

Distributed Multi-Robot Ergodic Coverage Control for Estimating Time-Varying Spatial Processes

Mattia Mantovani , *Student Member, IEEE*, Mattia Catellani , *Graduate Student Member, IEEE*,
and Lorenzo Sabattini , *Senior Member, IEEE*

Abstract—We present a fully distributed framework for multi-robot exploration and coverage of time-varying spatial processes in complex, non-convex environments. Building on heat-equation-driven adaptive coverage (HEDAC) and system ergodicity, the proposed approach enables robots to autonomously navigate arbitrary domains, reconstruct unknown spatial fields, and continuously balance exploration and coverage without centralized coordination. A temporal decay mechanism promotes adaptive monitoring by regulating the relevance of past observations. Simulation and real-world experiments demonstrate the effectiveness and robustness of the method.

Index Terms—Multi-robot systems, distributed robot systems.

I. INTRODUCTION

MULTI-ROBOT systems are increasingly used to explore and monitor environments where human operation is impractical. Central to this is coverage control, which allows robot teams to focus sensing resources on spatial regions of interest, typically described by a density distribution.

Traditional approaches, such as Voronoi-based coverage control [1], [2], position robots at the centroids of their Voronoi cells to optimize static coverage. These approaches, however, generally assume convex environments and strong connectivity, constraining their performance in cluttered or resource-limited scenarios. More recent solutions employing model predictive control (MPC) or graph neural networks (GNN) architectures have attempted to overcome these limitations, but they often incur high computational costs or require extensive ad hoc training [3], [4], [5].

Ergodic coverage control offers a dynamic alternative, aligning a robot's time-averaged trajectory with a spatial distribution of interest. The Spectral Multiscale Coverage (SMC) framework [6] introduced a metric-based formulation, later extended with time-optimal [7], [8] and receding-horizon methods [9]. Still, SMC assumes convexity, lacks inter-robot collision avoidance, and tends to emphasize global coverage at the expense of distant high-priority regions [10], [11].

Received 1 October 2025; accepted 19 January 2026. Date of publication 6 February 2026; date of current version 19 February 2026. This article was recommended for publication by Associate Editor D. Aksaray and Editor M. A. Hsieh upon evaluation of the reviewers' comments. (*Corresponding author: Mattia Mantovani.*)

The authors are with the Department of Sciences and Methods for Engineering (DISMI), University of Modena and Reggio Emilia, 42122 Reggio Emilia, Italy (e-mail: mattia.mantovani@unimore.it; mattia.catellani@unimore.it; lorenzo.sabattini@unimore.it).

Digital Object Identifier 10.1109/LRA.2026.3662641




Fig. 1. Snapshot from a real-world experiment demonstrating the proposed distributed strategy using a team of three Uvify IFO-S drones. Each drone runs the algorithm fully onboard and collaborates with neighboring agents via minimal communication. The image shows the experiment in progress, accompanied by an RViz visualization depicting the environment, robot trajectories, and spatial process estimation.

The Heat Equation Driven Area Coverage (HEDAC) framework [12] addresses some of these issues by embedding the ergodic metric in a stationary heat equation. This formulation naturally handles non-convex domains and encourages local coverage through heat diffusion. Obstacle modeling was further generalized via the Finite Element Method in [13]. However, HEDAC does not incorporate inter-robot safety and fails to guarantee a minimum distance between robots. A further limitation of prior work is the assumption of full knowledge of a static spatial process. Real-world applications often involve unknown, time-varying phenomena requiring online exploration, learning, and monitoring. While Gaussian Processes (GPs) are widely used for spatial modeling [14], [15], [16], their computational cost and sensitivity to redundant data limit scalability [17]. Recent sampling strategies have tackled this by retaining only informative data points [17], or decaying older samples to promote adaptiveness [18]. Yet these methods face trade-offs between memory efficiency and process tracking fidelity. Despite recent advances, to the best of the authors' knowledge, no existing method simultaneously integrates distributed ergodic control, responsiveness to time-varying spatial processes, and collision-aware motion in complex, non-convex environments. The only notable exception is the decentralized ergodic control framework of [19]. However, their approach relies on a pre-trained GP to model the spatial process, without explicitly leveraging GP uncertainty, and is tailored to target localization rather than coverage control. As such, it does not address coverage and estimation of time-varying processes, and therefore cannot serve

as a meaningful baseline for our problem. Other works have investigated multi-agent ergodic trajectory optimization in the context of exploration with time-varying sensor visibility [20], but these approaches likewise do not provide a solution to the persistent coverage problem.

To overcome these limitations, we propose a fully distributed framework for exploration and monitoring of unknown, time-varying spatial processes in complex environments using multi-robot teams with limited sensing and communication. Building on HEDAC [12] and recent time-varying GP models [18], our approach enables robots to collaboratively reconstruct and track dynamic spatial fields while continuously adapting to environmental changes. Our key contributions are:

- A fully distributed ergodic coverage framework for multi-robot exploration and adaptive monitoring of unknown, time-varying spatial processes;
- A novel ergodic metric based on a dynamic goal density function that accounts for both exploration and coverage of key areas in the environment.
- A novel temporal decay mechanism for continuous adaptation to evolving spatial phenomena;
- Extensive simulations and real-world experiments demonstrating robustness, scalability, and effectiveness in complex, non-convex environments.

