacle, guaranteeing that the agent will never collide with obstacles as long as the integration time step is sufficiently fine. However, the repulsive fields are only defined for circular obstacles, making it difficult to conduct a comparison study against other methods for more complex environments.

Alternatively, Ayvali *et. al.* [12] use stochastic trajectory optimization for ergodic coverage (STOEC) to compute coverage trajectories, which lends itself naturally to obstacle avoidance. During the stochastic sampling step, rejecting all the samples that lead to an obstacle intersection is an easy way to guarantee that the chosen trajectory will be collision-free. However, the drawback of this

method is implicit. To ensure good ergodic coverage, the duration of sampled trajectories must be long enough. However, the probability of randomly sampling a valid trajectory in a convoluted environment vanishes exponentially as the time horizon increases. In practice, STOEC struggles to find paths around large obstacles or maze-like environments.

Lerch et al. [14] propose a control barrier function to prevent the trajectory from entering the obstacle volumes, then use gradient-based optimization methods to numerically minimize the ergodic metric with respect to the trajectory and control inputs. Unfortunately, this method also struggles with convoluted maps since the computed trajectories get stuck behind large obstacles or inside small rooms.

2.2 Laplace-Beltrami basis functions for non-Euclidean spaces

The Fourier basis functions are commonly used for spectral analysis in $\mathbb{R}^n$. They appear in the ergodic metric to help us compare the spectral coefficients of the information distribution and the time-averaged robot statistics. However, it is well known that they are not appropriate when the underlying space is not square [19, 20], as evidenced by the necessity of Laplace's spherical harmonic functions. The most general solution for spectral analysis on an arbitrary manifold uses the eigenfunctions of the Laplace-Beltrami operator. It is used in computational geometry to summarize functions on meshes [15, 21], especially when the space has curvature or complicated boundaries. In a similar vein, we utilize it to summarize the spatial and temporal statistics of our information distribution in ergodic search, especially in the presence of complex obstacles.

3 Theory

We begin our discussion on ergodic theory with a compact metric space X, which for this paper we assume is a subset of $\mathbb{R}^n$ equipped with the Borel σ-algebra $\mathscr{B}$. In addition, we have some known information distribution represented by a probability measure $\mu : \mathscr{B} \to \mathbb{R}$. Together these objects form a measure space $(X, \mathscr{B}, \mu)$. We study the behavior of a robot on this space by treating the robot as a point $x \in X$ and its time evolution as a function $T^t : (X, \mathbb{R}) \to X$; the robot's configuration space is X and a robot starting at x will end up at $T^t(x)$ after t seconds. Thus the ergodic criterion is the convergence of the robot's trajectory statistics to the desired distribution, stated in (1).

$$\frac{1}{\tau}\int_0^\tau \delta_{[T^t(x)]}dt \to \mu \tag{1}$$

In this section, we detail the generation of ergodic trajectories by first studying measure preserving flows, an essential component of ergodic systems. Then we prove that following the flows results in an ergodic trajectory by showing that they are uniquely ergodic. Lastly, we discuss ergodic search on a fixed time horizon, using the randomly generated trajectory as an initial guess and minimizing the ergodic metric using gradient-based methods.

3.1 Measure-preserving flows

Fundamental to ergodic theory is the concept of *measure-preserving flows*, as measure-preserving is a prerequisite for ergodicity [22]. Thus it is useful to search for an ergodic flow within the space of all measure-preserving flows. Flows (Definition 1) are functions that simulate motion along a vector field, whether that field is time varying or not. A flow is called measure-preserving (Definition 2) with respect to a measure space $(X, \mathscr{B}, \mu)$ if it leaves the measure unchanged.

Definition 1 (Flow). *We define a* ***time-invariant flow*** *as a family of continuous functions $T^t : X \to X$ such that for any $s, t \in \mathbb{R}$, $T^t(T^s(x)) = T^{t+s}(x)$. Each flow is related to a vector field over X by the exponential map; $T = \exp(v)$ and $v = \log(T)$. Equivalently, the time-invariant flow can be defined as the flow along its vector field $v(x)$.*

On the other hand, a ***time-varying flow*** *is defined as the flow along a time-varying vector field $v(x, t)$.*

We note that our time parameter $t \in \mathbb{R}$ is a real number, thus admitting a negative time flow. For this paper, we assume our flows are invertible, and thus T^{-t} is the inverse map of T^t.

