Decentralized PSO-Based Guidance for Multi-UAV Wildfire Monitoring and Frontier Tracking

H. Chakraa, M. Feurgard, and M. Bronz
École Nationale de l'Aviation Civile (ENAC), Université de Toulouse, 31400 Toulouse, France.

http://www.imavs.org/

Abstract

This paper proposes a guidance framework for adaptive navigation and coverage control in autonomous Unmanned Aerial Vehicles (UAVs) swarms operating in rapidly evolving wildfire environments. In the navigation phase, each UAV computes the bearing to its nearest discovered frontier cell and advances along that heading, with an iterative separation to prevent inter-drone collisions. In the coverage phase, a decentralized Particle Swarm Optimization (PSO) based controller searches the velocity space independently for each drone. Its fitness simultaneously maximizes front-line coverage, rewards exploration of the pre-frontier region just ahead of the fire edge, and attracts each drone toward its assigned angular sector of the perimeter. A shared belief fire map, updated incrementally through exploration, drives both frontier extraction and the PSO fitness. We evaluate the framework with frontier and fire-body Intersection-over-Union (IoU), sensor overlap rate, and per-UAV decision latency. Finally, the proposed framework was experimentally validated through real-world flight tests using a fleet of DJI Mini 3 quadrotors.

1 Introduction

Wildfires rank among the most destructive natural disasters, causing irreversible ecological damage, economic losses, and severe threats to human lives and infrastructures [1]. The increasing frequency and intensity of wildfires highlight the urgent need for efficient monitoring technologies. Unmanned Aerial Vehicles (UAVs) have emerged as a promising solution, offering higher spatial and temporal resolution and enabling access to inaccessible terrain [2].

The central challenge is guidance in a dynamic and partially observed environment. A wildfire front evolves continuously under wind, terrain, and fuel conditions [3]. Effective fleet coordination involves two complementary subproblems: (i) *navigation*, guiding drones from initial positions to the active fire, and (ii) *coverage and tracking*, distributing drones along the frontier boundary to maximize observability, following the propagation (see Figure 1).

Emails: {hamza.chakraa, mael.feurgard, murat.bronz}@enac.fr

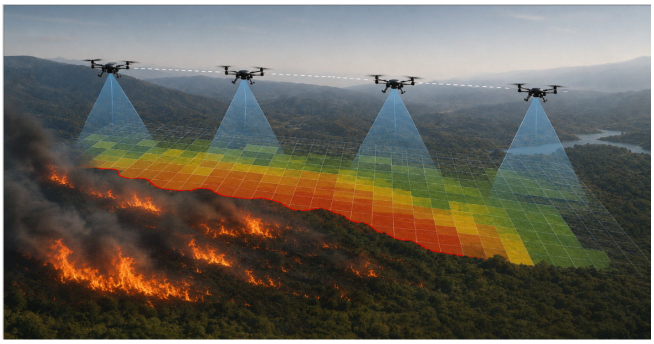

Figure 1: Illustration of a wildfire monitoring scenario. UAVs cooperatively track the evolving wildfire frontier. Each drone observes a local region within its onboard sensor range, and then information is constructed.

This paper addresses this problem through a unified guidance framework. The proposed approach operates in two sequential phases. In the navigation phase, each UAV is driven toward the evolving wildfire frontier using a heading-based controller designed for partial observability, where no prior map of the fire is available. The swarm incrementally builds the situation through onboard sensing, progressively revealing the fire structure as the mission progresses. Once sufficient proximity to the wildfire is achieved, the system transitions to the coverage phase, where the objective shifts from reaching the fire to spatially distributing UAVs along its boundary. This phase is formulated as a decentralized Particle Swarm Optimization (PSO) problem [4], enabling each UAV to independently optimize its velocity while remaining coordinated at the swarm level. PSO has been widely adopted in UAV deployment and coverage tasks due to its balance between computational efficiency and global search capability [5]. A key component of this work is a frontier extraction module based on OpenCV image-processing techniques, which continuously reconstructs the wildfire boundary from the shared belief fire map.

The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 presents the system overview and models. Section 4 details the multi-UAV guidance strategy. Section 5 reports simulation results and field demonstrations with DJI Mini 3 drones, and finally, Section 6 concludes.

http://www.imavs.org/

2 RELATED WORK

Satellite imagery and ground-based sensor networks provide large-scale coverage, but their low temporal resolution limits real-time response [5]. UAV platforms bridge this gap with on-demand deployment at higher spatial and temporal resolution, including inaccessible terrains to ground personnel [6, 7]. Bailon-Ruiz *et al.* [6] introduced an early multi-UAV replanning system for wildfire monitoring, later extended to continuous real-time perimeter tracking [8]. Heracleous *et al.* [7] demonstrated accurate frontier reconstruction from aerial imagery, and Seraj *et al.* [9] formalized cooperative coverage with quality of service and safety guarantees. Despite this progress, guidance under partial observability (where the fire boundary must be discovered incrementally) remains a central open challenge.

Classical Coverage Path Planning (CPP) problems assume a known, static target region and produce offline tours that guarantee complete coverage with bounded redundancy. The wildfire setting fundamentally breaks this assumption: the region of interest is the active fire perimeter, which is unknown *a priori*, partially observable, and continuously deforming under wind and fuel gradients [8]. Coverage must therefore be treated as an online dynamic problem in which the target itself moves and grows, sensor information arrives incrementally, and replanning occurs at every control step. This shift, from offline CPP over a fixed map to adaptive tracking of an evolving frontier, motivates dedicated dynamic formulations and distinguishes wildfire guidance from conventional multi-UAV area coverage.

