

Collective chemotactic search

Adam Wysocki^{1,*}, Hugues Meyer^{1,2,†} and Heiko Rieger^{1,‡}¹*Department of Theoretical Physics & Center for Biophysics, Saarland University, 66123 Saarbrücken, Germany*²*School of Physics and Astronomy, University of Nottingham, Nottingham NG7 2RD, United Kingdom*

(Received 6 September 2024; accepted 26 June 2026; published 18 August 2026)

We investigate collective search by self-propelled agents that are repelled by their own chemically produced trails, a minimal mechanism that simultaneously generates indirect interactions and memory. Using lattice and off-lattice models, we show that this mechanism enhances search efficiency through two distinct regimes. In a weak-memory regime, chemical cues are short-lived and interactions primarily promote spatial separation between agents. This reduces redundant exploration while preserving mobility, leading to an optimal trade-off between spatial order and persistence. In a strong-memory regime, long-lived chemical trails induce effective self-avoidance, strongly suppressing revisits and long search times. Here, optimal search occurs at finite memory strength: Permanently persistent trails lead to self-caging, while moderate memory enables efficient exploration. At higher densities, overlapping chemical trails give rise to a collective self-avoidance mechanism that yields substantial cooperative speedup without global spatial order. Together, these results show how chemically mediated memory and interactions can optimize collective search across distinct dynamical regimes.

DOI: [10.1103/qrhc-249z](https://doi.org/10.1103/qrhc-249z)

I. INTRODUCTION

Search processes, in which mobile agents explore an environment to locate targets, are ubiquitous across physics, biology, and robotics [1]. Physically relevant search processes are often stochastic in nature [2–4] and include chemical reaction kinetics [5–7], genetic transcription [8,9], bacteria searching for nutrients [10], immune cells searching for pathogenic cells [11], foraging animals [12–15], and swarming robots [16,17], to name but a few. In all these settings, search efficiency is a key performance metric and is commonly quantified by the mean first-passage time (MFPT) [3], the average time required to reach a target for the first time.

The problem of optimizing search has been studied extensively for single agents and for collections of noninteracting searchers. Optimal strategies are known for a wide range of stochastic processes, including Lévy flights [18], persistent [19] and intermittent random walks [20], random walks with n -step memory [21,22], random walks with resetting [23–27], and random walks interacting with the environment [28].

For multiple *independent* (i.e., noninteracting) agents, the MFPT typically decreases trivially with agent number, and the optimal strategy largely mirrors that of a single searcher [29–34]. In contrast, far less is known about how *interactions*

between agents can qualitatively alter collective search efficiency [35–37].

A central limitation of *collective search* is redundant exploration. Repeated visits to already explored regions slow down target discovery and diminish the benefits of parallelism. In principle, redundant exploration can be suppressed in two complementary ways. First, agents may retain *memory* of their past trajectories and avoid revisits. Second, interactions between agents may promote spatial separation, allowing the search domain to be effectively divided into nonoverlapping regions [35–37]. While memory is known to strongly enhance search efficiency at the single-agent level [21,22,38], its role in collective search—and its interplay with interagent interactions—remains poorly understood.

In this work, we address these questions using autochemotactic particles [39–44] as a minimal and universal model of collective search. Autochemotactic agents produce and respond to a chemical field, thereby interacting both with their own past trajectories and with chemical traces left by other agents. This single mechanism naturally combines *self-interaction*, which implements memory, with *indirect particle-particle interactions*, which mediate communication without direct contact. As such, autochemotaxis provides a unified framework to study how memory and interaction-induced coordination jointly shape collective search.

Beyond its theoretical simplicity, autochemotaxis is biologically well motivated. Chemical signaling is one of the most elementary communication mechanisms available to microorganisms and underlies processes such as bacterial navigation [10,45], immune cell migration [46,47], and collective cell dynamics during development [48]. Chemotactic interactions have also been realized in synthetic matter, like self-propelled microdroplets that communicate via chemorepulsive trails [49] or colloidal particles that leave

*Contact author: a.wysocki@lusi.uni-sb.de†Contact author: hugues.meyer@nottingham.ac.uk‡Contact author: heiko.rieger@uni-saarland.de

phase-change trails [50], mimicking tunable pheromone interactions as among ants. Synthetic chemotaxis in active matter systems generically leads to pattern formation like dynamic clusters and waves [40,51,52] and has recently been studied intensively [41–44,49,53]. Moreover, it is a natural example of nonreciprocal interactions, i.e., interactions that violate the actio=reactio principle and that became recently a major research focus [54–64].

From this perspective, autochemotactic repulsion represents a minimal but ubiquitous strategy by which searchers can share information about explored regions and suppress redundant exploration. Using this framework, we show that collective search efficiency is controlled by a nontrivial balance between memory, directional persistence, and spatial organization. We identify two qualitatively distinct regimes. In a weak-memory regime, efficient search emerges from interaction-induced spatial ordering that effectively partitions the search domain. In a strong-memory regime, long-lived but decaying chemical trails suppress revisits through collective self-avoidance, even in the absence of global order. In both cases, optimal search does not coincide with maximal memory or maximal spatial order. Instead, it arises from an intermediate regime in which memory and interactions are strong enough to reduce redundancy, yet sufficiently weak to avoid overconstraining motion.

All results are obtained numerically. Despite the simplicity of the microscopic rules, collective autochemotactic search is a history-dependent many-body problem with non-Markovian dynamics, which makes an analytical treatment intractable.

The paper is organized as follows. In Sec. II A, we introduce two microscopic models of autochemotaxis. Section II B defines the key observables used throughout the paper. In Sec. II C, we describe how to distinguish between weak- and strong-memory regimes. Section III presents results in the weak-memory regime, where the dynamics is dominated by effective long-range repulsive interactions and self-propulsion. In Sec. IV, we investigate the strong-memory regime, in which chemical trails persist over long times, and highlight the resulting differences in search strategies. Finally, Sec. V summarizes our findings and discusses their broader implications.

II. MODELS, OBSERVABLES, AND MEMORY SCALES

A. Models of autochemotaxis

Autochemotactic motion couples chemical production along particle trajectories with orientation response to chemical gradients. To describe search by such agents, we employ both an off-lattice model, which captures realistic active Brownian dynamics, and a lattice model, which can access much lower densities and strong-memory regimes.

1. Off-lattice model

We consider a two-dimensional off-lattice active Brownian particle (ABP) model in which each self-propelled disk secretes a diffusing chemical field that mediates autochemotactic repulsion [39,65]. Particle i with position $\mathbf{r}_i$ moves with constant propulsion speed v_0 along an intrinsic orientation $\mathbf{e}_i(t) = [\cos \varphi_i(t), \sin \varphi_i(t)]^\top$, undergoing rotational

diffusion. Its translational dynamics obey

$$\dot{\mathbf{r}}_i(t) = v_0 \mathbf{e}_i(t) + \sum_{\substack{j=1 \\ j \neq i}}^N \frac{\mathbf{f}_{ij}}{\gamma_t}, \quad (1)$$

where $\mathbf{f}_{ij}$ is a short-range repulsive force preventing overlap.

Each particle deposits chemical at its instantaneous location. The field $c(\mathbf{r}, t)$ diffuses with coefficient D_c and decays with rate α_c :

$$\dot{c} = D_c \nabla^2 c - \alpha_c c + h_c \sum_{i=1}^N \delta(\mathbf{r} - \mathbf{r}_i(t)). \quad (2)$$

The particles respond to gradients of their own collective field through a chemotactic torque,

$$\dot{\varphi}_i(t) = \beta [\nabla c(\mathbf{r}_i(t), t) \times \mathbf{e}_i(t)]_z + \sqrt{2D_r} \xi_i(t), \quad (3)$$

which rotates the particle orientation away from chemical gradients for $\beta > 0$ (chemorepulsion). Here, $\varphi_i(t)$ denotes the polar angle of the propulsion direction. The second term on the right-hand side is the usual rotational noise for ABPs. We simulate N particles in a square box of length L using periodic boundary conditions. Unless otherwise stated, lengths are measured in units of the particle radius a and time in units of the rotational diffusion time $1/D_r$. The packing fraction is defined as $\phi = \pi a^2 N / L^2$. In Appendix A, we provide additional details of the ABP implementation.

