Preprints of the 23rd IFAC World Congress
Busan, Republic of Korea, August 23-28, 2026

Information Density Matching Driven Optimal Transport Control for UAV Swarm Exploration in Spatiotemporal Fields

Hao Gao [*,***] **Yun Gao** [*,**] **Jiawen Wang** [****] **Siyi Zhou** [*,***]
Shiheng Zhang [*] **Jinni Zhou** [***] **Yiding Ji** [*]

[*] *Robotics and Autonomous Systems Thrust, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China (e-mail: {hgao906, szhou335, szhang593}@connect.hkust-gz.edu.cn; jiyiding@hkust-gz.edu.cn).*
[**] *Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong SAR, China (e-mail: y.gao@gaoyunailab.com).*
[***] *College of Future Technology, Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China (e-mail: eejinni@hkust-gz.edu.cn).*
[****] *School of Astronautics, Harbin Institute of Technology, Harbin 150001, China (e-mail: 2023110488@stu.hit.edu.cn).*

Abstract: Distributed exploration and field reconstruction with UAV swarms remain challenging due to stringent scalability constraints, limited onboard computation, and requirement for uncertainty-driven coordination. We propose an information density matching (IDM) framework where each UAV maintains a sparse Gaussian Process model of the unknown spatiotemporal field and generates a variance-based information density. A distributed optimal-transport control law then drives the swarm toward the target density, enabling them to migrate toward high-uncertainty regions while preserving spatial dispersion. The swarm density forms a Wasserstein gradient flow and a Lyapunov type dissipation result is established under bounded approximation errors. Simulations in static, multi-peak, and dynamic fields demonstrate that our approach achieves fast uncertainty reduction, competitive reconstruction accuracy, stable swarm distribution, and lower communication load compared to several baseline methods.

Keywords: UAV Swarms; Spatiotemporal Field; Information Density Matching; Optimal Transport; Autonomous Exploration.

1. INTRODUCTION

Autonomous coordination of UAV swarms has become a central topic in large-scale environmental monitoring (Javed et al., 2024), aerial exploration (Gao et al., 2025b), and spatiotemporal field estimation (Du et al., 2025). In such tasks, multi-agent systems operate in unstructured environments, acquire local noisy measurements, and infer an evolving field. The mismatch between high-dimensional environmental uncertainty and low-dimensional agent motion introduces challenges beyond classical distributed control (Gao et al., 2025a), which often assumes global observability, analytic field models, or static objectives (Wong et al., 2023). Scalable, uncertainty-aware, and analytically grounded swarm exploration therefore remains an open problem at the intersection of distributed control, information-driven sensing, and statistical learning (Jang et al., 2020; Santilli et al., 2022).

Existing approaches exhibit limitations in uncertainty-dominated spatiotemporal environments. Coverage-based methods such as Voronoi (Gao et al., 2026) and Lloyd formulations offer geometric guarantees but ignore the spatial distribution of estimation uncertainty, resulting in slow information gain and inefficient sensing allocation (Reily and Zhang, 2021). Heuristic and biomimetic strategies are robust to environmental variability but often rely on reactive rules without principled objectives or stability characterizations, which may cause clustering or redundant exploration (Liu et al., 2022; Huang et al., 2025). Learning-based methods address some of these issues but require centralized training or data aggregation and provide limited guarantees on convergence, scalability, or coherence (Chen et al., 2020; Bartolomei et al., 2021). Information-theoretic exploration frameworks (Wu et al., 2023) partially mitigate these challenges, yet many rely on centralized Gaussian Process (GP) models (Viseras et al., 2016) and do not yield tractable distributed control laws.

[⋆] This work is supported by National Natural Science Foundation of China grants 62303389 and 62373289; Guangdong Province Scientific Research Platform and Project Scheme grant 2024KTSCX039; Youth Talent Support Program of Guangdong Association for Science and Technology grant SKXRC2025463 and Guangdong Provincial Key Lab of Integrated Communication, Sensing and Computation for Ubiquitous Internet of Things (No.2023B1212010007). *(Corresponding authors: Jinni Zhou and Yiding Ji.)*

©2026 the authors. Accepted by IFAC for publication under a Creative Commons Licence CC-BY-NC-ND

These limitations indicate a structural gap between distributed control and uncertainty-aware exploration. A comprehensive framework should construct an uncertainty-driven objective from statistical field estimation (Dergachev and Yakovlev, 2025), design a distributed controller compatible with swarm-level density dynamics (Hsu et al., 2021), and provide analytical guarantees under local sensing and communication constraints (Newaz et al., 2023). To address this gap, we develop an information density matching framework in which each UAV maintains a sparse GP approximation and computes a variance-driven target density. A distributed optimal transport (OT) control law then drives the swarm density toward this information density, enabling self-organization toward epistemically uncertain regions while maintaining spatial dispersion. The closed-loop dynamics are interpreted as a Wasserstein gradient flow, and a Lyapunov function is used to show monotonic reduction of density mismatch under mild regularity conditions. Simulations in static, multi-peak, and dynamic fields demonstrate improved uncertainty reduction, reconstruction accuracy, and swarm stability compared with coverage, Lévy-flight (Ramachandran et al., 2020), and chemotaxis-based schemes (Yang et al., 2015).

