Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Dual-pheromone stigmergic framework for decentralized swarm search and rescue

Eashwar Sathyamurthy

eashwar.sathyamurthy@gmail.com

University of Maryland, College Park

Naitri Rajyaguru

University of Maryland, College Park

Yiannis Aloimonos

University of Maryland, College Park

Article

Keywords: swarm robotics, stigmergy, search and rescue, decentralized control, multi-robot coverage, bio-inspired coordination

Posted Date: November 5th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11601876/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Dual-pheromone stigmergic framework for decentralized swarm search and rescue

Eashwar Sathyamurthy[1*], Naitri Rajyaguru[2*] and Yiannis Aloimonos[1]

[1*]Mechanical Engineering, University of Maryland, College Park, MD, 20742, USA.
[2]University of Maryland Institute for Advanced Computer Studies, University of Maryland, College Park, MD, 20742, USA.

*Corresponding author(s). E-mail(s): eashwar.sathyamurthy@gmail.com; nrajyagu@umd.edu;
Contributing authors: jyaloimo@umd.edu;

Abstract

Robot teams search large areas faster by spreading out and avoiding repeated visits. Centralized planners coordinate this by assigning regions using information about robot positions and progress. However, maintaining these assignments becomes difficult when communication is unreliable. Decentralized control reduces this dependence, but robots lack shared records of where others have searched. We present STIGSAR, which uses stigmergy: robots leave environmental traces that others can sense. Temporary repulsive traces discourage revisits. When one robot discovers a target, it deposits a persistent attractive trace to guide others toward it. In simulation, STIGSAR approached centralized coverage performance and made 7–12× fewer visits per covered cell than random walk. Compared with single-trace variants that completed every mission, STIGSAR reduced median normalized mission cost by 38–45%. Robotarium experiments demonstrated recruitment with virtual traces. RoboMaster robots deposited and sensed physical repulsive trails, continuing search after robot loss without direct communication or online task allocation.

Keywords: swarm robotics, stigmergy, search and rescue, decentralized control, multi-robot coverage, bio-inspired coordination

1 Introduction

Multi-robot search is most effective when robots spread across an environment and avoid repeatedly covering the same regions. Centralized planners coordinate this process by using information about robot positions and search progress to divide the workspace and assign regions to individual robots [1, 2]. However, maintaining these assignments requires continuous updates as robots move or complete their tasks. When communication becomes unreliable or robots fail, this information may become outdated, leaving the planner with an inaccurate record of the remaining search area. Decentralized control avoids dependence on a central planner, but robots acting on local information may not know which regions others have already searched [3, 4]. The challenge is therefore to coordinate coverage without relying on continuous communication, while retaining information about the team's previous activity.

This need for shared information becomes more important when the team's objective changes during a mission. We consider a search task in which robots must locate targets whose positions are initially unknown. A target represents a location of interest, such as a hazardous site, where additional robots may be required after discovery. Before a target is found, robots should disperse to cover different regions. Once it is detected, however, additional robots may need to converge on that location [5]. We refer to this process as post-discovery recruitment, in which additional robots are guided toward a detected target until a prescribed number arrive. The team therefore needs to retain two kinds of information: where robots have recently searched and where a target has been discovered. Figure 1 illustrates this transition using a wildfire-response scenario.

Stigmergy offers a way to retain this information in the environment. In a stigmergic system, an agent modifies its surroundings, and other agents later respond to that modification [6–9]. Ant pheromone trails are a familiar biological example and



Fig. 1: Stigmergic swarm search and recruitment in a wildfire-response scenario.

have inspired distributed routing, exploration, coverage, and patrolling [10, 11]. In robotics, these traces need not be chemical; they can be represented by physical markers, RFID-based traces, or virtual fields that store information at specific locations in the workspace [12–16]. Because the information is associated with a location rather than an individual robot, other robots can access it when they encounter that location, without direct communication with the robot that left it.

Repulsive pheromone traces have been used to reduce repeated exploration [3], while RFID-based pheromones have supported target search and tracking with physical robots [16]. Earlier virtual-pheromone systems combined attraction and repulsion and described target-related signaling for disaster-response scenarios [17]. Multiple-pheromone platforms such as ColCOSΦ allow distinct signals to drive different robot responses [18]. Physical stigmergic trails have also recently been used to regulate group formation without direct inter-robot communication [19]. More recent methods extend pheromone-based coordination to dynamic coverage [20], hazardous exploration [21], and asynchronous swarm coordination [22]. SPACE combines attractive frontier guidance with a repellent exploration pheromone and a robot-density field for large-scale collision-aware exploration [23]. Although robot actions are computed locally, SPACE assumes a shared occupancy grid maintained by a collaborative mapping backend, and its objective is exploration rather than post-discovery recruitment. Recent pheromone-based search-and-rescue work has also used separate virtual fields for coverage, victim information, environmental risk, and communication quality within a learning framework [24]. These fields support decentralized victim search and detection within a multi-agent reinforcement-learning framework, rather than recruitment of additional robots after target discovery.

Although these studies demonstrate how environmental traces can coordinate robot behavior, search and recruitment place different demands on the information they carry. During search, traces should discourage robots from revisiting recently explored regions. However, this information must gradually lose influence so that robots can return to previously searched areas when necessary. After target discovery, a different signal is needed to attract robots toward the detected location. Unlike the search trace, this signal must remain available long enough for other robots to encounter and follow it. The two behaviors therefore require environmental memories that differ in both their function and persistence. This raises the question of whether separating these functions into distinct environmental memories improves coordination when a robot team transitions from search to recruitment.

We present STIGSAR, a decentralized stigmergic framework that uses two pheromone fields to retain these different kinds of information. During search, robots deposit a temporary repulsive field that records recent visits and discourages redundant coverage. When a robot discovers a target, it deposits a separate persistent attractive field along an outward spiral. This extends the signal into the surrounding area, where other searching robots can detect it and use attractive-pheromone measurements to move toward the target. Each robot independently switches between search, advertisement, and following based on target detection and available pheromone information, without online task allocation or direct robot-to-robot communication.

The main contributions are:

- A dual-pheromone stigmergic controller that uses environmental memory for decentralized search and recruitment. Temporary repulsive traces discourage repeated coverage, while persistent attractive traces guide robots toward discovered targets.
- Theoretical analysis of how pheromone evaporation influences local search decisions and sufficient conditions for finite-time target following. Controlled experiments compare the complete controller with designs that remove, merge, or reverse the functions of the two pheromone fields.
- Evaluation in simulation and on physical robots. Simulation experiments examine search performance, field design, recruitment, and robot loss across increasing workloads. Robotarium experiments validate decentralized coordination with virtual pheromone fields, while RoboMaster experiments demonstrate repulsive search and continued operation after robot loss using physically deposited and sensed traces.

2 Results

STIGSAR uses traces left in the environment to discourage repeated search and guide robots toward discovered targets. We evaluated whether this information improved search coverage, helped robots reach detected targets, and remained useful when robots stopped operating. Simulations allowed controlled comparisons, while experiments with physical robots tested how the same behaviors operated beyond simulation.

2.1 Experimental design and validation platforms

We used three experimental environments to examine different aspects of the framework. Simulation allowed us to change workspace size, team size, and target number, and to test what happened when one or both pheromone traces were modified. Robotarium [25] introduced physical robot motion while keeping the traces on a virtual grid linked to robot positions through overhead tracking. DJI RoboMaster robots [26] used UV-visible traces [27] that were deposited and sensed in the arena. This last platform tested the physical implementation of the repulsive search trace, rather than the full search-and-recruitment mechanism. Platform-specific sensing and motion are described in Section 4.1.1. The simulation and RoboMaster experiments were run on a computer equipped with an 11th-generation Intel Core i9-11950H processor (2.60 GHz) and 32 GB of RAM, while Robotarium experiments were conducted remotely on the Georgia Institute of Technology's Robotarium platform.

Table 1: **Experimental environments used for each behavioral evaluation.**

Evaluation	Simulation	Robotarium	RoboMaster
Search (Section 2.3)	✓	✓	✓ (Section 2.6)
Field ablation (Section 2.2)	✓	✓	
Recruitment (Section 2.4)	✓	✓	
Robot loss (Section 2.5)	✓	✓	✓ (Section 2.6)

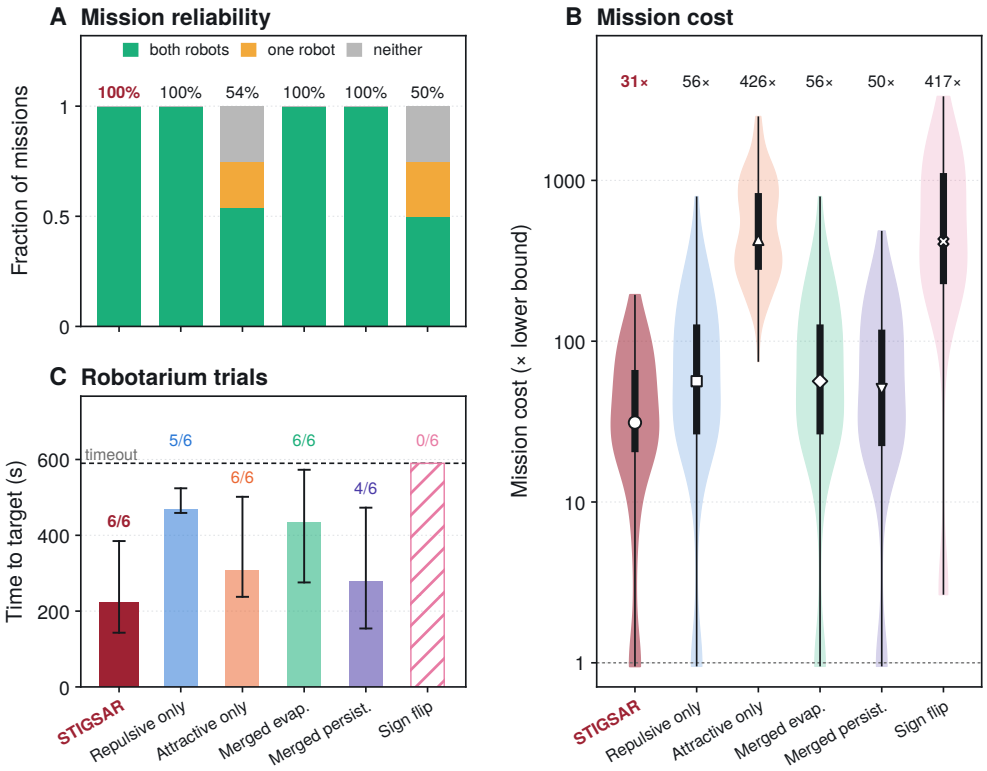


Fig. 2: **Field ablation.** **a**, Mission reliability across 240 simulations per variant (95% Wilson intervals). **b**, Mission cost normalized by the nearest initial robot–target Manhattan distance; timeouts are retained at the mission horizon, with medians and 95% bootstrap intervals shown. **c**, Robotarium time to target and success rate; labels show completed/total trials.

