

Asymptotically Optimal Ergodic Coverage on Generalized Motion Fields

Christian Hughes*, Yilang Liu*, Yanis Lahrach[¶], Julia Engdahl[†], Houston Warren[‡], Darrick Lee[§], Fabio Ramos[‡], Travis Miles[†], and Ian Abraham^{*‡}

*Yale University, [¶] Université Catholique de Louvain, [†]Rutgers University, [§]University of Edinburgh, [‡]University of Sydney




Fig. 1: Autonomous Migratory Tracking over the Gulf. We propose a novel coverage trajectory planner that reasons about coverage over abstract domains that evolves according to flow fields. Here, the search domain corresponds to the whale population that flows with the ocean currents. The autonomous uncrowded underwater system (UxS) equipped with our proposed method integrates the flow model into guaranteeing coverage and observational frequency despite its under-actuation.

Abstract—Autonomous robotic exploration in remote and extreme environments allows scientists to model complex transport phenomena and collective behaviors described by continuously deforming flow fields. Although these environments are naturally modeled as time-varying domains, most adaptive exploration methods assume static environments and fail to provide adequate coverage or satisfy any formal guarantees. This is especially the case in oceanography where autonomous underwater systems (UxS) have highly restrictive compute and payload requirements that necessitate path planning methods that yield robust data collection strategies in open-loop and underactuated settings.

In this work, to address the aforementioned issues, we propose to formulate adaptive search as an ergodic coverage problem and investigate certifying coverage in the ergodic sense over evolving domains with flow-induced dynamics. We expand upon recent work demonstrating maximum mean discrepancy (MMD) as a functional ergodic metric, and derive a flow-adaptive formulation that explicitly accounts for domain evolution within the coverage objective. We show that this approach preserves ergodic coverage guarantees in ambient flows and enables effective exploration in under-actuated, and even open-loop planning settings by integrating environment dynamics. Experiments validate that our method generalizes to diverse spatiotemporal processes including ocean exploration, and tracking human and cattle movement. Physical experiments on aerial and legged robotic platforms validate our ability to obtain ergodic coverage in non-convex, flow-restricted environments while respecting robot dynamics.

I. INTRODUCTION

Autonomous exploration remains a pivotal challenge in robotics, particularly in remote and extreme environments

This work is supported by the Office of Naval Research grant N000142512440 and the National Science Foundation under award NSF FRR 2238066. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or Office of Naval Research.

where planning for data collection is a long-term and arduous process. Classical approaches, such as random sampling [18] and greedy information-gain strategies [16], excel at quickly identifying high-value regions but fall short in incorporating long-term planning and providing coverage guarantees. Alternatively, traditional coverage-based planning methods, such as grid-based exploration [38, 5], have achieved moderate success at search tasks in structured and well-known environments, but lack consideration for environment changes that structurally alter the search domain. The core difficulty of exploring an evolving domain lies in the challenge of balancing efficient long-horizon visitation proportionality with short-term information gain. Furthermore, energy-constrained and underactuated systems, such as UxS, require stable and robust long-duration trajectories to avoid the high computational cost of iterative planning [31, 8]. Therefore, in this work, we seek to address the question: *can robotic systems perform effective coverage-based search and exploration over domains subjected to flows?*

Recent ergodic search methods have emerged as promising tools for exploration that balance short-term information gain with long-term coverage guarantees. Such methods construct exploration trajectories that maximize time-averaged visitation statistics proportional to a predefined information distribution [26], and, over sufficiently long time-horizons, allow robots to explore a domain such that their visitation asymptotically approaches a given search space’s utility distribution.

Unfortunately, traditional ergodic search methods assume a priori knowledge of the search domain and underlying information distribution [26, 37], and few consider time-varying definitions of information in the formulation [27], nor

Flow-Induced Domain

Inverted Flow Domain




Fig. 2: Illustration of Ergodic Trajectory Optimization Over a Flow-Induced Domain. The search domain is defined as a set of points uniformly drawn from the Stanford Dragon mesh. (Top) Shown is the evolution of the domain subject to a flow of a Duffing oscillator. Lighter green color denotes visitation by the trajectory rendered over time through the HSV colormap. (Bottom) Illustration of coverage fidelity on initial ambient domain. Our approach leverages a novel definition of ergodicity on the inverted flow map that facilitates computationally efficient open-loop trajectory optimization and establishes formal guarantees of coverage.

do they consider time-varying domains in the search problem (i.e. where the search space changes over time). While recent advancements have identified Maximum Mean Discrepancy (MMD), a two-sample statistical metric [9], as a kernel-based ergodic metric for exploration in arbitrary spatial domains [15, 21], its application has been largely restricted to stationary environments.

In this paper, we extend the ergodic MMD formulation to time-varying domains by defining ergodicity over discrete spatial samples that evolve according to a known flow field. Unlike spectral methods, MMD’s sample-based approach facilitates navigation in complex, non-convex geometries by decoupling the ergodic metric from a fixed coordinate structure. We derive a flow-adaptive variation of MMD that is capable of establishing asymptotic coverage guarantees in the reachable space over long-duration planning horizons. Our findings demonstrate that the proposed ergodic MMD framework effectively couples a robot’s control with the underlying flow physics of the domain. We demonstrate our approach on a number of robotic exploration problems that require traversal over evolving search landscapes, including under-actuated and, more critically, open-loop exploration problems. In summary, our contributions are:

- 1) A theoretical formulation of ergodicity that maintains coverage guarantees for information distributions transported by time-varying vector fields;
- 2) A novel MMD-based optimization method that enables ergodic exploration in non-convex and flow-constrained domains; and
- 3) Experimental demonstrations of our approach’s versatility in search problems subjected to fluidic currents, animal migration, and human crowd flows.

The remainder of this paper is structured as follows: Section II discusses prior and related work. Section III covers preliminary material on ergodicity, maximum-mean discrepancy, and ergodic control. Section IV then derives the proposed time-varying ergodic maximum-mean discrepancy method and provides theoretical analysis on the approach. Lastly, Sections V - VII present a series of experimental results, limitations,

and future research directions.

