SUPPLEMENTARY MATERIALS: Dual-pheromone stigmergic framework for decentralized swarm search and rescue

Eashwar Sathyamurthy[a], Naitri Rajyaguru[b], Yiannis Aloimonos[a]

[a] *Department of Mechanical Engineering, University of Maryland, College Park, Maryland, 20742, USA*
[b] *University of Maryland Institute for Advanced Computer Studies, University of Maryland, College Park, Maryland, 20742, USA*

1. Overview of the supplementary materials

This document provides the evaporation-rate sensitivity analysis supporting the choice of τ_{rep} in Section 4.1.3 of the main article, followed by the theoretical proofs and a summary of the experimental design. The sensitivity analysis is presented in Supplementary Section 2, the proofs in Supplementary Section 3, and the experimental design summary in Supplementary Section 4.

2. Sensitivity analysis of evaporation rate

The repulsive pheromone must persist long enough to discourage immediate revisits, but not so long that old trajectories prevent robots from re-entering previously searched regions. The base time constant was defined as

$$\tau_{\text{rep}} = \frac{H \times W}{N \times R}, \tag{S1}$$

where $H \times W$ is the number of grid cells, N is the number of robots, and R is the robot sensing and deposition radius. The sensitivity sweep evaluated $\tau_{\text{rep}} = m\tau_{\text{base}}$, where $\tau_{\text{base}} = (H \times W)/(N \times R)$ and $m \in \{0.5, 1, 2, 3, 4\}$. Thus, the curve labelled 1.0× Base in Fig. S1 is Eq. (S1), while the remaining curves shorten or lengthen its persistence using the indicated multiplier.

The experiment isolates the search phase of the framework so that the effect of repulsive-pheromone persistence can be evaluated without including target advertisement, attractive-pheromone following, or rescue time. At the start of each run, multiple robots were distributed randomly in the environment, and multiple targets were placed at random locations. The robots then executed the repulsive-pheromone search behavior until every target had been

Email addresses: eashwar.sathyamurthy@gmail.com (Eashwar Sathyamurthy), nrajyagu@umd.edu (Naitri Rajyaguru), jyaloimo@umd.edu (Yiannis Aloimonos)