Definition 2 (Measure-preserving). *Given a measure space $(X, \mathscr{B}, \mu)$ and a flow $T^t : X \to X$, T is* ***measure-preserving*** *on μ if for all subsets $A \subseteq X$, $\mu(A) = \mu(T^{-t}A)$.*

The converse is referred to as ***T-invariance****, where the same property indicates μ is T-invariant.*

To describe the space of measure-preserving flows, we can look at the differential change in measure that would result from moving a point in X along a vector field v. It turns out the space of all measure-preserving flows can be identified with the solutions to a partial differential equation (Theorem 1).

Theorem 1 (Measure-preserving flows). *Let $(X, \mathscr{B}, \mu)$ be a measure space (we assume $X \subseteq \mathbb{R}^n$). Let $p(x)$ be a continuous and differentiable probability density function associated with μ such that $\mu(A) = \int_A p(x)dx$.*

The set of all measure-preserving flows on this measure space can be described through their associated vector fields $T = \exp(\vec{v})$. Further, these vector fields form a linear subspace described by the following partial differential equation.

$$\nabla \cdot (p(x)\vec{v}(x)) = 0 \tag{2}$$

Proof. in Appendix A at (https://github.com/biorobotics/DARS24-Ergodic-Flows)

As the solutions to (2) form a linear subspace, we can approximate that subspace to a certain extent using a finite basis of vector fields, all of which satisfy the measure-preserving criterion. Our paper focuses primarily on an application to a 2D setting, and one method for finding measure-preserving vector fields on 2D spaces is to use the 2D curl. First, we choose a set of basis functions $\{u_i(x)\}$

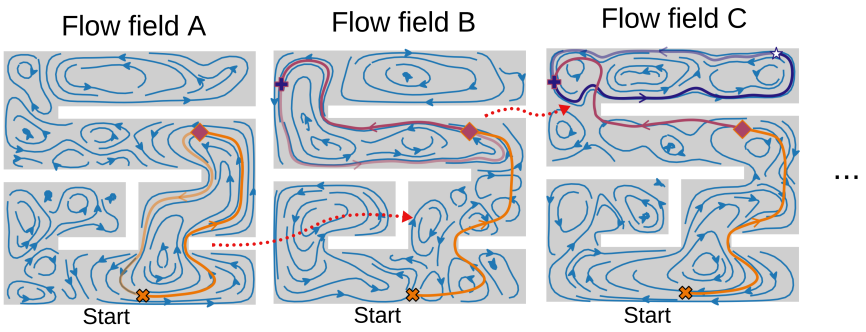


Fig. 2. Example measure-preserving flows for a uniform information distribution on a maze-like space, drawn in blue. (A) Starting from an initial condition marked by the brown X, the brown line represents a trajectory traced by following the flow field lines. Note that the flow lines form stationary limit cycles, as shown by the faded brown extension, so the robot will simply trace the same cycle repeatedly. (B) By cleverly switching the underlying flow field when the robot reaches the magenta diamond, we can escape the limit cycle and cover more space. (C) We can repeat this process when the robot reaches the purple plus, extending the trajectory to cover more of the map. This process can be repeated again when the robot reaches the white star ad infinitum.

as the eigenfunctions of the Laplacian, subject to a Dirichlet boundary condition $u(x) = 0$ for all $x \in \partial X$.

$$\nabla^2 u_i = \lambda_i u_i \tag{3}$$

Then we can construct vector fields using the 2D curl, where x and y denote the two dimensions. We can see that the vector fields in equation (4) always satisfy (2), and on the boundary ∂X, v will always point parallel to the boundary.

$$v_i(x,y) = \left[-\frac{1}{p(x,y)} \frac{\partial u_i(x,y)}{\partial y}, \frac{1}{p(x,y)} \frac{\partial u_i(x,y)}{\partial x} \right] \tag{4}$$

We note that following these basis flows, or any linear combination of them, will result in limit cycles and will *not* yield ergodic trajectories. This effect is visualized in Fig. 2 with the traced trajectory (solid and faded) for each respective flow field. Measure-preserving is a necessary condition for ergodicity, but is not sufficient. The process for generating ergodic trajectories using these measure-preserving flows requires time variance and is detailed in the next section.