For an ideal ABP, the *bare persistence length* is defined as $l_p = v_0 / (a D_r)$. Interactions may modify the orientational dynamics and, in principle, also the propulsion speed. Thus, we define an effective persistence length as

$$l_p^{\text{eff}} = \frac{v_0^{\text{eff}}}{a D_r^{\text{eff}}}, \quad (4)$$

where v_0^{eff} and D_r^{eff} denote the effective propulsion speed and rotational diffusion coefficient, respectively. The latter is obtained from the decay of the orientation autocorrelation function, $\langle \mathbf{e}(0) \cdot \mathbf{e}(t) \rangle \propto \exp(-D_r^{\text{eff}} t)$. In our autochemotactic model, the dominant effect of the trails is a modification of the orientational dynamics. Direct repulsive interactions can additionally affect the propulsion speed, particularly in dense or clustered states. However, the present work focuses on dilute and spatially homogeneous regimes, where persistent blocking of particles is absent and $v_0^{\text{eff}} \simeq v_0$. Consequently, variations of l_p^{eff} are primarily controlled by changes in D_r^{eff} .

This continuous model connects directly to the standard ABP model and, in the weak-memory regime, reduces effectively to an ABP with Yukawa-type repulsive interactions.

2. Lattice model

As a complementary approach to the continuous off-lattice model, we use a discrete autochemotactic random walk [44]. We place N walkers on the sites of a two-dimensional square lattice of size $L \times L$ with periodic boundary conditions. The state of a walker is described by the probability $\rho(i, t)$ of being at site i at time t . Its dynamics are governed by the master equation

$$\rho(i, t + \Delta t) = \sum_{j \in \mathcal{N}_i} p_{j \rightarrow i}(t) \rho(j, t), \quad (5)$$

where $\mathcal{N}_i$ is the set of nearest neighbors of site i , and the discrete time step is set to $\Delta t = 1$. Thus, lengths are in units of the lattice spacing and time in units of Δt .

The transition probability from site j to a neighboring site i reads

$$p_{j \rightarrow i}(t) = \frac{(1 + b_i) e^{-\beta c(i,t)}}{\sum_{k \in \mathcal{N}_j} (1 + b_k) e^{-\beta c(k,t)}}, \quad (6)$$

where $c(i, t)$ is the chemical concentration at site i and time t , and β is the chemotactic coupling strength. For $\beta > 0$, the walker preferentially moves toward sites with lower concentration, thus generating chemorepulsion, which is the case considered here. While no hard-core exclusion is implemented in the lattice model, repulsive chemotaxis dynamically suppresses multiple occupancy. We define the occupation fraction as $\varphi = N/L^2$. When comparing both models, we identify φ with the off-lattice packing fraction ϕ .

The concentration field evolves according to the discrete analog of Eq. (2) with diffusion constant D_c and decay rate α_c . At each time step, walkers deposit chemical on their current site. For $D_c = 0$ and $\alpha_c = 0$, deposited chemical persists indefinitely and the model reduces to a self-avoiding walk [66]. In this limit, the variance of the chemical field grows linearly in time, a signature commonly associated with aging effects [67]. However, within the parameter range and time windows investigated here, we did not observe a systematic dependence of the mean first-passage time on the equilibration time.

The parameter b_i encodes an *intrinsic* directional persistence: If the jump $j \rightarrow i$ continues in the same direction as the previous step, then $b_i = b$ (a constant), otherwise $b_i = 0$ [19]. In the absence of chemotaxis ($\beta = 0$), the mean number of consecutive steps in the same direction defines the *bare persistence length* $l_p = (4 + b)/3$. Chemotactic interactions modify this persistence: In the case of chemorepulsion ($\beta \neq 0$), walkers tend to maintain their direction for longer, leading to an *effective persistence length* $l_p^{\text{eff}} > l_p$.

The lattice and off-lattice models differ substantially in their microscopic dynamics. The purpose of considering both is not to establish quantitative equivalence, but to test the robustness of trail-mediated self-repulsion as a search mechanism. Agreement between the models should therefore be interpreted at a qualitative level.

B. Core observables

1. Collective speedup

To quantify the benefit of collective search, we compare the MFPT of N interacting agents with that of N noninteracting agents that couple only to their self-generated chemical field. We define the *collective speedup* as

$$\gamma_N = 1 - \frac{\bar{T}_N}{\bar{T}_N^{\text{ind}}}, \quad (7)$$

quantifying the relative gain in performance due to interactions. A value of $\gamma_N > 0$ indicates that interaction between agents leads to faster search than a purely individualistic strategy. Here, $\bar{T}_N$ is the MFPT for the first of N interacting agents to reach a randomly located target. The reference value $\bar{T}_N^{\text{ind}}$ is the MFPT for N independent searchers, which can

be obtained from single-agent search: Let $S_1(t)$ denote the survival probability of the target in a single-agent search, i.e., the probability that the target has not been found by time t . For N independent agents searching in parallel, the joint survival probability is $S_N(t) = [S_1(t)]^N$, and the corresponding MFPT becomes

$$\bar{T}_N^{\text{ind}} = \int_0^\infty [S_1(t)]^N dt. \quad (8)$$

In the special case where $S_1(t)$ decays exponentially, this reduces to $\bar{T}_N^{\text{ind}} = \bar{T}_1/N$. However, this simple scaling can break down for large N due to extreme-value effects associated with the closest searcher [68,69].

2. Spatial order parameter

Since the spatial arrangement of searchers affects collective search performance, we define a quantitative measure of spatial organization that distinguishes homogeneous from inhomogeneous particle distributions [70]. Given an instantaneous configuration of N searchers with positions $\{\mathbf{r}_i\}$, we compute the Voronoi tessellation of the full domain of size L^2 . Each searcher $\mathbf{r}_i$ is associated with a Voronoi cell of area A_i , containing all points closer to $\mathbf{r}_i$ than to any other particle. We normalize the cell areas by defining $\mathcal{A}_i = A_i/\langle A \rangle$, where the mean area is $\langle A \rangle = L^2/N$, and compute the standard deviation of the normalized areas as

$$\sigma_{\mathcal{A}} = \sqrt{\langle \mathcal{A}^2 \rangle - \langle \mathcal{A} \rangle^2}. \quad (9)$$

For randomly distributed particles (Poisson process), this value is known from numerical studies to be $\sigma_{\mathcal{A}}^{\text{ind}} = \sqrt{0.28} \approx 0.53$ in the limit of large N [71]. We define *spatial order parameter* as

$$\eta_{\mathcal{A}} = 1 - \frac{\sigma_{\mathcal{A}}}{\sigma_{\mathcal{A}}^{\text{ind}}}. \quad (10)$$

This quantity measures deviations from random spatial structure: $\eta_{\mathcal{A}} > 0$ indicates a more regular (uniform) spacing than in the random case, while $\eta_{\mathcal{A}} < 0$ indicates clustering. In the present work, we focus exclusively on parameter regimes with $\eta_{\mathcal{A}} \geq 0$. For example, if particles are uniformly assigned to N distinct subdomains as in a split search process, but randomly distributed within each subdomain, the spatial order is found to be $\eta_{\mathcal{A}}^{\text{split}} \approx 0.53$.

C. Memory regimes

The nature of the interaction between autochemotactic searchers is set by two characteristic lengths: a *chemical decay length*

$$\lambda_{\text{decay}} = \sqrt{\frac{D_c}{\alpha_c}}, \quad (11)$$

the typical distance a chemical molecule diffuses before decaying, and an *advective length*

$$\lambda_{\text{adv}} = \frac{2D_c}{v_0}, \quad (12)$$

the distance over which diffusion redistributes the chemical before the source moves away.