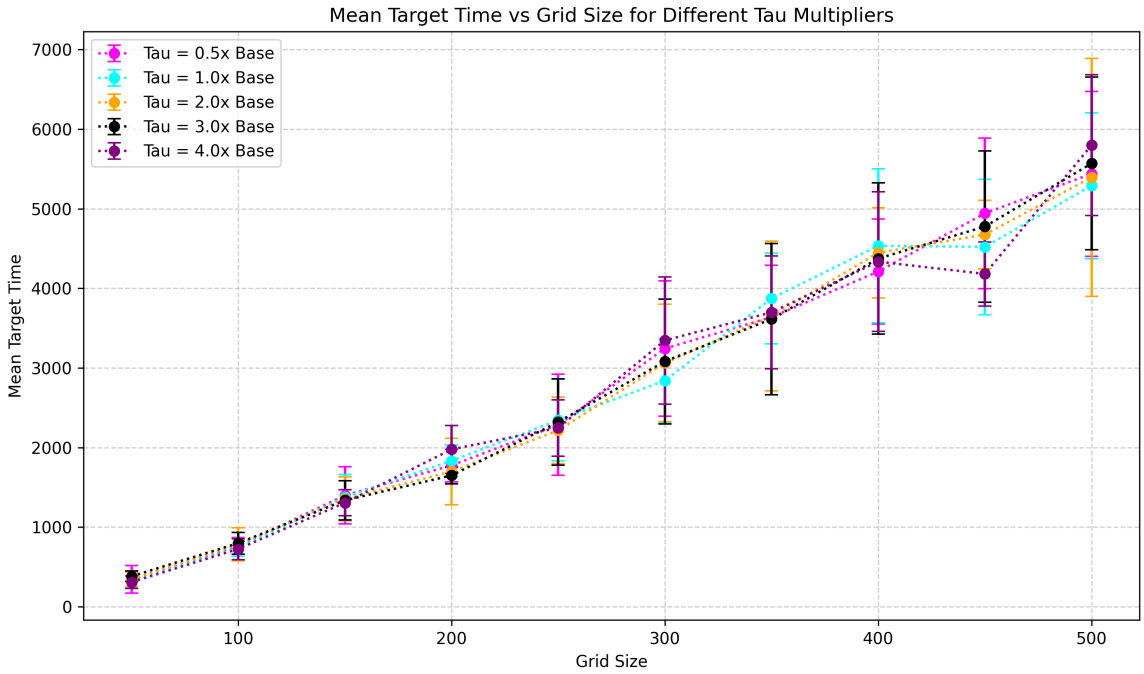


Figure S1: **Sensitivity of search-completion time to the repulsive-pheromone time constant.** Search-completion time is the number of simulation steps required for the randomly dispersed robots to find all targets. The base value is $\tau_{\text{base}} = (H \times W)/(N \times R)$. The five curves apply multipliers $m \in \{0.5, 1, 2, 3, 4\}$ to this base value. Points show the mean over ten runs, and error bars show one standard deviation.

found. A run ended when the final target was detected, and search-completion time was recorded as the number of simulation steps required to reach this condition.

The experiment used ten grid-size settings from 50 to 500. For each configuration, τ_{base} was recalculated using its grid dimensions, number of robots, and robot radius. This tests whether the same scaling rule can maintain consistent search performance as workspace size and collective search capacity change, rather than tuning one fixed evaporation constant for every environment. Each of the five multipliers was evaluated in ten independent runs at every grid size, giving 500 runs in total. The plotted marker is the mean search-completion time over the ten runs, and the error bar is one standard deviation.

Search-completion time increased with grid size for every multiplier, whereas the five multiplier curves remained close relative to the run-to-run variability. The multiplier with the lowest mean changed across grid sizes, so no tested alternative uniformly dominated the base value. When all grid sizes were pooled with equal numbers of runs, mean search-completion times were 2804.4, 2767.7, 2762.9, 2790.2, and 2792.5 steps for multipliers 0.5, 1, 2, 3, and 4, respectively. The full range among these means was approximately 1.5%, and the base result was approximately 0.2% above the lowest pooled mean. These results indicate that the base scaling preserves comparable search performance across the tested configurations without requiring a different multiplier for each grid size or swarm configuration.

Equation (S1) was therefore retained because it achieves near-minimum aggregate completion time without introducing an additional tuned multiplier. Its scaling also reflects the rate at which the swarm refreshes environmental information. A larger grid contains more cells and therefore increases τ_{rep}, allowing marks to persist over the longer traversal scale. Increasing either the number of robots or their sensing and deposition radius decreases τ_{rep} because the field is refreshed more frequently and over a larger region. In the field update, $\lambda_{\text{rep}} = \exp(-1/\tau_{\text{rep}})$, so increasing τ_{rep} produces slower evaporation, while decreasing it produces faster evaporation. The selected expression consequently adapts pheromone persistence to workspace size and collective sensing capacity.

3. Proofs

3.1. Proof of Theorem 1: Decay-Factor Sensitivity in Search

Theorem 1 (Decay-Factor Sensitivity in Search)**.** *Consider a robot r of radius R in* SEARCH *state at step t (Algorithm ??) comparing two candidate neighboring cells $\mathbf{c}, \mathbf{d} \in \mathcal{N}$ via their repulsive pheromone scores $u_1(\mathbf{c})$ and $u_1(\mathbf{d})$ under decay factor $\lambda_{\text{rep}} \in (0,1)$ (Eq. S4) as shown in Figure S2. Let $\mathcal{U}_\mathbf{c} = \mathcal{V}_R(\mathbf{c}) \setminus \mathcal{V}_R(\mathbf{d})$ and $\mathcal{U}_\mathbf{d} = \mathcal{V}_R(\mathbf{d}) \setminus \mathcal{V}_R(\mathbf{c})$ denote the non-overlapping portions of sensing regions of cells $\mathbf{c}$ and $\mathbf{d}$ respectively. For a fixed history of swarm visits up to time t, the following properties characterize how varying λ_{rep} affects local action preferences, without accounting for changes in subsequent robot trajectories or global coverage:*