The contributions of this work are summarized as follows:

(1) We formulate multi-UAV exploration as distributed information-density matching, where the target density is derived from GP posterior variance.
(2) We synthesize a distributed optimal transport controller whose density dynamics form a Wasserstein gradient flow, and establish Lyapunov stability of the resulting closed-loop system.
(3) Simulations in representative spatiotemporal fields demonstrate gains in coverage, density matching, communication efficiency, and cooperative behavior over common exploration baselines.

Notations. For a matrix M, $M^\top$ denotes its transpose. $\|\cdot\|$ is the Euclidean norm, I_n is the identity matrix, and $W_2(\cdot,\cdot)$ denotes the Wasserstein-2 distance. The swarm density, information density, local density estimate, and OT potential are denoted by ρ, $\rho^\star$, $\widehat{\rho}_i$, and ϕ_i, respectively.

2. MODELING AND PROBLEM FORMULATION

2.1 System Modeling

We consider N UAVs deployed in a bounded domain $\mathcal{X} \subset \mathbb{R}^d$, $d \in \{2,3\}$, to explore and reconstruct an unknown spatiotemporal field $f : \mathcal{X} \times \mathbb{R}_{\geq 0} \to \mathbb{R}$. The high-level translational motion of UAV i is modeled as

$$\dot{x}_i(t) = u_i(t) + w_i(t), \tag{1}$$

where $x_i(t) \in \mathcal{X}$ is the position, $u_i(t) \in \mathbb{R}^d$ is the high-level control input, and $w_i(t) \sim \mathcal{N}(0, \Sigma_w)$ is the process disturbance with covariance Σ_w. Each agent obtains

$$y_i(t) = f(x_i(t), t) + v_i(t), \tag{2}$$

where $y_i(t)$ is the scalar field measurement and $v_i(t) \sim \mathcal{N}(0, \sigma_v^2)$ is measurement noise with variance σ_v^2. Agents communicate through a connected undirected graph $\mathcal{G} = (\mathcal{A}, \mathcal{E})$, where $\mathcal{A} = \{1, \ldots, N\}$ is the agent set, $\mathcal{E}$ is the edge set, and $\mathcal{N}_i$ is the neighbor set of agent i.

To avoid ambiguity, we introduce the macroscopic swarm density before using it in the continuum model. For finite N, a smoothed empirical density is represented by

$$\rho_N^K(x,t) = \frac{1}{Z_N} \sum_{j=1}^{N} K_\rho(x - x_j(t)), \tag{3}$$

where $K_\rho(\cdot)$ is a nonnegative density-smoothing kernel and Z_N normalizes ρ_N^K such that $\int_{\mathcal{X}} \rho_N^K(x,t)dx = 1$. In the mean-field limit, ρ_N^K converges weakly to the macroscopic density $\rho(x,t)$. The corresponding continuum model is

$$\frac{\partial \rho}{\partial t}(x,t) + \nabla \cdot (\rho(x,t)u(x,t)) = 0. \tag{4}$$

Here $u(x,t)$ is the continuum velocity field induced by the agent-level controls.

Each UAV maintains a sparse GP posterior with mean $\mu_i(x,t)$ and variance $\sigma_i^2(x,t)$. Let $\sigma^2(x,t)$ denote the ideal centralized posterior variance obtained from all measurements. For any variance field $s(x,t) \in \{\sigma^2(x,t), \sigma_i^2(x,t)\}$, define the normalized information density as

$$\mathcal{R}[s](x,t) = \frac{s(x,t)}{\int_{\mathcal{X}} s(\xi,t)d\xi}. \tag{5}$$

The global and local information densities are then $\rho^\star = \mathcal{R}[\sigma^2]$ and $\rho_i^\star = \mathcal{R}[\sigma_i^2]$, respectively. The control objective is to drive $\rho(x,t)$ toward $\rho^\star(x,t)$ in a distributed manner using the local approximation $\rho_i^\star(x,t)$.

We adopt the following assumptions:

(A1) The communication graph $\mathcal{G}$ is connected.
(A2) The field $f(x,t)$ is continuous in x and piecewise continuous in t.
(A3) The distributed GP approximation satisfies $|\sigma^2(x,t) - \sigma_i^2(x,t)| \leq \varepsilon_{\mathrm{GP}}$ for all $x \in \mathcal{X}$, where $\varepsilon_{\mathrm{GP}} > 0$ is the GP variance approximation bound.

Assumption A2 allows finite temporal changes in the target field. Thus, after each field or GP update, $\rho^\star$ may be recomputed and the density error may exhibit small discontinuities, while the closed-loop motion remains continuous between consecutive updates.