FIG. 1. Simulation snapshots of the lattice autochemotactic model illustrating the weak- and strong-memory regimes. (Left) Weak-memory regime ($D_c = 2.5$, $\alpha_c = 0.1$, $2\sqrt{D_c\alpha_c} = 1$), where the chemical field is short-lived and does not form persistent trails. (Right) Strong-memory regime ($D_c = 0.01$, $\alpha_c = 0.001$, $2\sqrt{D_c\alpha_c} = 0.0063$), where persistent, asymmetric trails form along past trajectories and overlap with those of other walkers. In both panels, $N = 20$ walkers move in a box of size $L = 100$ ($\varphi = 0.002$) with coupling strength $\beta = 10$. The background color encodes the chemical concentration field and the red circles indicate walker positions.

To quantify the resulting asymmetry of the chemical profile, consider a walker moving at constant velocity v_0 . In its comoving frame, the steady-state concentration decays asymmetrically ahead and behind (see Appendix B). The ratio of the two length scales controls the fore-aft asymmetry of the profile. In the *zero-memory regime*, if

$$\lambda_{\text{decay}} \ll \lambda_{\text{adv}} \quad \text{or} \quad 2\sqrt{D_c\alpha_c} \gg v_0,$$

the chemical field is short-lived and effectively isotropic around each walker. Interactions are effectively instantaneous and correspond to a Yukawa-type potential in two dimensions. For intermediate values $\lambda_{\text{decay}} \simeq \lambda_{\text{adv}}$ or $2\sqrt{D_c\alpha_c} \simeq v_0$, the system enters a weak-memory regime in which the field retains a finite fore-aft asymmetry but remains short-lived in time, so that the particle primarily interacts with its recently emitted field rather than with a persistent trail. In contrast, in the *strong-memory regime*, if

$$\lambda_{\text{decay}} \gg \lambda_{\text{adv}} \quad \text{or} \quad 2\sqrt{D_c\alpha_c} \ll v_0,$$

the field becomes strongly asymmetric, forming long-lived trails along past trajectories. Walkers then interact with their own history, leading to trail-mediated self-avoidance and dynamically structured environments. For the lattice model, where $v_0 = 1$, by construction, the criterion reduces to $2\sqrt{D_c\alpha_c} \gg 1$ (zero memory) and $2\sqrt{D_c\alpha_c} \ll 1$ (strong memory). Figure 1 shows representative configurations of the chemical field in both regimes. The qualitative contrast between localized profiles and persistent trails anticipates the distinct search mechanisms analyzed below.

III. WEAK-MEMORY REGIME

A. Single searcher

In the weak-memory regime, $\lambda_{\text{decay}} \simeq \lambda_{\text{adv}}$, which lies between the zero-memory and the strong-memory regimes, the chemical field exhibits a finite fore-aft asymmetry but remains short-lived. The particle interacts primarily with the local

gradient of its recently emitted field rather than with a long-lived record of its past trajectory. Self-interaction, therefore, produces a change in persistence but does not introduce genuine long-term memory. We first consider the single-particle case ($N = 1$) to isolate self-interaction effects.

Several microscopic parameters influence the effective persistence length l_p^{eff} by affecting gradient strength or shape. In each model, we vary only one control parameter (β in the lattice model, D_c in the off-lattice model), noting that the results depend only on the resulting persistence length, not on the specific tuning mechanism.

In the lattice model, increasing the chemotactic coupling strength β enhances the response to the fore-aft asymmetry, biasing the motion to straighter trajectories and thereby increasing l_p^{eff} beyond its bare value l_p [see inset in Fig. 2(a)]. In the limit of very strong coupling, the particle approaches a ballistic regime.

In the off-lattice model, β is fixed, and instead we vary the chemical diffusion constant D_c . Changing D_c changes the shape of the self-generated field, which in turn modifies the effective rotational diffusion constant D_r^{eff} and thereby $l_p^{\text{eff}} = v_0/(aD_r^{\text{eff}})$. For small D_c , the chemical signal remains highly localized and the particle rarely overlaps with its own trail; for large D_c , the signal becomes broad and shallow. In both limits, $D_r^{\text{eff}} \rightarrow D_r$, and thus $l_p^{\text{eff}} \rightarrow l_p$. As a result, D_r^{eff} exhibits a minimum and correspondingly l_p^{eff} a maximum, at intermediate D_c [see inset of Fig. 2(b)].

This change in persistence directly impacts search efficiency. When the mean first-passage time $\bar{T}_1$ is plotted against l_p^{eff} [see Figs. 2(a) and 2(b)], the data collapse onto a single master curve in both models, matching the behavior of a persistent random walker and a simple ABP. In the lattice model, $\bar{T}_1$ diverges as $l_p^{\text{eff}} \rightarrow \infty$, consistent with persistent random walks [19]. Such a divergence does not occur for finite persistence in the off-lattice model, where orientations are continuous.

These results demonstrate that, in the weak-memory regime, autochemotactic self-interaction acts primarily as




FIG. 2. Single-agent search as a function of effective persistence length l_p^{eff} . Mean first-passage time $\bar{T}_1$ vs l_p^{eff} in (a) the lattice model, where l_p^{eff} is tuned by the chemotactic coupling strength β . Colors correspond to different values of the bare persistence length l_p . (b) Off-lattice model, where l_p^{eff} is tuned by the chemical diffusivity D_c . Colors indicate different values of D_c . Insets show how l_p^{eff} depends on the respective control parameter. For finite chemical decay, the data collapse onto a common master curve consistent with ABP behavior. Empty symbols in panel (b) correspond to $\alpha_c = 0$, where strong-memory effects lead to systematic deviations from the master curve. In all cases with finite chemical decay, the parameters lie in the finite-memory regime defined in Sec. IIC, where no persistent trail accumulation occurs.

a renormalization of persistence: Search efficiency is controlled by l_p^{eff} , independent of microscopic interaction details. However, in the absence of chemical decay ($\alpha_c = 0$), the equivalence with a simple ABP breaks down [see empty symbols in Fig. 2(b)]. The accumulating concentration field introduces strong long-term memory, and systematic deviations from the ABP master curve emerge. In particular, particles with small D_c can attain MFPTs below the ABP optimum and approach the ballistic limit $\bar{T}_1 = L^2/(4l_p)$. This shows that in the strong-memory regime, the self-generated trail does more than renormalize persistence but qualitatively modifies the search strategy. The collective consequences of this long-memory behavior are discussed in Sec. IV.




FIG. 3. (a) Effective persistence length l_p^{eff} of interacting autochemotactic lattice random walkers (RWs) (circles) compared to the single-particle value (dots) as a function of chemotactic coupling β . Collisions suppress the growth of l_p^{eff} at strong coupling. (b) Spatial order parameter η_A . (c) Collective MFPT $\bar{T}_N$ (circles) and the independent-search reference $\bar{T}_N^{\text{ind}}$ (dots). (d) Cooperative speedup γ_N . Parameters: bare persistence $l_p = 4/3$ and searcher occupation fraction $\varphi = 0.012$.

B. Collective search

Collective search performance emerges from the interplay between two competing ingredients: the tendency of agents to spread across space and avoid redundant exploration, and their ability to move persistently to efficiently scan unexplored regions. In autochemotactic systems, repulsive interactions promote spatial separation, while self-interaction with chemical trails enhances directional persistence. These mechanisms are intrinsically coupled, so that collective efficiency cannot be optimized by tuning either persistence or spatial order in isolation.

We first vary the chemotactic coupling β at fixed density to probe how this balance shapes collective dynamics in the lattice model (see Fig. 3). At weak coupling, increasing β primarily induces spatial order while leaving the effective persistence length l_p^{eff} unchanged [Figs. 3(a) and 3(b)]. A constant l_p^{eff} indicates that the initial performance gain is driven by spatial separation rather than persistence enhancement. Repulsion reduces redundant exploration, leading to a decrease of the MFPT [Fig. 3(c)] and a corresponding increase of the cooperative speedup γ_N [Fig. 3(d)]. At intermediate coupling, self-interaction enhances directional persistence. In contrast to the single-particle case, however, interparticle collisions




FIG. 4. Minimum MFPT $\min_{\beta, l_p}(\bar{T}_N \varphi)$ as a function of density φ in the lattice model. The corresponding optimal l_p^{eff} and η_A are shown. Here, l_p^∞ denotes the effective persistence length in the strong-coupling limit $\beta \rightarrow \infty$. Shaded regions indicate distinct density regimes with different optimal search strategies.

bound the growth of l_p^{eff} , which saturates at a finite value $l_p^\infty \equiv \lim_{\beta \rightarrow \infty} l_p^{\text{eff}}(\beta, \varphi)$ [Fig. 3(a)]. In this regime, the single-search MFPT is already close to its persistence-controlled optimum, and $\bar{T}_N \approx \bar{T}_N^{\text{ind}}$. Collective interactions therefore provide little additional advantage, and the speedup γ_N decreases and can even become negative. At strong coupling, a distinct lattice-specific effect appears. As $\beta \rightarrow \infty$, single-particle motion becomes ballistic and both $\bar{T}_1$ and $\bar{T}_N^{\text{ind}}$ diverge [see Fig. 2(a)]. In contrast, collisions regularize persistence in the interacting system, allowing particles to explore new regions so that the collective MFPT $\bar{T}_N$ remains finite. This leads again to an increase of γ_N at large β , even though neither spatial order nor persistence grows further.