Table 1 summarizes these evaluations; further experimental details are given in Section 4 and the Supplementary Information.

For search, we compared STIGSAR with centralized lawnmower coverage, which assigns regions using global information, and decentralized random walk, which has no shared record of previous visits. These methods provide references for planned coverage and uncoordinated exploration, respectively.

2.2 Field ablation

We first tested whether separating the two traces affected mission completion and efficiency. Robots searched for a target at an initially unknown location, and a trial was successful only when every robot's target encounter had been registered before the time limit. The simulations varied the number of robots and their initial distances from the target. Each approach was tested on the same 240 mission configurations, so differences in outcome could be attributed to the pheromone arrangement.

We compared six versions of the controller (Fig. 2). Full STIGSAR used both the temporary repulsive and persistent attractive traces. Two versions removed one trace: *repulsive only* retained the search trace, whereas *attractive only* retained the target-directed trace. Two further versions used a single trace for both behaviors, either allowing it to evaporate (*merged evaporative*) or keeping it persistent (*merged persistent*). The *sign-flip* version also used one trace but changed whether robots avoided or followed it after discovery.

In simulation, STIGSAR, repulsive only, and both merged versions each achieved 100% mission reliability across the 240 trials (Fig. 2a). Attractive only and sign flip reached 54% and 50% reliability, respectively (129/240 and 119/240 missions). Thus, eventual completion did not distinguish STIGSAR from the other three fully successful designs. The remaining distinction was the time required to complete the mission.

To account for differences in initial target distance, we divided the number of simulation steps until all target encounters were registered by the shortest initial Manhattan distance between any robot and the target. Runs that did not finish were retained at the mission time limit. STIGSAR had a median normalized cost of 31.2, compared with 56.4 for repulsive only, 56.4 for merged evaporative, and 50.5 for merged persistent (Fig. 2b). Among the versions that completed every simulated mission, separating the traces reduced median cost by 38–45%. Attractive only and sign flip had higher median costs of 426.5 and 416.7, respectively, with timeouts included.

The Robotarium results showed a similar difference between mission completion and the time required (Fig. 2c). STIGSAR, attractive only, and merged evaporative completed all six trials each, with reported median completion times of 224, 308, and 434 s, respectively. Repulsive only completed 5/6 trials with a completed-run median of 469 s, and merged persistent completed 4/6 with a completed-run median of 277 s. Sign flip completed 0/6 trials. STIGSAR therefore had the shortest reported median completion time among the three methods that completed every Robotarium trial.

2.3 Search

We evaluated how the temporary repulsive trace affected coverage, target discovery, and repeated visits while robots searched for targets distributed across the workspace. Each trial ended when all targets were found or when the allowed mission time, $T_{\max}$, was reached.

In the representative simulation, STIGSAR and centralized coverage approached complete coverage and found all 90 targets within approximately the first tenth of the allowed mission time (Fig. 3a,b). Random walk covered the workspace and found targets more slowly, reaching about 95% coverage only near the end of the time limit.

We then increased grid side length from 50 to 500 cells, team size from 5 to 50 robots, and the number of targets from 15 to 150 (Fig. 3c). STIGSAR remained close to the centralized reference in the steps needed to reach 95% coverage, whereas random walk required approximately an order of magnitude more steps in the largest configuration. Because grid size, robot number, and target number increased together, this comparison describes performance under increasing combined workloads rather than the effect of any single variable.

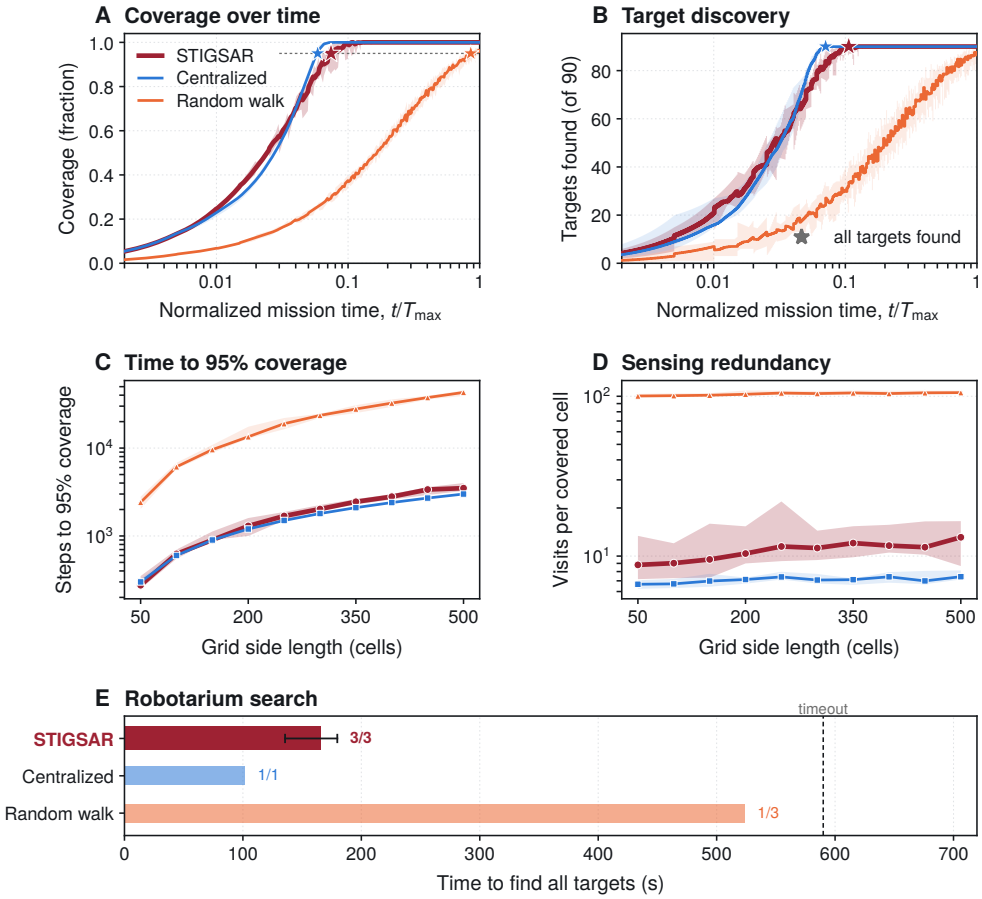


Fig. 3: **Search.** **a**, Coverage and **b**, target discovery over normalized mission time for a representative simulation configuration; lines show medians and shading the 2.5th–97.5th percentiles, with stars marking complete target discovery. **c**, Steps to 95% coverage and **d**, visits per covered cell across ten jointly scaled configurations of grid size, robot number, and target count; medians and 95% bootstrap intervals are shown. **e**, Robotarium time to find all targets; labels show completed/total trials and the dashed line marks the mission timeout.

The difference in repeated search was also substantial (Fig. 3d). The centralized planner made approximately 6–7 visits per covered cell, STIGSAR 8–14, and random walk about 100. STIGSAR therefore made approximately 7–12 times fewer visits per covered cell than random walk, although it remained more redundant than the globally coordinated sweep.

In Robotarium, STIGSAR found all ten targets in 3/3 trials, compared with 1/3 for random walk; the centralized reference completed its single trial (Fig. 3e). The figure reports elapsed time in seconds and retains incomplete trials at the time limit.