A complementary thread frames wildfire monitoring as an active perception problem, where agents maximize information gain about the fire state rather than cover fixed geometric regions. Julian and Kochenderfer [10] showed that Deep Reinforcement Learning (DRL) with information-driven objectives can outperform geometric coverage heuristics in decentralized surveillance. Wang and Bronz [11] proposed a Monte Carlo Tree Search (MCTS) belief-state planner for uncertainty estimation, fire tracking, and fuel-aware exploration. These methods show strong estimation performance but require either large training datasets or expensive online tree search. Evolutionary strategies such as Genetic Algorithms [12] and PSO offer a training-free alternative that remains competitive for deployment. PSO in particular performs strongly in decentralized multi-agent settings: its population-based search and support for multi-objective fitness make it well-suited for online re-optimization.

To the best of our knowledge, this is the first framework to apply per-drone decentralized PSO as a reactive velocity controller for wildfire frontier coverage. Unlike prior approaches that rely on discrete action sets [6, 8], offline distillation [11], or trained networks [13, 14], our framework is training-free, interpretable, and runs on a single CPU without a GPU or pre-collected data.

3 SYSTEM OVERVIEW

3.1 Wildfire propagation model

The wildfire environment is represented by a discrete two-dimensional grid $\mathcal{G} = \{1, \ldots, N\}^2$ where each cell corresponds to a spatial region of the terrain. The simulation evolves in discrete time steps. At each time step, The fire state is $\mathbf{X}_t \in [0,1]^{N\times N}$ evolving with a fuel map $\mathbf{F}_t \in \mathbb{R}_{\geq 0}^{N\times N}$. Wildfire propagation follows a Cellular Automata (CA) formulation on the Moore neighborhood. Full model details are in [15]. Fire is propagated every $N_f = 10$ drone control steps, yielding an effective ground-truth spread rate of approximately 0.5–$1\,\mathrm{m/s}$, consistent with measured rates for low-to-moderate intensity fires [8].

3.2 UAVs and observation model

We consider a fleet of M homogeneous UAVs operating over the fire grid. Each UAV i has a planar pose $\mathbf{p}_t^{(i)} = (x_t^{(i)}, y_t^{(i)}) \in \mathbb{R}^2$ and a heading $\theta_t^{(i)} \in (-\pi, \pi]$. The sensor footprint $\mathcal{D}_t^{(i)}$ of drone i at time t is a square field of view of side s cells centered at the grid projection of $\mathbf{p}_t^{(i)}$:

$$\mathcal{D}_t^{(i)} = \left\{ \tilde{\mathbf{p}} \in \mathcal{G} \;\middle|\; \left\| \left[\tilde{\mathbf{p}} - \mathrm{m2cell}(\mathbf{p}_t^{(i)}) \right] \right\|_\infty \leq \rho \right\} \tag{1}$$

where $\rho = \lfloor s/2 \rfloor$ is the FoV half-side in cells and $\mathrm{m2cell}(\cdot)$ denotes the coordinate transformation. Within $D_t^{(i)}$, drone i observes the true fire state $\mathbf{X}_t(\tilde{\mathbf{p}})$.

The team maintains a single shared data structure: the belief fire map $\hat{\mathbf{X}}_t \in [0,1]^{N\times N}$, initialized to zero. The belief map is updated incrementally as the fleet explores: after each control step, every cell inside the union of all drone FoV footprints is overwritten with its ground-truth fire value, while unobserved cells retain their previous belief:

$$\hat{\mathbf{X}}_{t+1}(\tilde{\mathbf{p}}) = \begin{cases} \mathbf{X}_t(\tilde{\mathbf{p}}) & \text{if } \tilde{\mathbf{p}} \in \bigcup_{i=1}^{M} \mathcal{D}_t^{(i)}, \\ \hat{\mathbf{X}}_t(\tilde{\mathbf{p}}) & \text{otherwise.} \end{cases} \tag{2}$$

This update rule models the physical sensing process: cells that no drone has yet overflown remain at their initialized value while observed cells hold the most recent ground-truth measurement. The belief frontier $\hat{\mathbf{R}}_t$ is extracted from $\hat{\mathbf{X}}_t$ after every observation update via the OpenCV pipeline described in Section 3.3.

3.3 Frontier extraction

The active fire frontier is extracted using a multi-step OpenCV pipeline designed to construct the estimated fire frontier, $\hat{\mathbf{R}}_t$, from the shared belief fire map $\hat{\mathbf{X}}_t$. The process begins by smoothing isolated burning pixels and suppressing artifacts that typically appear near the boundary. After this initial denoising stage, fragmented fire regions created by the blurring operation are reconnected using binary thresholding at $\tau = 60$ and morphological refinement. Residual artifact

http://www.imavs.org/

pixels are then removed using a median filter, while an accumulation pass fills small holes that would otherwise generate false interior contours.

Once a coherent fire mask has been obtained, the boundary is further regularized using Gaussian blur to reduce high-frequency irregularities before edge detection. Finally, the active fire frontier is extracted as a thin single-pixel contour using the Canny edge detector [16].

The resulting binary edge image $\mathcal{E}_t \in \{0,1\}^{N\times N}$ encodes detected frontier pixels, where $\mathcal{E}_t(\tilde{\mathbf{p}}) = 1$ indicates an edge and $\mathcal{E}_t(\tilde{\mathbf{p}}) = 0$ indicates background. The estimated frontier is then:

$$\hat{\mathbf{R}}_t = \{\tilde{\mathbf{p}} \in \mathcal{G} \mid \mathcal{E}_t(\tilde{\mathbf{p}}) = 1\} \tag{3}$$

The same pipeline is applied to the ground-truth fire map to obtain $\mathbf{R}_t$, used *exclusively* for metrics calculation (Section 5).