The β scan demonstrates that spatial order and effective persistence are intrinsically coupled: large values of both cannot be achieved simultaneously in the weak-memory regime (see Appendix C). Because efficient search requires both persistent motion and suppression of redundant exploration, the optimal strategy corresponds to a specific balance between them. Each pair of control parameters (l_p, β) therefore selects a distinct point in $(\eta_A, l_p^{\text{eff}})$ space. We now determine, for each density, the combination that minimizes the MFPT.

Figure 4 shows the minimum MFPT $\min_{\beta, l_p}(\bar{T}_N \varphi)$ as a function of density φ in the lattice model, together with the corresponding optimal values of l_p^{eff} and η_A . Three density regimes emerge, each characterized by a distinct optimal compromise between persistence and spatial separation. At low density, optimal search occurs at moderate coupling: Repulsion generates finite spatial order while persistence remains sufficiently large to ensure efficient exploration. At intermediate densities, strong coupling becomes optimal. Particles move nearly ballistically between collisions, spatial order




FIG. 5. Cooperative search speedup γ_N vs spatial order η_A for repulsive Yukawa ABPs, autochemotactic ABPs, and autochemotactic RWs. Color indicates the effective persistence length l_p^{eff} . The gray region corresponds to $\gamma_N > 0$ and $\eta_A > 0$, where cooperative search is enhanced and spatial order exceeds that of independent searchers. Colored curves show the Yukawa ABP reference data for different values of l_p^{eff} . While the models do not collapse onto a universal curve, they exhibit a common qualitative trend: Cooperative speedup is typically largest at intermediate spatial order and decreases again at high order. The lattice occupation fraction is $\varphi = 0.012$ and the off-lattice packing fraction is $\phi = 0.01$.

approaches its maximal attainable value, and the effective persistence saturates at l_p^∞ . At high density, however, strong coupling is detrimental. Frequent collisions suppress persistence, and the optimal strategy shifts toward weak coupling, $\beta \approx 0$, where interactions are minimal and persistence is controlled by the bare value l_p . Optimal search is achieved at an intermediate balance between persistence and spatial separation, rather than at maximal spatial order.

While the collective MFPT $\bar{T}_N$ quantifies absolute group performance, the cooperative speedup γ_N directly measures the efficiency gain due to interactions relative to independent searches. To test whether the relation between spatial order and cooperative gain depends on the microscopic interaction mechanism, we compare three distinct models: repulsive Yukawa ABPs, autochemotactic ABPs, and the lattice autochemotactic random-walker model. Figure 5 shows γ_N as a function of spatial order η_A . Although the microscopic dynamics differ substantially, all three models exhibit the same qualitative trend. A moderate degree of spatial order is generally associated with enhanced cooperative speedup, indicating that ordering suppresses redundant exploration and promotes a more efficient division of search effort. However, this trend is not monotonic. When the spatial order becomes too strong, particle motion becomes increasingly constrained, and the cooperative speedup can saturate or even decrease due to caging and overconstrained motion. The decrease of γ_N at large η_A occurs entirely within the homogeneous regime and is therefore not associated with clustering or band formation. The agreement between the models is qualitative rather than quantitative. While the data do not collapse onto a single universal curve, most autochemotactic data lie close to the reference curves of repulsive Yukawa ABPs. For a given l_p^{eff} , only a finite range of speedups is attainable, and optimal




FIG. 6. Mean first-passage time for repulsive ABPs. Minimum MFPT $\bar{T}_N^{\min}$, obtained by optimizing over the strength of the repulsive Yukawa interaction, as a function of the bare persistence length l_p (dots). For comparison, the dashed line shows the MFPT of independent ABPs $\bar{T}_N^{\text{ind}}$, while the dot-dashed line corresponds to MFPT $\bar{T}_N^{\text{split}}$ of an artificial domain-splitting strategy in which walkers explore nonoverlapping subdomains. The packing fraction is $\phi = 0.01$.

performance occurs at intermediate spatial order rather than maximal ordering, consistent with the density-dependent optimization shown in Fig. 4.

An even stronger test of mechanism independence is obtained by comparing soft repulsive interactions with an extreme limit in which spatial separation is enforced geometrically. Figure 6 shows the minimum MFPT for Yukawa-interacting active Brownian particles, optimized over interaction strength, together with an artificial domain-splitting strategy in which each walker explores a nonoverlapping subdomain. In practice, this is implemented by comparing N particles in a square box of side length L with a single particle in a box of side length $L/\sqrt{N}$, such that the area per searcher is identical in both cases. Despite the absence of explicit interactions in the latter scenario, the resulting MFPT reductions are quantitatively similar. This demonstrates that collective search enhancement does not depend on the microscopic form of repulsion, but rather on its effective role in allocating exploration space and regulating persistence.

Taken together, these results show that collective search is optimized by tuning interactions to enhance persistence without enforcing maximal spatial order.

IV. STRONG-MEMORY REGIME

A strong-memory regime arises when the chemical trail left by a walker is sufficiently sharp and asymmetric to encode spatially resolved information about its recent trajectory. This occurs when the advective screening length is much smaller than the chemical decay length, $\lambda_{\text{adv}} \ll \lambda_{\text{decay}}$ [see Eqs. (11) and (12)]. In our lattice model with propulsion speed $v_0 = 1$, this criterion reduces to $2\sqrt{D_c \alpha_c} \ll 1$. Except in limiting cases where diffusion is so large that chemical trails become spatially smeared, or decay is so fast that traces vanish almost instantly, this condition reliably predicts the onset of strong memory. In this regime, walkers remain sensitive to




FIG. 7. Representative trajectories of autochemotactic lattice random walkers of duration $T = 60$ at density $\phi = 0.005$ and coupling strength $\beta = 50$. (a) Weak-memory regime at $(D_c, \alpha_c) = (10^{-1}, 10^{-1})$, where chemical traces are short-lived and walkers do not effectively avoid previously explored regions, leading to frequent intersections between trajectories of different walkers. (b) Strong-memory regime at $(D_c, \alpha_c) = (10^{-3}, 10^{-3})$, where long-lived and localized chemical trails suppress crossings with both the walker's own past trajectory and the trails produced by others, resulting in pronounced collective self-avoidance.

chemical signals produced far in the past, and their dynamics is qualitatively distinct from that in the weak-memory limit. Representative trajectories of the two regimes are shown in Fig. 7.

An operational signature of strong memory is provided by the survival probability of targets. For sufficiently small D_c and α_c , the survival probability exhibits a compressed-exponential decay,

$$S_N(t) = \exp[-(t/T)^\xi], \quad (13)$$

with exponent $\xi > 1$, indicating a strong suppression of long search times. As D_c or α_c increase, the trail becomes less informative, ξ decreases toward unity, and the decay crosses over to a purely exponential form characteristic of weak memory. Across the explored parameter range, the condition $\xi > 1$ coincides with $2\sqrt{D_c \alpha_c} \ll 1$, providing a consistent identification of the strong-memory regime [see Fig. 8(a) and details in Appendix D].

In the strong-memory regime of our lattice model, chemical trails become narrow and localized, which facilitates frequent reorientations and can reduce the effective persistence length. Nevertheless, long-lived trails continue to repel the walker from previously visited regions, so that self-avoidance and the suppression of revisits persist even in the absence of large directional persistence [Fig. 7(b)].