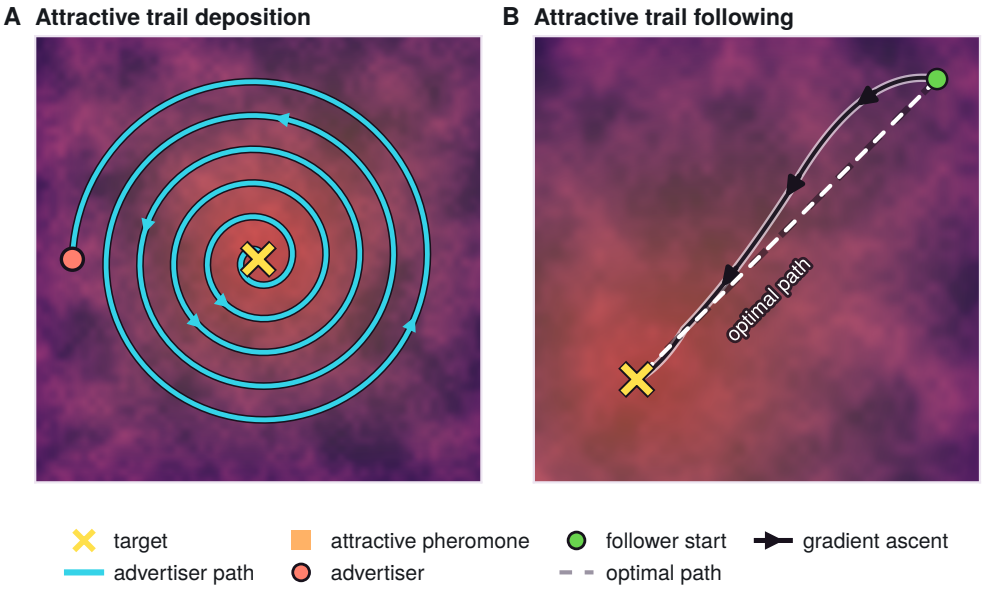


Fig. 4: **Recruitment. a**, The advertising robot deposits a persistent attractive field along an outward spiral after target discovery. **b**, A robot encountering the attractive field enters FOLLOW and moves toward the target using local field measurements, as shown by the black trajectory; the dashed line shows the Manhattan shortest path.

STIGSAR and the centralized reference completed their searches sooner than random walk, whose two unsuccessful trials reached the time limit. The single centralized trial serves as a deterministic reference rather than a statistical comparison.

2.4 Recruitment

We evaluated whether the attractive trace could guide a searching robot toward a detected target using local measurements. After finding a target, the first robot deposited the trace along an outward spiral so that other robots could encounter it beyond the target's immediate surroundings (Fig. 4a). A searching robot that detected the trace switched to FOLLOW and moved toward stronger attractive-pheromone measurements (Fig. 4b). To assess following rather than the time spent finding the trace, measurements began when the robot first detected it.

In Robotarium, the following robot reached the target in all six trials: 3/3 with an initial target separation of 30 grid cells and 3/3 at 40 cells. Figure 4b illustrates a representative following path alongside the Manhattan shortest-path reference.

Physical execution required a mean of 929 controller iterations in the 30-cell condition and 892 in the 40-cell condition. These are updates made while executing robot motion, not the number of grid cells travelled. They therefore describe physical execution rather than a distance-independent recruitment time.

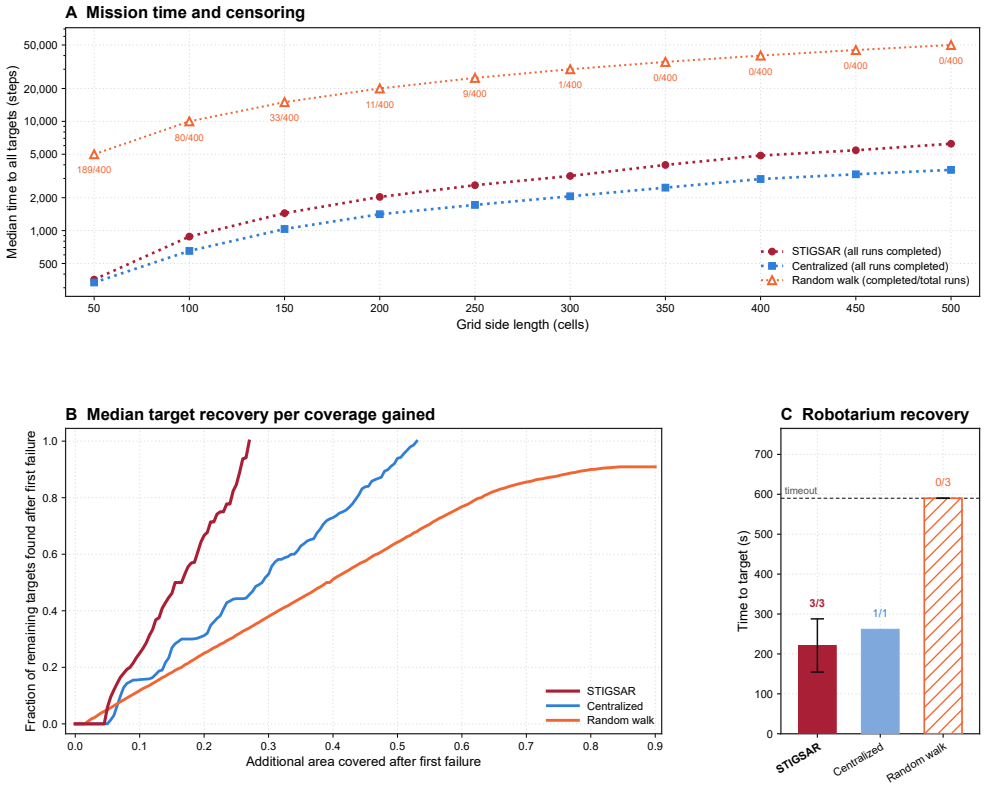


Fig. 5: **Robot loss. a**, Median mission time across ten jointly scaled simulation configurations after robot loss. Unfinished runs are retained at the mission horizon; labels show completed/total random-walk runs. **b**, Fraction of targets remaining at the first failure that were subsequently found, plotted against additional workspace coverage after failure. **c**, Robotarium time to find all targets after two of five robots were stopped. Bars show completion times, labels indicate completed/total trials, and the dashed line marks the timeout.

2.5 Robot loss

We also tested whether search could continue when some robots stopped without notifying the others. During the target-search task, selected robots stopped moving and depositing pheromone, but the remaining robots received neither a failure notification nor new search assignments. STIGSAR continued using environmental traces available at the time of failure and those deposited by the surviving robots.

Figure 5a shows the median mission time after robot loss across ten jointly scaled simulation configurations. STIGSAR and the centralized reference completed all tested runs, whereas random walk completed fewer missions as the workload increased. STIGSAR maintained mission completion without the reassignment used by the centralized reference.

Figure 5b compares the fraction of targets remaining at the first failure that were subsequently found as the surviving robots covered additional area. For a given amount of additional coverage, STIGSAR recovered a larger fraction of the remaining targets than the two reference methods in the plotted comparison. The curves describe target recovery relative to newly covered area, not additional coverage over time or total workspace coverage.

In Robotarium, two of the five robots were stopped during each search trial (Fig. 5c). STIGSAR still found all ten targets in 3/3 trials. Random walk completed none of its three trials, finding 2, 7, and 4 targets, respectively, while the centralized reference completed its single trial. These results demonstrate continued operation without explicit failure detection or reassignment, but do not isolate pheromone persistence as the sole reason for that behavior.

2.6 RoboMaster validation with physical pheromone trails

The preceding Robotarium experiments used virtual traces. To test whether robots could deposit and detect a trace directly in their surroundings, we used RoboMaster robots with UV-visible repulsive trails (Fig. 6a). The available setup was used to evaluate the search trace only; the attractive trace and recruitment were not physically implemented on this platform. Representative Robotarium recruitment and search trajectories are also shown in Fig. 6b,c.

We used two RoboMaster robots to test whether a small team could achieve broad coverage in a 4.0 m × 1.8 m arena, where the deposited trails remained detectable for approximately 15 s under our indoor conditions. For the robot-loss condition, three robots began searching and one was stopped after 60 s, leaving two to continue without notification. Each condition lasted 240 s and included three trials each for STIGSAR and random walk, together with one deterministic centralized reference trial.

In the search trials, STIGSAR covered an average of 78.8% (±5.5%) of the arena, compared with 60.6% (±5.7%) for random walk and 99.1% for the centralized reference. Figure 6d shows representative trials with 87.8% coverage for STIGSAR, 62.3% for random walk, and 99% for centralized coverage. Values are reported as mean ± sample standard deviation across three trials for STIGSAR and random walk; the centralized reference was evaluated once. The representative paths show broader exploration under STIGSAR than random walk, while centralized sweeping retained the highest coverage.

After one robot was stopped in the robot-loss trials, STIGSAR achieved 82.8% (±4.1%) mean final coverage, compared with 76.1% (±6.9%) for random walk and 83.5% for the centralized reference. STIGSAR reached 60% coverage after approximately 125 s (±7.3 s), compared with approximately 195 s (±52.2 s) for random walk. The robot-loss trials began with three robots whereas normal-search trials began with two; their coverage values therefore should not be interpreted as evidence that losing a robot improved performance.

These experiments establish that the repulsive search trace can be deposited and sensed by physical robots, including after one robot stops operating. The persistent attractive trace was evaluated with virtual fields in Robotarium, while its physical deposition and sensing were not evaluated in this study.