4 Multi-UAV Guidance Strategy

Before coverage can begin, each drone must reach the active fire frontier. The system operates under complete initial ignorance: $\hat{\mathbf{X}}_0 = \mathbf{0}_{N\times N}$ and $\hat{\mathbf{R}}_0 = \emptyset$, so no frontier is known at $t = 0$ and drones have no prior information about the fire location. Drones are initially dispatched to the four quadrant centroids of the operational area. As soon as any drone's FoV intersects a burning region and $\hat{\mathbf{R}}_t \neq \emptyset$, the transition to coverage occurs.

Once all UAVs have reached the frontier and the system switches to the coverage phase, each one runs an independent PSO instance at every control step to determine its velocity command. The only information shared between drones is their current positions $\{\mathbf{p}_t^{(j)}\}_{j\neq i}$ and the shared belief map $\hat{\mathbf{X}}_t$, both available to drone i. Before running the PSO, two preliminary operations are performed: sector assignment, which defines the portion of the frontier that drone i is responsible for, and the boundary check, which determines whether the PSO should run at all or whether the drone should instead return toward its assigned region. These are described in the following.

4.1 Sector assignment

To prevent fleet clustering, the belief frontier $\hat{\mathbf{R}}_t$ is partitioned into M equal angular sectors around its centroid $\bar{\mathbf{f}}_t$. Each frontier pixel is assigned to the sector whose angular slice contains its bearing from $\bar{\mathbf{f}}_t$, dividing the perimeter into M arcs of equal angular width $2\pi/M$. The sector centroid $\mathbf{s}^{(i)}$ for drone i is then the mean position of all frontier pixels in its assigned arc:

$$\mathbf{s}^{(i)} = \frac{1}{|\mathcal{S}^{(i)}|} \sum_{\mathbf{q}\in\mathcal{S}^{(i)}} \mathbf{q} \ \in \mathbb{R}^2 \tag{4}$$

where $\mathcal{S}^{(i)}$ is the set of positions of frontier pixels belonging to drone i's sector, falling back to $\bar{\mathbf{f}}_t$ if the sector is empty. Sector assignment is recomputed at every control step as the belief frontier evolves. If drone i drifts farther than the effective leash radius ℓ_{eff} from $\mathbf{s}^{(i)}$, the PSO is bypassed entirely and the drone returns directly toward its sector centroid. The leash radius adapts to the current frontier size so that drones are allowed to move farther as the fire grows:

$$\ell_{\text{eff}} = \max(\ell_0,\ 1.4 \cdot \bar{r}_{\mathcal{F}}) \tag{5}$$

where ℓ_0 is the base leash and $\bar{r}_{\mathcal{F}}$ is the mean distance from belief frontier pixels to $\bar{\mathbf{f}}_t$ in metres, serving as a proxy for the current frontier radius. A fixed leash would cause boundary-pinning as the frontier expands: drones would be repeatedly recalled before they can cover the outer portions of a large perimeter.

4.2 PSO update

A population of P particles evolves over K iterations, each particle k encoding a candidate drone velocity $\boldsymbol{\xi}_k = (v_x^k, v_y^k)^\top$ constrained to the speed disc $\mathcal{B}(V) = \{\mathbf{v} \in \mathbb{R}^2 \mid \|\mathbf{v}\|_2 \leq V\}$, which enforces a physical speed limit. The search first updates the particles internal search velocity $\boldsymbol{\nu}_k \in \mathbb{R}^2$:

$$\begin{aligned} \boldsymbol{\nu}_{k+1} = {} & w\,\boldsymbol{\nu}_k \\ & + c_1\mathbf{r}_1 \odot (\boldsymbol{\xi}_k^{\text{pb}} - \boldsymbol{\xi}_k) \\ & + c_2\mathbf{r}_2 \odot (\boldsymbol{\xi}^{\text{gb}} - \boldsymbol{\xi}_k) \end{aligned} \tag{6}$$

where $\boldsymbol{\xi}_k^{\text{pb}}$ is the personal best (the highest-fitness candidate this particle has visited across all previous iterations) and $\boldsymbol{\xi}^{\text{gb}}$ is the global best (the highest-fitness candidate found by any particle in the population). $\mathbf{r}_1$ and $\mathbf{r}_2$ are independent random vectors with each component drawn uniformly from $[0,1]$ at every iteration, and $\odot$ denotes element-wise multiplication. The second equation updates the candidate velocity itself:

$$\boldsymbol{\xi}_{k+1} = \text{proj}_{\mathcal{B}(V)}(\boldsymbol{\xi}_k + \boldsymbol{\nu}_{k+1}) \tag{7}$$

For temporal continuity, the population is warm-started at every control step: 80% of particles are initialized near the previous step's global best $\boldsymbol{\xi}^{\text{gb}}$ with additive Gaussian noise and the remaining 20% are sampled uniformly on $\mathcal{B}(V)$. This preserves knowledge of the previous solution while maintaining sufficient exploration to escape local optima.

4.3 Fitness function

The fitness Φ_i is evaluated independently for each of the P particles at every PSO iteration. For a particle k with candidate velocity $\boldsymbol{\xi}_k = (v_x^k, v_y^k)^\top$, the candidate position is:

$$\mathbf{p}'_k = \mathbf{p}_t^{(i)} + \boldsymbol{\xi}_k\,\mathrm{d}t \in \mathbb{R}^2 \tag{8}$$

and its grid projection is $\tilde{\mathbf{p}}'_k$. The fitness combines terms in two coordinate spaces: the coverage terms G_i and E_i use the grid projection $\tilde{\mathbf{p}}'_k$ to index into the belief map and count

http://www.imavs.org/

frontier pixels inside the FoV window, while the attraction A_i, separation C_i, and boundary B terms use the world position $\mathbf{p}'_k$ directly, since they involve physical distances for which the grid representation is inappropriate. The fitness score of a particle k with candidate velocity $\boldsymbol{\xi}_k$ is:

$$\begin{aligned}\Phi_i(\boldsymbol{\xi}_k) = &\underbrace{\lambda_G\, G_i(\tilde{\mathbf{p}}'_k)}_{\text{exclusive gain}} + \underbrace{\lambda_E\, E_i(\tilde{\mathbf{p}}'_k)}_{\text{exploratory gain}} \\ + &\underbrace{\lambda_A\, A_i(\mathbf{p}'_k)}_{\text{sector attraction}} - \underbrace{\lambda_C\, C_i(\mathbf{p}'_k)}_{\text{collisions}} - \underbrace{\lambda_B\, B(\mathbf{p}'_k)}_{\text{boundary}}\end{aligned} \tag{9}$$