3.2 Infinite-Limit Ergodicity from Random Flows

As mentioned in the previous section, tracing a 2D time-invariant flow will always result in a stationary limit cycle. To break out from the limit cycle, we need to follow a time-varying flow. This concept is illustrated in Fig. 2, where switching between flow fields A, B, and C at key points helps extend the trajectory to cover more of the space.

A random flow (Definition 3) is a simple method to generate a time-varying flow. In particular the random flow $\mathfrak{T}^t$ is a Markov process, and by the Krylov-Bogolyubov theorem [23] admits at least one invariant measure.

Definition 3 (Random flow). *Let $\mathfrak{T}^t$ be a time-invariant distribution of flows such that if $T^t \sim \mathfrak{T}^t, T^s \sim \mathfrak{T}^s$, then $T^t(T^s(\cdot)) \sim \mathfrak{T}^{t+s}$. $\mathfrak{T}^t$ is a Markov process. Note that the sampled flows are time-varying, but the overall stochastic process is time-invariant.*

Invariant measures are defined in terms of $\mathfrak{T}^t$'s push-forward measure $(\mathfrak{T}^t)_*$. μ is invariant for $\mathfrak{T}^t$ if $(\mathfrak{T}^t)_*\mu = \mu$.

Notably, if $\mathfrak{T}^t$ admits only one invariant measure, we can invoke *unique ergodicity* (Theorem 2) to prove that $\mathfrak{T}^t$ is ergodic. In other words, if there is exactly one measure μ invariant with respect to $\mathfrak{T}^t$, then $\mathfrak{T}^t$ must be ergodic with respect to μ. And therefore the spatial statistics of the trajectory traced by following T^t will converge to μ as $t \to \infty$.

Theorem 2 (Unique Ergodicity). *Let $(X, \mathscr{B}, \mu)$ be a measure space and let $\mathfrak{T}^t$ be a distribution of random flows. If μ is the only measure stationary with respect to $\mathfrak{T}^t$, then we say $\mathfrak{T}^t$ is uniquely ergodic. And in the infinite limit, trajectories sampled from $\mathfrak{T}^t$ will be ergodic.*

Proof. See Eisiedler's textbook [22], Chapter 4.3.

Using the basis vector fields $\{v_1, \cdots, v_n\}$ from (4), we can construct a distribution uniquely ergodic with respect to the given measure. With $U_1 \cdots U_n$ being n uniform random variables on the range $[-1, 1]$ and $u_i \sim U_i$, let $\mathfrak{T}^t$ be the random flow obtained by exponentiating the random vector field V^t.

$$\sum_{i=1}^{n} u_i v_i \sim V^t \tag{5}$$

Corollary 1 (Constructing a Uniquely Ergodic Random Flow). *Let $\mathfrak{T}^t$ be the random flow obtained from exponentiating (5). If the support of the desired measure μ has only one connected component, then $\mathfrak{T}^t$ is ergodic.*

Proof. Any measure invariant with V^t must be invariant for every basis vector field $\{v_i\}$. By construction, μ is invariant for all the $\{v_i\}$ and invariant with V^t.

On the other hand, there is no other measure ν invariant with every $\{v_i\}$, unless μ can be split (i.e. across multiple connected components). By contradiction, suppose $\nu \neq \mu$ is invariant with all $\{v_i\}$. The boundary of ν's support must be a shared limit cycle of all $\{v_i\}$. But if the support of μ has only one connected component, then the support of ν would have to match the support of μ and thus $\nu = \mu$, leading to a contradiction.