II. RELATED WORK

A. Information-Theoretic Exploration

Information-driven exploration has emerged as a crucial strategy for autonomous systems to efficiently reduce environmental uncertainty by optimizing trajectories according to the predicted information gain [4, 3]. Early foundations of information maximization used probabilistic models to minimize uncertainty across the search space [39], while later approaches use Gaussian Processes to model information gain as the inverse of predicted variance [17, 24]. These methods are highly effective for static mapping and reconstruction tasks where sensor data may be sparse or noisy.

Recent advancements have shown improvements upon traditional information-theoretic methods using sparse environment sampling [14, 6]. These methods have also been modified to enable flow-tracking, which allows for exploration in dynamic systems [25]. However, while successful in short-term tasks or those in which high-value regions are the sole priority, these methods often lack the long-horizon coverage guarantees necessary for comprehensive exploration. In this paper, we propose a flow-adaptive ergodic planner that provides theoretical guarantees for comprehensive coverage of the reachable space in non-stationary domains over long-horizon planning problems.

B. Ergodic Trajectory Optimization

Ergodic trajectory planning methods seek to minimize the discrepancy between a robot’s time-averaged visitation and a target spatial utility distribution. Early work in ergodic trajectory optimization demonstrated theoretical guarantees of full coverage of continuous-space domains over infinite time-horizons by minimizing a spectral ergodic metric based on Fourier basis functions [2, 26, 32]. However, these spectral methods are largely restricted to known, bounded Euclidean domains, which limits their utility in environments with non-convex geometries [15, 26].

Recent work has introduced Maximum Mean Discrepancy (MMD) as a functional ergodic metric for arbitrarily-structured

domains [15]. While this formulation enables exploration of complex and non-convex search spaces, current developments of this approach assume that the exploration space is static. We extend this approach by proposing a modified MMD metric that accounts for domain evolutions subject to approximated or known flow fields to enable simultaneous planning for current and future utility landscapes.

C. Flow Tracking

When interacting with environments whose topography or information measure are time-variant, it is crucial that autonomous systems possess the ability to track movements within their exploration manifold. Lagrangian Coherent Structures (LCS) provide a principled framework for analyzing time-dependent flows by identifying material surfaces that guide optimal motion in environments such as multi-agent swarms and underwater currents [34, 13, 12, 10, 33, 29].

Common approaches to manifold tracking involve approximating LCS by either estimating fluid velocity or tracking discrete points within the flow [23, 11]. Because LCS are extensions of both stable and unstable manifolds that generate time-dependent flows, LCS inherently denote regions of their represented flow where escape events are more likely to occur [28]. These unstable manifolds have also been shown to be approximable through local sampling of relevant regions [1]. However, full flow-approximation of environments under complex flows is often computationally expensive or requires high-density sampling. Alternatively, some strategies restrict exploration to easily-approximated regions within the flow manifold to reduce information requirements [20], but these methods risk missing important regions within the flow. In this paper, we leverage knowledge of the flow field to optimize trajectories that are robust to minor approximation irregularities and exploit ambient dynamics to maximize global coverage.

III. PRELIMINARIES

In this section, we review the fundamental principles of ergodic search techniques, and outline the maximum mean discrepancy metric that we use to establish the time-varying domain coverage problem.

A. Ergodicity

We define ergodicity by first considering a robot's state in discrete time as $x_t \in \mathcal{X} \subseteq \mathbb{R}^n$. Additionally, let us define $\mathbf{x} = \{x_0, \dots, x_{T-1}\}$ as the trajectory of the robot for time-horizon $T \in \mathbb{N}$ following some control input $u_t \in \mathcal{U} \subseteq \mathbb{R}^m$ for $t = 0, \dots, T-1$ and dynamics $x_{t+1} = f(x_t, u_t)$ from some initial condition x_0 , and $f : \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{X}$. Next, let us define the closed and bounded operational search space Ω where $\omega_t \in \Omega$ is a point in the operational search space. The function $g : \mathcal{X} \rightarrow \Omega$ then projects states $x_t \rightarrow \omega_t$, that is, g defines where the robot is in the operational search space.

Definition 1. (Time-Averaged Trajectory Distribution) Assuming deterministic dynamics, the time-averaged trajectory

distribution, e.g., the visitation histogram of a trajectory $\mathbf{x}$ in Ω , is given as

$$\rho_{\mathbf{x},T}(\omega) = \frac{1}{T} \sum_{t=0}^{T-1} \delta[\omega - g(x_t)], \quad (1)$$

where δ is a Dirac-delta function.

Definition 2. (Ergodicity) [40, Theorem 6.14] Let $\mu : \Omega \rightarrow \mathbb{R}^+$ be a probability measure. A trajectory $\mathbf{x} = \{x_t\}_{t=0}^{T-1}$ in Ω is said to be ergodic with respect to μ if its time-averaged statistics $\rho_{\mathbf{x},T}(\omega)$ at the limit $T \rightarrow \infty$ converge weakly to μ . In other words,

$$\lim_{T \rightarrow \infty} \int_{\Omega} \phi(\omega) d\rho_{\mathbf{x},T}(\omega) = \int_{\Omega} \phi(\omega) d\mu(\omega) \quad (2)$$

for all continuous functions $\phi \in \mathcal{C}(\Omega)$, where by definition,

$$\int_{\Omega} \phi(\omega) d\rho_{\mathbf{x},T}(\omega) = \frac{1}{T} \sum_{t=0}^{T-1} \phi(g \circ x_t) \quad (3)$$

and the equality becomes

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \phi(g \circ x_t) = \int_{\Omega} \phi(\omega) d\mu(\omega). \quad (4)$$

Intuitively, this statement of ergodicity suggests that a trajectory visits regions in Ω with probability proportional to the expected value of μ distributed over Ω . Establishing the ergodic property is ideal for information gathering and exploration that requires a robot to visit each place in the search domain. In general, measuring ergodicity for a robotic system does not occur over infinite time horizons and we typically consider problems where $T \ll \infty$ which we denote as sub-ergodic.