2.2 Gaussian Process Field Representation

Let $\mathcal{D}_i(t) = \{(x_i^\ell, y_i^\ell)\}_{\ell=1}^{n_i(t)}$ denote the local measurement history of agent i, where $n_i(t)$ is the number of stored measurements. Let $\mathcal{I}_i(t) = \{z_{i,\ell}\}_{\ell=1}^{m_i(t)}$ be the local inducing set, with $m_i(t)$ inducing points. For any query x, let $k_i(x) = [k(x, z_{i,1}) \cdots k(x, z_{i,m_i})]^\top$, where $k(\cdot,\cdot)$ is the GP covariance kernel. Let $K_i \in \mathbb{R}^{m_i \times m_i}$ be the kernel matrix over $\mathcal{I}_i(t)$, and let $\mathbf{y}_i$ be the vector of measurements or compressed pseudo-observations associated with the inducing representation. The resulting posterior is

$$\begin{aligned} \mu_i(x,t) &= k_i(x)^\top \left(K_i + \sigma_v^2 I_{m_i}\right)^{-1} \mathbf{y}_i, \\ \sigma_i^2(x,t) &= k(x,x) - k_i(x)^\top \left(K_i + \sigma_v^2 I_{m_i}\right)^{-1} k_i(x). \end{aligned} \tag{6}$$

Sparse GP communication transmits compressed posterior parameters instead of raw histories, preserving scalability. If $k(\cdot,\cdot)$ is Lipschitz, then the posterior variance and the induced target density are Lipschitz, which ensures a well-posed OT control problem.

2.3 Optimal Transport Preliminaries

Let $\mathcal{P}_2(\mathcal{X})$ denote the set of probability densities on $\mathcal{X}$ with finite second moments. For densities $\rho_1, \rho_2 \in \mathcal{P}_2(\mathcal{X})$, the squared 2-Wasserstein distance is

$$W_2^2(\rho_1, \rho_2) = \inf_{\gamma \in \Gamma(\rho_1, \rho_2)} \int_{\mathcal{X}\times\mathcal{X}} \|x - y\|^2 d\gamma(x, y), \quad (7)$$

where $\Gamma(\rho_1, \rho_2)$ denotes all transport couplings with marginals ρ_1 and ρ_2. We use the density-matching Lyapunov functional

$$V(\rho) = W_2^2(\rho, \rho^\star), \quad (8)$$

for which the Wasserstein gradient-flow velocity is

$$u(x, t) = -\nabla\phi(x, t), \quad (9)$$

where the global Kantorovich potential $\phi(x, t)$ satisfies

$$\nabla \cdot (\rho(x, t)\nabla\phi(x, t)) = \rho(x, t) - \rho^\star(x, t). \quad (10)$$

This OT formulation is not a greedy pointwise ascent on the uncertainty map. In the ideal continuum setting, the only zero-dissipation equilibrium of (10) is $\rho = \rho^\star$, so the density-level flow avoids spurious local minima. Theorem 3 then treats finite-agent sampling, actuator limits, and numerical approximation as bounded perturbations.

2.4 Problem Formulation

The objective is to design a distributed control law that reconstructs $f(x, t)$ while reducing epistemic uncertainty.

Problem 1. *Given a tam of UAVs with local measurements $\mathcal{D}_i(t)$ and communication graph $(\mathcal{A}, \mathcal{G})$, compute a distributed Gaussian process posterior that satisfies*

$$\sup_{x\in\mathcal{X}} |\sigma^2(x, t) - \sigma_i^2(x, t)| \le \varepsilon_{\mathrm{GP}}. \quad (11)$$

Problem 2. *Given the global target density $\rho^\star$ and its local approximation $\rho_i^\star$, design distributed control law $u_i(t)$ such that the swarm density governed by (4) satisfies*

$$\lim_{t\to\infty} W_2(\rho(\cdot, t), \rho^\star(\cdot, t)) = 0. \quad (12)$$

Remark 1. *With bounded Gaussian process and optimal transport errors, we are also going to prove that*

$$\frac{d}{dt}W_2^2(\rho, \rho^\star) \le -cW_2^2(\rho, \rho^\star) + \varepsilon_{\mathrm{app}}, \quad (13)$$

where $c > 0$ and $\varepsilon_{\mathrm{app}} = O(\varepsilon_{\mathrm{GP}})$ denotes the combined approximation error.

3. INFORMATION DENSITY DRIVEN CONTROL AND PROPERTY ANALYSIS

3.1 Distributed GP Mapping and Information Density

Each UAV maintains the sparse Gaussian process (GP) posterior in (6), prunes its inducing set to keep bounded complexity, and exchanges compressed inducing summaries with neighbors. The posterior variance induces normalized local target density $\rho_i^\star$ through (5). The local swarm density estimate explicitly includes the agent itself

$$\widehat{\rho}_i(x, t) = \frac{1}{Z_i} \sum_{j\in\mathcal{N}_i\cup\{i\}} K_\rho(x - x_j(t)), \quad (14)$$

where Z_i is the local normalization constant satisfying $\int_{\mathcal{X}} \widehat{\rho}_i(x, t)dx = 1$. Adding x_i is critical for physical consistency, since the realized density should represent all available local masses, including the agent's own contribution.