Thus μ is the only invariant measure, and $\mathfrak{T}^t$ is ergodic by Theorem 2. □

3.3 Ergodic metric on arbitrary 2D metric space

According to Corollary 1, randomly sampling a trajectory from $\mathfrak{T}^t$ as time $t \to \infty$ guarantees an ergodic trajectory, covering our space according to the given information distribution. However, there is no mention of how quickly the convergence happens or how well the trajectory covers the space for a finite time horizon.

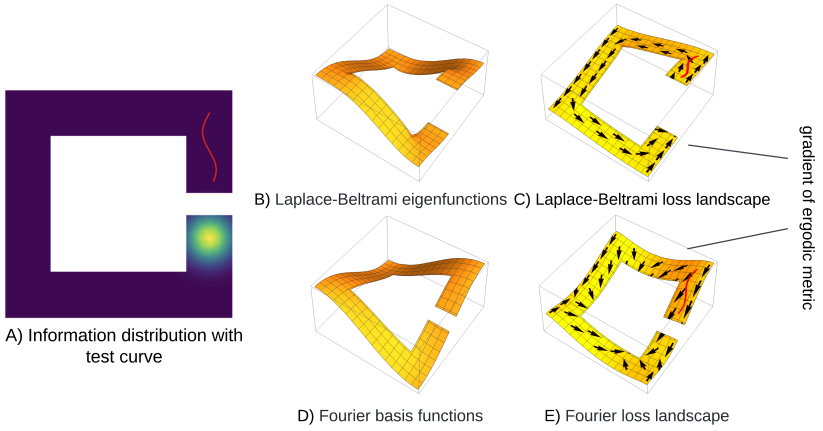


Fig. 3. Effect of modifying the ergodic metric to use the Laplace-Beltrami eigenfunctions. (A) A C-shaped map with an information peak (yellow) and test curve (red) are separated by a gap. Both basis function choices depend only on the geometry of the obstacle map and can be compared in figures B and D. The Laplace-Beltrami eigenfunctions (B) have the correct sense of continuity influenced by the boundaries of the free space, while the ordinary Fourier basis functions (D) try to maintain continuity within the gap. As a result, the Laplace-Beltrami loss landscape (C) smoothly pushes the test curve around the central obstacle, while the Fourier loss landscape (E) gets trapped in a local minima.

Thus we turn to the ergodic metric [7] used in many other ergodic search formulations to maximize coverage in a finite time horizon. However, rather than the standard Fourier basis functions for a square map $\phi_{ij} = \cos(k_i \pi x)\cos(k_j \pi y)$, we must instead find the spatial harmonic basis functions for our arbitrary metric space X using the eigenfunctions of the Laplace-Beltrami operator [24] (6).

$$\nabla^2 \phi_k = \lambda_k \phi_k \tag{6}$$

Because our space is digitally represented as a triangular mesh, we use the Finite Element Method (FEM) to compute mesh basis functions on the graph Laplace-Beltrami operator [17, 25, 26]. This yields the basis functions ϕ_k, with integrals approximated as the inner product $\int_X fg dx \approx f^\top M g = \langle f, g \rangle_M$. The kth spectral coefficient of the information distribution $p(x)$ is $\hat{\xi}_k = \langle p, \phi_k \rangle_M$ and the kth spectral coefficient of the agent trajectory is given as

$$\xi_k = \frac{1}{T} \sum_{t=1}^{T} \langle \phi_k, \delta_{x_t} \rangle_M = \frac{1}{T} \sum_{t=1}^{T} \phi_k(x_t) \tag{7}$$

The ergodic metric itself is the squared error between the information distribution and trajectory's spectral coefficients, weighted to place more importance on the low-frequency coefficients (8). The weighting formula $(1 + \sqrt{\lambda_k})^{-2}$ is chosen to ensure consistent behavior on an obstacle-free square map when compared to prior literature. We will henceforth refer to our modified ergodic metric as

the Laplace-Beltrami ergodic metric to distinguish it from the previous Fourier ergodic metric.

$$\mathcal{E} = \sum_{k=0}^{\infty} \left(1 + \sqrt{\lambda_k}\right)^{-2} \left(\xi_k - \hat{\xi}_k\right)^2 \tag{8}$$