B. Ergodic Control

The concept of ergodicity can be rigorously quantified by defining a metric that evaluates the difference between a trajectory's time-averaged statistics and a given target distribution. More specifically, the distance to ergodicity (2) is calculated by the following metric

$$\mathcal{E}_{\mu}(\mathbf{x}) = \left\| \mathcal{F}[\rho_{\mathbf{x},T}] - \mathcal{F}[\mu] \right\|_{\mathcal{F}}^2 \quad (5)$$

where $\mathcal{F}$ is a functional transform applied to the distributions ρ and μ , and $\|\cdot\|_{\mathcal{F}}$ is the associated norm in the transformed space.

Prior work has employed the Fourier transform as $\mathcal{F}$ [26, 22] and more recent approaches adopt kernel-based techniques to approximate the L^2 ergodic metric [37]. These metrics maintain the core property that $\mathcal{E}_{\mu} \rightarrow 0$ as $T \rightarrow \infty$ if and only if the trajectory $\mathbf{x}$ is ergodic with respect to the distribution μ .

Given the definition of a robot, i.e., its dynamic constraints, the goal of ergodic control is to generate trajectories that are ergodic in a defined closed and bounded operational search space Ω . To achieve ergodicity, the robot trajectory and controls are solved through an optimization problem that

minimizes the ergodic metric. A general optimization problem can be expressed as:

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{u}} \quad & \mathcal{E}_\mu(\mathbf{x}) + \sum_{t=0}^{T-1} \ell(x_t, u_t) \\ \text{subject to} \quad & h_1(x_t, u_t) = 0 \quad \forall t \in [0, T-1] \\ & h_2(x_t, u_t) \leq 0 \quad \forall t \in [0, T-1] \end{aligned} \quad (6)$$

where $\mathbf{u}$ represents the control input sequence, ℓ is a running cost function, and the constraints h_1 and h_2 encapsulate dynamics, kinematic relations, and boundary conditions.

C. Maximum Mean Discrepancy

The ergodic metric in (5) is obtained as a special case of a generic metric between distributions, where one of the distributions is the time-averaged trajectory distribution. Maximum mean discrepancy (MMD) is a two-sample statistical metric that measures the similarity between samples of two probability distributions using kernel mean embeddings in a Reproducing Kernel Hilbert Space (RKHS). Beyond expressivity, the advantage of MMD is that one does not necessarily need access to either underlying density function to compute the metric, only samples.

To define the MMD metric, let $k : \Omega \times \Omega \rightarrow \mathbb{R}^{\geq 0}$ be a positive-definite kernel, with $\mathcal{H}$ as the associated RKHS. The kernel mean embedding for a probability distribution p is given by:

$$p \mapsto \mu_p = \mathbb{E}_{\omega \sim p}[k(\omega, \cdot)]. \quad (7)$$

When this mapping is injective, the kernel is called *characteristic* [35], which allows for a unique representation of p in $\mathcal{H}$. The MMD between distributions p and q is expressed as

$$\text{MMD}_k^2(p, q) = \|\mu_p - \mu_q\|_{\mathcal{H}}^2 \quad (8)$$

which can be expanded to

$$\begin{aligned} \text{MMD}_k^2(p, q) = & \mathbb{E}_p[k(x, x')] \\ & - 2\mathbb{E}_{p,q}[k(x, y)] + \mathbb{E}_q[k(y, y')]. \end{aligned} \quad (9)$$

such that $x, x' \sim p$ and $y, y' \sim q$. For finite samples $\mathbf{x} = \{x_i\}_{i=1}^N$ and $\mathbf{y} = \{y_i\}_{i=1}^M$, the empirical approximation of the MMD can be calculated as

$$\begin{aligned} \text{MMD}_k^2(p, q) \approx & \overline{\text{MMD}}_k^2(\mathbf{x}, \mathbf{y}) = \frac{1}{N^2} \sum_{i=1}^N \sum_{i'=1}^N k(x_i, x_{i'}) \\ & - \frac{2}{NM} \sum_{i=1}^N \sum_{j=1}^M k(x_i, y_j) + \frac{1}{M^2} \sum_{j=1}^M \sum_{j'=1}^M k(y_j, y_{j'}). \end{aligned} \quad (10)$$

In this work, we seek to use the MMD metric to define *time-varying* search domains that allow robots to adapt to their search problem.

IV. ERGODIC SEARCH ON FLOWS

A. Ergodicity Over Flows

Here, we derive the notion of ergodicity for time-varying domains. We assume a compact static domain $\Omega_0 \subset E$ which is enclosed in a larger ambient space $E = \mathbb{R}^n$. Furthermore, we consider a probability measure $\mu_0 \in \mathcal{P}(E)$ which is supported on Ω_0 . In this work, we consider discrete-time varying domains as those that evolve by some induced *discrete flow*, i.e., a function $\omega_{t+1} = F(\omega_t) \forall \omega \in \Omega$.

Definition 3. (Discrete Flow Function) A discrete flow function is a function $\zeta : \mathbb{Z}^+ \rightarrow C(E, E)$ which satisfies:

- $\zeta_0 = \text{Id}_E$ is the identity on E ;
- $\zeta_s \circ \zeta_t = \zeta_{s+t}$.

We say that the discrete flow is invertible if $\zeta_t \in \text{Homeo}(E)$ for all t (a function $f : E \rightarrow E$ is a homeomorphism if f is continuous and has a continuous inverse f^{-1}).

The idea is that the flow induces the changes in both the domain and the underlying information measure. In particular, we set $\Omega_t = \zeta_t(\Omega_0)$, and define the *push-forward measure* $\mu_t := (\zeta_t)_* \mu_0$ by

$$\mu_t(A) = \mu_0(\zeta_t^{-1}(A)) \quad (11)$$

for any Borel set $A \subset E$. Note that μ_t is a probability measure supported on Ω_t . We will now consider the notion of ergodicity with respect to such time-varying measures induced by a flow, beginning with the invertible setting, and then the general setting.

Invertible Flows. We can define ergodicity for discrete invertible flows as follows.

Definition 4. (Time-Varying Ergodicity for Invertible Flows) Suppose ζ is a discrete invertible flow. We say that a trajectory $\mathbf{x} = \{x_t\}_{t=0}^{T-1}$ is ergodic with respect to the time-varying measure μ_t if

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \phi(\zeta_t^{-1} \circ g \circ x_t) = \int_{\Omega_0} \phi(\omega) d\mu_0(\omega) \quad (12)$$

for all $\phi \in C(\Omega_0)$.