3.2 OT-Based Distributed Exploration Control

Each UAV solves the local OT potential equation

$$\nabla \cdot (\widehat{\rho}_i\nabla\phi_i) = \widehat{\rho}_i - \rho_i^\star, \quad (15)$$

where $\phi_i(x, t)$ is the local Kantorovich potential. The distributed control input is

$$u_i(t) = -\nabla\phi_i(x_i(t), t). \quad (16)$$

Substituting $u = -\nabla\phi$ into (4) gives

$$\frac{\partial\rho}{\partial t} = \nabla \cdot (\rho\nabla\phi), \quad (17)$$

which is the Wasserstein gradient flow of V. The corresponding Lyapunov dissipation is

$$\frac{d}{dt}V(\rho(t)) = -2\int_{\mathcal{X}} \rho(x, t)\|\nabla\phi(x, t)\|^2 dx \le 0. \quad (18)$$

Thus, the swarm density is monotonically transported toward the information density.

3.3 Robustness and Distributed Implementation

In practice, GP approximation, local density estimation, discretized OT, and low-level tracking produce errors. Let

$$\|\widehat{\rho}_i - \rho\|_{L^1} + \|\rho_i^\star - \rho^\star\|_{L^1} \le \varepsilon_{\mathrm{app}}, \quad (19)$$

where $\varepsilon_{\mathrm{app}} > 0$ denotes the combined density-approximation error. The implemented velocity can be written as

$$\widehat{u}_i = -\nabla\phi_i = -\nabla\phi + \eta_i, \qquad \|\eta_i\| = O(\varepsilon_{\mathrm{app}}), \quad (20)$$

where η_i is the local implementation perturbation. This leads to an input-to-state stability type bound

$$\frac{d}{dt}W_2^2(\rho, \rho^\star) \le -cW_2^2(\rho, \rho^\star) + C\varepsilon_{\mathrm{app}}, \quad (21)$$

for constants $c, C > 0$. The distributed procedure is summarized in Algorithm 1.

Algorithm 1 Distributed Information Density Matching

1: Initialize inducing set $\mathcal{I}_i$, potential ϕ_i, period Δt, inducing cap M_{ind}.
2: **for** $t = 1, 2, \ldots, T_{\max}$ **do**
3: Measure $y_i(t)$ at $x_i(t)$ and update $\mathcal{D}_i$.
4: Update sparse GP posterior (μ_i, σ_i^2) and prune $\mathcal{I}_i$ if $|\mathcal{I}_i| > M_{\mathrm{ind}}$.
5: Exchange compressed GP summaries with neighbors.
6: Compute $\rho_i^\star$ from σ_i^2 and $\widehat{\rho}_i$ from $\mathcal{N}_i \cup \{i\}$.
7: Solve $\nabla \cdot (\widehat{\rho}_i\nabla\phi_i) = \widehat{\rho}_i - \rho_i^\star$.
8: Apply $u_i(t) = -\nabla\phi_i(x_i, t)$ and propagate the UAV state.
9: **end for**

Theorem 1. *Let each UAV evolve according to*

$$\dot{x}_i(t) = -\nabla\phi_i(x_i(t), t), \quad (22)$$

where ϕ_i solves (15). Define the empirical measure

$$\rho_N^\delta(x, t) = \frac{1}{N}\sum_{i=1}^{N} \delta_D(x - x_i(t)). \quad (23)$$

Assume ρ_N^δ admits a weak limit $\rho(x, t)$ as $N \to \infty$, and that $\|\widehat{\rho}_i - \rho\|_{L^1} \le \varepsilon_{\mathrm{app}}$. Then the limiting density satisfies

$$\partial_t\rho(x, t) = \nabla \cdot (\rho(x, t)\nabla\phi(x, t)), \quad (24)$$

where ϕ solves

$$\nabla \cdot (\rho\nabla\phi) = \rho - \rho^\star. \quad (25)$$

Proof. The control law induces a flow map F_t satisfying

$$\frac{d}{dt}F_t(x_0) = -\nabla\phi(F_t(x_0), t). \quad (26)$$

The empirical measure is transported by this flow

$$\rho_N^\delta(\cdot,t) = F_{t\sharp}\rho_N^\delta(\cdot,0), \tag{27}$$

where $F_{t\sharp}$ denotes the push-forward operator. Weak convergence and uniform boundedness allow passage to the limit, giving $\rho(\cdot,t) = F_{t\sharp}\rho(\cdot,0)$. Any measure transported by a differentiable flow satisfies

$$\partial_t\rho + \nabla\cdot(\rho(-\nabla\phi)) = 0, \tag{28}$$

or equivalently $\partial_t\rho = \nabla\cdot(\rho\nabla\phi)$. Since the distributed elliptic equations converge to the global equation under bounded L^1 perturbations in the density coefficients, the limit potential satisfies $\nabla\cdot(\rho\nabla\phi) = \rho - \rho^\star$.