II. BACKGROUND & NOTATION

We denote by $\mathbb{N}$, $\mathbb{R}$, $\mathbb{R}_{\geq 0}$, and $\mathbb{R}_{> 0}$ the set of natural, real, real non-negative, and real positive numbers. Let us define with $\|\cdot\|_n$ the n -dimensional norm. We define with $\mathbf{R}(\theta)$ the rotation matrix that aligns the local coordinate system according to a direction angle θ . Although real-world coverage problems are typically two-dimensional, the proposed method can be extended to n -dimensional problems. For clarity, we consider a two-dimensional domain $\Omega \subset \mathbb{R}^2$ with a Lipschitz continuous boundary. An arbitrary point in the domain is described with $\mathbf{x} \in \Omega$. An integrable probability density function $\phi : \Omega \rightarrow \mathbb{R}_{\geq 0}$ is defined to represent the areas of the environment Ω of highest relevance, such that $\int_{\Omega} \phi(\mathbf{x}) d\mathbf{x} = 1$. Let us define the robots' positions with $\mathbf{z}_i \in \Omega$ and their trajectories with $\Gamma_i : [0, t] \rightarrow \mathbb{R}^2$, for $i = 1, \dots, R$, where R is the total number of robots. The neighbor set of robot i within radius r is denoted by $N_{i,r}$.

A. Heat-Driven Area Coverage Framework (HEDAC)

In ergodic control formulations, the ergodic metric is a fundamental concept, typically defined as the difference between the empirical distribution of the robots' trajectories Γ_i , and the probability distribution defined onto the environment $\phi(\mathbf{x})$. In the original HEDAC formulation [12], the empirical distribution of the robots' trajectories is approximated using Radial Basis Functions (RBFs). This approach smooths the trajectory representation, replacing the standard SMC formulation, which traditionally represents the empirical distribution Γ_i using Dirac delta functions. The use of RBFs provides a more continuous and analytically manageable representation, enhancing the stability and interpretability of the method.

Therefore, the scalar field representing the spatial distribution of the error is defined as

$$e(\mathbf{x}, t) = c_{\phi}(\mathbf{x}) - c_{\gamma}(\mathbf{x}, t), \quad (1)$$

where $c_{\phi}(\mathbf{x})$ represents the coverage density from the robots' trajectories and $c_{\gamma}(\mathbf{x}, t)$ is the smoothed goal density. Finally, the ergodic metric is given by the L_2 norm of the error $E(t) = \|e(\mathbf{x}, t)\|_2$, which we aim to minimize over time, so that the system is ergodic, that is, $\lim_{t \rightarrow \infty} E(t) = 0$. In the HEDAC framework, a stationary heat equation is used to form a potential field to guide robots toward target areas, supporting ergodic coverage in non-convex domains with obstacles [13]. The solution of the heat equation $u(\mathbf{x}, t)$ is used to compute the robot input. We refer the reader to [12] for full details.

B. Robot Model & Control

We model each robot as a point mass with second-order planar dynamics and a limited field of view (FoV) defined by depth d_{fov} and angle θ_{fov} , capturing typical sensing constraints such as those from onboard cameras [21], [22].

The i -th robot state at time t is

$$\mathbf{x}_i(t) = \begin{bmatrix} \mathbf{z}_i(t) \\ \theta_i(t) \end{bmatrix}, \quad (2)$$

where $\mathbf{z}_i(t)$ is the planar position and $\theta_i(t) \in [0, 2\pi)$ is the heading angle. Given a known field $u(\mathbf{x}, t)$ from the heat equation, we define the desired velocity as the normalized gradient of the field $u(\mathbf{x}, t)$ evaluated for $\mathbf{x} = \mathbf{z}_i$, scaled by a proportional gain $k_v > 0$,

$$\mathbf{v}_d(\mathbf{z}_i, t) = k_v \frac{\nabla u(\mathbf{z}_i, t)}{\|\nabla u(\mathbf{z}_i, t)\|_2}. \quad (3)$$

Normalizing the field gradient enforces constant-speed motion and avoids stalls from eventual vanishing accelerations, provided $\nabla u \neq 0$. The acceleration dynamics of the i -th robot are

$$\ddot{\mathbf{x}}_i(t) = \begin{bmatrix} \ddot{\mathbf{z}}_i(t) \\ \ddot{\theta}_i(t) \end{bmatrix} = \begin{bmatrix} \min(k_p(\mathbf{v}_d(\mathbf{z}_i, t) - \dot{\mathbf{z}}_i(t)), a_{\max}) \\ \alpha_i(\theta_i, t) \end{bmatrix}, \quad (4)$$

where $k_p > 0$ controls acceleration responsiveness. The angular acceleration $\alpha_i(\theta, t)$ is a PD controller on heading error

$$\alpha_i(\theta_i, t) = \text{sat}_{\alpha_{\max}} \left(k_{\theta} e_{\theta_i} + d_{\theta} \frac{de_{\theta_i}}{dt} \right), \quad (5)$$

with heading error $e_{\theta_i} = \theta_{d_i} - \theta_i$, where the desired heading θ_{d_i} aligns with the potential field gradient. Gains $k_{\theta} > 0$, $d_{\theta} \geq 0$ regulate convergence, and $\text{sat}_{\alpha_{\max}}(\cdot)$ bounds angular acceleration to $\pm \alpha_{\max}$. The proposed strategy is model-agnostic, as its vector field can be tracked by many systems using suitable low-level controllers.

C. Gaussian Processes

GPs provide a principled framework for inferring latent functions from sparse data while quantifying predictive uncertainty [23], [24]. In this work, we consider an underlying

task-dependent spatial field $\phi(\mathbf{x})$ representing a quantity of interest distributed over the environment and sensed by the robots (e.g., temperature, gas concentration, or a vision-based detection likelihood). We model this field as $\phi(\mathbf{x}) \sim \text{GP}(\mu(\mathbf{x}), \Sigma(\mathbf{x}, \mathbf{x}'))$, where $\mu(\mathbf{x})$ is the mean and $\Sigma(\mathbf{x}, \mathbf{x}')$ is the covariance function. Assuming spatial smoothness, we use the squared exponential kernel

$$\kappa_{\text{SE}}(\mathbf{x}, \mathbf{x}') = \sigma_c^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{2l^2}\right), \quad (6)$$

with σ_c and l controlling variance and correlation length, respectively. Observations are noisy measurements $y(\mathbf{x}) = \phi(\mathbf{x}) + \mathcal{N}(0, \sigma_n^2)$, where σ_n^2 is the sensor noise. This observation model abstracts the sensing process by assuming pointwise access to the field; more complex or indirect sensor models can be incorporated through an appropriate measurement mapping. Let $\mathcal{D}_t = \{\mathbf{X}_t, \mathbf{y}_t, \mathbf{t}_t\}$ be the dataset of N samples collected up to time t , with inputs $\mathbf{X}_t = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, outputs $\mathbf{y}_t = \{y_1, \dots, y_N\}$, and timestamps $\mathbf{t}_t = \{t_1, \dots, t_N\}$.