The invertibility of the flow allows us to map our trajectory $g(x_t)$ back to this original domain. Therefore, this motivates one definition of an MMD metric by mapping the trajectory statistics back to Ω_0 , and using the standard MMD metric for μ_0 . In particular, we have

$$\begin{aligned} \overline{\text{MMD}}_{k, \zeta}^2(\rho_{\mathbf{x}}, \mu) & \\ = & \frac{1}{T^2} \sum_{t, t'=0}^{T-1} k(\zeta_t^{-1} \circ g \circ x_t, \zeta_{t'}^{-1} \circ g \circ x_{t'}) \\ & - \frac{2}{TM} \sum_{t, j} k(\zeta_t^{-1} \circ g \circ x_t, \omega_j) + \frac{1}{M^2} \sum_{j, j'=1}^M k(\omega_j, \omega_{j'}), \end{aligned} \quad (13)$$

In simpler terms, this is just the information measure pushed forward onto the domain at time t .




Fig. 3: Long-term Modeling of Gulf Stream Ocean Currents. The utility contour represents areas where flow approximation models disagree. Using the proposed method, an underwater robot visits regions proportional to their utility. The clearing utility contours in sampled regions indicate the robot's success in resolving localized flow model ambiguities.

where all the ω_j 's are sampled from μ_0 .

General Flows. In general, if we consider a non-invertible flow, the heuristic idea of time-varying ergodicity is that the trajectory statistics $\rho_{\mathbf{x},T}$ should weakly converge to μ_T as $T \rightarrow \infty$. However, $\rho_{\mathbf{x},T}$ contains samples collected at different times, and we need to map these samples forward to time T along the flow in order to compare with the measure μ_T . In particular, we define the *trajectory distribution with respect to the forward flow ζ at time $T-1$* as

$$\rho_{\mathbf{x},T}^{\zeta,f} := \frac{1}{T} \sum_{t=0}^{T-1} \delta[\omega - \zeta_{T-1-t} \circ g(x_t)]. \quad (14)$$

Now, we can give the general definition of time-varying ergodicity.

Definition 5. (General Time-Varying Ergodicity) Suppose ζ is a discrete flow. We say that a trajectory $\mathbf{x} = \{x_t\}_{t=0}^{T-1}$ is ergodic with respect to the time-varying measure μ_t if

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \phi(\zeta_{T-1-t} \circ g \circ x_t) - \int_{\Omega_T} \phi(\omega) d\mu_T(\omega) = 0 \quad (15)$$

for all compactly supported $\phi \in C_c(E)$.

Note that with this definition of ergodicity, what matters is what the measure μ_T looks like as $T \rightarrow \infty$. In particular, if the flow is not injective, and identifies two points $x, y \in \Omega_0$; for instance $\zeta_t(x) = \zeta_t(y)$, then this definition of ergodicity cannot distinguish between which of the two points have been visited. Note that in the Lipschitz invertible setting, the two definitions are equivalent.

Proposition 1. Suppose ζ is a discrete invertible flow where each $\zeta_t : E \rightarrow E$ is a Lipschitz continuous function with Lipschitz constant $\|\zeta_t\|_{\text{Lip}} \leq 1$. Then, Definitions 4 and 5 of time-varying ergodicity are equivalent.

This provides an alternate definition of a time-varying MMD metric which does not involve a backwards map, but rather

samples at a later time,

$$\begin{aligned} \overline{\text{MMD}}_{k,\zeta}^2(\rho_{\mathbf{x}}, \mu) & \\ = \frac{1}{T^2} \sum_{t,t'=0}^{T-1} k(\zeta_{T-t-1} \circ g \circ x_t, \zeta_{T-t'-1} \circ g \circ x_{t'}) & \\ - \frac{2}{TM} \sum_{t,j} k(\zeta_{T-t-1} \circ g \circ x_t, \omega_j) + \frac{1}{M^2} \sum_{j,j'=1}^M k(\omega_j, \omega_{j'}), & \end{aligned} \quad (16)$$

where now, the ω_j 's are sampled from μ_{T-1} . Recall that we do not need to measure ergodicity as $T \rightarrow \infty$ and instead only consider finite-time horizon problems. Thus, we can use the definitions of MMD to define two formulations of ergodic control for time-varying domains: 1) when the flow of the domain is given over a finite time; and 2) when the transform $\zeta_{t'}$ of the domain at some time t' to its original configuration Ω_0 can be measured (e.g., when we are tracking an object and know where it is at the current time).

While this formulation enables exploration over moving targets, the discrete samples that move according to a known flow field also can be viewed as a non-stationary distribution where the samples move according to max-likelihood. Hence, the target distribution is inherently non-stationary and assumed to be evolving subject to some flow. This allows practitioners to plan for arbitrary transient distributions by constructing a virtual flow field that relocates sample densities over time.

B. Ergodic Trajectory Optimization on Flows

We now formulate the optimal control problem for generating ergodic trajectories on flows. First, let us define the ergodic metric given a general flow ζ (either forward Def. 5 or backward Def. 4) as

$$\mathcal{E}_{\mu}^k(\mathbf{x}, \zeta) = \overline{\text{MMD}}_{k,\zeta}^2(\rho_{\mathbf{x}}, \mu) \quad (17)$$

where ζ encodes some measurable transform or known flow at the current time t that projects the initial information measure μ_0 towards μ_t . Given state and control constraints h_1, h_2 on the robotic system, e.g., initial condition, control saturation, the optimization problem is defined as

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{u}} \mathcal{E}_{\mu}^k(\mathbf{x}, \zeta) & \\ \text{subject to } h_1(x_t, u_t) = 0 \ \forall t \in [0, T-1] & \\ h_2(x_t, u_t) \leq 0 \ \forall t \in [0, T-1] & \end{aligned} \quad (18)$$

for constraints $h_1(x_t, u_t)$ and $h_2(x_t, u_t)$. Note that in this formulation, we only need information about the flow map (which specifies whether we are looking at future samples from information density μ_t or pushing back samples over the original space). An illustration of the approach is shown in Fig. 2.