3.4 Ergodic metric minimization for finite time horizon

To minimize the ergodic metric (8) on a finite time horizon, we formulate the problem as a minimization problem over the state x_a^t for every agent at every time and vector field coefficients $u_i(t)$. The first constraint comes from forcing the trajectory to follow the measure-preserving flows (5), and the second constraint enforces a maximum velocity.

$$\begin{aligned} \min_{x_a^t, u_i(t)} \quad & \mathcal{E}(x_a^t) \\ \text{s.t.} \quad & \dot{x}_a^t = \sum_i u_i(t) v_i(x_a^t) \qquad \forall\ t \in [0, T] \\ & \|\dot{x}_a^t\| \leq v_{\max} \qquad \forall\ t \in [0, T] \end{aligned} \tag{9}$$

Picking some time discretization Δt and writing $x_a^{t+1} = f(x_a^t, u_i(t))$ as the flow function for following $\dot{x}_a^t$ for one timestep, we can rewrite the constraints of (9) as $x_a^{t+1} = f(x_a^t, u_i(t))$ and $\|x_a^{t+1} - x_a^t\| \leq v_{\max} \Delta t$.

We use gradient descent to solve this problem, forcing the Δx and Δu updates to respect the linearized constraints at every iteration. η is gradient step size and κ is a penalty term for the change in vector field.

$$\begin{aligned} \min_{\Delta x, \Delta u} \quad & \|\eta \nabla_x \mathcal{E} - \Delta x\|^2 + \kappa \|\Delta u\|^2 \\ \text{s.t.} \quad & x_a^{t+1} + \Delta x_a^{t+1} = f(x_a^t, u_i(t)) + \partial_x f(\cdot) \Delta x_a^t + \partial_u f(\cdot) \Delta u \\ & \|x_a^{t+1} + \Delta x_a^{t+1} - x_a^t - \Delta x_a^t\| \leq v_{\max} \Delta t \end{aligned} \tag{10}$$

With the linear approximation of the flow constraints, we notice that (10) is a second-order cone problem (SOCP), and we minimize ergodicity iteratively by running a SOCP solver to find the Δx and Δu_i updates.

4 Results and Discussion

In order to evaluate the efficacy of our approach in complex environments, we present results from experiments on three maps (see Fig. 4) with uniform and nonuniform information distributions. We also compare the performance of various algorithms in single and multi-agent settings. In particular, we compare our results against



Fig. 4. Obstacle maps used for testing. From left to right, we have the obstacle-free map, the Maze map, and the Rooms map.

STOEC [12] and the control barrier function method [14], observing at least an order of magnitude improvement in performance as the obstacles in the environment get more complex. Note that the experiments are run with a max velocity of 1 unit/s, with the single-agent case running for a duration of $t = 5$s. The multi-agent comparison cases use a total of 7 agents, with each agent having a trajectory of $t = 1$s.

4.1 Obstacle-free Map Experiment Results

The obstacle-free map is used as a baseline to ensure that our method matches the performance of control barrier function-based implementation [14]. The experiment results for the single-agent uniform (Left), single-agent nonuniform (Center), and multi-agent nonuniform (Right) cases are shown in Fig. 5. The final ergodic metric scores for the four experiments, i.e., uniform or non-uniform (gaussian) coverage using single or multiple agents, are reported in Table 1, with F being the Fourier ergodic metric and LB being our modified Laplace-Beltrami ergodic metric. We observe that our method performs comparably to the baseline. Note that for the square map, the Fourier and Laplace-Beltrami ergodic metrics are identical. The numerical difference in the table is solely due to approximation error.

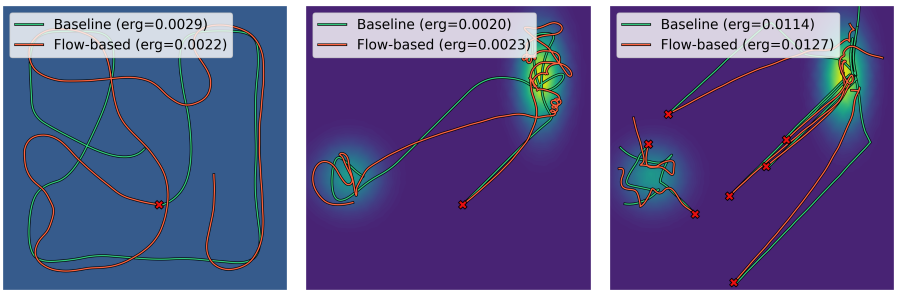


Fig. 5. Results for obstacle-free square map comparing our and the baseline methods. From left to right, we have the trajectories generated for the uniform, nonuniform, and multi-agent nonuniform test cases. The ergodic metric value for each trajectory is shown in the figure legend.