III. PROBLEM STATEMENT

We consider the problem of enabling a team of mobile robots to autonomously explore, learn, and monitor unknown time-varying spatial processes in complex, non-convex environments. The primary challenges are:

- 1) Autonomous exploration and reconstruction of a dynamic spatial process.
- 2) Persistent monitoring of high-interest regions, balancing exploration and coverage in real time.
- 3) Fully distributed operation, with local coordination and collision avoidance guarantees to improve coverage efficiency.

The problem is grounded in the following constraints:

- *Limited sensing*: each robot perceives the environment only through its onboard sensors (e.g., camera, LiDAR), resulting in a narrow, camera-like field of view and sparse observations.
- *Limited communication*: communication is modeled as range-limited, and only robots within range may share poses and local dataset samples, while those outside the range cannot exchange information.
- *Sparse data acquisition*: due to the limited FoV, each robot collects only partial, spatially sparse observations at each timestep.

Problem 1: Design a fully distributed strategy that enables a multi-robot team to explore and estimate an unknown, time-varying spatial process in a complex environment, while continuously balancing exploration and perpetual coverage of key areas.

IV. DISTRIBUTED ERGODIC COVERAGE CONTROL

In this section, we introduce our proposed fully distributed ergodic coverage control approach within the HEDAC framework from [12]. An overview of the proposed closed-loop system architecture is represented in Fig. 2.




Fig. 2. System architecture. Blue blocks denote control and dynamical components, orange blocks denote learning-based modules, and gray blocks denote data exchange and filtering.

A. Distributed Gaussian Process

Unlike prior ergodic coverage methods that assume full knowledge of a static spatial process, we employ time-varying GPs to enable distributed exploration and persistent monitoring of dynamic environments. Each robot maintains its own local GP and collaborates with neighbors through limited dataset sharing.

Following [18] and the formulation in Section II-C, predictions are computed at target locations $\mathbf{X}^* = \{\mathbf{x}_1^*, \dots, \mathbf{x}_M^*\}$ using spatio-temporal decay, such that older or spatially distant samples contribute less to the posterior $(\phi^* | \mathcal{D}_t) \sim \mathcal{N}(\mu(\mathbf{X}^*), \Sigma(\mathbf{X}^*))$, with

$$\begin{aligned} \mu(\mathbf{X}^*) &= \tilde{\kappa}(\mathbf{X}^*, \mathbf{X}_t)^\top [\tilde{\kappa}(\mathbf{X}_t, \mathbf{X}_t) + \sigma_n^2 \mathbf{I}]^{-1} \mathbf{y}_t, \\ \Sigma(\mathbf{X}^*) &= \kappa(\mathbf{X}^*, \mathbf{X}^*) - \tilde{\kappa}(\mathbf{X}^*, \mathbf{X}_t)^\top \\ &\quad \cdot [\tilde{\kappa}(\mathbf{X}_t, \mathbf{X}_t) + \sigma_n^2 \mathbf{I}]^{-1} \tilde{\kappa}(\mathbf{X}^*, \mathbf{X}_t). \end{aligned} \quad (7)$$

The decayed covariance matrices are

$$\begin{aligned} \tilde{\kappa}(\mathbf{X}_t, \mathbf{X}_t) &= \kappa_{\text{SE}}(\mathbf{X}_t, \mathbf{X}_t) \odot \mathbf{D}_t \odot \mathbf{T}^D, \\ \tilde{\kappa}(\mathbf{X}^*, \mathbf{X}_t) &= \kappa_{\text{SE}}(\mathbf{X}^*, \mathbf{X}_t) \odot \mathbf{d}_t^\top \odot \mathbf{T}^d, \end{aligned} \quad (8)$$

with $\odot$ denoting element-wise multiplication. The temporal decay components are

$$\begin{aligned} [\mathbf{T}^d]_i &= T_\tau(\Delta t_i), \quad \Delta t_i = t - t_i, \\ [\mathbf{T}^D]_{ij} &= \begin{cases} [\mathbf{T}^d]_i \cdot [\mathbf{T}^d]_j & \text{if } i \neq j, \\ 1 & \text{otherwise.} \end{cases} \end{aligned} \quad (9)$$

As in [18], an exponential decay is used, defined as $T_\tau(\Delta t) = \exp(-\Delta t/\lambda_t)$, where λ_t controls how quickly information loses relevance. However, the spatial decay in [18] is based on sample index proximity rather than true spatial distance, causing inconsistencies, i.e., samples that are far in space but close in index may be incorrectly treated as nearby. Alternative approaches implement decay by augmenting the noise variance model with a time-dependent term that grows over time [25]. However, this mechanism does not account for the spatial correlations

between samples and relies solely on a heuristic aging scheme. To address this, we define spatial decay directly in terms of Euclidean distance

$$\begin{aligned} [\mathbf{D}_t]_{ij} &= \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|_2^2 / \lambda_s), \\ [\mathbf{d}_t]_i &= \exp(-\|\mathbf{x}_i - \mathbf{z}_i\|_2^2 / \lambda_s), \end{aligned} \quad (10)$$

where $\mathbf{x}_i, \mathbf{x}_j \in \mathbf{X}_t$ are data samples, $\mathbf{z}_i$ is the current robot position, and λ_s sets the spatial decay rate. By embedding spatial and temporal decay in the kernel, measurement influence decreases as data become older or farther from the robot's current location. The GP therefore assigns higher uncertainty to predictions based on outdated or distant data. The decay parameters λ_t and λ_s control how quickly information loses relevance over time and space, providing a principled way to downweight past measurements while preserving spatial correlations.