C. Computational Complexity

Per optimization step, the computational complexity of the flow-adaptive MMD metric is $O(T^2 + TM + M^2)$. Because the target distribution’s self-similarity term is constant throughout the optimization, we can eliminate this term to reduce the computational scaling to $O(T^2 + TM)$. This scaling indicates that, for shorter time-horizons, the optimization’s performance is primarily sensitive to the spatial sample resolution M . Conversely, for longer time-horizons, the computational overhead is quadratically dominated by the number of trajectory control knots T . The summation-based structure of MMD also makes it highly compatible with GPU parallelization, which can reduce the computational scaling to approximately linear with respect to the planning horizon.

V. RESULTS

This section details simulated and physically implemented results of the proposed flow-adaptive ergodic planner’s approach in a variety of scenarios. We compare the effectiveness of our method against an existing flow-adaptive information-theoretic trajectory optimizer [36] and a flow-modified information-maximization approach. Additionally, we examine the variation in trajectory behavior between identical flow-adaptive ergodic algorithms when provided with varying levels of actuation. Finally, we showcase the success of our method in real-world scenarios including cattle behavior analysis and indoor human search tasks. Specifically, we are interested in addressing the following questions:

- To what degree does flow-adaptive ergodic planning mitigate the limitations of under-actuated control in high-velocity fields?
- What types of flows can the proposed method cover?
- Can ergodic planning methods yield robustness to variations in approximated flow fields?

A. Comparison to Existing Methods

Here, we demonstrate the proportional coverage effectiveness of our flow-adaptive ergodic planner by comparing our approach to two information-theoretic baselines:

- **Flow-Modified Information Maximization** [30] is a naive implementation of information maximization with integrated flow dynamics that recognizes visitation on moving samples.
- **Energy-Optimal Planning** [19, 36] is a dynamic programming approach to navigating over uncertain currents towards high-value regions by solving a grid-structured partial differential equation.

Each solver is provided the same optimization parameters, including velocity control bounds of 1.74 knots, and tasked




Fig. 4: Illustrative Trajectories for Eddy Exploration. A robot is tasked with exploring an eddy subjected to a flow (Fig. 3) to refine a flow model. Each panel displays a robot’s trajectory (red line) and the resulting reduction in model uncertainty. Our method leverages the underlying currents to maximize coverage of high-utility regions.

with exploring an eddy within the Gulf of America for one month over 6 hour control intervals (see Fig. 3 for eddy location). The utility grid shown in Figure 4 is a normalized distribution of the discrepancy between the true Gulf flow data for the month of October, 2025, against a single-day flow sample collected from August, 2025. This data is intended to mimic the discrepancies between existing flow-approximation methods used in the Gulf, and the velocity field used for the robot dynamics is similarly collected for the month of October and changes every 6 hours. Given this data, the robot seeks to minimize the overall flow discrepancy by sampling the true flow field every 6 hours during the mission and resolving ambiguities in the flow, as shown by the reduction of contour intensity in regions near the robot.

As shown in Figure 4, the information maximization approach performs moderately well, achieving an overall flow discrepancy reduction of 12.9%. However, the lack of emphasis on wide-spread coverage in the info-max algorithm results in revisitation of previously-seen areas, resulting in sub-optimal performance. The energy-optimal solver performed mildly worse than the info-max approach due to its numerical stability’s reliance on the satisfaction of the Courant-Friedrichs-Lewy condition, which cannot be guaranteed in all flow-subjected environments [7]. Additionally, the feedback policy of this method defines its displacement as a full thrust along the policy gradient, leading to coarse trajectories that may overshoot and miss high-value regions [19]. In contrast, the flow-adaptive ergodic planner produces smooth, dynamically-consistent trajectories resulting in a total flow discrepancy reduction of 18.3%. The flow-adaptive ergodic objective’s simultaneous balance between local information gain and global visitation proportionality allows the robot to anticipate the advection of high-utility features instead of reacting to them, which yields a 42% improvement over the best-performing baseline.

TABLE I: Eddy Model Refinement Comparison

Method	Mean Model Error Reduction
Flow-Adaptive EMD (Ours)	18.28 ± 3.80 %
Flow-Modified Info-Max	12.87 ± 0.65 %
Energy-Optimal Planning	11.47 ± 7.51 %

B. Energy-Constrained Exploration

Next, we evaluate the variation in our methods’ coverage effectiveness when provided with varying levels of actuation

capacity. We construct an environment with 75 randomly-placed samples and generate a vortex flow field that influences a robot and these samples. The robot is simulated for 30 different random sample sets and provided with identical optimization parameters except for a varied velocity control bound between 0 m/s and 0.5 m/s. When placed in the vortex, where the maximum flow speed is 3.46 m/s and the average flow speed is 2.56 m/s, the robot is tasked with maximizing the number of points visited within 10 seconds.




Fig. 5: Ablation Study on Coverage Effectiveness Under Actuation Constraints. (Top) Snapshots indicate an initialization of randomly-placed samples within a vortex field with a peak speed of 3.46 m/s and a mean speed of 2.56 m/s. Dark blue color denotes visitation by the trajectory. (Bottom) Performance as a function of velocity control limits. Even at lower control limits, the proposed method is able to leverage the flow field to assist with visiting particles that evolve on the flow.

When provided with the maximum velocity control bound of 0.5 m/s, approximately 19% of the average flow speed, the flow-adaptive ergodic planner enables the robot to visit up to 98.67% of spatial samples within 10 seconds (see Fig. 5). As shown in in Figure 5, the coverage quality of the robot’s trajectories steadily increases as the robot’s velocity bound increases and the robot gains the ability to maneuver through the flow. While the flow-adaptive ergodic planner is inherently limited by a given robot’s ability to traverse a given flow field, even in under-actuated systems, our method plans trajectories that leverage the flow field and the robot’s limited actuation to work with the flow and maximize visitation.