A more microscopic characterization of strong memory is provided by the suppression of self-crossings [72]. Long-lived chemical trails effectively repel the walker from its own previously visited locations, strongly reducing the number of times its trajectory intersects its past path [Fig. 7(b)]. To quantify this effect, we measure the average number of crossings $\bar{N}_{\text{cross}}$ between a walker's trajectory and previously deposited chemical trails during the time until the target was found, and compare it to the corresponding value $\bar{N}_{\text{cross}}^{\text{ind}}$ for a random walker with the same effective persistence length. We define




FIG. 8. Strong-memory effects on collective search in the lattice model. (a) Location of the optimal mean first-passage time $\min(\bar{T}_N \varphi)$ (filled circles) and of the maximal spatial order $\max(\eta_A)$ (open diamonds) in the (D_c, α_c) plane for increasing searcher density φ . The upper-right shaded region corresponds to weak memory, where $2\sqrt{D_c \alpha_c} > 1$ and chemotactic memory effects are negligible. The colored contours delimit the strong-memory regime ($\xi > 1$); the enclosed region expands with increasing searcher density φ . Colors encode the searcher density φ and are used consistently across all panels; representative densities are indicated in panel (a). The single-searcher case is indicated by a black outlined symbol and corresponds to $\varphi = 10^{-4}$ in a system of size $L = 100$. (b) Optimal search cost $\min(\bar{T}_N \varphi)$ as a function of density φ . (c) Path-crossing reduction parameter χ_{cross} , averaged over the strong- and weak-memory regimes. (d) Cooperative speedup γ_N , averaged over the strong- and weak-memory regimes. (e) Average spatial order η_A within the strong-memory regime compared to the global maximum of $\max(\eta_A)$. The chemotactic coupling strength is $\beta = 50$ and the bare persistence length is $l_p = 4/3$.

the normalized crossing reduction

$$\chi = 1 - \frac{\bar{N}_{\text{cross}}}{\bar{N}_{\text{ind}}}, \quad (14)$$

so that $\chi > 0$ indicates reduced self-crossing and enhanced self-avoidance. While the exponent ξ characterizes the global survival statistics, the crossing reduction χ provides a local, trajectory-level measure of memory-induced self-avoidance.

Strong memory substantially enhances search efficiency, but remarkably, infinite memory does not. In the limit $D_c = \alpha_c = 0$, corresponding to a true self-avoiding walk [66], chemical trails never decay and permanently exclude previously visited regions. While this completely suppresses revisits at short times, it ultimately becomes detrimental: closed loops of chemical traces form effective barriers that either trap the walker within its own trail network (self-caging) [49] or temporarily isolate unexplored regions from the accessible domain. Both effects reduce the effective search space and increase the MFPT. Efficient search therefore emerges at *finite* memory strength: Chemical traces must persist long enough to prevent frequent revisits, yet decay on a timescale that allows reentry into regions that were previously enclosed. This competition explains why the minimal MFPT lies inside the strong-memory regime but away from the self-avoiding limit [Fig. 8(a)], highlighting the importance of optimal forgetting. While the ideal space-covering lower bound $\bar{T}_1 = L^2/2$ cannot be surpassed, finite memory allows MFPT values that are much closer to this bound than those obtained with permanently self-avoiding trajectories [Fig. 8(b)].

To isolate the role of memory from chemotactic interactions, we also considered a minimal lattice model in which self-avoidance is imposed explicitly over a finite temporal

window (a “snake” walk) (see Appendix E for details). Figure 9 shows the mean first-passage time as a function of the snake length l_{snake} . In all cases, the MFPT exhibits a clear minimum at finite memory. When rescaled by system size, the optimal memory length collapses as $l_{\text{snake}}^* \sim L^2$, corresponding to the time required to explore the domain once. This minimal model reproduces the same trade-off observed in autochemotactic search: Insufficient memory leads to frequent revisits, whereas excessive memory causes self-caging and impedes exploration.




FIG. 9. Mean first-passage time for the snake lattice model as a function of the normalized memory length l_{snake}/L^2 for different system sizes L . All curves exhibit a clear minimum at finite memory, with optimal values collapsing at $l_{\text{snake}} \sim L^2$.

To further isolate the role of memory from a mere reduction of accessible search volume, we also investigated a minimal model in which lattice sites stochastically switch between accessible and inaccessible states, independently of the walker trajectory. While this introduces a fluctuating obstacle field that constrains motion, it does not systematically suppress redundant revisits, since site accessibility carries no information about the walker's past exploration. Thus, this model fails to reproduce the strong reduction of the MFPT observed for snake walks or autochemotactic trails. This confirms that efficient search in the strong-memory regime relies on trajectory-dependent memory that actively penalizes revisits rather than on a passive reduction of the accessible search space.

The same trail-mediated mechanism governs collective behavior. As the searcher density increases, the location of the minimal MFPT shifts toward larger values of D_c and α_c [see Fig. 8(a)]. At high density, very strong memory becomes counterproductive: The superposition of many long-lived chemical trails leads to an oversuppression of trajectory crossing, which can hinder access to remaining unexplored regions and slow down the search. Increasing D_c or α_c weakens memory and partially restores access to these regions, at the cost of allowing some revisits. The optimal search strategy therefore shifts toward intermediate memory strength, balancing revisit suppression against collective overavoidance. The same collective overavoidance mechanism also explains why the operational boundary defined by $\xi > 1$ shifts toward larger (D_c, α_c) as the searcher density increases.

Within the strong-memory regime, the resulting collective speedup γ_N is significantly larger than in the weak-memory regime [see Fig. 8(d)]. Crucially, the origin of this speedup differs qualitatively from that discussed in Sec. III. While efficient search at weak memory relies on spatial ordering and an effective partitioning of the search domain, spatial order remains negligible throughout the strong-memory regime [Fig. 8(e)]. The order parameter η_A stays small, and its global maximum appears only at larger (D_c, α_c) , outside the strong-memory region [see Fig. 8(a)]. Instead, walkers in the strong-memory regime avoid not only their own long-lived chemical traces but also those produced by others, thereby reducing redundant exploration without inducing global alignment. We observe that path crossing is highly reduced and that the normalized crossing suppression χ is systematically enhanced as compared to the weak-memory regime [see Fig. 8(c)]. The resulting network of overlapping chemical traces forms a dynamically evolving constraint landscape, reminiscent of an entangled polymer melt, with walkers acting as mobile end points. It is this trail-mediated collective self-avoidance—rather than spatial order—that underlies the pronounced cooperative speedup observed at strong memory.

Taken together, these results establish the strong-memory regime as a distinct mode of collective search. Long-lived but decaying chemical trails regulate exploration by suppressing revisits at both the single-particle and collective level. Optimal search performance is achieved when memory is strong enough to encode past motion, yet sufficiently transient to avoid self-caging. This mechanism is fundamentally distinct from the order-driven enhancement operative in the weak-memory regime and identifies environmentally

mediated memory, encoded in the self-generated chemical field, as an independent and powerful control parameter for collective search.

We note that other collective states known in chemotactic systems, such as traveling bands, can severely degrade search performance by concentrating exploration at the band front; a scaling analysis is provided in Appendix F.

V. CONCLUSIONS AND OUTLOOK

In this work, we have shown that efficient collective search does not arise from spatial order alone, but from the mechanism by which redundant exploration is suppressed. Using autochemorepulsive searchers as a minimal model that combines indirect interactions and memory, we identified two qualitatively distinct regimes of collective optimization, governed by the lifetime of self-generated chemical trails.

In the weak-memory regime, chemical cues are short-lived and primarily act to mediate repulsive interactions between agents. Here, efficient search relies on a balance between directional persistence and spatial separation: Moderate ordering distributes searchers across the domain while preserving sufficient mobility. In this regime, excessive ordering or overly strong interactions reduce performance by constraining motion and promoting collective structures that do not contribute effectively to exploration.