Theorem 2. *Let $V(\rho) = W_2^2(\rho,\rho^\star)$. Under the density dynamics in Theorem 1, ρ evolves as the Wasserstein gradient flow of V, and*

$$\frac{d}{dt}V(\rho(t)) = -2\int_{\mathcal{X}} \rho(x,t)\|\nabla\phi(x,t)\|^2 dx \le 0. \tag{29}$$

Proof. The Benamou–Brenier formulation writes

$$W_2^2(\rho,\rho^\star) = \inf_{\varrho,\upsilon}\int_0^1\int_{\mathcal{X}} \varrho(x,s)\|\upsilon(x,s)\|^2 dxds, \tag{30}$$

subject to the continuity equation $\partial_s\varrho + \nabla\cdot(\varrho\upsilon) = 0$. The optimal transport velocity has the form $\upsilon^\star = -\nabla\phi$, where ϕ satisfies the OT potential equation. Hence, the dynamics in Theorem 1 coincide with the steepest descent of V in the Wasserstein metric. Thus, the variation satisfies

$$\frac{d}{dt}V = 2\int_{\mathcal{X}} \phi\,\partial_t\rho\,dx. \tag{31}$$

Using $\partial_t\rho = \nabla\cdot(\rho\nabla\phi)$ and integrating by parts yields

$$\frac{d}{dt}V = -2\int_{\mathcal{X}} \rho\|\nabla\phi\|^2 dx. \tag{32}$$

The integrand is nonnegative and vanishes only when $\nabla\phi \equiv 0$, which implies $\rho = \rho^\star$. Thus V is a Lyapunov function for density matching.

Theorem 3. *Assume the approximation errors satisfy*

$$\|\widehat{\rho} - \rho\|_{L^1} + \|\widehat{\rho}^\star - \rho^\star\|_{L^1} + \|\nabla\widehat{\phi} - \nabla\phi\| \le \varepsilon_{\text{app}}, \tag{33}$$

where $\widehat{\rho}$, $\widehat{\rho}^\star$, and $\widehat{\phi}$ denote the implemented approximations of ρ, $\rho^\star$, and ϕ, respectively. Let the implemented velocity be $\widehat{u} = -\nabla\widehat{\phi}$. Then, under a standard transport-Poincaré inequality on the compact domain, there exist constants $c, C > 0$ independent of ε_{app} such that

$$\frac{d}{dt}W_2^2(\rho,\rho^\star) \le -cW_2^2(\rho,\rho^\star) + C\varepsilon_{\text{app}}. \tag{34}$$

Consequently,

$$\limsup_{t\to\infty} W_2(\rho(t),\rho^\star) = O(\sqrt{\varepsilon_{\text{app}}}). \tag{35}$$

Proof. Write the approximate control as

$$\widehat{u} = -\nabla\phi + \eta, \qquad \|\eta\| = O(\varepsilon_{\text{app}}), \tag{36}$$

where η is the aggregate velocity perturbation. The density evolves as

$$\partial_t\rho = \nabla\cdot(\rho\nabla\phi) + \nabla\cdot(\rho\eta). \tag{37}$$

Differentiating $V(\rho) = W_2^2(\rho,\rho^\star)$ gives

$$\frac{d}{dt}V = -2\int_{\mathcal{X}} \rho\|\nabla\phi\|^2 dx - 2\int_{\mathcal{X}} \rho\langle\nabla\phi,\eta\rangle dx. \tag{38}$$

By Cauchy–Schwarz and Young's inequalities,

$$\left|\int_{\mathcal{X}} \rho\langle\nabla\phi,\eta\rangle dx\right| \le \frac{1}{2}\int_{\mathcal{X}} \rho\|\nabla\phi\|^2 dx + C\varepsilon_{\text{app}}. \tag{39}$$

Therefore, the following inequlity holds

$$\frac{d}{dt}V \le -\int_{\mathcal{X}} \rho\|\nabla\phi\|^2 dx + C\varepsilon_{\text{app}}. \tag{40}$$

The transport-Poincaré inequality gives

$$\int_{\mathcal{X}} \rho\|\nabla\phi\|^2 dx \ge cW_2^2(\rho,\rho^\star), \tag{41}$$

which yields the ISS-type inequality. Solving it gives $W_2^2(\rho,\rho^\star) = O(\varepsilon_{\text{app}})$ and hence $W_2(\rho,\rho^\star) = O(\sqrt{\varepsilon_{\text{app}}})$.

4. SIMULATION RESULTS

4.1 Simulation Setup

Simulations are conducted in a bounded $100 \times 100 \times 30$ m^3 workspace. The horizontal plane is the exploration domain, and the vertical dimension is retained for the dynamics of the quadrotor. Unless otherwise stated, $N = 12$ UAVs are initialized randomly, the control period is $\Delta t = 0.5$ s, and each run lasts 300 steps. The swarm model (1) is used as a high-level kinematic layer. Each control command u_i is sent as a velocity reference to a cascaded attitude–velocity controller of a full 6-DoF quadrotor. Then the realized translational motion is written as

$$\dot{x}_i(t) = u_i(t) + e_i(t) + w_i(t), \tag{42}$$

where $e_i(t)$ denotes bounded tracking error caused by attitude dynamics, saturation, and delay. This error enters the perturbation term η_i in Theorem 3; therefore, bounded low-level tracking mismatch does not destroy density-level stability but may enlarge the final tracking neighborhood. The maximum speed and acceleration are $v_{\max} = 4$ m/s and $a_{\max} = 6$ m/s^2, with cruise altitude $h_c = 12$ m.