Exclusive coverage gain G_i: This term rewards the candidate position for covering frontier pixels that no teammate currently observes. The gain counts frontier pixels visible from $\mathbf{p}'_k$ that no teammate sees:

$$G_i(\tilde{\mathbf{p}}'_k) = \frac{1}{s^2}\left|\mathcal{D}(\tilde{\mathbf{p}}'_k) \cap \hat{\mathbf{R}}_t \cap \left(\mathcal{G} \setminus \bigcup_{j\neq i} \mathcal{D}_t^{(j)}\right)\right| \tag{10}$$

where $|\cdot|$ denotes set cardinality and $\mathcal{D}(\tilde{\mathbf{p}}'_k)$ denote the cell FoV window centered at the candidate grid position $\tilde{\mathbf{p}}'_k$. This single term is self-driving: a stationary drone exhausts its exclusive cells and is pushed outward; entering the fire body yields zero gain; only tangential motion along the ring continuously reveals new exclusive cells.

Exploratory gain E_i: When the frontier is fully covered, the exploratory gain extends each drone's attention beyond the current fire edge and into a pre-frontier zone $\mathcal{P}_t$ (the unburned cells that the fire has not yet reached but will reach next). It counts the cells in $\mathcal{P}_t$ visible from $\tilde{\mathbf{p}}'_k$ that no other UAV currently covers:

$$E_i(\tilde{\mathbf{p}}'_k) = \frac{1}{s^2}\left|\mathcal{D}(\tilde{\mathbf{p}}'_k) \cap \mathcal{P}_t \cap \left(\mathcal{G} \setminus \bigcup_{j\neq i} \mathcal{D}_t^{(j)}\right)\right| \tag{11}$$

and the pre-frontier zone is obtained by morphological expansion of the belief frontier followed by exclusion of already-burning cells:

$$\begin{aligned}\mathcal{P}_t = &\left\{\tilde{\mathbf{p}} \in \mathcal{G} \;\middle|\; \min_{\tilde{\mathbf{p}}' \in \hat{\mathbf{R}}_t} \|\tilde{\mathbf{p}} - \tilde{\mathbf{p}}'\|_\infty \leq n_s\right\} \\ &\setminus \left(\hat{\mathbf{R}}_t \cup \{\tilde{\mathbf{p}} \in \mathcal{G} \mid \hat{\mathbf{X}}_t(\tilde{\mathbf{p}}) > 0.5\}\right)\end{aligned} \tag{12}$$

The expansion radius n_s cells corresponds to next fire propagation steps giving drones enough anticipatory range to observe cells before they ignite.

Sector attraction A_i: When both G_i and E_i are zero (which occurs when the frontier is fully covered and the pre-frontier buffer is also saturated), the fitness Φ_i receives no signal from the coverage terms. In this condition, the PSO has no basis for preferring one direction over another, and drones may drift arbitrarily or cluster together. The sector attraction term prevents this by providing a continuous, direction-dependent signal that pulls each drone toward its assigned sector centroid $\mathbf{s}^{(i)}$:

$$A_i(\mathbf{p}'_k) = \frac{1}{1 + \dfrac{\|\mathbf{p}'_k - \mathbf{s}^{(i)}\|_2}{D}} \tag{13}$$

where D is the diagonal of the operational area in meters, used as a normalization factor.

Collision penalty C_i: This term penalizes candidate positions that would bring UAV i too close to any other one, discouraging the PSO from selecting velocities that create near-collisions. For each drone $j \neq i$, the penalty is zero when the candidate distance $\|\mathbf{p}'_k - \mathbf{p}_t^{(j)}\|_2$ exceeds the minimum separation d_{sep}:

$$C_i(\mathbf{p}'_k) = \sum_{j\neq i}\left[\max\Big(0,\; d_{\text{sep}} - \|\mathbf{p}'_k - \mathbf{p}_t^{(j)}\|_2\Big)\right]^2 \tag{14}$$

This means the penalty increases rapidly as the violation grows. The sum over all other agents means the penalty also scales with the number of simultaneous violations, strongly discouraging positions where drone i would be close to multiple drones at once.

Boundary repulsion B: This term penalizes candidate positions that approach the edges of the operational area. If drone i flies within one FoV half-width of any grid edge, part of its sensor footprint extends outside the modeled area, wasting sensing capacity. The penalty activates linearly as soon as the candidate enters this margin along any edge and increases toward the boundary.