To manage the $\mathcal{O}(N^3)$ computational cost of GPs, we adopt threshold-based data filtering inspired by [17], [18]. New samples are added when their predictive standard deviation satisfies $\sigma_{u,i}(\mathbf{x}') \geq \sigma_{\text{in}}$, with $\sigma_{u,i}(\mathbf{x}) = \sqrt{\Sigma_i(\mathbf{x})}$, while existing samples are removed once $\sigma(\mathbf{x}_t) \geq \sigma_{\text{out}}$, with thresholds defined as $\sigma_{\text{in}} = e_{\text{in}}/z$ and $\sigma_{\text{out}} = e_{\text{out}}/z$, and $e_{\text{in}} > e_{\text{out}}$ to prevent premature data removal. This selective data management bounds computation and communication costs while preserving adaptability in dynamic environments. For additional information, readers are referred to [17], [18].

B. Dynamic Goal Density

We introduce a novel dynamic goal density that adaptively integrates information from the GP, facilitating both exploration and targeted monitoring of the environment. From each robot's GP_i , the estimated spatial process $\mu_i(\mathbf{x})$ and the corresponding uncertainty $\sigma_{u,i}(\mathbf{x}) = \sqrt{\Sigma_i(\mathbf{x})}$ are computed, $\forall \mathbf{x} \in \Omega$, from (7). This uncertainty representation effectively distinguishes between explored and unexplored regions, ensuring an informed and adaptive coverage strategy. For every robot, we then combine these two components into a straightforward unified spatial distribution, $\tilde{F}_i(\mathbf{x})$, which highlights both areas that require monitoring and those that need further exploration. This distribution is defined as

$$\tilde{F}_i(\mathbf{x}) = (\exp(\mu_i(\mathbf{x})) - 1) + (\exp(\sigma_{u,i}(\mathbf{x})) - 1). \quad (11)$$

This formulation directs robots to prioritize exploration when the predictive uncertainty $\sigma_{u,i}(\mathbf{x})$ exceeds the mean estimate $\mu_i(\mathbf{x})$, while favoring the monitoring of high-priority regions otherwise. To eliminate contributions when $\mu_i(\mathbf{x}) \approx 0$ or $\sigma_{u,i}(\mathbf{x}) \approx 0$, we offset both terms by subtracting 1 from $\exp(\mu_i(\mathbf{x}))$ and $\exp(\sigma_{u,i}(\mathbf{x}))$. The resulting normalized unified spatial distribution is then defined as

$$F_i(\mathbf{x}) = \frac{\tilde{F}_i(\mathbf{x})}{\int_{\Omega} \tilde{F}_i(\mathbf{x}) d\mathbf{x}}. \quad (12)$$

Here, $\mu_i(\mathbf{x})$ and $\sigma_{u,i}(\mathbf{x})$ are normalized to ensure that both mean and uncertainty contribute proportionally to the final distribution. This proposed combination effectively highlights areas that require both monitoring and further exploration, even in complex environments. At each time step, the goal density is

updated dynamically as the robot refines its knowledge of the evolving spatial process through the time-varying GP, either via exploration or by incorporating data shared by its neighbors. This update follows the formulation provided in [12] yielding a novel local dynamic goal density computed by the i -th robot as

$$g_i(\mathbf{x}) = \langle \varphi_{\gamma}, F_i \rangle_{\Omega} = \int_{\Omega} \varphi_{\gamma}(\mathbf{x}) F_i(\mathbf{x}) d\mathbf{x}. \quad (13)$$

Here, $\varphi_{\gamma}(\mathbf{x})$ is an RBF acting as a smoothing with shape γ defined as in [12]. As the robot explores, it continually refines its understanding of the environment through the time-varying GPs that facilitate adaptiveness to the dynamic nature of the underlying monitored process.

C. Distributed Ergodic Control

To enable scalable, fully distributed collaboration, robots must avoid globally sharing their entire coverage histories, which is both unnecessary and computationally inefficient. Since a robot's decisions primarily depend on avoiding redundant coverage with nearby robots, only local information exchange is required. Each robot therefore maintains a local coverage density that includes its own past observations and those of nearby neighbors within a communication range r . This localized strategy significantly reduces communication overhead and acts as a spatiotemporal memory horizon, promoting coordination while relaxing bandwidth demands.

We define the local coverage density of robot i as

$$c_i(\mathbf{x}, t) = \frac{1}{t} \int_0^t p_i(\mathbf{r}_i^{\tau}) d\tau, \quad (14)$$

where $p_i(\mathbf{r}_i^{\tau})$ is the instantaneous sensing footprint centered at $\mathbf{r}_i^{\tau} = \mathbf{R}(\theta_i(\tau))(\mathbf{z}_i(\tau) - \mathbf{x}) \in \mathbb{R}^2$, where $\mathbf{x}$ is the point of interest. For simplicity, we assume all robots are equipped with identical camera-like sensors, leaving the consideration of heterogeneous capabilities for future work. Furthermore, at time t , we define the sensor detection function as an RBF centered within the sensor's FoV, such that

$$p_i(\mathbf{r}_i^t) = \exp(-(\epsilon \mathbf{r}_i^t)^2), \quad (15)$$

here, $\epsilon > 0$ controls the shape of the RBF and is a tunable parameter. The explicit inclusion of time in the notation accounts for dynamic changes in the detection function. In complex environments, external factors such as obstacles can partially occlude the sensor's FoV, altering its effective coverage.