The scalar field is synthesized on the horizontal plane using spatiotemporal random Fourier features, with correlation lengths sampled from $l_s \in [15, 25]$ m and $l_t \in [20, 60]$ s. Each UAV maintains at most 32 inducing points and exchanges compressed GP summaries with neighbors within communication radius $r_c = 30$ m. The plane is discretized into a 100×100 grid. The coverage footprint radius is $r_{\text{cov}} = 16$ m, and the kernel bandwidth for the density estimate is $b_\rho = 12$ m. Then the main metrics are

$$\begin{aligned} \text{RMSE} &= \sqrt{\frac{1}{N_g}\sum_{q=1}^{N_g}\Big(\widehat{f}(x_q) - f(x_q)\Big)^2}, \\ \text{Stability} &= \text{Var}\left\{\|p_i - p_j\|^2 : i \ne j\right\}, \\ \text{CommLoad}(t) &= \frac{\sum_{\ell=1}^{N_{\text{pkt}}(t)} \texttt{packet_bytes}_\ell}{\Delta t}. \end{aligned} \tag{43}$$

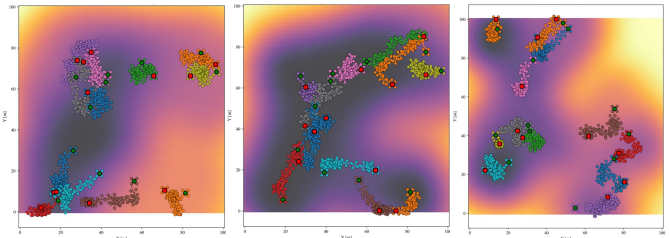
Here N_g is the number of grid points, x_q is the q-th grid point, $\widehat{f}$ is the GP-reconstructed field, p_i is the UAV position used for the formation-stability metric, and $N_{\text{pkt}}(t)$ is the number of communication packets transmitted during the simulation step at time t. Lower RMSE, density error, stability variance, and communication load are preferred.

In the proposed method, $k_{\text{explore}} = 4.0$, $k_{\text{repulse}} = 2.0$, and a low-pass filter with coefficient $\alpha = 0.4$ smooths the commanded velocities. Here k_{explore} and k_{repulse} are the information-seeking and repulsion gains, respectively, and α is the velocity-filter coefficient. The apparent discontinuities in the density-error curves occur when the spatiotemporal field, GP posterior, or target density is

periodically updated. This behavior is consistent with Assumption A2, which permits piecewise continuity in time; between updates, the controller evolves continuously.

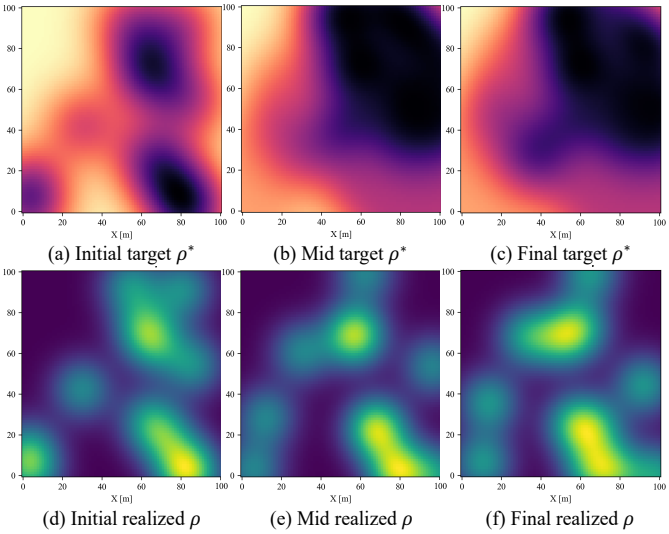
4.2 Results and Discussion

Fig. 1 shows that IDM generates smooth and space-filling trajectories that visit high-uncertainty regions. Fig. 2 further illustrates how the realized density reshapes toward the information density. Since $\rho^\star$ changes when the GP posterior is updated, a nonzero transient density error is expected, especially in dynamic fields.



(a) (b) (c)

Fig. 1. UAV trajectories of the proposed IDM controller overlaid on the final GP variance map.



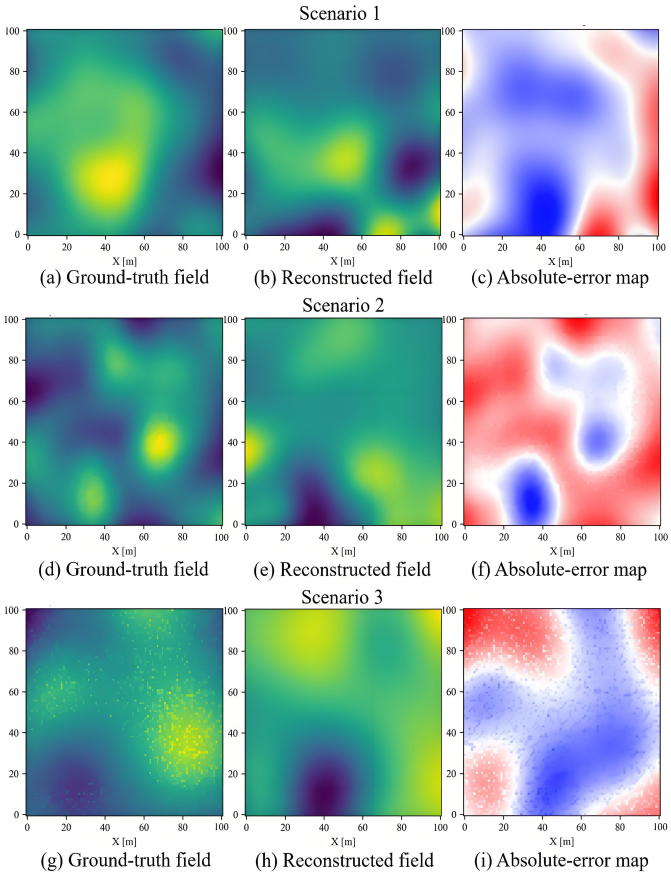
(a) Initial target ρ^* (b) Mid target ρ^* (c) Final target ρ^*