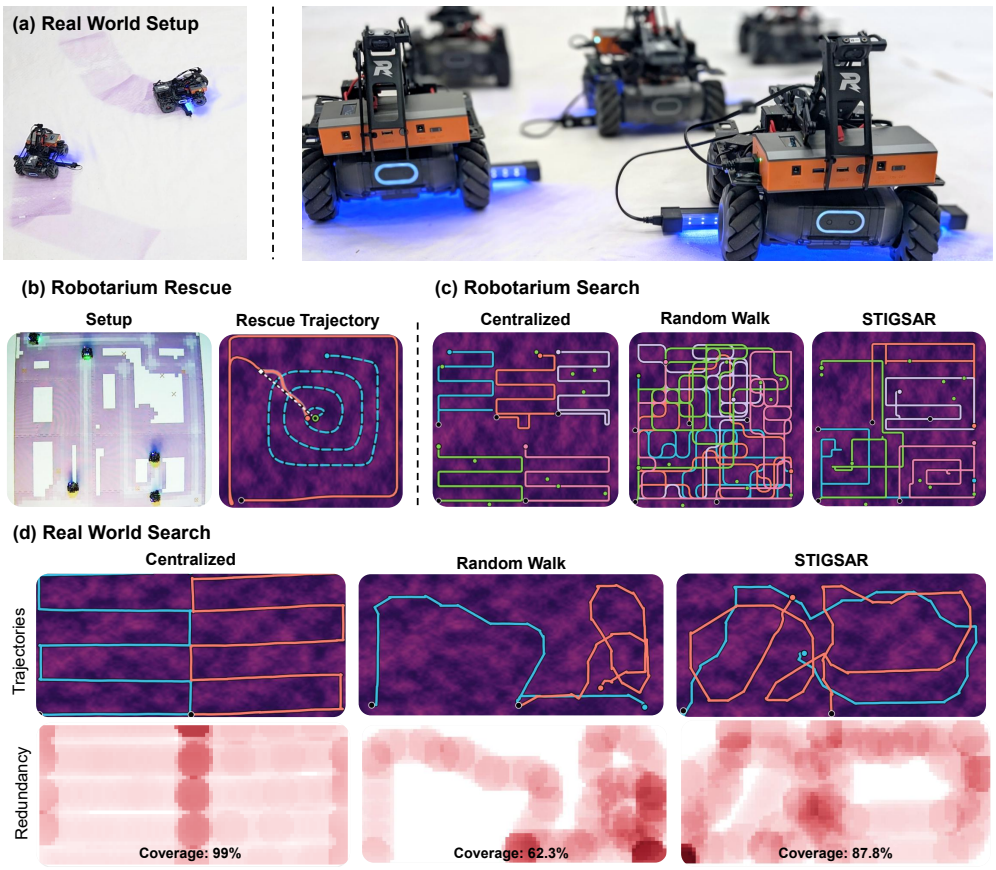


Fig. 6: **Physical platforms. a**, RoboMaster setup and UV-visible trail. **b**, Representative Robotarium target-response setup and trajectory. **c**, Representative Robotarium search trajectories for centralized coverage, random walk, and STIGSAR. **d**, Representative RoboMaster trajectories and visitation maps for the same three methods; percentages report covered area in the shown trials.

3 Discussion

The experiments show that environmental memory can support decentralized coordination, but its usefulness depends on the information retained and how long that information remains available. During search, recent visits should discourage other robots from covering the same regions. After a target is found, its location must remain available while additional robots move toward it. In the controlled comparisons, STIGSAR reduced median normalized mission cost by 38–45% relative to the alternative designs that completed all simulated missions. Because these alternatives also completed every mission, the main difference was the time required to complete the search-and-recruitment task. The result supports separating the two memories

under the tested conditions, rather than showing that coordination with a single trace is impossible.

Earlier robot systems have demonstrated attractive and repulsive pheromones, multiple signalling channels, and coordination through environmental traces [17–19, 23, 24]. STIGSAR builds on these mechanisms by testing their roles within a mission that requires robots to first spread across a workspace and then reach a common location after a discovery. The contribution is therefore not the use of multiple pheromones itself, but experimental evidence that separating recent search history from persistent target information can improve the efficiency of this change in collective behavior.

The robot-loss experiments show a further consequence of storing coordination information outside individual robots. When robots stopped without notifying the team, the others continued searching without receiving replacement assignments. Traces deposited before a failure could remain available to guide subsequent decisions, so the team did not need to reconstruct its search history through direct communication. These observations do not establish that trace persistence alone caused the observed tolerance to robot loss. They also do not imply that STIGSAR consistently outperforms centralized planning, which remained an effective coverage reference. Rather, they show that useful coordination can continue without explicit failure detection or online reassignment in the tested environments.

The physical experiments distinguish decentralized decision-making from the means used to store and sense environmental information. In Robotarium, robots made local decisions using virtual pheromone fields registered through overhead tracking. Although no centralized search assignments were computed, this implementation still depended on external localization and a maintained virtual field. The RoboMaster experiments instead used a UV-visible trace deposited and sensed directly in the arena, where it remained detectable for approximately 15 s. By comparison, Salman *et al.* demonstrated photochromic stigmergy using eight e-puck robots and traces that faded in approximately 50 s [27]. These studies used different tasks and experimental conditions, so their trail lifetimes and performance should not be treated as a controlled comparison. Our physical results establish that the repulsive search trace can support coverage and continued operation after robot loss; they do not yet validate a physically deposited attractive recruitment trace.

Several limitations remain. The experiments used bounded, connected, and largely obstacle-free environments, and the number of physical trials was small. The finite-time analysis of target following applies only when the attractive signal remains detectable and the robot's local movement choices satisfy the stated conditions; it is not an unconditional guarantee in cluttered settings. Obstacles, localization error, uneven deposition, and imperfect sensing could all change how reliably robots interpret environmental traces. Physical photochromic systems also require a prepared surface and appropriate illumination, while marker technologies more generally differ in sensing requirements, information capacity, and deployment cost. Realizing both STIGSAR signals physically would require traces that can coexist, be detected independently, and retain different persistence times.

The wildfire-response example in Fig. 1 illustrates a possible need for robots to disperse during search and converge after a discovery; the experiments do not model wildfire dynamics or suppression. The underlying coordination problem also arises in tasks such as marking hazards or previously inspected regions. Whether the same benefits extend to those applications remains to be established. Future experiments should therefore examine obstacle-rich environments, uncertain localization and trail sensing, and a physical implementation of both memories on the same substrate. These findings suggest that the information left in the environment, and how long it lasts, can shape how a decentralized team changes its collective behavior.

4 Methods

STIGSAR coordinates search and post-discovery recruitment through environmental traces rather than direct robot-to-robot messages or online task assignments. To explain how the controller uses this shared information, we first define what each robot knows and how the two pheromone fields are updated. We then describe how robots search, advertise a detected target, and follow a trace toward it, before presenting the theoretical analysis and experimental protocols.

4.1 Problem formulation and operating assumptions

We consider a team of N robots searching a bounded two-dimensional workspace for targets at initially unknown locations. Robots first explore to discover these targets. When a target is detected, the discovering robot advertises its location so that additional robots can find it. We study the full search-and-recruitment controller in grid-based simulation and on physical robots in Robotarium [25], using virtual pheromone fields. We separately test physical deposition and sensing of the repulsive search trace on DJI RoboMaster robots [26, 27]; this platform does not implement attractive recruitment.

For the grid-based controller in simulation and Robotarium, let the workspace be represented by an $H \times W$ rectangular grid,

$$\mathcal{C} = \{1, \dots, H\} \times \{1, \dots, W\}, \tag{1}$$

where $\mathbf{c} \in \mathcal{C}$ denotes a cell. Simulated robots move between neighboring cells, while Robotarium robots move continuously toward physical waypoints associated with these cells. RoboMaster robots also move continuously, but select their movements using color-pixel observations from an onboard camera rather than the grid-based pheromone controller. A separate spatial grid is used to evaluate RoboMaster coverage, not to navigate. All tested workspaces were bounded and contained no static obstacles. We describe each platform's sensing and movement rules below.

The virtual-field controllers in simulation and Robotarium update at discrete steps $t \in \mathbb{Z}_{\geq 0}$. Robots use a sensing radius R, measured in grid cells, to detect targets and attractive pheromone; the search rule also accesses repulsive-field values around candidate next positions, as explained below. RoboMaster instead senses the physical trace through a forward-facing camera. In the virtual controller, a searching robot refreshes

repulsive pheromone at its occupied cell. After detecting a target, an advertising robot assigns attractive pheromone to cells in its surrounding sensing region (Algorithm 2).

Each robot in simulation and Robotarium maintains a private binary memory of previously visited cells. For robot r, this memory is defined as

$$M_r(\mathbf{c}, t) = \begin{cases} 1, & \text{robot } r \text{ occupied cell } \mathbf{c} \text{ at or before time } t, \\ 0, & \text{otherwise.} \end{cases} \tag{2}$$

This memory is private to robot r. It records where that robot has been, whereas the repulsive field retains information deposited by other searching robots as well. The two records therefore serve different purposes when the robot selects its next movement. Private visit memory is used in simulation and Robotarium; RoboMaster instead responds to the physical trace observed by its camera.

4.1.1 Robot states, local sensing, and actions

Although the platforms evaluate the same principle of environmental memory, they differ in what a robot can sense and how it moves. Table 2 compares these interfaces, and Fig. 7 illustrates their sensing regions and available actions.

In simulation and Robotarium, the local sensing region of robot r at time t is defined as

$$\mathcal{V}_R(\mathbf{p}_r(t)) = \{\mathbf{c} \in \mathcal{C} \mid \|\mathbf{c} - \mathbf{p}_r(t)\|_\infty \leq R\}, \tag{3}$$

where $\mathbf{p}_r(t)$ is the robot's current grid position and R is its sensing radius in cells. The region contains up to $(2R+1)^2$ cells and is used for target detection and attractive-pheromone measurements. When selecting a search movement, however, the robot evaluates repulsive pheromone around each candidate next cell. These candidate regions can extend beyond $\mathcal{V}_R(\mathbf{p}_r(t))$, so their values are read directly from the virtual field in simulation and Robotarium. Thus, the decisions are decentralized, but the virtual search rule is not restricted to measurements inside the robot's current sensing region. The repulsive and attractive deposition rules are defined in Section 4.1.2 and Section 4.4, respectively.

In the simulation, the state of robot r is represented by its grid position at time step t,

$$\mathbf{p}_r(t) = \begin{bmatrix} x_r(t) & y_r(t) \end{bmatrix}^\mathsf{T}.$$

The available robot actions are moving one cell north, south, east, or west per time step. The simulation treats robots as points and enforces separation, so it does not use an additional collision-avoidance controller.