Table 1: PSO fitness terms and weights

Term	**Symbol**	**Weight**
Individual gain	G_i	$\lambda_G = 1.00$
Exploratory gain	E_i	$\lambda_E = 0.35$
Sector attraction	A_i	$\lambda_A = 0.40$
Collision penalty	C_i	$\lambda_C = 0.25$
Boundary repulsion	B	$\lambda_B = 0.04$

Fitness weights: The five weights are summarized in Table 1. The coverage terms G_i and E_i carry the highest weights as

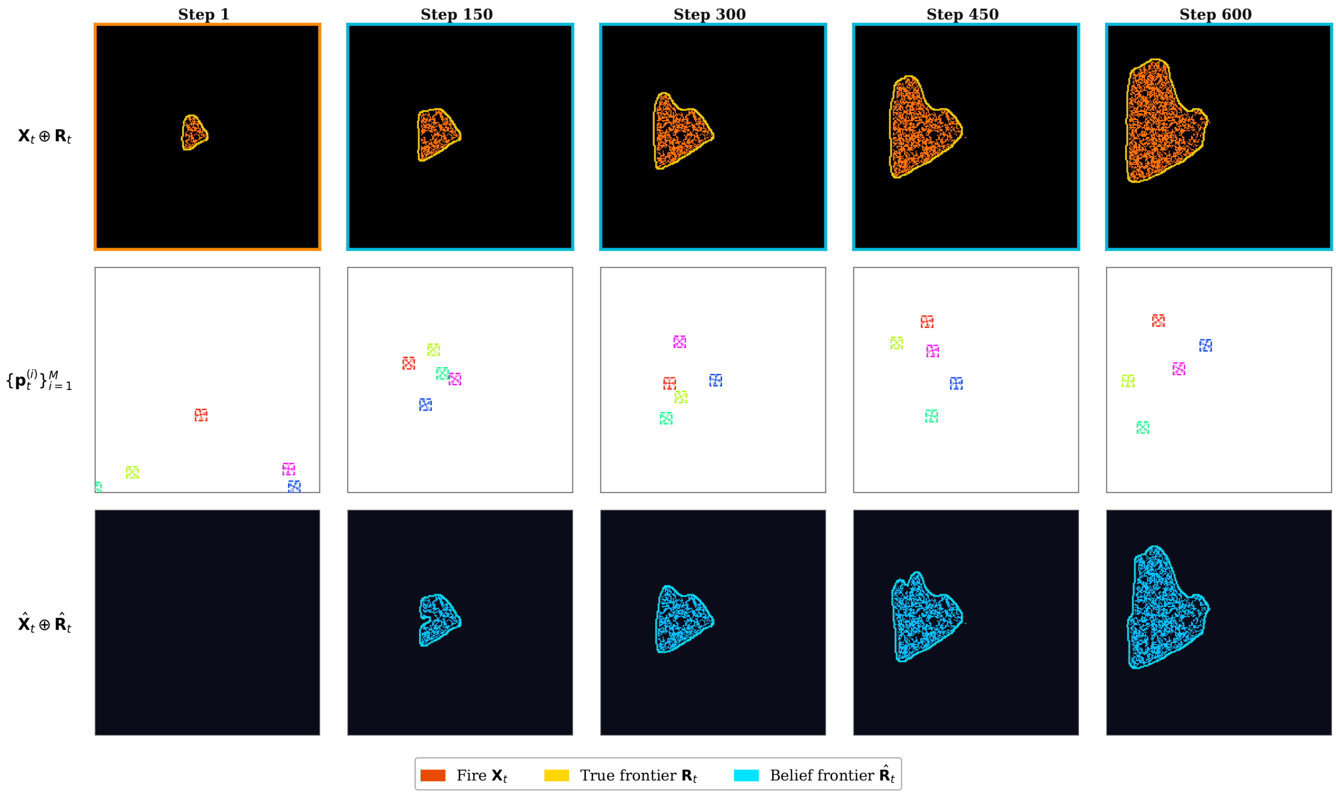

Figure 2: Simulation snapshot shown at five selected timesteps. *Row 1*: true fire state $\mathbf{X}_t$ overlaid with the ground-truth frontier $\mathbf{R}_t$ (yellow). *Row 2*: UAV positions $\{\mathbf{p}_t^{(i)}\}_{i=1}^M$ with their square FoV footprints. *Row 3*: shared belief map $\hat{\mathbf{X}}_t$ overlaid with the belief frontier $\hat{\mathbf{R}}_t$ (cyan), built incrementally from FoV observations. At $t = 0$, the belief map is entirely dark, reflecting complete initial ignorance $\hat{\mathbf{X}}_0 = \mathbf{0}$. As drones overfly burning regions, observed cells are overwritten with ground-truth values and the belief frontier grows, progressively narrowing the gap with the true frontier. The progression illustrates how the fleet collectively reconstructs the fire perimeter through exploration alone, with no prior map of the environment.

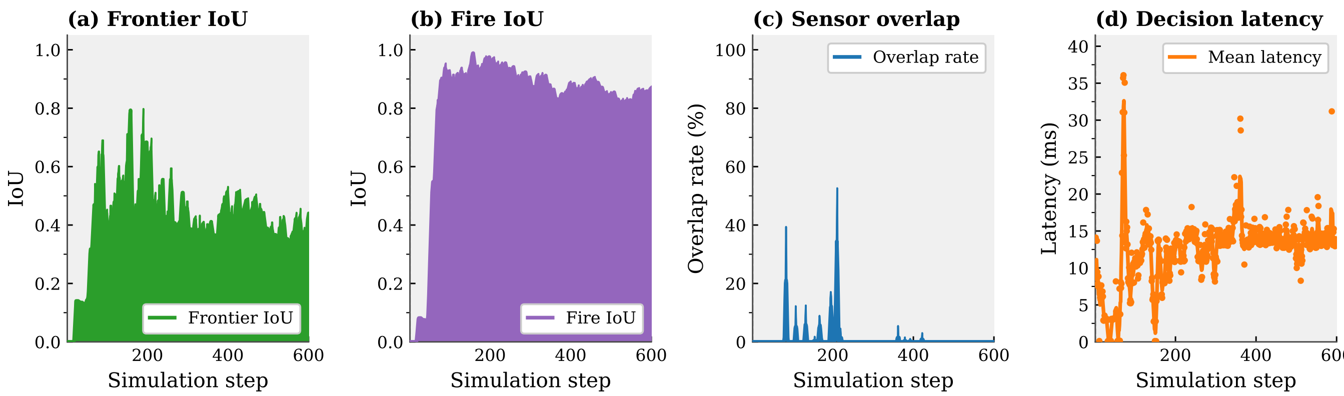

Figure 3: Per-step metrics over 600 simulation steps (200×200 grid). *(a)* Frontier IoU over the full run. *(b)* Fire IoU, reflecting coverage of the burning body. *(c)* Sensor overlap rate: low sustained value confirms efficient fleet distribution. *(d)* Per-drone PSO decision latency, well within the $1\,\mathrm{s}$ budget.

http://www.imavs.org/

http://www.imavs.org/

they are the primary drivers of cooperative behavior. The sector attraction A_i is weighted at $\lambda_A = 0.40$ to maintain distribution when coverage saturates, but also weak enough to be overridden by large G_i or E_i.