(d) Initial realized ρ (e) Mid realized ρ (f) Final realized ρ

Fig. 2. Evolution of the information density $\rho^\star$ and realized swarm density ρ in Scenario 1.

Fig. 3 evaluates field reconstruction. The visible discrepancies in Scenario 2 mainly occur near high-gradient or moving regions, where sparse measurements and temporal evolution make the GP posterior harder to align with the instantaneous ground truth. The maximum absolute error can exceed the maximum reconstructed value when the reconstructed field is locally small but the ground truth contains sharp peaks. This does not contradict a low global RMSE, which averages errors over the grid. For real deployment, such local errors are important for threshold-based tasks, so IDM should be used together with posterior variance or confidence indicators rather than relying only on the posterior mean.

Fig. 4 and Table 1 compare IDM with Lloyd, frontier, Lévy, chemotaxis, and random exploration. IDM achieves the lowest or tied-lowest density error, the most stable formation, and the lowest communication load, while maintaining competitive coverage and RMSE. It is not



(a) Ground-truth field (b) Reconstructed field (c) Absolute-error map

(d) Ground-truth field (e) Reconstructed field (f) Absolute-error map

(g) Ground-truth field (h) Reconstructed field (i) Absolute-error map

Fig. 3. Field reconstruction in three scenarios using IDM. Rows from left to right represent ground truth, GP reconstruction, and absolute error, respectively.

uniformly best for every metric: frontier achieves the lowest RMSE in this particular setting, and random exploration obtains slightly higher coverage. However, IDM provides the most balanced performance because it jointly considers uncertainty reduction, density matching, communication efficiency, and swarm coherence.

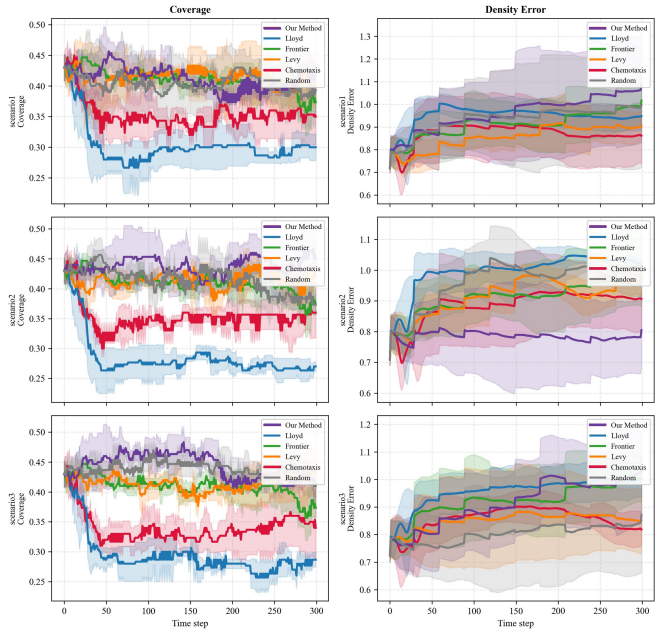


Fig. 4. Benchmarking against baseline exploration policies in terms of coverage and density-matching error.

The parameters r_{cov} and b_ρ affect evaluation and density estimation, respectively. Reducing r_{cov} lowers absolute coverage for all methods but does not change the controller itself. In contrast, changing b_ρ modifies the smoothness of $\widehat{\rho}_i$: too small a bandwidth produces noisy local density estimates, whereas too large a bandwidth oversmooths local swarm structure. Communication radius and bandwidth affect the GP error ε_{GP} and the density-estimation error in Theorem 3; lower communication increases the final residual neighborhood but does not invalidate the stability mechanism. Therefore, the method is expected to remain effective under moderate parameter changes, although severe communication loss or overly sparse sensing will reduce reconstruction quality.

Table 1. Comparison of exploration policies.

Policy	Cov.	Dens. Err.	RMSE	Stab.	Comm.
IDM	0.42	**0.84**	1.25	**827.84**	**28.52**
Lloyd	0.27	1.01	1.30	1552.16	33.48
Frontier	0.40	**0.84**	**1.13**	829.17	35.39
Lévy	0.41	0.94	1.59	970.40	36.87
Chemotaxis	0.27	0.87	1.57	991.50	33.33
Random	**0.43**	0.87	1.37	833.08	37.63

5. CONCLUSION

This paper presented an information density matching framework that unifies sparse GP field estimation with OT-based swarm coordination for UAV exploration in spatiotemporal environments. Each UAV constructs a variance-driven information density from its local GP posterior and applies a distributed OT control law to move the swarm density toward this target. The resulting closed-loop density dynamics admit a Wasserstein gradient-flow interpretation, and Lyapunov analysis establishes monotonic density-mismatch reduction with robustness to bounded GP, OT, and tracking errors. Simulations across static, multi-peak, and dynamic fields show that IDM achieves balanced performance in coverage, density matching, reconstruction accuracy, communication load, and formation stability compared with representative baselines. Future work will focus on hardware experiments, richer sensitivity analysis, and extensions to heterogeneous teams under intermittent communication.