In contrast, the strong-memory regime is controlled by long-lived and localized chemical trails that encode information about past trajectories. Search efficiency is then enhanced not by global spatial order, which remains negligible, but by trail-mediated self-avoidance. At both the single-agent and collective level, long-lived trails suppress revisits and strongly reduce long search times, leading to compressed-exponential survival statistics and pronounced cooperative speedup. Importantly, optimal search is achieved at finite memory strength: Permanently persistent trails induce self-caging and fragment the accessible search space, while moderate decay allows reentry into previously enclosed regions and maintains global connectivity.

These results demonstrate that collective search efficiency can be regulated by qualitatively different physical mechanisms depending on memory strength. Weak memory favors order-driven domain sharing, whereas strong memory gives rise to collective self-avoidance without alignment or large-scale ordering. This distinction highlights memory as an independent and powerful control parameter for collective search, rather than a secondary correction to persistence or interaction strength.

Beyond autochemotactic systems, our findings point to a broader principle: Interactions enhance collective search only insofar as they suppress redundant exploration without overconstraining motion. The similarity between chemotactic repulsion, soft pairwise interactions, and even enforced domain partitioning underscores that the microscopic origin of avoidance is less important than its effective impact on exploration dynamics.

Several open directions remain. Alternative forms of self-generated gradients, including attractant sinks [45,73,74] or chemokinetic responses [75,76], may realize different trade-offs between memory and mobility. More complex cost

functions that account for signal production or energetic constraints could modify optimal strategies [77], particularly in dense or competitive environments. Finally, extending these ideas to adaptive or learning agents [78–81] may help bridge stochastic search theory, active-matter physics, and collective decision-making in biological and artificial systems.

Taken together, our work establishes autochemotactic repulsion as a minimal and versatile framework for studying collective search and reveals how tuning memory and interactions can switch the dominant mechanism by which groups explore space efficiently.

ACKNOWLEDGMENTS

We acknowledge financial support from DFG through the Collaborative Research Center SFB 1027. This work was supported by the Medical Research Council (Grant No. MR/Y003845/1). Parts of the manuscript were revised and edited using ChatGPT (OpenAI) to improve clarity and expression. We thank Reza Shaebani for helpful discussions.

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

APPENDIX A: OFF-LATTICE MODEL DETAILS

Particles are disks of radius a interacting through a harmonic repulsion $V(r) = \frac{k}{2}(2a - r)^2$ for $r \leq 2a$, with stiffness k . The force on particle i due to particle j is $\mathbf{f}_{ij} = -\nabla_{\mathbf{r}_i} V(r_{ij})$, where $r_{ij} = |\mathbf{r}_i(t) - \mathbf{r}_j(t)|$ is the distance between particles i and j .

The source term in Eq. (2) is regularized using a Gaussian profile centered at $\mathbf{r}_i - a\mathbf{e}_i$ to reduce trivial self-interaction, and the local gradient in Eq. (3) is sampled at $\mathbf{r}_i + a\mathbf{e}_i$.

Rotational noise $\xi_i(t)$ in Eq. (3) is Gaussian white noise with $\langle \xi_i(t) \rangle = 0$ and $\langle \xi_i(t) \xi_j(t') \rangle = \delta_{ij} \delta(t - t')$. We assume that the translational and rotational friction constants correspond to those of a sphere with radius a , i.e., $\gamma_r/\gamma_t = 4/3 a^2$.

We use a simulation box of size $L/a = 100$ with periodic boundary conditions. The system is first equilibrated before starting any measurements.

APPENDIX B: STATIONARY CONCENTRATION PROFILE

Consider a two-dimensional diffusion equation with a single source moving at constant speed v_0 :

$$\frac{\partial c}{\partial t}(\mathbf{r}, t) = D_c \nabla^2 c(\mathbf{r}, t) - \alpha_c c(\mathbf{r}, t) + h_c \delta(x - v_0 t, y). \quad (\text{B1})$$

Introducing the comoving field $\tilde{c}(x, y, t) = c(x + v_0 t, y, t)$, the steady-state profile $c_\infty(x, y) = \lim_{t \rightarrow \infty} \tilde{c}(x, y, t)$ satisfies

$$\nabla^2 c_\infty + 2\lambda_{\text{adv}}^{-1} \frac{\partial c_\infty}{\partial x} - \lambda_{\text{decay}}^{-2} c_\infty + \frac{h_c}{D_c} \delta(x, y) = 0, \quad (\text{B2})$$

where $\lambda_{\text{adv}} = 2D_c/v_0$ and $\lambda_{\text{decay}} = \sqrt{D_c/\alpha_c}$. Taking the Fourier transform gives

$$\hat{c}_\infty(k_x, k_y) = \frac{h_c/D_c}{k_x^2 + k_y^2 - 2ik_x\lambda_{\text{adv}}^{-1} + \lambda_{\text{decay}}^{-2}}. \quad (\text{B3})$$

Transforming back to real space yields

$$c_\infty(x, y) = \frac{h_c}{2\pi D_c} e^{-x\lambda_{\text{adv}}^{-1}} K_0(\kappa \sqrt{x^2 + y^2}), \quad (\text{B4})$$

where K_0 is the modified Bessel function of the second kind and $\kappa^2 = \lambda_{\text{adv}}^{-2} + \lambda_{\text{decay}}^{-2}$. For large distances $\kappa r \gg 1$, with $r = \sqrt{x^2 + y^2}$, the asymptotic form

$$K_0(u) \simeq \sqrt{\frac{\pi}{2u}} e^{-u} \quad (u \rightarrow \infty)$$

gives along the x axis

$$c_\infty(x, 0) \approx \frac{h_c}{D_c \sqrt{8\pi\kappa|x|}} \begin{cases} e^{-(\kappa - \lambda_{\text{adv}}^{-1})|x|}, & \text{if } x < 0, \\ e^{-(\kappa + \lambda_{\text{adv}}^{-1})|x|}, & \text{if } x > 0. \end{cases} \quad (\text{B5})$$

The decay length therefore differs ahead and behind the particle, leading to the fore-aft asymmetry discussed in the main text.

For $\kappa \gg \lambda_{\text{adv}}^{-1}$ or, equivalently, for $\lambda_{\text{decay}} \ll \lambda_{\text{adv}}$, the profile becomes isotropic,

$$c_\infty(r) \sim \frac{h_c}{D_c \sqrt{4\pi\kappa r}} e^{-\kappa r}, \quad (\text{B6})$$

recovering the two-dimensional Yukawa form.

APPENDIX C: ORDER-PERSISTENCE RELATION IN THE WEAK-MEMORY REGIME

In the weak-memory regime, spatial order and effective persistence cannot be tuned independently. Varying the bare persistence l_p and the chemotactic coupling β simultaneously modifies both the effective persistence length l_p^{eff} and the spatial order parameter η_A . Figure 10 shows the resulting trajectories in the $(\eta_A, l_p^{\text{eff}})$ plane at fixed density. Two regimes are observed. For moderate β , spatial order increases while $l_p^{\text{eff}} \approx l_p$ remains nearly unchanged, indicating that repulsive separation dominates. For larger β , the data collapse onto a boundary curve, revealing a trade-off: For a given l_p^{eff} , spatial order is bounded, and conversely, high spatial order cannot be achieved at arbitrarily large persistence. This constraint originates from the finite interaction range λ_{decay} . Long-range velocity correlations would be required to achieve both high persistence and strong ordering simultaneously, but chemotactic alignment is limited when λ_{decay} is shorter than the persistence length or the typical interparticle spacing $1/\sqrt{\varphi}$. As a result, the maximal attainable order decreases with increasing l_p . In the strong-coupling limit $\beta \rightarrow \infty$, for all values of l_p , the system converges to a common point $(\eta_A^\infty, l_p^\infty)$. In this regime, particles follow the steepest descent of the concentration field and move ballistically until redirected by collisions. The typical trajectory length is therefore set by a mean free path $l_{\text{MFP}} = 1/(\sqrt{8}\varphi d)$, where $d \simeq \lambda_{\text{decay}}/3$ is the effective interaction diameter. This confirms that in the strong-coupling regime, collective persistence is collision limited.