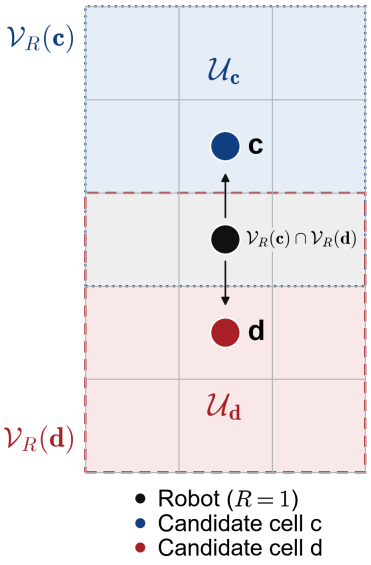


Figure S2: Sensing regions of candidate cells $\mathbf{c}$ and $\mathbf{d}$.

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 visited.*
2. *Equal numbers of cells last visited 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 visited cells but fewer cells visited 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 visited 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.*

Proof. The proof uses two operating assumptions from SEARCH in Section 4 of the main manuscript: (i) the repulsive pheromone field is initially zero everywhere ($\Phi_{\text{rep}}(\mathbf{q}, 0) = 0$ for all $\mathbf{q} \in \mathcal{C}$), and (ii) at each time step, any cell occupied by at least one robot has its repulsive pheromone reset to 1, while all other cells are multiplied by the decay factor $\lambda_{\text{rep}} \in (0,1)$. We

hold the history of robot positions up to step t fixed and examine how the score comparison between two valid neighboring cells $\mathbf{c}, \mathbf{d} \in \mathcal{N}$ depends on λ_{rep}.

To prove the first property, we write out the repulsive pheromone scores evaluated by Algorithm 1 at step t:

$$u_1(\mathbf{c}) = \sum_{\mathbf{q}\in\mathcal{V}_R(\mathbf{c})} \Phi_{\text{rep}}(\mathbf{q}, t), \qquad u_1(\mathbf{d}) = \sum_{\mathbf{q}\in\mathcal{V}_R(\mathbf{d})} \Phi_{\text{rep}}(\mathbf{q}, t). \tag{S2}$$

As illustrated in Figure **??**, the sensing region $\mathcal{V}_R(\mathbf{c})$ can be split into the cells shared with $\mathcal{V}_R(\mathbf{d})$ and the non-overlapping cells $\mathcal{U}_\mathbf{c} = \mathcal{V}_R(\mathbf{c}) \setminus \mathcal{V}_R(\mathbf{d})$. Similarly, $\mathcal{V}_R(\mathbf{d})$ splits into the same shared region $\mathcal{V}_R(\mathbf{c}) \cap \mathcal{V}_R(\mathbf{d})$ and $\mathcal{U}_\mathbf{d} = \mathcal{V}_R(\mathbf{d}) \setminus \mathcal{V}_R(\mathbf{c})$. Subtracting $u_1(\mathbf{d})$ from $u_1(\mathbf{c})$ cancels the identical sum over the shared cells $\mathcal{V}_R(\mathbf{c}) \cap \mathcal{V}_R(\mathbf{d})$:

$$\begin{aligned} u_1(\mathbf{c}) - u_1(\mathbf{d}) &= \left(\sum_{\mathbf{q}\in\mathcal{V}_R(\mathbf{c})\cap\mathcal{V}_R(\mathbf{d})} \Phi_{\text{rep}}(\mathbf{q}, t) + \sum_{\mathbf{q}\in\mathcal{U}_\mathbf{c}} \Phi_{\text{rep}}(\mathbf{q}, t) \right) \\ &\quad - \left(\sum_{\mathbf{q}\in\mathcal{V}_R(\mathbf{c})\cap\mathcal{V}_R(\mathbf{d})} \Phi_{\text{rep}}(\mathbf{q}, t) + \sum_{\mathbf{q}\in\mathcal{U}_\mathbf{d}} \Phi_{\text{rep}}(\mathbf{q}, t) \right) \\ &= \sum_{\mathbf{q}\in\mathcal{U}_\mathbf{c}} \Phi_{\text{rep}}(\mathbf{q}, t) - \sum_{\mathbf{q}\in\mathcal{U}_\mathbf{d}} \Phi_{\text{rep}}(\mathbf{q}, t). \end{aligned} \tag{S3}$$