REFERENCES

Bartolomei, L., Teixeira, L., and Chli, M. (2021). Semantic-aware active perception for uavs using deep reinforcement learning. In *IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*, 3101–3108.

Chen, F., Martin, J.D., Huang, Y., Wang, J., and Englot, B. (2020). Autonomous exploration under uncertainty via deep reinforcement learning on graphs. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 6140–6147.

Dergachev, S. and Yakovlev, K. (2025). Decentralized uncertainty-aware multi-agent collision avoidance with model predictive path integral. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 12456–12463.

Du, Z., Luo, C., Min, G., Wu, J., Luo, C., Pu, J., and Li, S. (2025). A survey on autonomous and intelligent swarms of uncrewed aerial vehicles (UAVs). *IEEE Trans. on Intel. Transportation Systems*, 26(10), 14477 – 14500.

Gao, Y., Gao, H., Ji, Y., Zhou, J., and Shi, Y. (2025a). A dual calibration framework for exploring environments using heterogeneous robot swarms. In *Annual Conference of the IEEE Industrial Electronics Society*, 1–7.

Gao, Y., Gao, H., Ma, W., Xiong, H., and Ji, Y. (2026). Learning load-balanced distributed coverage for robot swarms via graph attention networks. In *IEEE Intl. Conference on Robotics and Automation*, 4222–4229.

Gao, Y., Zhou, Z., Gao, H., Zhang, S., and Ji, Y. (2025b). Hybrid control for robot swarms to detect critical nodes in heterogeneous sensor networks. In *IEEE Intl. Conf. on Automation Science and Engineering*, 3462–3467.

Hsu, K.C., Rubies-Royo, V., Tomlin, C.J., and Fisac, J.F. (2021). Safety and liveness guarantees through reach-avoid reinforcement learning. In *17th Robotics: Science and Systems*.

Huang, Y., Wang, H., Bai, X., Cai, X., Yu, H., and Ren, Y. (2025). Biomimetic multi-UAV swarm exploration with U2U communications under resource constraints. *IEEE Trans. on Vehicular Technology*, 74(6), 9750 – 9766.

Jang, D., Yoo, J., Son, C.Y., Kim, D., and Kim, H.J. (2020). Multi-robot active sensing and environmental model learning with distributed gaussian process. *IEEE Robotics and Automation Letters*, 5(4), 5905–5912.

Javed, S., Hassan, A., Ahmad, R., Ahmed, W., Ahmed, R., Saadat, A., and Guizani, M. (2024). State-of-the-art and future research challenges in uav swarms. *IEEE Internet of Things Journal*, 11(11), 19023–19045.

Liu, J., Lv, Y., Yuan, Y., Chi, W., Chen, G., and Sun, L. (2022). An efficient robot exploration method based on heuristics biased sampling. *IEEE Transactions on Industrial Electronics*, 70(7), 7102–7112.

Newaz, A.A.R., Alsayegh, M., Alam, T., and Bobadilla, L. (2023). Decentralized multi-robot information gathering from unknown spatial fields. *IEEE Robotics and Automation Letters*, 8(5), 3070–3077.

Ramachandran, R.K., Kakish, Z., and Berman, S. (2020). Information correlated lévy walk exploration and distributed mapping using a swarm of robots. *IEEE Transactions on Robotics*, 36(5), 1422–1441.

Reily, B. and Zhang, H. (2021). Team assignment for heterogeneous multi-robot sensor coverage through graph representation learning. In *IEEE International Conference on Robotics and Automation*, 838–844.

Santilli, M., Mukherjee, P., Williams, R.K., and Gasparri, A. (2022). Multirobot field of view control with adaptive decentralization. *IEEE Transactions on Robotics*, 38(4), 2131–2150.

Viseras, A., Wiedemann, T., Manss, C., Magel, L., Mueller, J., Shutin, D., and Merino, L. (2016). Decentralized multi-agent exploration with online-learning of gaussian processes. In *IEEE International Conference on Robotics and Automation*, 4222–4229.

Wong, A., Bäck, T., Kononova, A.V., and Plaat, A. (2023). Deep multiagent reinforcement learning: Challenges and directions. *Artificial Intelligence Review*, 56(6), 5023–5056.

Wu, Y., Fu, Y., and Wang, S. (2023). Information-theoretic exploration for adaptive robotic grasping in clutter based on real-time pixel-level grasp detection. *IEEE Trans. on Industrial Elect.*, 71(3), 2683–2693.

Yang, B., Ding, Y., Jin, Y., and Hao, K. (2015). Self-organized swarm robot for target search and trapping inspired by bacterial chemotaxis. *Robotics and Autonomous Systems*, 72, 83–92.