FIG. 10. Effective persistence length l_p^{eff} vs spatial order η_A for fixed density $\varphi = 0.012$ of the autochemotactic lattice model. For each bare persistence length l_p , increasing β traces out a distinct trajectory in the $(\eta_A, l_p^{\text{eff}})$ plane that converges to a common strong-coupling state $(\eta_A^\infty, l_p^\infty)$.

APPENDIX D: IDENTIFICATION OF MEMORY REGIMES FROM FIRST-PASSAGE STATISTICS

To identify the crossover between weak- and strong-memory behavior, we fitted the survival probability $S_N(t)$ to the stretched-exponential form $S_N(t) = \exp[-(t/T)^\xi]$. Figure 11(a) shows representative survival probabilities for a single autochemotactic random walker ($N = 1$) together with the corresponding fits. For large α_c , the stretching exponent remains close to $\xi \simeq 1$, corresponding to approximately exponential first-passage statistics. As α_c decreases, ξ increases above unity, indicating compressed-exponential survival. The resulting dependence of ξ on α_c for two different densities φ is shown in Fig. 11(b). We use the crossover from $\xi \approx 1$ to $\xi > 1$ to identify the onset of the strong-memory regime.

APPENDIX E: FINITE-MEMORY SELF-AVOIDANCE: SNAKE MODEL

To disentangle the effect of memory from the specific chemotactic interaction mechanism, we introduce a minimal lattice model in which self-avoidance is imposed explicitly over a finite temporal window. This “snake” model serves as a conceptual reference that isolates the role of memory while preserving the essential trade-off between revisit suppression and self-caging discussed in Sec. IV.

We consider a single walker on a two-dimensional square lattice of size $L \times L$ with periodic boundaries. Each lattice site i carries a clock variable τ_i that records the time elapsed since the site was last visited. At each time step, the walker attempts to jump to a neighboring site j with probability

$$p_j \propto \exp[-\beta \Theta(\tau_j - l_{\text{snake}})], \quad (\text{E1})$$

where l_{snake} defines the memory length and Θ is the Heaviside function. Sites visited within the last l_{snake} steps are therefore disfavored, while older sites are revisited without penalty. After each move, the clock field is updated as $\tau_i \rightarrow \tau_i + 1$ on all sites except the current position, where τ_i is reset to zero.




FIG. 11. (a) Single-particle survival probability $S_1(t)$ in the lattice model for $\varphi = 10^{-4}$ and $D_c = 10^{-4}$ at two values of the decay rate α_c . Lines show stretched-exponential fits $S_1(t) = \exp[-(t/T)^\xi]$. (b) Stretching exponent ξ as a function of α_c for two densities φ at fixed $D_c = 10^{-4}$. The crossover from $\xi \approx 1$ to $\xi > 1$ marks the transition from weak- to strong-memory behavior.

For $l_{\text{snake}} = 0$, the dynamics reduces to an unbiased random walk. In the opposite limit $l_{\text{snake}} \rightarrow \infty$, revisits are permanently forbidden and the trajectory approaches that of a true self-avoiding walk. The parameter l_{snake} thus provides a direct control over the memory depth of the searcher.

Figure 9 shows the mean first-passage time to a randomly placed target as a function of the normalized memory length l_{snake}/L^2 for several system sizes. In all cases, the MFPT exhibits a pronounced minimum at finite memory. When rescaled by the system area, the optimal memory length collapses as

$$l_{\text{snake}}^* \sim L^2, \quad (\text{E2})$$

corresponding to the typical time required to explore the domain once.

For memory lengths shorter than l_{snake}^* , frequent revisits dominate and search efficiency is reduced. For longer memory lengths, self-caging effects emerge as permanently excluded regions fragment the accessible space, again increasing the MFPT. This minimal model therefore reproduces the same qualitative trade-off observed in autochemotactic search: Efficient exploration requires memory that is strong enough to suppress redundant revisits, yet sufficiently transient to avoid topological trapping.

Importantly, the snake model contains neither chemotactic interactions nor collective effects. Its agreement with the autochemotactic results demonstrates that the existence of an




FIG. 12. Example of a band formed by autochemotactic lattice random walkers. The background color shows the chemical concentration field, while white circles mark the walkers. The band spans the system horizontally and propagates downward at approximately constant speed.

optimal, finite memory length is a generic consequence of constrained exploration dynamics rather than a model-specific feature.

APPENDIX F: INEFFICIENT SEARCH BY BANDS

Autochemotactic particles are known to display a rich collective phenomenology beyond homogeneous states [40–44,49,51–53]. In particular, autochemorepulsive interactions can induce effective alignment, which at sufficiently high particle density and for slowly diffusing chemical cues leads to the spontaneous formation of traveling bands [40,44]. Such banded states are also observed in our lattice simulations [see Fig. 12]. A typical band spans the system along one direction and propagates ballistically in the transverse direction, often containing a large fraction of all searchers.

While band formation represents a robust emergent pattern of chemotactic interactions, it is strongly disadvantageous for collective search. Because only walkers located at the leading

edge of the band encounter unexplored territory, particles in the bulk repeatedly rescan regions that have already been visited. As a result, the search efficiency is effectively controlled by the motion of the band front alone.

This observation allows for a simple estimate of the MFPT in the banded state. For a system of linear size L in which a band propagates ballistically at speed v_0 , the MFPT scales as

$$\bar{T}_{\text{band}} \simeq \frac{L}{2v_0}, \quad (\text{F1})$$

independent of the total number of searchers N . Here, dimensional units are used for clarity.

For comparison, N independent persistent searchers homogeneously distributed across the domain achieve an MFPT,

$$\bar{T}_{\text{hom}} \simeq \frac{L^2}{4Nv_0a}, \quad (\text{F2})$$

where a is the particle radius. This expression applies in the ballistic regime $l_p \gg a$ and represents the optimal scaling for independent active Brownian particles. The factor a reflects the effective transverse width $2a$ swept out by each particle.

Comparing the two estimates yields

$$\frac{\bar{T}_{\text{hom}}}{\bar{T}_{\text{band}}} = \frac{L}{2Na}, \quad (\text{F3})$$

which is smaller than unity as soon as the number of searchers exceeds the minimal number required to span the system. Physically, this reflects the fact that in a homogeneous population, all particles contribute to parallel exploration, whereas in a banded state, only the leading edge explores new territory.

For less persistent walkers, the corresponding MFPT estimate for homogeneous search acquires logarithmic corrections of order $\ln L$, but the qualitative conclusion remains unchanged. In that case, one finds $\bar{T}_{\text{hom}}/\bar{T}_{\text{band}} \sim L/(Nl_p)$ for $l_p \ll L$. An analogous scaling argument can be formulated for persistent random walkers.

In chemotactic systems, bands typically contain far more particles than the minimal number required to span the system. Equation (F3) therefore implies that the MFPT in the banded state is substantially larger than in the homogeneous state. We thus conclude that, despite being a common outcome of chemotactic interactions, band formation suppresses the advantages of collective search by concentrating exploration at a single moving front. Banding should therefore be regarded as a failure mode of collective search rather than an optimized strategy.

[1] D. Grebenkov, R. Metzler, and G. Oshanin, *Target Search Problems* (Springer, Cham, 2024).
 [2] N. G. van Kampen, *Stochastic Processes in Physics and Chemistry* (North-Holland, Amsterdam, 1992).
 [3] S. Redner, *A Guide to First-Passage Processes* (Cambridge University Press, Cambridge, 2001).
 [4] P. C. Bressloff, *Stochastic Processes in Cell Biology* (Springer, Heidelberg, 1992).
 [5] M. Smoluchowski, Versuch einer mathematischen theorie der koagulationskinetik kolloider lösungen, *Z. Phys. Chem.* **92**, 129 (1917).

[6] F. C. Collins and G. E. Kimball, Diffusion-controlled reaction rates, *J. Colloid Sci.* **4**, 425 (1949).
 [7] S. Rice, *Diffusion-Limited Reactions* (Elsevier, Amsterdam, 1985).
 [8] O. G. Berg, R. B. Winter, and P. H. Von Hippel, Diffusion-driven mechanisms of protein translocation on nucleic acids. 1. Models and theory, *Biochemistry* **20**, 6929 (1981).
 [9] J. Gorman and E. C. Greene, Visualizing one-dimensional diffusion of proteins along DNA, *Nat. Struct. Mol. Biol.* **15**, 768 (2008).