In the Robotarium experiments, the physical state includes position and heading, while $\mathbf{p}_r(t)$ continues to denote the two-dimensional grid location used by the pheromone controller,

$$\mathbf{s}_r(t) = \begin{bmatrix} x_r(t) & y_r(t) & \theta_r(t) \end{bmatrix}^\mathsf{T}.$$

The robot selects a neighboring grid cell and converts it into a physical waypoint. Its position determines the waypoint, while its heading $\theta_r(t)$ is used by the unicycle motion controller to generate linear and angular velocity commands. Overhead tracking provides robot positions and maintains the virtual pheromone field,

Table 2: Controller implementation in the three experimental environments.

Component	Simulation	Robotarium	DJI RoboMaster
Robot state	2D Position	2D Position and heading	2D Position and heading
Actions	North, south, east, west	Adjacent-cell waypoints	Left, straight, right
Local sensing region	Cells within the sensing radius	Cells within the sensing radius	Forward 120° field of view
Collision avoidance	Not applicable	Collision radius	Forward 120° field of view
Pheromone representation	Pheromone-cell intensity	Virtual pheromone grid	UV-visible physical trail

but does not assign search regions. Robotarium's control-barrier-function safety certificates [28] provide collision avoidance by enforcing a minimum separation between robots (Fig. 7b).

RoboMaster does not use fixed grid directions. Instead, it divides the forward-facing 120° camera view into three sectors and uses detected trail-colored pixels to select a left, straight, or right movement. As the robot moves, its UV source activates the photochromic substrate along its path, leaving a trace that other robots can sense. The selected movement is executed through turning and forward commands with both forward collision checks and arena-boundary checks (Fig. 7c). Position and heading are estimated through wheel odometry. These experiments evaluate search coverage and continued operation after robot loss, not physical target detection or recruitment.

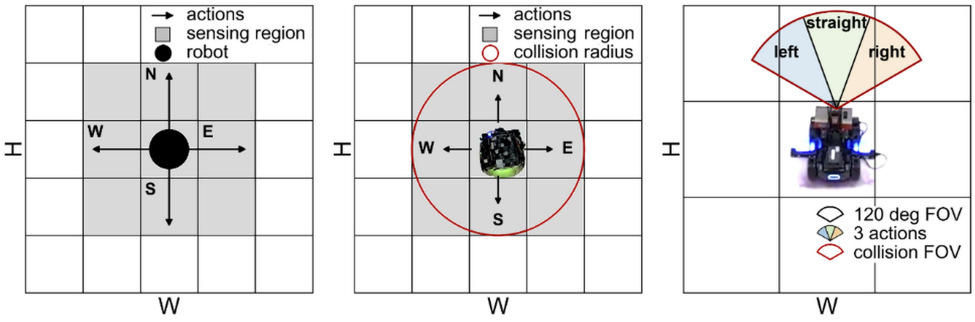


Fig. 7: Local sensing and actions across experimental environments. a, Simulated robot with four cardinal actions and local sensing region. **b**, Robotarium robot with four neighboring waypoints and collision radius. **c**, RoboMaster with left, straight, and right actions within a 120° field of view and forward collision checking.

The common feature is that deposited information influences subsequent search decisions. We now define the two virtual fields used together in simulation and Robotarium; RoboMaster implements only the physical repulsive-trace counterpart.

4.1.2 Dual-pheromone fields

In simulation and Robotarium, the two types of information are represented by separate pheromone fields. The repulsive field stores recent search activity, while the attractive field retains information leading to a detected target. Each field assigns a normalized concentration to grid cell $\mathbf{c}$ at time t,

$$\Phi_{\text{rep}}(\mathbf{c}, t) \in [0, 1], \qquad \Phi_{\text{att}}(\mathbf{c}, t) \in [0, 1],$$

where zero indicates the absence of pheromone.

To record search activity, each robot in SEARCH refreshes repulsive pheromone at its occupied cell. Let $\mathcal{R}_{\text{Search}}(t)$ be the set of robots currently searching. The cells refreshed at step t are therefore

$$\mathcal{D}_{\text{rep}}(t) = \bigcup_{r \in \mathcal{R}_{\text{Search}}(t)} \{\mathbf{p}_r(t)\}, \tag{4}$$

Each refreshed cell is assigned a concentration of one. Where no robot refreshes the field, its concentration decays from the preceding step according to

$$\Phi_{\text{rep}}(\mathbf{c}, t+1) = \begin{cases} 1, & \mathbf{c} \in \mathcal{D}_{\text{rep}}(t), \\ \lambda_{\text{rep}} \Phi_{\text{rep}}(\mathbf{c}, t), & \mathbf{c} \notin \mathcal{D}_{\text{rep}}(t). \end{cases} \tag{5}$$

The decay factor λ_{rep} is defined by

$$\lambda_{\text{rep}} = \exp\left(-\frac{1}{\tau_{\text{rep}}}\right) \tag{6}$$

where $\tau_{\text{rep}} > 0$ controls how long the repulsive pheromone persists after deposition. Larger values produce slower decay, retaining information about recent robot visits for longer.

A fixed decay time would not account for changes in workspace size, robot number, or sensing radius. We therefore selected the decay timescale empirically as

$$\tau_{\text{rep}} = \frac{HW}{NR}, \tag{7}$$

where H and W are the grid dimensions, N is the number of robots, and R is the sensing radius in grid cells. This empirically selected rule increases pheromone persistence with workspace area and decreases it with team size or sensing radius. Its sensitivity is evaluated in Supplementary Section 2.

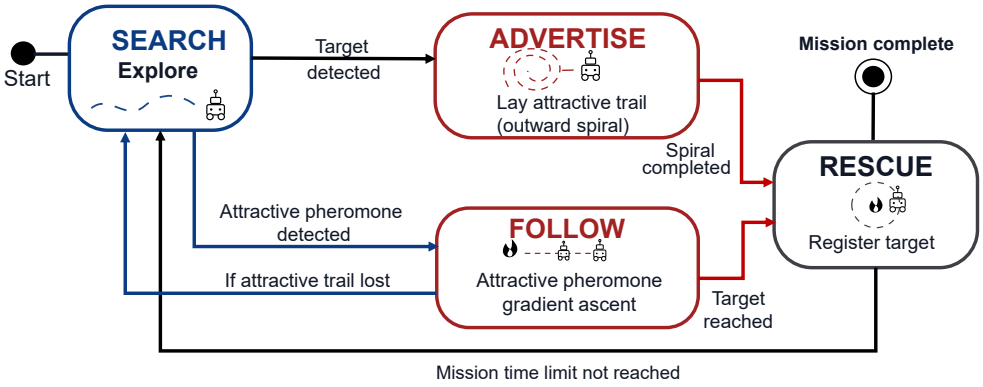


Fig. 8: Local state transitions governing search, target advertisement, convergence, rescue, and mission termination.

Unlike the repulsive field, the attractive field guides robots toward detected targets and remains persistent during recruitment. Its decay factor is therefore set to

$$\lambda_{\text{att}} = 1. \tag{8}$$

These settings let search traces lose influence while target-related information remains available for recruitment. The RoboMaster repulsive trace also fades, but its approximately 15-second detectability is a property of the physical material and is not generated by the virtual decay equation.

4.2 Dual-pheromone stigmergic behavior

The virtual-field controller follows the state machine in Fig. 8. All robots begin in SEARCH, using repulsive pheromone and private visit memory to explore (Section 4.3). When one robot detects a target, it changes to ADVERTISE and extends an attractive trace outward from that location (Section 4.4). The other searching robots do not need to know where the target is in advance: a robot that encounters the attractive trace changes to FOLLOW and moves toward the target (Section 4.5). For mission evaluation, detecting a target within the prescribed target-detection radius counts as arrival. A follower therefore enters RESCUE and registers its encounter upon detection. The discovering robot registers its earlier encounter only when advertisement ends, without having to return to the target. Simulation tracks these records for evaluation, and recruitment is complete when the prescribed number of robots has registered the target. A run then terminates, or ends at the maximum mission horizon if recruitment is incomplete. The robots select their behaviors from sensed information rather than from new task assignments.

4.3 Repulsive-pheromone search behavior

The robot first checks whether its sensing region contains a target. If so, it stops searching and enters ADVERTISE to make the discovery available to the others. If no target is detected but an attractive trace is present, it enters FOLLOW. Otherwise, the robot remains in SEARCH and chooses its next movement using repulsive pheromone and its private visit memory (Algorithm 1). Target detection therefore takes priority over following, and following takes priority over further exploration.

To compare possible search movements, the robot first identifies its valid neighboring cells $\mathbf{c} \in \mathcal{N}$. For each candidate, it sums the repulsive pheromone in the region around that cell at time t:

$$u_1(\mathbf{c}) = \sum_{\mathbf{q} \in \mathcal{V}_R(\mathbf{c})} \Phi_{\text{rep}}(\mathbf{q}, t),$$

together with the local visit-memory value

$$u_2(\mathbf{c}) = M_r(\mathbf{c}, t).$$

The next cell is selected using lexicographic minimization,

$$\mathbf{next} = \operatorname*{lexmin}_{\mathbf{c} \in \mathcal{N}} \left(u_1(\mathbf{c}), u_2(\mathbf{c}) \right).$$

The repulsive score is considered first, favoring regions with less recent search activity. If two candidates have equal repulsive scores, the robot prefers one it has not visited before. When all valid neighboring cells have already been visited, Algorithm 1 instead selects a random neighbor for that step. The robot checks the same conditions again at its next search step.

Before moving, the robot marks its current cell in private memory and refreshes the repulsive trace there. Pheromone elsewhere decays according to Eq. 5. As a result, recently visited regions have a stronger influence on subsequent search decisions than regions visited earlier, without permanently excluding either from exploration.