C. Experimental Validation

Real-World Exploration Using Crazyflie Drone. First, we demonstrate our system’s compatibility to robot dynamics and sparse information distributions by performing a monitoring




Fig. 6: Experimental Deployment in a Simulated Flow-Field. (Top) Simulated flow-following trajectories for a fleet of 8 Spheros. (Bottom) Corresponding hardware implementation using a Crazyflie 2.1 drone (blue trail) and a Sphero fleet (red trails) over 40 seconds. The proposed method effectively manages the drone’s flight dynamics to maintain ergodic coverage over target distributions subjected to time-varying flow patterns and sample sparsity.

task on a fleet of 8 Spheros using a Crazyflie 2.1 drone. The Spheros are distributed within a 3m x 3m space and instructed to follow trajectories that imitate each Sphero’s movement through a flow field for 40 seconds (shown in Fig. 6). A Crazyflie drone is then instructed to traverse the space, dynamically independent of the flow, to observe each of the Spheros and revisit Spheros once each has been seen. Our method successfully plans a coverage trajectory that visits each Sphero within 30 seconds while respecting drone dynamics. The MMD-based ergodic metric’s use of discrete spatial samples to represent a continuous distribution allows for trajectory optimization at varying levels of sample density, even in systems of extreme sparsity (e.g., 8 samples).

Simulated Use-Case of Cattle Behavior Inspection. When obtaining new cattle, farmers often need to observe the interactions between newly-acquired cattle and the existing herd. A common point of contention for more aggressive cattle is during feeding time, when the herd movement is approximated according to a given cow’s proximity to each feeding spot.

We evaluate our method’s effectiveness at anomaly detection by monitoring aggressive interactions during cattle feeding. As shown in Fig. 7, we conduct 30 simulations involving 30 cows randomly distributed across a 30m x 30m ranch with three feeding troughs (flow attractors). In each trial, two cows are selected as aggressive and a drone with a 8m sensing radius is tasked with identifying conflicts involving aggressive cows.

The reliance of the energy-based planner on the Courant-Friedrichs-Lewy (CLF) condition, which requires that the nu-




Fig. 7: Validation of Ergodic Coverage In a Behavioral Flow Environment. (Left) Thirty cows are shown in a fenced-in region at $t = 0$, where gray vectors indicate the directed vector field towards the nearest feeding trough to each cow. (Right) A drone executes an ergodic trajectory designed to prioritize regions of high cattle density to locate potential conflicts between cattle. Our method successfully positions the drone over problematic cows (marked with exclamation points) during feeding time. We show that the proposed method generalizes beyond fluid currents to any motion-predictable system.

merical wave speed of the PDE is at least as fast as the physical wave speed, results in unstable numerical solutions during the path-planning process due to the high-velocity gradients required to model the herd’s behavioral flow. In contrast, our flow-adaptive ergodic planner successfully guides the drone to aggressive interactions with a 50.8% success rate in an open-loop setting (see Tab. II). While information maximization tends to see coverage improvements in search problems with sparse data (e.g., 30 cows), our method balances prioritization of high-utility visitation and wide-spread coverage, resulting in a 42.8% improvement in performance. The success of this use-case demonstrates that the flow-adaptive ergodic planner enables exploration over predictable behavioral flows in addition to fluid flow.

TABLE II: Open-Loop Cattle Skirmish Detection Comparison

Method	Cows Seen	Success (%)
Total Cattle Fights	59	
Flow-Adaptive EMD (Ours)	30	50.847
Flow-Modified Info-Max [30]	21	35.593
Energy-Optimal Planning [19]	✗	✗

Indoor Human Search on Quadraped Robot. In architecturally-constrained environments, such as hallways, human motion is similarly approximable as a flow. To test our flow-adaptive ergodic planner on motion patterns with less adherence to the instructed flow, we deploy a Unitree Go2 quadraped robot with an approximated flow for 4 people walking through a floor of an office building. Each person is instructed to follow a movement pattern (shown in Fig. 8) which is provided to the Go2’s trajectory planner to perform uniform search over the 4 people to locate the person wearing an all-black outfit. To test the robustness of our solver, each person was instructed to introduce stochastic deviations, such as lateral drifting or mid-path stops, rather than following the path perfectly. For additional implementation details, please see Appx. A.

As shown in Figure 8, the Go2 successfully maneuvered

the indoor non-convex domain to observe each person on the floor and locate the target, with the robot successfully locating the target during the 2-minute run-time. The kernel-based spatial integration inherent to MMD provides a form of stochastic smoothing that yields resilience to minor flow model deviations, such as those made by humans following a path. With this validation, we show that our flow-adaptive planner successfully plans trajectories that adhere to high-dimensional robot dynamics while maintaining sufficient robustness to locate a human target who performs minor deviations from their instructed movement pattern. The success of this experiment validates that architecturally-imposed motion constraints (e.g., hallways) can create pseudo-deterministic flows that are sufficiently structured for ergodic coverage planning.

Simulated Whale Tracking on Real-World Gulf Data.

Finally, we perform a simulated validation of our method in the Gulf stream where an underwater robot with 1.16 knots of actuation ability is assigned to inspect 6 whales traveling with the ocean’s current. As shown in Fig. 1, our approach successfully observes each whale over the course of a 1-month migration while leveraging the ambient flow to ensure visitation before the pod separates. This validates our method’s consideration of the appropriate time to visit each sample in the search space by sometimes delaying short-term visitation to maximize long-term proportionality.

VI. LIMITATIONS

The primary limitation of the proposed approach is its dependence on the fidelity of the environmental flow model in open-loop. While this reliance is negligible in domains that evolve through known flows, large stochastic environments, such as oceanic currents or unconstrained crowds, introduce inherent model uncertainties. This model uncertainty presents two distinct errors in our trajectory planning: coverage degradation and state-estimation error, which lead to short-term performance sensitivities.

Theoretically, ergodic search methods are robust to modeling inaccuracies over infinite time-horizons, but would require excessive computation. A recursive online adaptation, like MPC with flow model correction, could resolve this issue, but its impact on coverage guarantees and how one would measure success in a MPC-like setting require significantly more analysis beyond the scope of this work.

VII. CONCLUSION

In this paper, we present a novel formulation of the coverage-based exploration problem as an ergodic trajectory optimization solved over time-varying flows. We derive the optimization objective for ergodic exploration over general flows using the maximum mean discrepancy (MMD) metric and show that coverage guarantees over the reachable search space are maintained despite the influence of flow dynamics in open-loop planning. Our results demonstrate that our approach outperforms traditional information-theoretic baselines by at least 42% while maintaining high-fidelity coverage in under-actuated problem settings without the need for replanning.