Thus, the shared sensing region has no effect on the difference between the two scores. Next, consider the pheromone value $\Phi_{\text{rep}}(\mathbf{q}, t)$ of any individual cell $\mathbf{q}$. Under the update rule in Eq.5 of main manscript given below

$$\Phi_{\text{rep}}(\mathbf{q}, t+1) = \begin{cases} 1, & \mathbf{q} \in \mathcal{D}_{\text{rep}}(t), \\ \lambda_{\text{rep}} \Phi_{\text{rep}}(\mathbf{q}, t), & \mathbf{q} \notin \mathcal{D}_{\text{rep}}(t). \end{cases} \tag{S4}$$

if $\mathbf{q}$ has never been visited by any other robot up to step t, then $\Phi_{\text{rep}}(\mathbf{q}, t) = 0$. If $\mathbf{q}$ was last visited k steps ago (that is, at time step $t-k$, where $0 \le k \le t$), its pheromone value was set to 1 at that visit and multiplied by λ_{rep} during each of the k subsequent steps without another visit, giving

$$\Phi_{\text{rep}}(\mathbf{q}, t) = \lambda_{\text{rep}}^k. \tag{S5}$$

Any earlier visits to cell $\mathbf{q}$ before step $t-k$ were overwritten when its concentration was reset to 1. For each $k \in \{0, 1, \ldots, t\}$, let $n_\mathbf{c}(k)$ be the number of cells in $\mathcal{U}_\mathbf{c}$ that were last visited k steps ago, and let $n_\mathbf{d}(k)$ be the number of cells in $\mathcal{U}_\mathbf{d}$ that were last visited k steps ago. Grouping the cells in Eq. (S3) by k yields

$$\sum_{\mathbf{q}\in\mathcal{U}_\mathbf{c}} \Phi_{\text{rep}}(\mathbf{q}, t) = \sum_{k=0}^{K} n_\mathbf{c}(k) \lambda_{\text{rep}}^k, \qquad \sum_{\mathbf{q}\in\mathcal{U}_\mathbf{d}} \Phi_{\text{rep}}(\mathbf{q}, t) = \sum_{k=0}^{K} n_\mathbf{d}(k) \lambda_{\text{rep}}^k, \tag{S6}$$

where $K \le t$ is the largest number of steps since the most recent visit to any visited cell in $\mathcal{U}_\mathbf{c} \cup \mathcal{U}_\mathbf{d}$ (taking $K = 0$ if no cell in either region has been visited). Subtracting the two sums in Eq. (S6) and defining the visit-count difference at each step k as $b_k = n_\mathbf{c}(k) - n_\mathbf{d}(k)$ gives

$$D_{\mathbf{cd}}(\lambda_{\text{rep}}) = u_1(\mathbf{c}) - u_1(\mathbf{d}) = \sum_{k=0}^{K} \big(n_\mathbf{c}(k) - n_\mathbf{d}(k)\big) \lambda_{\text{rep}}^k = \sum_{k=0}^{K} b_k \lambda_{\text{rep}}^k. \tag{S7}$$

This is a finite polynomial in λ_{rep} whose coefficients b_k depend only on when each cell in $\mathcal{U}_\mathbf{c}$ and $\mathcal{U}_\mathbf{d}$ was last visited, which establishes Property 1.

To prove the second property, we examine the sign of the polynomial in Eq. (S7) for $\lambda_{\text{rep}} \in (0,1)$. First, suppose $\mathcal{U}_{\mathbf{c}}$ and $\mathcal{U}_{\mathbf{d}}$ have equal numbers of cells last visited k steps ago for every $k \in \{0,\dots,K\}$. Then $n_{\mathbf{c}}(k) = n_{\mathbf{d}}(k)$, so $b_k = 0$ for all k. Every term in Eq. (S7) vanishes, giving $u_1(\mathbf{c}) - u_1(\mathbf{d}) = 0$, or $u_1(\mathbf{c}) = u_1(\mathbf{d})$ for all $\lambda_{\text{rep}} \in (0,1)$. Second, suppose $\mathcal{U}_{\mathbf{c}}$ has no more such cells than $\mathcal{U}_{\mathbf{d}}$ at any step k ($n_{\mathbf{c}}(k) \le n_{\mathbf{d}}(k)$, so $b_k \le 0$ for all k) and strictly fewer at least at one step $k^\star$ ($n_{\mathbf{c}}(k^\star) < n_{\mathbf{d}}(k^\star)$, so $b_{k^\star} < 0$). Because $\lambda_{\text{rep}} \in (0,1)$, every power λ_{rep}^k is strictly positive. Therefore, every term $b_k \lambda_{\text{rep}}^k$ in Eq. (S7) is nonpositive, and the $k^\star$ term is strictly negative:

$$u_1(\mathbf{c}) - u_1(\mathbf{d}) = \sum_{k=0}^{K} b_k \lambda_{\text{rep}}^k \le b_{k^\star} \lambda_{\text{rep}}^{k^\star} < 0. \tag{S8}$$

Thus $u_1(\mathbf{c}) < u_1(\mathbf{d})$ for every $\lambda_{\text{rep}} \in (0,1)$, which proves Property 2.

To prove the third property, suppose $\mathcal{U}_{\mathbf{c}}$ contains more recently visited cells than $\mathcal{U}_{\mathbf{d}}$ (so $b_k > 0$ for some smaller indices k) while $\mathcal{U}_{\mathbf{d}}$ contains more cells visited earlier (so $b_k < 0$ for some larger indices k). Because Algorithm 1 prefers the candidate with the lower repulsive pheromone score u_1, the robot prefers $\mathbf{d}$ when $D_{\mathbf{cd}}(\lambda_{\text{rep}}) > 0$ and prefers $\mathbf{c}$ when $D_{\mathbf{cd}}(\lambda_{\text{rep}}) < 0$. For small λ_{rep} near 0, higher powers λ_{rep}^k decay rapidly, so the polynomial in Eq. (S7) is dominated by its lowest-order nonzero term (the most recent visits), where $b_k > 0$ and $\mathbf{d}$ has the lower score. As λ_{rep} approaches 1, all powers λ_{rep}^k approach 1 and older visits retain their weight; if the negative coefficients from older visits outweigh the positive coefficients so that $\sum_{k=0}^{K} b_k < 0$, then $D_{\mathbf{cd}}(\lambda_{\text{rep}})$ becomes negative near $\lambda_{\text{rep}} = 1$, reversing the robot's preference from $\mathbf{d}$ to $\mathbf{c}$.

Every switch in preference by the robot whether to move next to either cell $\mathbf{c}$ or cell $\mathbf{d}$ as λ_{rep} varies over $(0,1)$ requires $D_{\mathbf{cd}}(\lambda_{\text{rep}})$ to change sign and therefore to pass through a root $\lambda_{\text{rep}} \in (0,1)$ where $D_{\mathbf{cd}}(\lambda_{\text{rep}}) = 0$. By Descartes' rule of signs [1], the number of positive real roots of the polynomial $D_{\mathbf{cd}}(\lambda_{\text{rep}}) = \sum_{k=0}^{K} b_k \lambda_{\text{rep}}^k$ (and hence the number of roots in the sub-interval $(0,1)$) cannot exceed the number of sign changes in the nonzero sequence of visit-count differences $(b_0, b_1, \dots, b_K)$. Furthermore, to obtain a tighter bound specific to the open interval $\lambda_{\text{rep}} \in (0,1)$, we can map $(0,1)$ one-to-one onto $(0,\infty)$ via the substitution

$$z = \frac{1-\lambda_{\text{rep}}}{\lambda_{\text{rep}}} > 0 \quad \Longleftrightarrow \quad \lambda_{\text{rep}} = \frac{1}{1+z}. \tag{S9}$$