Table 1. The Fourier (F) and Laplace-Beltrami (LB) ergodic metric values for the trajectory produced by each method on the obstacle-free square map. The ergodicity values between these approaches are comparable in the absence of obstacles.

Obstacle-free map	Single Agent				Multi-Agent			
	Uniform		Gaussians		Uniform		Gaussians	
	F	LB	F	LB	F	LB	F	LB
Baseline [14]	0.0028	0.0029	0.0019	0.0020	0.0006	0.0005	0.0118	0.0114
Flow-based (ours)	0.0022	0.0022	0.0025	0.0023	0.0004	0.0004	0.0131	0.0127

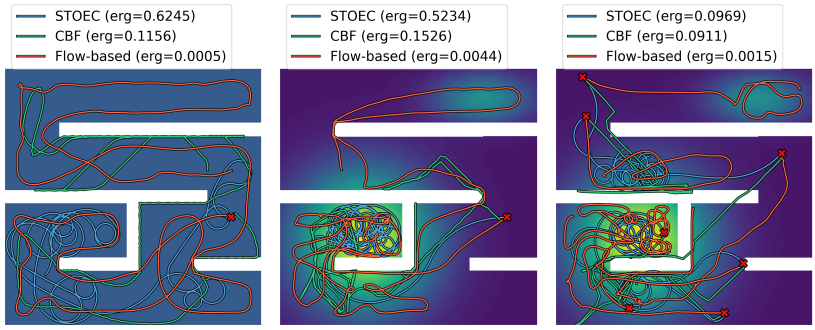


Fig. 6. Results for maze map comparing our and prior methods. From left to right, we have the trajectories generated for the uniform, nonuniform, and multi-agent nonuniform test cases. The ergodic metric value for each trajectory is shown in the figure legend, calculated using the Laplace-Beltrami ergodic metric.

4.2 Maze Map Experiment Results

The maze map poses a path-finding problem in addition to the search problem, and the challenge lies in the large distance between points on opposite sides of a wall and the back-and-forth snaking required to navigate the maze. The trajectories computed from various approaches are illustrated in Fig. 6 and their respective ergodicity values are summarized in Table 2. The maze map results demonstrate that our flow-based method is not only able to avoid the obstacles, but can also navigate around large obstacles. We observe that although STOEC [12] also successfully avoids wall collisions, it gets trapped in some corner of the maze, unable to explore the entire map. The control barrier function approach [14], on the other hand, hugs the walls very closely and occasionally navigates through the wall leading to infeasible trajectories. In summary, we observe that our method is able to better explore the maze in all cases, evidenced by the lower ergodic metric value.

Table 2. The Fourier (F) and Laplace-Beltrami (LB) ergodic metric values for the trajectory produced by each method on the maze map. When compared with two prior works, our approach shows at least an order of magnitude improvement in the ergodicity values.

Maze map	Single Agent				Multi-Agent			
	Uniform		Gaussians		Uniform		Gaussians	
	F	LB	F	LB	F	LB	F	LB
STOEC [12]	0.0658	0.6245	0.0480	0.5234	0.0095	0.0552	0.0194	0.0976
Control barrier [14]	0.0067	0.0710	0.0392	0.1610	0.0011	0.0160	Collision	
Flow-based (ours)	0.0004	0.0005	0.0022	0.0044	0.0001	0.0001	0.0007	0.0015

4.3 Rooms Map Experiment Results

Lastly, the rooms map connects two circular rooms with two narrow corridors. Robots must move in a straight, single-file line through the corridors to avoid collision. The trajectories computed from various approaches are illustrated in