To incorporate neighbor information, each robot aggregates only the spatially relevant contributions using a weighting function $w_r(\mathbf{x}, \mathbf{z}_i(t))$ centered at its current location $\mathbf{z}_i(t)$. This yields the distributed coverage density

$$\tilde{c}_i(\mathbf{x}, t) = c_i(\mathbf{x}, t) + \sum_{j \in N_{i,r}} w_r(\mathbf{x}, \mathbf{z}_j(t)) c_j(\mathbf{x}, t), \quad (16)$$

where $N_{i,r}$ denotes the set of neighbors within range r , and w_r may be chosen as a hard indicator or a smooth decay kernel (e.g., Gaussian) to enforce locality with a support radius r that is the communication range. Building on this, similar to (1), we propose a new distributed error field that is used by each robot

to guide its trajectory, defined as

$$\tilde{e}_i(\mathbf{x}, t) = g_i(\mathbf{x}) - \tilde{c}_i(\mathbf{x}, t), \quad (17)$$

where $g_i(\mathbf{x})$ is the dynamic goal density computed as in (13). As in [12], the distributed ergodic metric minimized by each robot according to (17) is $\mathcal{E}_i = \|\tilde{e}_i(\mathbf{x}, t)\|_2$. This formulation enables distributed operation through neighbor-limited information sharing.

This newly defined ergodic metric serves as the basis for the local source term that each robot independently computes to solve its own heat equation. Following [12], we define the local source for each robot as

$$s_i(\mathbf{x}, t) = \begin{cases} \tilde{e}_i(\mathbf{x}, t)^2 & \text{if } \tilde{e}_i(\mathbf{x}, t) \geq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (18)$$

In our formulation, each robot autonomously solves a local heat equation using only its own data and information shared by nearby neighbors. Building on the framework from [12], the distributed heat equation is defined as

$$\begin{aligned} \rho \Delta u(\mathbf{x}, t) &= \beta u(\mathbf{x}, t) - s_i(\mathbf{x}, t), \\ \frac{\partial u}{\partial \mathbf{n}} &= 0, \quad \text{on } \delta\Omega, \end{aligned} \quad (19)$$

where $\mathbf{n}$ is the outward unit normal on the boundary $\delta\Omega$. The term $\rho \Delta u(\mathbf{x}, t)$ governs spatial diffusion, with larger ρ producing wider influence. Domain-wide decay is controlled by $\beta u(\mathbf{x}, t)$, where $\beta = \beta_0/|\Omega|$ ensures scale invariance.

Unlike [12], which relies on localized cooling to reduce robot overlap without enforcing guarantees, we explicitly integrate a potential-field-based collision avoidance mechanism, as introduced in [26], to maintain a minimum safety distance r_s between robots. The method in [26] provides theoretical guarantees of maintaining the minimum safety distance, under the mild assumption that the initial deployment does not violate it. This ensures safe, distributed coverage while preserving the ergodic exploration behavior.

D. Stability Analysis

We study the stability of the distributed scheme induced by the dynamic goal density (13). For smoothly time-varying goal densities with bounded temporal derivative, the ergodic error remains uniformly bounded, while in the static case asymptotic convergence is recovered as in classical HEDAC. Instantaneous changes in the spatial process induce discontinuities in the Lyapunov function; however, between such events the closed-loop system either exhibits exponential convergence (in static intervals) or remains uniformly bounded (under smooth variations).

Following classic Lyapunov stability theory, let

$$V_i(t) = \frac{1}{2} \|g_i(x, t) - \tilde{c}_i(x, t)\|_2^2 = \frac{1}{2} \int \tilde{e}_i(x, t)^2 dx \quad (20)$$

be the Lyapunov functional associated with the distributed ergodic error of robot i . In general, we note that $V_i(t) \geq 0$ for all t , and $V_i(t) = 0$ if and only if $\tilde{c}_i(x, t) = g_i(x, t)$.

We consider the more general case where the spatial process evolves smoothly in time. Here, asymptotic convergence cannot

be guaranteed because the target density is changing; instead, we derive a uniform ultimate bound that depends on the process's rate of change.

Theorem 1 (Bounded tracking for smoothly varying processes): Consider a multi-robot system composed of R robots, where each robot $i \in \{1, \dots, R\}$ evolves according to the dynamics in (4) and applies the control input defined in (3). Each robot maintains its own local dynamic goal density g_i , which it updates based on its individual GP as well as information received from neighboring robots. Then, the local ergodic error remains bounded with an asymptotic tight bound proportional to $D_{\max}$.

Proof: To prove the statement, we show that the local Lyapunov functional $V_i(t)$ introduced in (20) is uniformly ultimately bounded, and, in particular, $\lim_{t \rightarrow \infty} V_i(t) \leq D_{\max}^2/2\alpha^2$. When $\partial g_i/\partial t \neq 0$, the Lyapunov derivative includes the nominal HEDAC term and a disturbance term:

$$\dot{V}_i(t) = -2\alpha V_i(t) + \int_{\Omega} \tilde{e}_i(x, t) \frac{\partial g_i}{\partial t}(x, t) dx.$$

Applying the Cauchy–Schwarz inequality gives

$$\int_{\Omega} \tilde{e}_i(x, t) \frac{\partial g_i}{\partial t}(x, t) dx \leq \|\tilde{e}_i(\cdot, t)\|_2 \left\| \frac{\partial g_i}{\partial t}(\cdot, t) \right\|_2,$$

and since $\|\tilde{e}_i(\cdot, t)\|_2 = \sqrt{2V_i(t)}$, define

$$D_i(t) := \left\| \frac{\partial g_i}{\partial t}(\cdot, t) \right\|_2.$$

To simplify the inequality, set $W_i(t) = \sqrt{V_i(t)}$, so $V_i = W_i^2$ and

$$\dot{W}_i(t) = \frac{\dot{V}_i(t)}{2\sqrt{V_i(t)}} \leq -\alpha W_i(t) + \frac{1}{\sqrt{2}} D_i(t). \quad (i)$$