5 RESULTS AND FIELD DEPLOYMENT

We evaluate the framework in two scenarios: a large-scale simulation on a 200×200 grid, with 5 UAVs and real flights deployment on the ENAC football field with DJI Mini 3 quadrotors on a 100×100 grid with 3 UAVs.

5.1 Simulation

Scenario: We consider a 200×200 grid (cell size $\Delta = 5\,\mathrm{m}$, total area $1\,\mathrm{km}^2$) with $M = 5$ UAVs operating at $V = 5\,\mathrm{m/s}$ for simulation. Fire is ignited at the grid center and propagated every $N_f = 10$ steps, yielding an effective spread rate of approximately 0.5–1 $\mathrm{m/s}$, consistent with low to moderate intensity fires [8]. Full parameters are listed in Table 2.

Table 2: Simulation parameters

Type	Parameter	Value
Environment	Grid size $N \times N$	200×200
	Cell resolution Δ	$5\,\mathrm{m}$
	Fire update period N_f	10 steps
UAV	Cruise speed V	$5\,\mathrm{m/s}$
	Field of view s	10 cells
	Fleet size M	5
PSO guidance	Population size P	40
	Iterations K	30
	Inertia weight w	0.60
	Acceleration $c_1 = c_2$	1.50
	Separation distance d_{sep}	$40\,\mathrm{m}$

Results: Figure 2 shows the evolution of the simulation across selected timesteps. *Row 1* shows the true fire state $\mathbf{X}_t$ overlaid with the ground-truth frontier. *Row 2* shows the UAV positions $\{\mathbf{p}_t^{(i)}\}$ with FoV footprints. *Row 3* shows the shared belief map $\hat{\mathbf{X}}_t$ with the belief frontier in cyan, built incrementally from FoV observations. At $t = 0$ the belief map is entirely dark, reflecting complete initial ignorance $\hat{\mathbf{X}}_0 = \mathbf{0}$. As drones fly over the fire, observed cells are updated and the belief frontier narrows the gap with the ground-truth frontier.

Figure 3 reports four per-step metrics over the simulation steps. We evaluate the framework along four complementary dimensions. **Frontier Intersection over Union (IoU)** measures pixel-level agreement between the belief frontier $\hat{\mathbf{R}}_t$ and the ground-truth fire edge $\mathbf{R}_t$; it directly quantifies how accurately the fleet tracks the active fire boundary. **Fire IoU** measures the overlap between the belief fire body $\hat{\mathbf{X}}_t$ and the true burning area $\mathbf{X}_t$. It reflects how well the interior of the fire has been observed, complementing the boundary-focused frontier IoU. Further, **Sensor overlap rate** is the fraction of currently covered frontier pixels observed by two or more drones simultaneously, a low value indicates that the fleet is spatially distributed with minimal redundant sensing, which is the goal of the sector assignment mechanism. Finally, **Decision latency** is the wall-clock time spent by each drone computing its PSO velocity command, it is a direct indicator of real-time feasibility on a ground station without GPU acceleration.

5.2 Field deployment

Scenario: The framework was demonstrated on a fleet of $M = 3$ DJI Mini 3 quadrotors flying on the ENAC football field ($\approx 116\,\mathrm{m} \times 65\,\mathrm{m}$). The fire environment was simulated on a ground station using a 100×100 grid ($\Delta = 0.55\,\mathrm{m}$, with a total area of $55\,\mathrm{m} \times 55\,\mathrm{m}$), positioned as a square sub-region centred within the rectangular field boundary, as illustrated in Figure 4(c–d). Fire is ignited at the grid center and propagated every $N_f = 10$ steps. Drone sensing was simulated by the square FoV footprint $\mathcal{D}_t^{(i)}$ centred on each drone's GPS-reported position, so the shared belief map was updated from real flight telemetry. Velocity commands were forwarded to the flying drones via the WildBridge interface [17], which exposes the DJI Mobile SDK via HTTP and RTSP, following the architecture validated in prior work [18]. Guidance parameters for this scenario are $V = 2.5\,\mathrm{m/s}$, $s = 12$ cells ($6.6\,\mathrm{m}$), and $d_{\text{sep}} = 15\,\mathrm{m}$.

Results: This setup validates the complete guidance pipeline under real flight conditions. On the field, the navigation phase was extremely short, drones encountered the simulated fire frontier within the first few seconds of flight and PSO coverage activated almost immediately. Figure 4 shows the drones flying on the ENAC football field during the demonstration, the DroneViz visualization (developed in ENAC) displays real-time drone positions, alongside a snapshot of the simulated fire state and the first 80 steps of GPS trajectories. Figure 5 shows the corresponding metrics.

The small operational area plays a decisive role in the observed performance. The FoV footprints collectively cover a substantial fraction of the grid at every step, so the belief map saturates rapidly and frontier IoU approaches unity within a few steps. Similarly, fire IoU reaches ≈ 1.0 almost immediately since the entire burning body falls within the observable area. These near-perfect scores therefore reflect the scale of the deployment rather than a fundamental improvement over the large-scale simulation.