4.4 Attractive pheromone advertiser behavior

After detecting a target, the discovering robot switches from search to ADVERTISE. Its role is now to leave a signal that other searching robots can encounter. To extend this signal beyond the target's immediate surroundings, it moves outward along a spiral (Fig. 4a and Algorithm 2). The spacing between spiral arms is $g = R - 1$ cells, which is positive for $R > 1$. Unlike repulsive pheromone, which is refreshed at occupied cells, attractive pheromone is assigned to the sensing region surrounding the advertiser. The assigned intensity decreases with distance from the detected target $\mathbf{t}_{\text{loc}}$, making stronger values available nearer the target:

$$S(\mathbf{c}) = \frac{1}{1 + \|\mathbf{c} - \mathbf{t}_{\text{loc}}\|_2},$$

Algorithm 1 Repulsive-Pheromone Search Step for Robot r

Require: Position $\mathbf{p}_r$, repulsive pheromone field Φ_{rep}, attractive pheromone field Φ_{att}, visit memory M_r, sensing radius R, time step t

if TargetDetected($\mathcal{V}_R(\mathbf{p}_r(t))$) **then**
  return Advertise
end if
if $\max_{\mathbf{q}\in\mathcal{V}_R(\mathbf{p}_r(t))} \Phi_{\text{att}}(\mathbf{q}, t) > 0$ **then** ▷ If attractive pheromone detected
  return Follow
end if
$\mathcal{N} \leftarrow$ GetValidNeighbors($\mathbf{p}_r$)
for each neighbor $\mathbf{c} \in \mathcal{N}$ **do**
  $u_1(\mathbf{c}) \leftarrow \sum_{\mathbf{q}\in\mathcal{V}_R(\mathbf{c})} \Phi_{\text{rep}}(\mathbf{q}, t)$
  $u_2(\mathbf{c}) \leftarrow M_r(\mathbf{c}, t)$
end for
if $u_2(\mathbf{c}) > 0$ for all $\mathbf{c} \in \mathcal{N}$ **then** ▷ All neighboring cells visited
  next $\leftarrow$ random cell from $\mathcal{N}$ ▷ Select random neighbor
else
  next $\leftarrow \text{lexmin}_{\mathbf{c}\in\mathcal{N}}\,(u_1(\mathbf{c}), u_2(\mathbf{c}))$
end if
$M_r(\mathbf{p}_r, t) \leftarrow 1$
$\Phi_{\text{rep}}(\mathbf{p}_r, t) \leftarrow 1$
$\mathbf{p}_r \leftarrow$ **next**
return Search

The function equals one at the target and decreases farther away. To retain the strongest signal already written to a cell, the field is updated as

$$\Phi_{\text{att}}(\mathbf{c}, t) \leftarrow \max(\Phi_{\text{att}}(\mathbf{c}, t), S(\mathbf{c}))\,,$$

Taking the maximum retains the strongest attractive value previously assigned to each cell. The advertiser continues extending the trace while the intensity at its current position exceeds $I_{\min}$. After the intensity reaches or falls below this threshold, it completes the current update and enters Rescue. Only then is its earlier target detection registered and counted toward the required number of robots. Simulation stores the detected target location, while Robotarium records detection status and time step. Because the attractive field does not evaporate, other robots can still encounter and follow the deposited trace toward the target (Section 4.5).

The recruitment experiments consider a single target and end upon completion or timeout. The use of persistent attractive cues for repeated or simultaneous recruitment to multiple targets was not evaluated here.

Algorithm 2 Attractive-Pheromone Advertisement Step for Robot r

Require: Position $\mathbf{p}_r$, target location $\mathbf{t}_{\text{loc}}$, attractive pheromone field Φ_{att}, sensing radius R, minimum intensity $I_{\min}$, time t

$g \leftarrow R - 1$ ▷ Set spacing between adjacent spiral arms
$(\Delta x, \Delta y) \leftarrow \textsc{NextSpiralStep}(g)$
next $\leftarrow \mathbf{p}_r(t) + (\Delta x, \Delta y)$
$\mathbf{p}_r(t) \leftarrow$ **next**
for each $\mathbf{c} \in \mathcal{V}_R(\mathbf{p}_r(t))$ **do**
    $S(\mathbf{c}) \leftarrow \dfrac{1}{1 + \|\mathbf{c} - \mathbf{t}_{\text{loc}}\|_2}$
    $\Phi_{\text{att}}(\mathbf{c}, t) \leftarrow \max(\Phi_{\text{att}}(\mathbf{c}, t), S(\mathbf{c}))$
end for
$S(\mathbf{p}_r(t)) \leftarrow \dfrac{1}{1 + \|\mathbf{p}_r(t) - \mathbf{t}_{\text{loc}}\|_2}$
if $S(\mathbf{p}_r(t)) \leq I_{\min}$ **then**
    return Rescue
end if
return Advertise

4.5 Attractive pheromone follow behavior

A robot that encounters the advertised trace enters Follow. Its task is to move toward the detected target, not simply to trace the spiral deposited by the advertiser. Selecting only the strongest adjacent cell could cause it to follow the spiral instead of approaching the target. The follower therefore uses all positive attractive measurements in its sensing region to estimate a preferred direction (Algorithm 3 and Fig. 4b):

$$\mathbf{v}_r(t) = \sum_{\substack{\mathbf{c} \in \mathcal{V}_R(\mathbf{p}_r(t)) \\ \Phi_{\text{att}}(\mathbf{c}, t) > 0}} \Phi_{\text{att}}^2(\mathbf{c}, t)\,(\mathbf{c} - \mathbf{p}_r(t))\,. \tag{9}$$

Stronger attractive measurements receive greater weight because their intensity is squared. If these contributions cancel or the vector magnitude falls below $\varepsilon_v = 10^{-6}$, the robot returns to Search. Otherwise, the normalized vector $\widehat{\mathbf{v}}_r(t) = \mathbf{v}_r(t)/\|\mathbf{v}_r(t)\|$ gives a preferred direction. The robot then scores its neighboring cells by both alignment with this direction and their local attractive intensity; the selected move is not necessarily closer to the target:

$$\mathbf{next} = \operatorname*{argmax}_{\substack{\mathbf{c} \in \mathcal{N} \\ \mathbf{c} \notin \mathcal{H}_r}} \left[(\mathbf{c} - \mathbf{p}_r(t)) \cdot \widehat{\mathbf{v}}_r(t) + \lambda_{\text{local}} \Phi_{\text{att}}(\mathbf{c}, t)\right] \tag{10}$$

The weight $\lambda_{\text{local}} = 0.1$ determines how strongly local intensity influences this choice in simulation and Robotarium. To avoid repeatedly moving between the same cells, the follower also excludes locations in its short-term history $\mathcal{H}_r$, initialized empty and limited to six cells. If every neighbor is excluded, the history is cleared so that a movement can still be selected (Algorithm 3). Following continues until the robot detects the target within the prescribed radius, whereupon it enters Rescue and

Algorithm 3 Attractive-Pheromone Follow Step for Robot r

Require: Position $\mathbf{p}_r(t)$, attractive pheromone field Φ_{att}, position history $\mathcal{H}_r$, sensing radius R, weight λ_{local}, time t

if TargetDetected$(\mathcal{V}_R(\mathbf{p}_r(t)))$ **then**
  return Rescue
end if
$\mathbf{v}_r(t) \leftarrow \mathbf{0}$
for each $\mathbf{c} \in \mathcal{V}_R(\mathbf{p}_r(t))$ such that $\Phi_{\text{att}}(\mathbf{c}, t) > 0$ **do**
  $w(\mathbf{c}) \leftarrow (\Phi_{\text{att}}(\mathbf{c}, t))^2$ ▷ Weight stronger pheromone
  $\mathbf{v}_r(t) \leftarrow \mathbf{v}_r(t) + w(\mathbf{c})(\mathbf{c} - \mathbf{p}_r(t))$
end for
if $\|\mathbf{v}_r(t)\| \leq \varepsilon_v$ **then**
  return Search ▷ No clear direction
end if
$\widehat{\mathbf{v}}_r(t) \leftarrow \mathbf{v}_r(t)/\|\mathbf{v}_r(t)\|$
$\mathcal{N} \leftarrow$ GetValidNeighbors$(\mathbf{p}_r(t))$
if $\mathcal{N} \subseteq \mathcal{H}_r$ **then**
  $\mathcal{H}_r \leftarrow \emptyset$ ▷ Allow revisiting cells
end if
next $\leftarrow \underset{\substack{\mathbf{c}\in\mathcal{N} \\ \mathbf{c}\notin\mathcal{H}_r}}{\text{argmax}} [(\mathbf{c} - \mathbf{p}_r(t)) \cdot \widehat{\mathbf{v}}_r(t) + \lambda_{\text{local}}\Phi_{\text{att}}(\mathbf{c}, t)]$
$\mathbf{p}_r(t+1) \leftarrow$ **next**
Append $\mathbf{p}_r(t+1)$ to $\mathcal{H}_r$
return Follow

registers its arrival. Simulation stores these records in private robot memory, while Robotarium uses a shared record for experimental evaluation. The run ends when recruitment is complete or the mission time limit is reached. The conditions under which following is guaranteed to make progress are stated in Theorem 2.

4.6 Theoretical analysis

The two rules require different analyses. For Search, we ask whether changing the repulsive decay factor can change a robot's preference between two candidate movements under the same deposition history. For Follow, we identify conditions that make each chosen step reduce the distance to the target. The first result concerns a local choice under fixed history; the second gives a conditional convergence guarantee. Neither establishes complete coverage or target arrival in arbitrary environments.