Assume the temporal derivative of the spatial process is uniformly bounded, i.e., $D_i(t) \leq D_{\max}$. Consider the linear comparison system

$$\dot{y}(t) = -\alpha y(t) + \frac{1}{\sqrt{2}} D_{\max}, \quad y(0) = W_i(0),$$

whose solution is

$$y(t) = e^{-\alpha t} W_i(0) + \frac{1}{\alpha\sqrt{2}} D_{\max} (1 - e^{-\alpha t}).$$

Since $W_i(t)$ satisfies the inequality (i) and $y(t)$ satisfies the corresponding equality, the comparison lemma ensures $W_i(t) \leq y(t)$ for all $t \geq 0$. Hence,

$$W_i(t) \leq e^{-\alpha t} W_i(0) + \frac{1}{\alpha\sqrt{2}} D_{\max} (1 - e^{-\alpha t}).$$

Taking the limit as $t \rightarrow \infty$ yields

$$\lim_{t \rightarrow \infty} W_i(t) \leq \frac{D_{\max}}{\alpha\sqrt{2}}.$$

Since $V_i(t) = W_i(t)^2$, the Lyapunov function, and thus the ergodic error, is bounded as follows

$$\lim_{t \rightarrow \infty} V_i(t) \leq \frac{D_{\max}^2}{2\alpha^2},$$

TABLE I
SUMMARY OF PARAMETERS USED IN THE EXPERIMENTS

Parameter	Description	Value
<i>Robot sensing and communication</i>		
d_{fov}	Sensor field-of-view depth	5 m
θ_{fov}	Sensor field-of-view angle	90°
r	Communication range	5 m
r_s	Minimum safety separation	1 m
<i>Robot dynamics and control</i>		
k_p	Linear velocity proportional gain	3
k_θ	Heading proportional gain	1
d_θ	Heading derivative gain	$2\sqrt{k_\theta}$
a_{max}	Max linear acceleration	2.5 m/s
α_{max}	Max angular acceleration	2.5 rad/s ²
<i>Heat equation and coverage parameters</i>		
γ	RBF shape parameter (goal density smoothing)	$1/d_{fov}$
ϵ	RBF shape parameter (sensor footprint)	$1/d_{fov}$
ρ	Heat diffusion coefficient	1
β	Heat equation decay coefficient	1
<i>Gaussian Process parameters</i>		
$\lambda_{s,t}$	Decay length-scale	100
σ_n	Measurement noise std. deviation	0.01
e_{in}	GP inclusion threshold	0.90
e_{out}	GP removal threshold	0.70

and this result completes the proof. $\square$

This result establishes a formal tracking guarantee: slowly varying environments lead to small steady-state errors, while faster variations result in larger, but still finite, errors. The static case arises as a special instance of the theorem with $D_{max} = 0$, for which the bound collapses to zero, thereby recovering the asymptotic convergence guarantees of the classical HEDAC framework.

Overall, the analysis shows that the proposed distributed controller preserves classical HEDAC stability in static environments, provides bounded tracking performance under smooth temporal variations, and remains well behaved between instantaneous changes, while a full global characterization under arbitrary jump sequences is left for future work.

V. EXPERIMENTAL VALIDATION

In this section, we conduct a comprehensive set of experiments to evaluate the proposed approach. Simulations were conducted on a PC equipped with an *Intel Core i7* CPU, an *NVIDIA GeForce RTX 4080* GPU, and 16 GB of RAM.

The experiments consider key parameters, including those governing the heat equation, the time-varying GPs used by the robots, and their intrinsic capabilities. Each robot's GP parameters are updated online by fitting the streaming data collected as the robots explore and exchange data, rather than being fixed in advance. We fix all robot-level and heat-equation parameters as reported in Table I, to ensure a consistent comparison across experiments. The parameters in Table I combine standard practice with task-specific intuition. Robot-level parameters (sensing and communication ranges, controller gains, dynamic limits) are set to reflect typical platforms and ensure smooth, stable coverage. HEDAC constants follow common choices in ergodic-control work, yielding a well-conditioned balance between coverage and trajectory smoothness. GP model parameters follow prior literature [17], [18]: spatial and temporal decay rates encode assumptions on process evolution, while the pruning thresholds e_{in} and e_{out} regulate online updates to retain




Fig. 3. Evolution of the ergodic metric over time in a static spatial process. Four robots, each running the proposed distributed strategy, monitor the environment with randomized initial positions and key area locations. The black line shows the oracle baseline from [12], where all robots share a centralized ergodic metric and immediately converge to the key area, without exploration or sustained monitoring. The proposed method, using a time-varying GP and dynamic goal density, balances exploration (higher metric) and focused coverage (lower metric) for effective long-term monitoring.

only informative measurements. These settings were robust across all experiments, but reducing manual tuning and deriving principled guidelines, especially for GP filtering and HEDAC weights, remains future work, including establishing theoretical guarantees on boundedness, convergence, and stability of the GP estimates.

We evaluate the proposed method in two scenarios. First, we consider a static spatial process, where baseline approaches typically converge to key regions and cease exploration. In contrast, our strategy is designed to sustain an explore-cover-repeat behavior, reflected by oscillations in the ergodic metric. Second, we assess adaptiveness by comparing against the oracle HEDAC framework [12] in a scenario where the spatial process undergoes an abrupt change in location, requiring renewed exploration. In both cases, the region of interest is modeled as a single Gaussian peak in a complex environment.

This setup highlights the distinguishing behavior of the proposed distributed strategy, which explicitly accounts for sensing constraints such as limited range and restricted field of view, unlike the oracle baseline. Finally, real-world experiments validate the applicability of the approach under practical sensing and operational conditions.