Fig. 8: Implementation of Flow-Adaptive Ergodic Coverage to Indoor Search Tasks. A Unitree Go2 quadruped robot (red trajectory) is tasked with locating a person of interest (wearing all black) among a group of four people. The motion of each person is modeled as a constrained social flow. The robot's camera view confirms the detection of the target during exploration in loose approximations of social flows.

The effectiveness of our proposed approach is demonstrated in highly restrictive oceanographic applications that are unable to perform onboard computation. Additionally, our approach allows autonomous platforms to leverage the ambient currents to improve mobility and intercept high-value features that are advected through the search space. Last, we show that the proposed method generalizes to behavioral flows, such as general human traffic patterns and herd movement, by treating architectural constraints or universal attractors as pseudo-deterministic flow fields.

REFERENCES

- [1] *An Experimental Testbed for Multi-Robot Tracking of Manifolds and Coherent Structures in Flows*, Dynamic Systems and Control Conference, 10 2013. doi: 10.1115/DSCC2013-3745. URL <https://doi.org/10.1115/DSCC2013-3745>.
- [2] Ian Abraham, Ahalya Prabhakar, and Todd D. Murphy. Active area coverage from equilibrium. In Marco Morales, Lydia Tapia, Gildardo Sánchez-Ante, and Seth Hutchinson, editors, *Algorithmic Foundations of Robotics XIII*, pages 284–300, Cham, 2020. Springer International Publishing. ISBN 978-3-030-44051-0.
- [3] Francesco Amigoni and Vincenzo Caglioti. An information-based exploration strategy for environment mapping with mobile robots. *Robotics and Autonomous Systems*, 58(5):684–699, 2010. ISSN 0921-8890. doi: <https://doi.org/10.1016/j.robot.2009.11.005>. URL <https://www.sciencedirect.com/science/article/pii/S0921889009002024>.
- [4] Frédéric Bourgault, Alexei Makarenko, Stefan Williams, Ben Grocholsky, and Hugh Durrant-Whyte. Information based adaptive robotic exploration. volume 1, pages 540 – 545 vol.1, 02 2002. ISBN 0-7803-7398-7. doi: 10.1109/IRDS.2002.1041446.
- [5] Rik Bähneemann, Nicholas Lawrance, Jen Jen Chung, Michael Pantic, Roland Siegwart, and Juan Nieto. *Re-visiting Boustrophedon Coverage Path Planning as a Generalized Traveling Salesman Problem*, page 277–290. Springer Singapore, 2021. ISBN 9789811594601. doi: 10.1007/978-981-15-9460-1_20. URL http://dx.doi.org/10.1007/978-981-15-9460-1_20.
- [6] Jeffrey A. Caley and Geoffrey A. Hollinger. Environment prediction from sparse samples for robotic information gathering. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10577–10583, 2020. doi: 10.1109/ICRA40945.2020.9197263.
- [7] R. Courant, K. Friedrichs, and H. Lewy. On the partial difference equations of mathematical physics. *IBM J. Res. Dev.*, 11(2):215–234, March 1967. ISSN 0018-8646. doi: 10.1147/rd.112.0215. URL <https://doi.org/10.1147/rd.112.0215>.
- [8] E.L. Creed, J. Kerfoot, C. Mudgal, and H. Barrier. Transition of slocum electric gliders to a sustained operational system. In *Oceans '04 MTS/IEEE Techno-Ocean '04 (IEEE Cat. No.04CH37600)*, volume 2, pages 828–833 Vol.2, 2004. doi: 10.1109/OCEANS.2004.1405565.
- [9] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012. URL <http://jmlr.org/papers/v13/gretton12a.html>.
- [10] G. Haller. Distinguished material surfaces and coherent structures in three-dimensional fluid flows. *Physica D: Nonlinear Phenomena*, 149(4):248–277, 2001. ISSN 0167-2789. doi: [https://doi.org/10.1016/S0167-2789\(00\)00199-8](https://doi.org/10.1016/S0167-2789(00)00199-8). URL <https://www.sciencedirect.com/science/article/pii/S0167278900001998>.
- [11] G. Haller. Lagrangian coherent structures from approximate velocity data. *Physics of Fluids*, 14(6):1851–1861, 06 2002. ISSN 1070-6631. doi: 10.1063/1.1477449. URL <https://doi.org/10.1063/1.1477449>.
- [12] George Haller. Lagrangian coherent structures. *Annual Review of Fluid Mechanics*, 47(Volume 47, 2015):137–162, 2015. ISSN 1545-4479. doi: <https://doi.org/10.1146/annurev-fluid-010313-141322>. URL <https://www.annualreviews.org/content/journals/10.1146/annurev-fluid-010313-141322>.
- [13] Imran Hayat, Ryan T. Black, and George Ilhwan Park. Reference map technique for lagrangian exploration of coherent structures. *Journal of Fluid Mechanics*, 1000, December 2024. ISSN 1469-7645. doi: 10.1017/jfm.2024.1037. URL <http://dx.doi.org/10.1017/jfm.2024.1037>.
- [14] Geoffrey A. Hollinger and Gaurav S. Sukhatme. Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9):1271–1287, 2014. doi: 10.1177/0278364914533443. URL <https://doi.org/10.1177/0278364914533443>.
- [15] Christian Hughes, Houston Warren, Darrick Lee, Fabio