Theorem 1 (Decay-Factor Sensitivity in Search) *Consider a robot r in* Search *comparing two neighboring cells $\mathbf{c}, \mathbf{d} \in \mathcal{N}$ at step t (Algorithm 1). Their repulsive scores, $u_1(\mathbf{c})$ and $u_1(\mathbf{d})$, depend on the decay factor $\lambda_{\text{rep}} \in (0, 1)$ (Eq. 5). As illustrated in Fig. 9, let $\mathcal{U}_\mathbf{c} = \mathcal{V}_R(\mathbf{c}) \setminus \mathcal{V}_R(\mathbf{d})$ and $\mathcal{U}_\mathbf{d} = \mathcal{V}_R(\mathbf{d}) \setminus \mathcal{V}_R(\mathbf{c})$ be the non-overlapping parts of their sensing regions. Holding the swarm's deposition history fixed up to step t, we obtain the following properties of the pairwise score comparison:*

1. *The score difference* $u_1(\mathbf{c}) - u_1(\mathbf{d})$ *is independent of the shared sensing region* $\mathcal{V}_R(\mathbf{c}) \cap \mathcal{V}_R(\mathbf{d})$ *and reduces to a finite polynomial in* λ_{rep} *determined solely by when each cell in the non-overlapping portions of sensing regions* $\mathcal{U}_\mathbf{c}$ *and* $\mathcal{U}_\mathbf{d}$ *was last refreshed with repulsive pheromone.*
2. *Equal numbers of cells last refreshed with repulsive pheromone at every prior time step in* $\mathcal{U}_\mathbf{c}$ *and* $\mathcal{U}_\mathbf{d}$ *give equal scores* ($u_1(\mathbf{c}) = u_1(\mathbf{d})$). *When one region (say,* $\mathcal{U}_\mathbf{c}$*) has no more such cells at any time step and fewer at least one, its score remains strictly lower* ($u_1(\mathbf{c}) < u_1(\mathbf{d})$). *Both results hold for all* $\lambda_{\text{rep}} \in (0, 1)$.
3. *When* $\mathcal{U}_\mathbf{c}$ *contains more recently refreshed cells but fewer cells refreshed earlier than* $\mathcal{U}_\mathbf{d}$*, changing* λ_{rep} *may reverse the robot's preference between* $\mathbf{c}$ *and* $\mathbf{d}$ *based on their repulsive pheromone values. Let* $b_k = n_\mathbf{c}(k) - n_\mathbf{d}(k)$ *denote the difference in the numbers of cells last refreshed* k *steps ago in the two regions, and let* K *be the largest such elapsed time. The number of strict preference reversals over* $\lambda_{\text{rep}} \in (0, 1)$ *is bounded above by the number of sign changes in the nonzero sequence* $(b_0, \ldots, b_K)$*, and hence by* K.

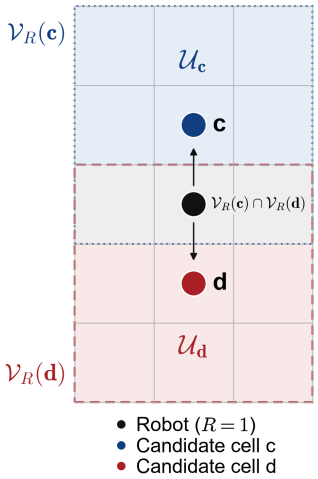


Fig. 9: Sensing regions of candidate cells $\mathbf{c}$ and $\mathbf{d}$.

Theorem 2 (Finite-time convergence of trail following) *Consider a robot* r *at position* $\mathbf{p}_r$ *that enters* FOLLOW *(Algorithm 3) at time* t_0 *and selects next actions according to Eq. 10. Suppose that at each time step before target detection, at least one unvisited neighboring cell* $\mathbf{c}^+ \in \mathcal{N} \setminus \mathcal{H}_r$ *reduces the Manhattan distance* d_1 *to the target* $\mathbf{t}_{\text{loc}}$*. This eligibility condition is assumed rather than guaranteed by the short-term history update rule. If the cell* $\mathbf{c}^+$*'s directional alignment with* $\widehat{\mathbf{v}}_r(t)$ *exceeds that of every non-progress neighbor* $\mathbf{m}$ *by more than the weighted local pheromone variation,*

$$(\mathbf{c}^+ - \mathbf{m}) \cdot \widehat{\mathbf{v}}_r(t) > \lambda_{\text{local}} \max_{\mathbf{c},\mathbf{q} \in \mathcal{N} \setminus \mathcal{H}_r} |\Phi_{\text{att}}(\mathbf{c}, t) - \Phi_{\text{att}}(\mathbf{q}, t)|\,, \tag{11}$$

then the robot moves closer to the target at every time step and reaches it in at most

$$T_{\text{reach}} \leq d_1(\mathbf{p}_r(t_0), \mathbf{t}_{\text{loc}}) \tag{12}$$

time steps, unless the target is detected earlier and the robot enters RESCUE.

Theorem 1 shows how the repulsive decay factor λ_{rep} can change the ranking of candidate search movements when deposition history is fixed. Because the repulsive score u_1 is considered before private memory u_2, the robot may prefer a previously visited cell even when an unvisited neighbor is available. The theorem characterizes this local dependence on pheromone persistence; it does not establish how changing the factor would alter the full search trajectory. Theorem 2 gives a different result: when its directional and eligibility conditions hold at every step, the follower makes progress toward the target and reaches it in finite time.

The full proofs and their additional assumptions are given in Supplementary Section 3.

4.7 Experimental design and baselines

To assess whether the environmental traces improved search, we compared STIGSAR with centralized lawnmower coverage and decentralized random walk. For matched comparisons, the methods used the same workspace size, team size, initial robot and target positions, sensing radius, and mission horizon. Initial positions were generated with NumPy's random-number generator using fixed seeds so that the comparisons could be reproduced.

For the centralized lawnmower baseline, the workspace was partitioned into equal-sized convex sections according to the team size, and each robot executed a deterministic back-and-forth (boustrophedon) sweep within its assigned section. For the decentralized random-walk baseline, each robot selected uniformly at random among its valid movement actions at every step without maintaining or reading shared environmental memory. In the field-ablation study, the five reduced or merged variants modified only the pheromone update and state-selection rules while preserving the same sensing radius, movement primitives, and target-detection radius as the complete STIGSAR controller.

4.8 Simulation experiments

To examine how search performance changed with workload, we increased workspace size, team size, and target count together. Across ten configurations, the square-grid side length was $L \in \{50, 100, \ldots, 500\}$, the team size was $N \in \{5, 10, \ldots, 50\}$, and each workspace contained $3N$ randomly placed targets. Each configuration used a horizon of $T_{\max} = \lceil 10L^2/N \rceil$ steps, with 100 independent runs each for STIGSAR and random walk and 10 for the centralized reference.

The robot-loss experiments used the same ten configurations to test whether search could continue when robots stopped without warning. In each run, 20% of the robots were immobilized according to one of four failure schedules: early, middle, late, or uniform. Each schedule had 100 runs for STIGSAR and random walk and 10 for the centralized baseline. Stopped robots no longer moved or deposited pheromone. The centralized reference reassigned their unfinished regions, whereas STIGSAR's remaining robots continued with their private visit records and the available pheromone field, without receiving new assignments.

To test whether separating the two fields improved recruitment, the field-ablation and recruitment simulations used matched starting conditions. They included 50×50 and 100×100 grids with one target, team sizes $N \in \{1, 2, 3\}$, and initial robot-to-target Manhattan distance shells of 20, 30, 40, and 50 cells. Ten seeded runs per combination gave 240 runs per condition. For field ablation, each run's normalized mission cost was

$$J_{\text{mission}} = \frac{T_{\text{all}}}{d_{\min}}, \tag{13}$$

where T_{all} is the step at which every robot's target encounter has been registered (or $T_{\max}$ if the run timed out). Target detection within the prescribed radius counts as arrival, but the discovering robot is registered only after it finishes advertising. The

reference distance is the shortest initial Manhattan distance between any robot and the target:

$$d_{\min} = \min_{r} d_1(\mathbf{p}_r(0), \mathbf{t}_{\text{loc}}) \tag{14}$$

In the recruitment suite, trajectory segments were analyzed from the time step at which a searching robot first detected the attractive field and entered FOLLOW, measuring step-by-step Manhattan distance d_1 to the target for comparison with Theorem 2.

4.9 Robotarium experiments

Robotarium provided physical motion while retaining the virtual pheromone controller. Trials took place in a $3.2\,\text{m} \times 2.0\,\text{m}$ arena, with a $1.9\,\text{m} \times 1.9\,\text{m}$ active workspace represented by a 50×50 virtual pheromone grid. High-level commands were issued at a nominal 30 Hz, waypoints were updated at approximately 6 Hz, and trials ended after 590 s of wall-clock time if incomplete.

Search and robot-loss trials used five robots and ten fixed target locations sampled uniformly across the grid and shared across all methods; in the robot-loss condition, two of the five robots were stopped mid-mission. Field-ablation trials used two robots with a single target placed at the center cell $(25, 25)$, evaluated across initial robot-to-target separations of 30 and 40 grid cells, with three trials per separation and six trials per variant.

Recruitment trials placed the target and advertising robot at $(25, 25)$ and initialized one to three searching robots at separations of 30 and 40 cells (three trials per separation).

4.10 DJI RoboMaster experiments

To examine physical environmental memory rather than a virtual field, we used RoboMaster robots to deposit and detect a UV-visible repulsive trail in a $4.0\,\text{m} \times 1.8\,\text{m}$ arena. This experiment tested search only; attractive recruitment was not implemented on RoboMaster. STIGSAR, centralized lawnmower coverage, and random walk were compared through trajectories, visitation maps, and covered area. Coverage was measured on a separate $0.05\,\text{m} \times 0.05\,\text{m}$ grid of 80×36 cells. This grid did not guide the robots, whose movements and odometry-based localization remained continuous.