A. Simulations: Static Process

We first evaluate the proposed approach in a baseline scenario with a static spatial process. A team of four robots applies the distributed strategy using time-varying GPs and dynamic goal densities to monitor a static Gaussian field. Fig. 3 reports the mean and standard deviation of the ergodic metric for each robot over multiple simulation runs, along with the HEDAC baseline [12]. Initial robot positions and the Gaussian peak location are randomized across runs.

Although each robot optimizes a modified ergodic metric with respect to its own dynamic goal density (17), performance is evaluated against the ground-truth spatial distribution to enable consistent comparison across methods. This metric quantifies how effectively robots cover the key regions of interest.




Fig. 4. A simulation run showing one robot's behavior while monitoring a spatial process represented as a heatmap (color denotes the value of the true underlying process, with red indicating higher intensity). The robot's trajectory is color-coded by its instantaneous ergodic metric; brighter colors indicate lower metric values (i.e., proximity to key regions), while darker colors denote exploratory phases. Black circles denote initial positions and blue dots final positions; other robots' trajectories are shown in gray. This illustrates the robot's ability to dynamically alternate between exploration and focused coverage.

As shown in Fig. 3, the method exhibits the expected explore-repeat behavior: lower ergodic values correspond to focused coverage of the key region, while higher values indicate exploratory motion. Fig. 4 illustrates a representative run, showing the environment, the spatial process, and the trajectory of a single robot color-coded by its instantaneous ergodic metric.

The visualization highlights the adaptive nature of the approach: the ergodic metric increases during exploration of less-visited areas and decreases as the robot returns to the key region. The GP decay mechanism enables gradual forgetting of past observations, promoting re-exploration even in this static scenario.

B. Simulations: Time-Varying Process

To evaluate our distributed strategy in a dynamic setting, we consider a challenging scenario where the spatial process abruptly shifts location within a complex environment. A team of four robots must explore, learn, and monitor the process using our method. To increase difficulty, the shift occurs instantaneously at a predefined time step.

Fig. 5 shows the ergodic metric (mean and standard deviation) over time across multiple runs, with robot trajectories color-coded. Since distributed ergodic control in time-varying scenarios remains largely unexplored, we compare our approach against the HEDAC framework [12] (purple line). At time step 5000, marked by the dashed black line, the spatial process changes location. Our distributed strategy, leveraging time-varying GPs and dynamic goal densities, maintains stable ergodic values and rapidly adapts to the new target. The decay mechanism fosters continuous exploration, enabling fast detection, coverage redirection, and efficient learning of the updated process. By contrast, while the oracle achieves lower initial ergodic values through a globally shared metric, it lacks adaptability: after the process shifts, coverage remains biased toward outdated regions, requiring substantial re-exploration to adapt. Fig. 6 illustrates a representative run, focusing on one robot's trajectory and the corresponding ergodic metric before and after the process change. Initially, the robot balances exploration with




Fig. 5. Ergodic metric over time during multiple simulation runs, where a team of four robots monitors a time-varying spatial process in a complex environment using the proposed distributed strategy. Each robot's metric is shown in a different color, while the purple line represents the HEDAC method [12]. At time step 5000 (dashed black line), the spatial process abruptly changes location. Our method enables adaptive exploration and coverage through time-varying GPs and dynamic goal densities, while the centralized baseline exhibits lower initial metrics but lacks responsiveness to process changes.




Fig. 6. Simulation run showing one robot's behavior before and after a sudden change in the spatial process (visualized as a heatmap; color indicates the process intensity, with red denoting higher values). Each subplot displays 5000 time steps. The robot's trajectory is color-coded by its instantaneous ergodic metric (bright colors correspond to lower metric values). Black circles mark initial positions, red dots indicate robot positions at the moment of the process change, and blue dots denote final positions; other robots' trajectories appear in gray.

monitoring of the first key area (red). After the shift at time step 5000, it detects the change, explores, and reallocates coverage to the new key region. The ergodic metric highlights periods when the robot effectively monitors the relevant area, both before and after the shift.

C. Real-World Experiments

We validate the proposed distributed strategy in real-world experiments using three *Uvify IFO-S* drones, configured as in simulation. Each drone runs the full algorithm onboard on an

Nvidia Jetson Nano (4 GB) and interfaces with a *PX4* flight controller. The ergodic controller computes acceleration commands that are internally integrated into velocity setpoints and sent to *PX4* for low-level stabilization and constraint handling. Robots communicate over WiFi as independent ROS nodes, exchanging only minimal information (positions and selected dataset samples) with nearby neighbors. A ground station hosts the *ROS core* and *QGroundControl* for supervision and data logging, but performs no inference, control, or coordination.

Each robot generates onboard synthetic scalar measurements to validate the distributed inference-control loop independently of perception. Despite real-world effects such as GPS inaccuracies, IMU drift, estimator latency, and WiFi packet loss, the system remains robust due to its low-bandwidth, neighbor-only communication. Integrating real perception with realistic FoV and sensing-noise models is left for future work.

Fig. 1 shows a representative snapshot, with a full recording provided in the accompanying video.

VI. CONCLUSIONS AND FUTURE WORKS

In conclusion, we presented a distributed multi-robot framework for autonomous exploration and adaptive monitoring of time-varying spatial processes in complex environments. Each robot maintains an online time-varying Gaussian Process model from local measurements and selectively shares information with nearby teammates to improve inference without centralized coordination. A dynamic goal density and distributed coverage mechanism balance exploration and focused monitoring, while GP temporal decay ensures persistent revisiting of critical regions. Safe and coordinated motion is achieved through a heat-equation-based controller augmented with minimum-distance constraints.