6 CONCLUSION

In this work, we presented a multi-UAV guidance framework for autonomous wildfire frontier tracking. The proposed approach combines a navigation strategy with a decentralized velocity-space PSO controller designed for dynamic

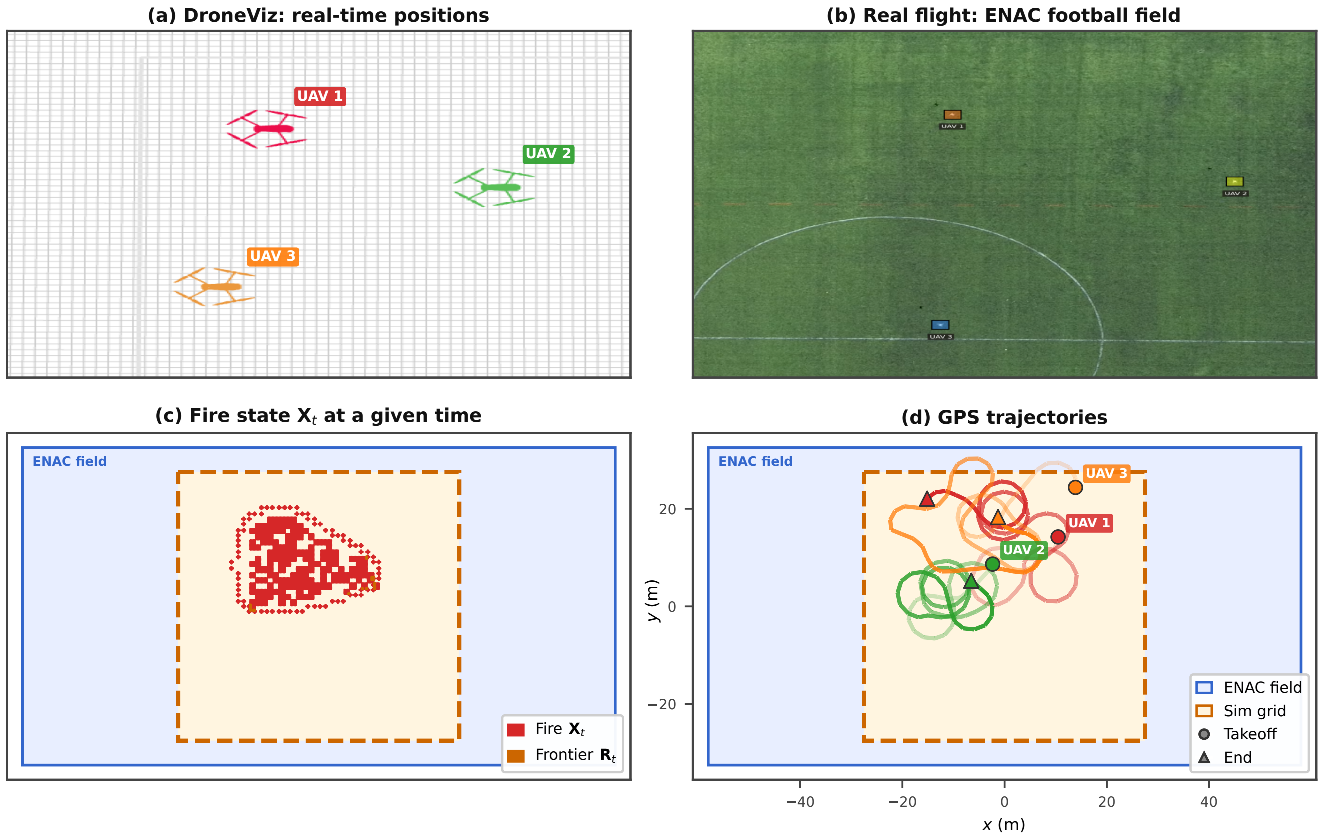

Figure 4: Field demonstration with $M = 3$ DJI Mini-3 quadrotors on the ENAC football field. Velocity commands from the guidance algorithm were forwarded via WildBridge framework. The fire environment and drone sensing were simulated on the ground station using a 100×100 grid ($\Delta = 0.55\,\mathrm{m}$).

http://www.imavs.org/

coverage under partial observability. The fitness function integrates exclusive frontier gain, exploratory sensing, sector attraction, collision avoidance, and boundary-awareness terms to promote efficient and coordinated swarm behavior. In addition, the framework's interfaces enable straightforward deployment on real UAV systems without additional command translation layers. Future work will include extensive comparisons with state-of-the-art planning and learning-based approaches, robustness analyses under communication and sensing uncertainties. Finally, an extension to support heterogeneous teams combining fixed-wing and quadrotors will be considered.

CODE AND DATA AVAILABILITY

Code is available at: `https://github.com/drones-fireflies/pso-guidance`

ACKNOWLEDGMENTS

The authors acknowledge the support of the French National Research Agency (Agence Nationale de la Recherche, ANR) under reference ANR-23-IAS2-0002-01 for the FireFlies TSIA ANR Project.

REFERENCES

[1] Chi Yuan, Youmin Zhang, and Zhixiang Liu. A survey on technologies for automatic forest fire monitoring, detection, and fighting using unmanned aerial vehicles and remote sensing techniques. *Canadian Journal of Forest Research*, 45(7):783–792, 2015.

[2] Eladio Altamira-Colado, Daniel Cuevas-González, Marco A. Reyna, Juan Pablo García-Vázquez, Roberto L. Avitia, and Alvaro R. Osornio-Vargas. Drone-Assisted Particulate Matter Measurement in Air Monitoring: A Patent Review. *Atmosphere*, 15(5):515, 2024.

[3] Yuqin Wang, Fengsen Gao, and Minghui Li. Probabilistic Path Planning for UAVs in Forest Fire Monitoring: Enhancing Patrol Efficiency through Risk Assessment. *Fire*, 7(7):254, 2024.