- Ramos, and Ian Abraham. Ergodic trajectory optimization on generalized domains using maximum mean discrepancy, 2024. URL <https://arxiv.org/abs/2410.10599>.
- [16] Stefan Ivić, Bojan Crnković, Luka Grbčić, and Lea Matleković. Multi-uav trajectory planning for 3d visual inspection of complex structures. *Automation in Construction*, 147:104709, March 2023. ISSN 0926-5805. doi: 10.1016/j.autcon.2022.104709. URL <http://dx.doi.org/10.1016/j.autcon.2022.104709>.
- [17] Maani Ghaffari Jadidi, Jaime Valls Miró, Rafael Valencia, and Juan Andrade-Cetto. Exploration on continuous gaussian process frontier maps. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6077–6082, 2014. doi: 10.1109/ICRA.2014.6907754.
- [18] L.E. Kavraki, P. Svestka, J.-C. Latombe, and M.H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996. doi: 10.1109/70.508439.
- [19] Matthias Killer, Marius Wiggert, Hanna Krasowski, Manan Doshi, Pierre F.J. Lermusiaux, and Claire J. Tomlin. Maximizing seaweed growth on autonomous farms: A dynamic programming approach for underpowered systems operating in uncertain ocean currents. *IEEE Robotics and Automation Letters*, 10(10):10745–10752, 2025. doi: 10.1109/LRA.2025.3604727.
- [20] Dhanushka Kularatne and M. Hsieh. Tracking attracting manifolds in flows. *Autonomous Robots*, 41, 12 2017. doi: 10.1007/s10514-017-9628-y.
- [21] Yanis Lahrach, Christian Hughes, and Ian Abraham. Search at scale: Improving numerical conditioning of ergodic coverage optimization for multi-scale domains, 2025. URL <https://arxiv.org/abs/2512.05229>.
- [22] Darrick Lee, Cameron Lerch, Fabio Ramos, and Ian Abraham. Stein variational ergodic search, 2024. URL <https://arxiv.org/abs/2406.11767>.
- [23] Kun Li, Chiya Savari, and Mostafa Barigou. Computation of lagrangian coherent structures from experimental fluid trajectory measurements in a mechanically agitated vessel. *Chemical Engineering Science*, 254:117598, 2022. ISSN 0009-2509. doi: <https://doi.org/10.1016/j.ces.2022.117598>. URL <https://www.sciencedirect.com/science/article/pii/S0009250922001828>.
- [24] Roman Marchant and Fabio Ramos. Bayesian optimisation for informative continuous path planning. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6136–6143, 2014. doi: 10.1109/ICRA.2014.6907763.
- [25] Roman Marchant, Fabio Ramos, and Scott Sanner. Sequential bayesian optimisation for spatial-temporal monitoring. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, UAI’14, page 553–562, Arlington, Virginia, USA, 2014. AUAI Press. ISBN 9780974903910.
- [26] George Mathew and Igor Mezic. Metrics for ergodicity and design of ergodic dynamics for multi-agent systems. *Physica D: Nonlinear Phenomena*, 240, 02 2011. doi: 10.1016/j.physd.2010.10.010.
- [27] A. Mavrommati, Emmanouil Tzorakoleftherakis, I. Abraham, and T. D. Murphey. Real-time area coverage and target localization using receding-horizon ergodic exploration. *IEEE Transactions on Robotics*, 34(1):62–80, Feb. 2018. doi: 10.1109/TRO.2017.2757959.
- [28] James Paulos, Nick Eckenstein, Tarik Tosun, Jungwon Seo, Jay Davey, Jonathan Greco, Vijay Kumar, and Mark Yim. Automated self-assembly of large maritime structures by a team of robotic boats. *IEEE Transactions on Automation Science and Engineering*, 12(3):958–968, 2015. doi: 10.1109/TASE.2015.2416678.
- [29] J. PENG and J. O. DABIRI. Transport of inertial particles by lagrangian coherent structures: application to predator–prey interaction in jellyfish feeding. *Journal of Fluid Mechanics*, 623:75–84, 2009. doi: 10.1017/S0022112008005089.
- [30] Brent Schlotfeldt, Nikolay Atanasov, and George J. Pappas. Maximum information bounds for planning active sensing trajectories. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4913–4920, 2019. doi: 10.1109/IROS40897.2019.8968147.
- [31] Oscar Schofield, Josh Kohut, David Aragon, Liz Creed, Josh Graver, Chip Haldeman, John Kerfoot, Hugh Roarty, Clayton Jones, Doug Webb, and Scott Glenn. Slocum gliders: Robust and ready. *Journal of Field Robotics*, 24(6):473–485, 2007. doi: <https://doi.org/10.1002/rob.20200>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.20200>.
- [32] Sherry E. Scott, Thomas C. Redd, Leonid Kuznetsov, Igor Mezić, and Christopher K.R.T. Jones. Capturing deviation from ergodicity at different scales. *Physica D: Nonlinear Phenomena*, 238(16):1668–1679, 2009. ISSN 0167-2789. doi: <https://doi.org/10.1016/j.physd.2009.05.003>. URL <https://www.sciencedirect.com/science/article/pii/S0167278909001523>.
- [33] Carmine Senatore and Shane D. Ross. Fuel-efficient navigation in complex flows. In *2008 American Control Conference*, pages 1244–1248, 2008. doi: 10.1109/ACC.2008.4586663.
- [34] Cody R. Short and Kathleen C. Howell. Lagrangian coherent structures in various map representations for application to multi-body gravitational regimes. *Acta Astronautica*, 94(2):592–607, 2014. ISSN 0094-5765. doi: <https://doi.org/10.1016/j.actaastro.2013.08.020>. URL <https://www.sciencedirect.com/science/article/pii/S0094576513003408>.
- [35] Bharath K. Sriperumbudur, Kenji Fukumizu, and Gert R. G. Lanckriet. Universality, Characteristic Kernels and RKHS Embedding of Measures. 12(70):2389–2410. ISSN 1533-7928.
- [36] Deepak N. Subramani, Patrick J. Haley Jr., and Pierre F. J. Lermusiaux. Energy-optimal path planning in the

coastal ocean. *Journal of Geophysical Research: Oceans*, 122(5):3981–4003, 2017. doi: <https://doi.org/10.1002/2016JC012231>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2016JC012231>.

- [37] Muchen Sun, Ayush Gaggar, Peter Trautman, and Todd Murphey. Fast ergodic search with kernel functions. *arXiv preprint arXiv:2403.01536*, 2024.
- [38] Yoonchang Sung, Zhiang Chen, Jnaneshwar Das, and Pratap Tokekar. A survey of decision-theoretic approaches for robotic environmental monitoring, 2023. URL <https://arxiv.org/abs/2308.02698>.
- [39] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. A probabilistic approach for concurrent map acquisition and localization for mobile robots. 12 1997.
- [40] Peter Walters. *An Introduction to Ergodic Theory*. Springer Science & Business Media, 2000. ISBN 978-0-387-95152-2.