Several limitations motivate future work. First, both the HEDAC-based coverage and time-varying GP filtering rely on parameters that could benefit from more principled, less parametric designs to reduce tuning and improve robustness. Second, while local stability guarantees are provided for instantaneous changes in the spatial process, a full global analysis under rapid or repeated jumps requires hybrid-systems tools such as dwell-time conditions. Third, explicitly accounting for bandwidth constraints through communication-aware strategies is essential for large-scale deployments. Finally, extending the hardware validation to larger robot teams and integrating real perception with realistic FoV and sensing-noise models are key steps toward practical real-world monitoring applications.

REFERENCES

- [1] J. Cortes, S. Martinez, and F. Bullo, "Spatially-distributed coverage optimization and control with limited-range interactions," *ESAIM: Control, Optimisation Calculus Variations*, vol. 11, no. 4, pp. 691–719, Oct. 2005.
- [2] F. Pratioli, B. Capelli, and L. Sabattini, "On coverage control for limited range multi-robot systems," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 9957–9963.
- [3] A. Carron and M. N. Zeilinger, "Model predictive coverage control," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 6107–6112, 2020.
- [4] M. Catellani, M. Gabbi, and L. Sabattini, "HMPCC: Human-aware model predictive coverage control," in *Proc. IEEE Int. Symp. Multi-Robot Multi-Agent Syst.*, Singapore, 2025, pp. 1–7.
- [5] W. Gosrich et al., "Coverage control in multi-robot systems via graph neural networks," in *Proc. Int. Conf. Robot. Automat.*, 2022, pp. 8787–8793.
- [6] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D: Nonlinear Phenomena*, vol. 240, no. 4–5, pp. 432–442, 2011.
- [7] L. M. Miller and T. D. Murphey, "Trajectory optimization for continuous ergodic exploration," in *Proc. Amer. Control Conf.*, 2013, pp. 4196–4201.
- [8] M. Sun, A. Gaggari, P. Trautman, and T. Murphey, "Fast ergodic search with kernel functions," *IEEE Trans. Robot.*, vol. 41, pp. 1841–1860, 2025.
- [9] D. E. Dong, I. Berger, and Abraham, "Time-optimal ergodic search: Multiscale coverage in minimum time," *Int. J. Robot. Res.*, Sage Publications: London, England, vol. 44, no. 10–11, 2025, pp. 1664–1683.
- [10] E. Ayvaci, H. Salman, and H. Choset, "Ergodic coverage in constrained environments using stochastic trajectory optimization," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2017, pp. 5204–5210.
- [11] C. Lerch, D. Dong, and I. Abraham, "Safety-critical ergodic exploration in cluttered environments via control barrier functions," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 10205–10211.
- [12] S. Ivić, B. Crnković, and I. Mezić, "Ergodicity-based cooperative multi-agent area coverage via a potential field," *IEEE Trans. Cybern.*, vol. 47, no. 8, pp. 1983–1993, Aug. 2017.
- [13] S. Ivić, A. Sikirica, and B. Crnković, "Constrained multi-agent ergodic area surveying control based on finite element approximation of the potential field," *Eng. Appl. Artif. Intell.*, vol. 116, 2022, Art. no. 105441.
- [14] M. Santos, U. Madhushani, A. Benevento, and N. E. Leonard, "Multi-robot learning and coverage of unknown spatial fields," in *Proc. Int. Symp. Multi-Robot Multi-Agent Syst.*, 2021, pp. 137–145.
- [15] K.-C. Ma, L. Liu, and G. S. Sukhatme, "Informative planning and online learning with sparse Gaussian processes," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 4292–4298.
- [16] L. Wei, A. McDonald, and V. Srivastava, "Multi-robot Gaussian process estimation and coverage: Deterministic sequencing algorithm and regret analysis," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 9080–9085.
- [17] M. Mantovani, F. Pratioli, and L. Sabattini, "Distributed coverage control for spatial processes estimation with noisy observations," *IEEE Robot. Automat. Lett.*, vol. 9, no. 5, pp. 4431–4438, May 2024.
- [18] F. Pratioli, M. Mantovani, A. Prorok, and L. Sabattini, "Distributed coverage control for time-varying spatial processes," *IEEE Trans. Robot.*, vol. 41, pp. 1602–1617, 2025.
- [19] I. Abraham and T. D. Murphey, "Decentralized ergodic control: Distribution-driven sensing and exploration for multiagent systems," *IEEE Robot. Automat. Lett.*, vol. 3, no. 4, pp. 2987–2994, Oct. 2018.
- [20] E. Wittemyer, A. Rao, I. Abraham, and H. Choset, "Multi-agent ergodic exploration under smoke-based time-varying sensor visibility constraints," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2025, pp. 5452–5458.
- [21] S. Ivić, A. Andrejč UK, and S. Družeta, "Autonomous control for multi-agent non-uniform spraying," *Appl. Soft Comput.*, vol. 80, pp. 742–760, 2019.
- [22] S. Ivić, "Motion control for autonomous heterogeneous multiagent area search in uncertain conditions," *IEEE Trans. Cybern.*, vol. 52, no. 5, pp. 3123–3135, May 2022.
- [23] C. E. Rasmussen, *Gaussian Processes in Machine Learning*. Berlin, Heidelberg: Springer, 2004, pp. 63–71.
- [24] E. Schulz, M. Speekenbrink, and A. Krause, "A tutorial on Gaussian process regression: Modelling, exploring, and exploiting functions," *J. Math. Psychol.*, vol. 85, pp. 1–16, 2018.
- [25] V. Gómez-Verdejo, E. Parrado-Hernández, and M. Martínez-Ramón, "Adaptive sparse Gaussian process," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 11, pp. 16383–16395, Nov. 2024.
- [26] I. I. Hussein and D. M. Stipanovic, "Effective coverage control for mobile sensor networks with guaranteed collision avoidance," *IEEE Trans. Control Syst. Technol.*, vol. 15, no. 4, pp. 642–657, Jul. 2007.