[4] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of ICNN95 - International Conference on Neural Networks*, pages 1942–1948, 1995.

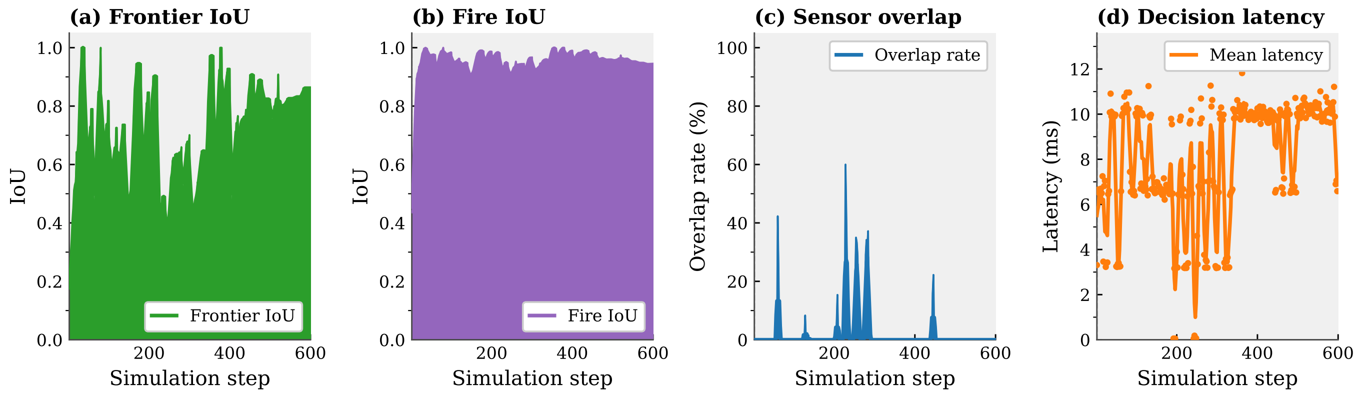

Figure 5: Per-step metrics in the field scenario with DJI Mini 3 quadrotors. *(a)* Frontier IoU over the full run. *(b)* Fire IoU, reflecting coverage of the burning body. *(c)* Sensor overlap rate. *(d)* Per-drone PSO decision latency

[5] Alessandro Giuseppi, Roberto Germanà, Federico Fiorini, Francesco Delli Priscoli, and Antonio Pietrabissa. UAV Patrolling for Wildfire Monitoring by a Dynamic Voronoi Tessellation on Satellite Data. *Drones*, 5(4):130, 2021.

[6] Rafael Bailon-Ruiz, Simon Lacroix, and Arthur Bit-Monnot. Planning to Monitor Wildfires with a Fleet of UAVs. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4729–4734, 2018.

[7] Constantinos Heracleous, Panayiotis Kolios, and Christos G. Panayiotou. UAV-based system for real-time wildfire perimeter propagation tracking. In *2023 31st Mediterranean Conference on Control and Automation (MED)*, pages 13–18, 2023.

[8] Rafael Bailon-Ruiz, Arthur Bit-Monnot, and Simon Lacroix. Real-time wildfire monitoring with a fleet of UAVs. *Robotics and Autonomous Systems*, 152:104071, 2022.

[9] Esmaeil Seraj, Andrew Silva, and Matthew Gombolay. Multi-UAV planning for cooperative wildfire coverage and tracking with quality-of-service guarantees. *Autonomous Agents and Multi-Agent Systems*, 36(2):39, 2022.

[10] Kyle D. Julian and Mykel J. Kochenderfer. Distributed Wildfire Surveillance with Autonomous Aircraft Using Deep Reinforcement Learning. *Journal of Guidance, Control, and Dynamics*, 42(8):1768–1778, 2019.

[11] Liyang Wang and Murat Bronz. Efficient decentralized multi-uav wildfire monitoring via mcts-distilled diffusion policies. In *14th International Conference on Unmanned Aircraft Systems (ICUAS)*, Corfu, Greece, 2026.

[12] Josy John and Suresh Sundaram. Genetic Algorithm-based Routing and Scheduling for Wildfire Suppression using a Team of UAVs. In *2024 IEEE Congress on Evolutionary Computation (CEC)*, pages 1–8, 2024.

[13] Sahand Khoshdel, Qi Luo, and Fatemeh Afghah. PyroTrack: Belief-Based Deep Reinforcement Learning Path Planning for Aerial Wildfire Monitoring in Partially Observable Environments. In *2024 American Control Conference (ACC)*, pages 601–607, 2024.

[14] Maxime Collignon, Adolfo Perrusquía, Antonios Tsourdos, and Weisi Guo. Search and rescue operations in wildfires using unmanned aerial vehicles: A multi-agent deep reinforcement learning approach. *Neurocomputing*, 653:131211, 2025.

[15] Hamza Chakraa and Murat Bronz. Accelerating Wildfire Spread Prediction for Guidance: A ConvLSTM Deep Learning Approach. In *AIAA SCITECH 2026 Forum*, Orlando, FL, 2026.

[16] John Canny. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):679–698, 1986.

[17] Edouard Rolland, Kilian Meier, Murat Bronz, Aditya Shrikhande, Tom Richardson, Ulrik Pagh Schultz Lundquist, and Anders Christensen. WildBridge: Ground station interface for lightweight multi-drone control and telemetry on dji platforms. *Proceedings of the 13th International Conference on Robot Intelligence Technology and Applications (RiTA 2025)*, 2025.

[18] Hamza Chakraa, Matthieu Verdoucq, João Machado, and Murat Bronz. A preliminary study on smoke plume observation with drone swarms. In *14th International Conference on Unmanned Aircraft Systems (ICUAS)*, Corfu, Greece, 2026.

http://www.imavs.org/