Multiplying $D_{\mathbf{cd}}\big(1/(1+z)\big)$ by $(1+z)^K > 0$ yields the transformed polynomial

$$Q_{\mathbf{cd}}(z) = (1+z)^K D_{\mathbf{cd}}\left(\frac{1}{1+z}\right) = \sum_{k=0}^{K} b_k (1+z)^{K-k} = \sum_{j=0}^{K} c_j z^{K-j}, \tag{S10}$$

where expanding $(1+z)^{K-k} = \sum_{i=0}^{K-k} \binom{K-k}{i} z^i$ and collecting powers of z^{K-j} gives the coefficients

$$c_j = \sum_{k=0}^{j} \binom{K-k}{j-k} b_k, \qquad j = 0,1,\dots,K. \tag{S11}$$

Because $\lambda_{\text{rep}} = 1/(1+z)$ is a strictly decreasing bijection between $\lambda_{\text{rep}} \in (0,1)$ and $z \in (0,\infty)$, the roots of $D_{\mathbf{cd}}(\lambda_{\text{rep}})$ in $(0,1)$ correspond one-to-one with the positive roots $z > 0$ of $Q_{\mathbf{cd}}(z)$. Applying Descartes' rule of signs to $Q_{\mathbf{cd}}(z)$ shows that the number of preference switches on $(0,1)$ is bounded above by the number of sign changes in $(c_0, c_1, \dots, c_K)$, which is determined entirely by the visit-count differences $b_k = n_{\mathbf{c}}(k) - n_{\mathbf{d}}(k)$ (and never exceeds K). This completes the proof of Property 3 and of Theorem **??**. □

3.2. Proof of Theorem 2: Finite-time convergence of trail following

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 equation given below (Eq 10 of main document):*

$$\mathbf{next} = \underset{\substack{\mathbf{c}\in\mathcal{N}\\ \mathbf{c}\notin\mathcal{H}_r}}{\operatorname{argmax}} \left[(\mathbf{c}-\mathbf{p}_r(t))\cdot\widehat{\mathbf{v}}_r(t) + \lambda_{\text{local}}\Phi_{\text{att}}(\mathbf{c},t)\right] \tag{S12}$$

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,

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

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

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

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

Proof. The proof relies on the following assumptions stated in the theorem for each time step $t\ge t_0$ while robot r is in the FOLLOW state prior to detecting the target: (i) the robot moves on a four-neighbor grid where each step changes the Manhattan distance $d_1(\mathbf{p}_r(t),\mathbf{t}_{\text{loc}}) = \|\mathbf{p}_r(t)-\mathbf{t}_{\text{loc}}\|_1$ by ±1; (ii) the sensed attractive-pheromone vector satisfies $\|\mathbf{v}_r(t)\|>0$ so that the unit direction $\widehat{\mathbf{v}}_r(t)=\mathbf{v}_r(t)/\|\mathbf{v}_r(t)\|$ is well defined and the robot remains in FOLLOW; (iii) the set of feasible (unvisited) neighboring cells $\mathcal{F}_r(t)=\mathcal{N}\setminus\mathcal{H}_r$ contains at least one progress cell $\mathbf{c}^+$ satisfying

$$d_1(\mathbf{c}^+,\mathbf{t}_{\text{loc}}) = d_1(\mathbf{p}_r(t),\mathbf{t}_{\text{loc}}) - 1; \tag{S13}$$

and (iv) for every feasible neighbor $\mathbf{m}\in\mathcal{F}_r(t)$ that does not reduce the Manhattan distance to the target ($d_1(\mathbf{m},\mathbf{t}_{\text{loc}})\ge d_1(\mathbf{p}_r(t),\mathbf{t}_{\text{loc}})$), the directional alignment of $\mathbf{c}^+$ exceeds that of $\mathbf{m}$ by more than the maximum weighted difference in local attractive pheromone across $\mathcal{F}_r(t)$:

$$(\mathbf{c}^+-\mathbf{m})\cdot\widehat{\mathbf{v}}_r(t) > \lambda_{\text{local}}\Delta\Phi_{\max}(t), \tag{S14}$$

where $\lambda_{\text{local}}>0$ and

$$\Delta\Phi_{\max}(t) = \max_{\mathbf{c},\mathbf{q}\in\mathcal{F}_r(t)} \left|\Phi_{\text{att}}(\mathbf{c},t)-\Phi_{\text{att}}(\mathbf{q},t)\right|. \tag{S15}$$