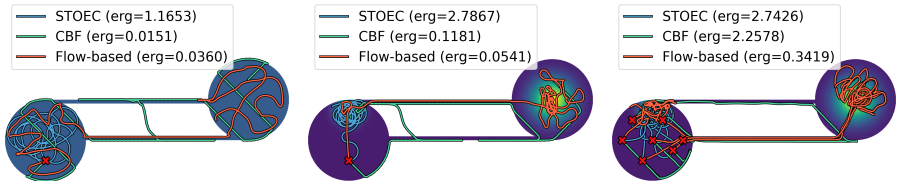


Fig. 7. Results for the rooms map comparing our and prior approaches. From left to right, we have the trajectories generated for the uniform, nonuniform, and multi-agent nonuniform test cases. The ergodic metric value for each trajectory is shown in the figure legend, calculated using the Laplace-Beltrami ergodic metric.

Fig. 7 and their respective ergodicity values are summarized in Table 3. The rooms map results show that our method is able to traverse the long corridors and explore the full area. We observe that STOEC [12] is unable to cross from the left room to the right, as the corridors are too narrow. And while the control barrier approach [14] is able to cross over, it often gets stuck with parts of the trajectory in forbidden regions. In summary, our method not only explores both rooms successfully in all the test cases but also minimizes the ergodic metric, as evidenced by the lower ergodicity values.

Table 3. The Fourier (F) and Laplace-Beltrami (LB) ergodic metric values for the trajectory produced by each method on the rooms map. When compared with two prior works, our approach shows at least an order of magnitude improvement in the ergodicity values.

Rooms map	Single Agent				Multi-Agent			
	Uniform		Gaussian		Uniform		Gaussian	
	F	LB	F	LB	F	LB	F	LB
STOEC [12]	0.1805	1.1653	0.4869	2.7867	0.1465	1.1422	0.4507	2.7426
Control barrier [14]	Collision		Collision		0.0255	0.2786	0.2958	2.2578
Flow-based (ours)	0.0024	0.0358	0.0189	0.0541	0.0116	0.1140	0.0441	0.3419

5 Conclusions and Future Work

In this work, we introduce the use of measure-preserving flows to construct feasible ergodic trajectories for point-size robots in a convoluted environment. We support this approach with mathematical proofs and show that a randomly time-varying combination of the underlying measure-preserving flows is guaranteed to yield an ergodic trajectory in the infinite time limit. Further, we present the Laplace-Beltrami ergodic metric to handle complex environments. Simulated results with single and multi-agent systems on maps that represent interior hallways and corridors illustrate the robustness of our approach. Empirical results indicate an order of magnitude improvement in ergodicity values and show success in scenarios where other approaches fail. Future directions for this work will focus on preventing collisions between finite-sized robots. Additionally, it is worth incorporating robot dynamics and physical parameters in the generation of these vector fields to help deploy generated trajectories on physical robotic systems and confirm the theoretical guarantees in the presence of physical constraints.