Search trials used two DJI RoboMaster S1 robots. Robot-loss trials began with three robots; one robot was stopped after 60 sec at its current position, leaving two robots to continue without receiving a failure notification. Search and robot-loss conditions each comprised three trials for STIGSAR and random walk and one deterministic centralized reference trial. Each trial ran for 240 sec. Robot position and heading were obtained from the RoboMaster S1 wheel odometry. For STIGSAR and random walk, left, straight, and right actions corresponded to heading changes of $+30°$, $0°$, and $-30°$, respectively, followed by a nominal 0.25 m forward movement. Actions did not have fixed durations; rotation and translation continued under wheel-odometry feedback until the commanded heading and distance were reached, with low-level commands refreshed every 0.05 s.

The physical repulsive trail was produced using photochromic powder mixed with an acrylic binder. A 360 nm UV source activated the material while the robot travelled at $0.8\,\mathrm{m\,s^{-1}}$. The activated trail remained detectable for approximately 15 sec under indoor lighting without direct sunlight, independently of the numerical decay factor λ_{rep} used in the virtual pheromone fields. Trail detection used color-pixel thresholding in HSV space, with hue, saturation, and value ranges of $[110, 140]$, $[24, 85]$, and $[148, 174]$, respectively. Matching pixels were counted within the left, straight, and right regions of the forward-facing camera view.

4.11 Metrics and statistical analysis

For field ablation (Fig. 2), mission reliability is the fraction of runs in which all robots register a target encounter by $T_{\max}$; normalized mission cost J_{mission} is defined in Eq. (13). Robotarium completion is reported using wall-clock time in seconds and controller steps.

In the search experiments (Fig. 3), coverage is the fraction of workspace cells observed within the sensing region of at least one robot, and target discovery tracks the cumulative number of detected targets over normalized mission time $t/T_{\max}$. Search speed across scaling configurations is measured by the number of simulation steps required to reach 95% coverage. A cell visit is counted whenever a robot occupies a grid cell, and visits per covered cell (total cell visits divided by the number of unique observed cells) quantify exploration redundancy, where lower values indicate less repeated visitation.

In the recruitment experiments (Fig. 4), convergence is evaluated from the step a searching robot enters FOLLOW by tracking its Manhattan distance $d_1(\mathbf{p}_r(t), \mathbf{t}_{\mathrm{loc}})$ to the target cell.

For the robot-loss experiments (Fig. 5), we report median mission time across ten simulation configurations, with unfinished runs assigned the mission horizon and completion counts reported separately. Post-failure target recovery is measured as the fraction of remaining targets subsequently found against additional workspace coverage, while Robotarium results report time to find all targets and completion counts within the timeout.

All statistical comparisons are descriptive rather than based on null-hypothesis significance tests. Simulation outcomes are summarized by medians, with empirical 2.5th to 97.5th percentile bands for time-series curves and 95% bootstrap intervals (computed by resampling runs) for scalar summaries. Binary mission-reliability proportions use 95% Wilson intervals. Robotarium results are reported using completed-trial medians where noted, while RoboMaster results are summarized using means and standard deviations across repeated trials. Representative physical-robot trajectories are shown separately.

Data availability. All data, experimental materials, and videos needed to evaluate and reproduce the paper's conclusions are publicly available at https://github.com/Eashwar-S/Dual_pheromone_stigmergic_controller.

Code availability. The code used to implement STIGSAR and reproduce the reported analyses is publicly available at https://github.com/Eashwar-S/Dual_pheromone_stigmergic_controller.

References

[1] Burgard, W., Moors, M., Stachniss, C. & Schneider, F. E. Coordinated multi-robot exploration. *IEEE Transactions on Robotics* **21**, 376–386 (2005).

[2] Gerkey, B. P. & Matarić, M. J. A formal analysis and taxonomy of task allocation in multi-robot systems. *The International Journal of Robotics Research* **23**, 939–954 (2004).

[3] Kuyucu, T., Tanev, I. & Shimohara, K. Superadditive effect of multi-robot coordination in the exploration of unknown environments via stigmergy. *Neurocomputing* **148**, 83–90 (2015).

[4] Zedadra, O., Jouandeau, N., Seridi, H. & Fortino, G. Multi-agent foraging: State-of-the-art and research challenges. *Complex Adaptive Systems Modeling* **5**, 3 (2017).

[5] Ebert, J. T., Berlinger, F., Haghighat, B. & Nagpal, R. A hybrid PSO algorithm for multi-robot target search and decision awareness 11520–11527 (2022).

[6] Grassé, P.-P. La reconstruction du nid et les coordinations inter-individuelles chez Bellicositermes natalensis et Cubitermes sp. La théorie de la stigmergie. *Insectes Sociaux* **6**, 41–80 (1959).

[7] Bonabeau, E., Dorigo, M. & Theraulaz, G. *Swarm Intelligence: From Natural to Artificial Systems* (Oxford University Press, New York, 1999).

[8] Şahin, E. *Swarm robotics: From sources of inspiration to domains of application* Vol. 3342 (Springer, Berlin, Heidelberg, 2005).

[9] Brambilla, M., Ferrante, E., Birattari, M. & Dorigo, M. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence* **7**, 1–41 (2013).

[10] Dorigo, M., Maniezzo, V. & Colorni, A. Ant system: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* **26**, 29–41 (1996).

[11] Dorigo, M. & Stützle, T. *Ant Colony Optimization* (MIT Press, 2004).

[12] Payton, D. W., Daily, M. J., Estowski, R., Howard, M. & Lee, C. Pheromone robotics. *Autonomous Robots* **11**, 319–324 (2001).

[13] Parunak, H. V. D., Purcell, M. & O'Connell, R. Digital pheromones for autonomous coordination of swarming UAVs 3446 (2002).

[14] Parunak, H. V. D., Brueckner, S. A. & Sauter, J. *Digital Pheromones for Coordination of Unmanned Vehicles* Vol. 3374 of *Lecture Notes in Artificial Intelligence* (Springer, Berlin, Heidelberg, 2004).

[15] Salman, M. *et al.* Phormica: Photochromic pheromone release and detection system for swarm robotics. *Frontiers in Robotics and AI* **7**, 591402 (2020).

[16] Tang, Q., Xu, Z., Yu, F., Zhang, Z. & Zhang, J. Dynamic target searching and tracking with swarm robots based on stigmergy mechanism. *Robotics and Autonomous Systems* **120**, 103251 (2019).

[17] Payton, D., Estkowski, R. & Howard, M. Compound behaviors in pheromone robotics. *Robotics and Autonomous Systems* **44**, 229–240 (2003).

[18] Liu, T., Sun, X., Hu, C., Fu, Q. & Yue, S. A multiple pheromone communication system for swarm intelligence. *IEEE Access* **9**, 148721–148737 (2021).

[19] Jagla, T., Schulte, D., Wieseke, N. & Middendorf, M. Robotic group formation with phosphorescent trails. *Autonomous Robots* **50**, 35 (2026).

[20] Tran, V. P., Garratt, M. A., Kasmarik, K. & Anavatti, S. G. Dynamic frontier-led swarming: Multi-robot repeated coverage in dynamic environments. *IEEE/CAA Journal of Automatica Sinica* **10**, 646–661 (2023).

[21] Ardiny, H. & Beigzadeh, A. M. Enhancing radioactive environment exploration with bio-inspired swarm robotics: A comparative analysis of lévy flight and stigmergy methods. *Robotics and Autonomous Systems* **181**, 104794 (2024).

[22] Tinoco, F. D., da Silva, M. & Saska, M. Pherocom: Decentralised and asynchronous robot swarm communication with pheromone-based marking. *arXiv* (2022).

[23] Que, H. *et al.* SPACE: Swarm pheromone fields for adaptive collision-aware exploration. *arXiv preprint arXiv:2606.29372* (2026).

[24] Yacoub, P., Kaouach, M. M., Khatab, E. & Shalash, O. Directional pheromone gradient observations for decentralized multi-agent reinforcement learning in swarm drone search and rescue. *Drones* **10**, 648 (2026).

[25] Wilson, S., Glotfelter, P., Wang, L., Cortés, J. & Egerstedt, M. The robotarium: Globally impactful opportunities, challenges, and lessons learned in remote-access, distributed control of multirobot systems. *IEEE Control Systems Magazine* **40**, 26–44 (2020).

[26] DJI. RoboMaster S1: Educational Robot. https://store.dji.com/product/robomaster-s1 (2019). Accessed: 2026-07-11.

[27] Salman, M., Garzón Ramos, D. & Birattari, M. Automatic design of stigmergy-based behaviours for robot swarms. *Communications Engineering* **3**, 30 (2024).

[28] Wang, L., Ames, A. D. & Egerstedt, M. Safety barrier certificates for collisions-free multirobot systems. *IEEE Transactions on Robotics* **33**, 661–674 (2017).

Acknowledgements. This study received no funding.

Author contributions.. E.S. and N.R. developed the methodology and designed the experiments. E.S. conceived the study, implemented the framework and conducted the experiments, analyzed the results, developed and drafted the theoretical proofs and drafted the manuscript. N.R. contributed to the experimental analysis, visualization and figure preparation, and manuscript development. Y.A. reviewed the manuscript and provided critical feedback. All authors reviewed and approved the final manuscript.

Competing interests.. The authors declare no competing interests.

Supplementary information. Supplementary Information contains full proofs and assumptions, the τ sensitivity sweep, additional configurations and maps, and Supplementary Videos of the physical-robot experiments can be accessed from here.

Ethics approval and consent to participate.. Not applicable.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- supplementary.pdf
- STIGSARsummaryvideo.mp4