At step t, Algorithm 3 selects the next cell in $\mathcal{F}_r(t)$ that maximizes the objective function from Eq. (S12):

$$J_t(\mathbf{c}) = (\mathbf{c}-\mathbf{p}_r(t))\cdot\widehat{\mathbf{v}}_r(t) + \lambda_{\text{local}}\Phi_{\text{att}}(\mathbf{c},t). \tag{S16}$$

To show that the robot is guaranteed to select a distance-reducing move, we compare the score $J_t(\mathbf{c}^+)$ of the progress neighbor $\mathbf{c}^+$ against the score $J_t(\mathbf{m})$ of any feasible non-progress neighbor $\mathbf{m}\in\mathcal{F}_r(t)$. (If $\mathcal{F}_r(t)$ contains no non-progress neighbors, then every available move reduces the distance to the target and the claim holds immediately.) Subtracting $J_t(\mathbf{m})$ from $J_t(\mathbf{c}^+)$ gives

$$\begin{aligned} J_t(\mathbf{c}^+)-J_t(\mathbf{m}) &= \big((\mathbf{c}^+-\mathbf{p}_r(t))-(\mathbf{m}-\mathbf{p}_r(t))\big)\cdot\widehat{\mathbf{v}}_r(t) + \lambda_{\text{local}}\big(\Phi_{\text{att}}(\mathbf{c}^+,t)-\Phi_{\text{att}}(\mathbf{m},t)\big)\\ &= (\mathbf{c}^+-\mathbf{m})\cdot\widehat{\mathbf{v}}_r(t) + \lambda_{\text{local}}\big(\Phi_{\text{att}}(\mathbf{c}^+,t)-\Phi_{\text{att}}(\mathbf{m},t)\big). \end{aligned} \tag{S17}$$

By the definition of $\Delta\Phi_{\max}(t)$ in Eq. (S15), the pheromone difference between $\mathbf{c}^+$ and $\mathbf{m}$ is at most $\Delta\Phi_{\max}(t)$ in magnitude, which implies

$$\Phi_{\text{att}}(\mathbf{c}^+, t) - \Phi_{\text{att}}(\mathbf{m}, t) \geq -\Delta\Phi_{\max}(t). \tag{S18}$$

Multiplying Eq. (S18) by $\lambda_{\text{local}} > 0$ and substituting it into Eq. (S17) yields

$$J_t(\mathbf{c}^+) - J_t(\mathbf{m}) \geq (\mathbf{c}^+ - \mathbf{m}) \cdot \widehat{\mathbf{v}}_r(t) - \lambda_{\text{local}}\Delta\Phi_{\max}(t). \tag{S19}$$

By the condition in Eq. (S14), the right-hand side of Eq. (S19) is strictly positive, so

$$J_t(\mathbf{c}^+) > J_t(\mathbf{m}). \tag{S20}$$

Because $\mathbf{c}^+$ achieves a strictly higher score than every feasible non-progress neighbor $\mathbf{m}$, no non-progress neighbor can maximize $J_t(\mathbf{c})$. Consequently, the cell $\mathbf{p}_r(t{+}1) = \text{argmax}_{\mathbf{c}\in\mathcal{F}_r(t)} J_t(\mathbf{c})$ chosen by Algorithm 3 must be a progress neighbor (either $\mathbf{c}^+$ itself or another neighbor that also reduces d_1 by 1).

Therefore, at every step $t \geq t_0$ before target detection, the Manhattan distance to the target decreases by 1:

$$d_1\big(\mathbf{p}_r(t+1), \mathbf{t}_{\text{loc}}\big) = d_1\big(\mathbf{p}_r(t), \mathbf{t}_{\text{loc}}\big) - 1. \tag{S21}$$

Applying Eq. (S21) over k consecutive steps from t_0 gives

$$d_1\big(\mathbf{p}_r(t_0+k), \mathbf{t}_{\text{loc}}\big) = d_1\big(\mathbf{p}_r(t_0), \mathbf{t}_{\text{loc}}\big) - k. \tag{S22}$$