References

1. Robin R Murphy, Satoshi Tadokoro, Daniele Nardi, Adam Jacoff, Paolo Fiorini, Howie Choset, and Aydan M Erkmen. Search and rescue robotics.
2. Matthew Dunbabin and Lino Marques. Robots for environmental monitoring: Significant advancements and applications. *IEEE Robotics & Automation Magazine*, 19(1):24–39, 2012.
3. Shivang Patel, Senthil Hariharan, Pranav Dhulipala, Ming C Lin, Dinesh Manocha, Huan Xu, and Michael Otte. Multi-agent ergodic coverage in urban environments. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8764–8771. IEEE, 2021.
4. Howie Choset. Coverage for robotics – a survey of recent results. *Annals of Mathematics and Artificial Intelligence*, 31:113–126, 2001.
5. Enric Galceran and Marc Carreras. A survey on coverage path planning for robotics. *Robotics and Autonomous Systems*, 61(12):1258–1276, 2013.
6. Ercan U Acar, Howie Choset, Alfred A Rizzi, Prasad N Atkar, and Douglas Hull. Morse decompositions for coverage tasks. *The international journal of robotics research*, 21(4):331–344, 2002.
7. George Mathew and Igor Mezić. Metrics for ergodicity and design of ergodic dynamics for multi-agent systems. *Physica D: Nonlinear Phenomena*, 240(4-5):432–442, 2011.
8. Alice Hubenko, Vladimir Fonoberov, George Mathew, and Igor Mezic. Multiscale adaptive search. *IEEE Trans Syst Man Cybern B Cybern*, 41(4):1076–87, 2011 Aug.
9. Lauren M. Miller and Todd D. Murphey. Trajectory optimization for continuous ergodic exploration. In *2013 American Control Conference*, pages 4196–4201, 2013.
10. 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, 2015.
11. Zhongqiang Ren, Akshaya Kesarimangalam Srinivasan, Howard Coffin, Ian Abraham, and Howie Choset. A local optimization framework for multi-objective ergodic search. *arXiv preprint arXiv:2207.02923*, 2022.
12. Elif Ayvali, Hadi Salman, and Howie Choset. Ergodic coverage in constrained environments using stochastic trajectory optimization. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5204–5210. IEEE, 2017.
13. Hadi Salman, Elif Ayvali, and Howie Choset. Multi-agent ergodic coverage with obstacle avoidance. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 27, pages 242–249, 2017.
14. Cameron Lerch, Dayi Dong, and Ian Abraham. Safety-critical ergodic exploration in cluttered environments via control barrier functions. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10205–10211. IEEE, 2023.
15. Bruno Lévy. Laplace-beltrami eigenfunctions towards an algorithm that" understands" geometry. In *IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06)*, pages 13–13. IEEE, 2006.
16. Martin Reuter, Silvia Biasotti, Daniela Giorgi, Giuseppe Patanè, and Michela Spagnuolo. Discrete laplace–beltrami operators for shape analysis and segmentation. *Computers & Graphics*, 33(3):381–390, 2009.
17. Bruno Vallet and Bruno Lévy. Spectral geometry processing with manifold harmonics. In *Computer Graphics Forum*, volume 27, pages 251–260. Wiley Online Library, 2008.

18. Zhongqiang Ren, Akshaya Kesarimangalam Srinivasan, Bhaskar Vundurthy, Ian Abraham, and Howie Choset. A pareto-optimal local optimization framework for multiobjective ergodic search. *IEEE Transactions on Robotics*, 39(5):3452–3463, 2023.
19. William Elwood Byerly. *An elemenatary treatise on Fourier's series, and spherical, cylindrical, and ellipsoidal harmonics, with applications to problems in mathematical physics*. Dover Publicatiions, 1893.
20. Thomas Murray MacRobert. Spherical harmonics: An elementary treatise on harmonic functions, with applications. *(No Title)*, 1967.
21. Justin Solomon, Keegan Crane, and Etienne Vouga. Laplace-beltrami: The swiss army knife of geometry processing. In *Symposium on Geometry Processing Graduate School (Cardiff, UK, 2014)*, volume 2, 2014.
22. Manfred Einsiedler. *Ergodic theory*. Springer, 2011.
23. Nicolas Kryloff and Nicolas Bogoliouboff. La théorie générale de la mesure dans son application à l'étude des systèmes dynamiques de la mécanique non linéaire. *Annals of mathematics*, pages 65–113, 1937.
24. Uwe Graichen, Roland Eichardt, Patrique Fiedler, Daniel Strohmeier, Frank Zanow, and Jens Haueisen. Sphara-a generalized spatial fourier analysis for multisensor systems with non-uniformly arranged sensors: Application to eeg. *PloS one*, 10(4):e0121741, 2015.
25. Ramsay Dyer, Hao Zhang, Torsten Möller, and Andrew Clements. An investigation of the spectral robustness of mesh laplacians. Technical report, Simon Fraser University, 2007.
26. Hao Zhang, Oliver van Kaick, and Ramsay Dyer. Spectral methods for mesh processing and analysis. In *Proceedings of Eurographics State-of-the-art Report*, volume 122, pages 1865–1894. Eurographics Association Prague, 2007.