Since $d_1\big(\mathbf{p}_r(t_0+k), \mathbf{t}_{\text{loc}}\big) \geq 0$, the distance reaches 0 (meaning the robot reaches $\mathbf{t}_{\text{loc}}$) in at most $k = d_1\big(\mathbf{p}_r(t_0), \mathbf{t}_{\text{loc}}\big)$ steps. Because Algorithm 3 checks whether the target lies within the sensing region $\mathcal{V}_R(\mathbf{p}_r(t))$ at the beginning of each step and transitions to RESCUE as soon as the target is detected (which occurs when $\|\mathbf{p}_r(t) - \mathbf{t}_{\text{loc}}\|_\infty \leq R$), the robot either enters RESCUE earlier or reaches $\mathbf{t}_{\text{loc}}$ in at most $T_{\text{reach}} \leq d_1\big(\mathbf{p}_r(t_0), \mathbf{t}_{\text{loc}}\big)$ time steps. □

4. Experimental Design Summary

The experimental design spans simulation, Robotarium, and RoboMaster with increasing physical realism. Search, field ablation, recruitment, and robot-loss experiments correspond to Sections 2.1–2.5 of the main manuscript, while Section 2.6 reports the RoboMaster physical-trail validation. Table S1 summarizes the configurations and repetitions used in each evaluation.

Within each matched experiment, compared methods used the same workspace configuration, robot and target initialization, mission horizon, and failure schedule where applicable. Fixed deterministic seeds were used throughout the reported experiments, based on a common base seed of 42, so that matched methods were evaluated under consistent initial conditions. For RoboMaster experiments, an explicit seed was supplied rather than the default wall-clock initialization.

Centralized lawnmower coverage provides a globally coordinated reference, whereas random walk provides a decentralized baseline without environmental memory. The field-ablation experiments isolate the contribution of the two pheromone fields by comparing the

Table S1: **Summary of the experimental design.** Simulation provides large-scale evaluation, Robotarium introduces physical robot dynamics with virtual pheromone fields, and RoboMaster evaluates search and robot loss using physically realized UV-visible trails.

Evaluation	Experimental configuration	Repetitions
Simulation		
Search	Ten joint workload levels with square grids from 50 by 50 to 500 by 500 cells, 5–50 robots, and 15–150 targets. Grid size, swarm size, and target count increase together.	Centralized: 100 STIGSAR: 1,000 Random walk: 1,000
Robot loss	Same ten search configurations with 1–10 failed robots and four failure-time conditions: early, middle, late, and mixed. Failed robots stop without notifying the remaining robots.	Centralized: 400 STIGSAR: 4,000 Random walk: 4,000
Recruitment	Two grid sizes (50 by 50 and 100 by 100), 1–3 searching robots, four initial distance shells (20, 30, 40, and 50 cells), and 10 repetitions per combination.	240 runs
Field ablation	Same factorial design as recruitment, evaluated for repulsive-only, attractive-only, evaporative merged-field, persistent merged-field, sign-flip, and dual-pheromone STIGSAR.	240 per variant 1,440 total
Robotarium		
Search	Centralized lawnmower, STIGSAR, and random walk evaluated with physical differential-drive robots and virtually projected pheromone fields.	1 centralized 3 STIGSAR 3 random walk
Robot loss	Same search comparison with two robots stopped during execution without notifying the remaining robots.	1 centralized 3 STIGSAR 3 random walk
Recruitment	Attractive-pheromone following from initial separations of 30 and 40 grid cells.	3 per separation 6 total
Field ablation	Six controller variants evaluated at initial separations of 30 and 40 grid cells.	12 per variant 72 total
RoboMaster		
Search	Centralized lawnmower, STIGSAR, and random walk evaluated using physically sensed UV-visible trails.	1 centralized 3 STIGSAR 3 random walk
Robot loss	Same physical-trail comparison with robot loss introduced during search.	1 centralized 3 STIGSAR 3 random walk

full STIGSAR controller with five reduced or merged-field variants. Search performance is characterized in the main manuscript using target discovery, coverage, steps to 95% coverage, and visits per covered cell; recruitment uses target-reaching success and convergence time; field ablation uses mission completion and normalized completion cost; and robot-loss experiments compare performance with the corresponding no-failure condition.

References

[1] Anderson, B., Jackson, J., & Sitharam, M. (1998). Descartes' rule of signs revisited. *The American Mathematical Monthly*, *105*(5), 447–451.