- [10] H. C. Berg, *E. Coli in Motion* (Springer, Heidelberg, 2005).
- [11] M. F. Krummel, F. Bartumeus, and A. Gérard, T cell migration, search strategies and mechanisms, *Nat. Rev. Immunol.* **16**, 193 (2016).
- [12] M. Hassell and T. Southwood, Foraging strategies of insects, *Annu. Rev. Ecol. Syst.* **9**, 75 (1978).
- [13] J. F. Traniello, Foraging strategies of ants, *Annu. Rev. Entomol.* **34**, 191 (1989).
- [14] A. M. Edwards, R. A. Phillips, N. W. Watkins, M. P. Freeman, E. J. Murphy, V. Afanasyev, S. V. Buldyrev, M. G. da Luz, E. P. Raposo, H. E. Stanley, *et al.*, Revisiting Lévy flight search patterns of wandering albatrosses, bumblebees and deer, *Nature (London)* **449**, 1044 (2007).
- [15] T. T. Hills, C. Kalff, and J. M. Wiener, Adaptive Lévy processes and area-restricted search in human foraging, *PLoS One* **8**, e60488 (2013).
- [16] M. Senanayake, I. Senthoooran, J. C. Barca, H. Chung, J. Kamruzzaman, and M. Murshed, Search and tracking algorithms for swarms of robots: A survey, *Robot. Auton. Syst.* **75**, 422 (2016).
- [17] D. S. Drew, Multi-agent systems for search and rescue applications, *Curr. Robot. Rep.* **2**, 189 (2021).
- [18] M. A. Lomholt, K. Tal, R. Metzler, and K. Joseph, Lévy strategies in intermittent search processes are advantageous, *Proc. Natl. Acad. Sci. USA* **105**, 11055 (2008).
- [19] V. Tejedor, R. Voituriez, and O. Bénichou, Optimizing persistent random searches, *Phys. Rev. Lett.* **108**, 088103 (2012).
- [20] O. Bénichou, C. Loverdo, M. Moreau, and R. Voituriez, Intermittent search strategies, *Rev. Mod. Phys.* **83**, 81 (2011).
- [21] H. Meyer and H. Rieger, Optimal non-Markovian search strategies with n -step memory, *Phys. Rev. Lett.* **127**, 070601 (2021).
- [22] A. Klimek and R. R. Netz, Optimal non-Markovian composite search algorithms for spatially correlated targets, *Europhys. Lett.* **139**, 32003 (2022).
- [23] L. Kusmierz, S. N. Majumdar, S. Sabhapandit, and G. Schehr, First order transition for the optimal search time of Lévy flights with resetting, *Phys. Rev. Lett.* **113**, 220602 (2014).
- [24] A. Chechkin and I. M. Sokolov, Random search with resetting: A unified renewal approach, *Phys. Rev. Lett.* **121**, 050601 (2018).
- [25] P. C. Bressloff, Search processes with stochastic resetting and multiple targets, *Phys. Rev. E* **102**, 022115 (2020).
- [26] A. Stanislavsky and A. Weron, Optimal non-Gaussian search with stochastic resetting, *Phys. Rev. E* **104**, 014125 (2021).
- [27] O. Tal-Friedman, A. Pal, A. Sekhon, S. Reuveni, and Y. Roichman, Experimental realization of diffusion with stochastic resetting, *J. Phys. Chem. Lett.* **11**, 7350 (2020).
- [28] A. Altshuler, O. L. Bonomo, N. Gorohovsky, S. Marchini, E. Rosen, O. Tal-Friedman, S. Reuveni, and Y. Roichman, Environmental memory facilitates search with home returns, *Phys. Rev. Res.* **6**, 023255 (2024).
- [29] S. Redner and P. L. Krapivsky, Capture of the lamb: Diffusing predators seeking a diffusing prey, *Am. J. Phys.* **67**, 1277 (1999).
- [30] G. Oshanin, O. Vasilyev, P. L. Krapivsky, and J. Klafter, Survival of an evasive prey, *Proc. Natl. Acad. Sci. USA* **106**, 13696 (2009).
- [31] C. Mejía-Monasterio, G. Oshanin, and G. Schehr, First passages for a search by a swarm of independent random searchers, *J. Stat. Mech.: Theory Exp.* (2011) P06022.
- [32] A. Nandi and D. Das, Capture of a diffusive prey by multiple predators in confined space, *Phys. Rev. E* **102**, 062109 (2020).
- [33] D. S. Grebenkov, R. Metzler, and G. Oshanin, From single-particle stochastic kinetics to macroscopic reaction rates: Fastest first-passage time of n random walkers, *New J. Phys.* **22**, 103004 (2020).
- [34] D. S. Grebenkov, Depletion of resources by a population of diffusing species, *Phys. Rev. E* **105**, 054402 (2022).
- [35] N. P. Tani, A. Blatt, D. A. Quint, and A. Gopinathan, Optimal cooperative searching using purely repulsive interactions, *J. Theor. Biol.* **361**, 159 (2014).
- [36] S. Ro, J. Yi, and Y. W. Kim, Target searches of interacting Brownian particles in dilute systems, *Phys. Rev. E* **107**, 064143 (2023).
- [37] R. Martínez-García, J. M. Calabrese, T. Mueller, K. A. Olson, and C. López, Optimizing the search for resources by sharing information: Mongolian gazelles as a case study, *Phys. Rev. Lett.* **110**, 248106 (2013).
- [38] J. d'Alessandro, A. Barbier-Chebbah, V. Cellierin, O. Benichou, R. Mège, R. Voituriez, and B. Ladoux, Cell migration guided by long-lived spatial memory, *Nat. Commun.* **12**, 4118 (2021).
- [39] J. Taktikos, V. Zaburdaev, and H. Stark, Modeling a self-propelled autochemotactic walker, *Phys. Rev. E* **84**, 041924 (2011).
- [40] B. Liebchen, D. Marenduzzo, and M. E. Cates, Phoretic interactions generically induce dynamic clusters and wave patterns in active colloids, *Phys. Rev. Lett.* **118**, 268001 (2017).
- [41] W. T. Kranz and R. Golestanian, Trail-mediated self-interaction, *J. Chem. Phys.* **150**, 214111 (2019).
- [42] Z. Mokhtari, R. I. A. Patterson, and F. Höfling, Spontaneous trail formation in populations of auto-chemotactic walkers, *New J. Phys.* **24**, 013012 (2022).
- [43] A. Barbier-Chebbah, O. Bénichou, and R. Voituriez, Self-interacting random walks: Aging, exploration, and first-passage times, *Phys. Rev. X* **12**, 011052 (2022).
- [44] H. Meyer and H. Rieger, Alignment interaction and band formation in assemblies of autochemorepulsive walkers, *Phys. Rev. E* **108**, 034604 (2023).
- [45] X. Fu, S. Kato, J. Long, H. H. Mattingly, C. He, D. C. Vural, S. W. Zucker, and T. Emonet, Spatial self-organization resolves conflicts between individuality and collective migration, *Nat. Commun.* **9**, 2177 (2018).
- [46] B. A. David and P. Kubes, Exploring the complex role of chemokines and chemoattractants *in vivo* on leukocyte dynamics, *Immunol. Rev.* **289**, 9 (2019).
- [47] R. H. Insall, P. Paschke, and L. Tweedy, Steering yourself by the bootstraps: How cells create their own gradients for chemotaxis, *Trends Cell Biol.* **32**, 585 (2022).
- [48] E. Donà, J. D. Barry, G. Valentin, C. Quirin, A. Khmelinskii, A. Kunze, S. Durdu, L. R. Newton, A. Fernandez-Minan, W. Huber, M. Knop, and D. Gilmour, Directional tissue migration through a self-generated chemokine gradient, *Nature (London)* **503**, 285 (2013).
- [49] B. V. Hokmabad, J. Agudo-Canalejo, S. Saha, R. Golestanian, and C. C. Maass, Chemotactic self-caging in active emulsions, *Proc. Natl. Acad. Sci. USA* **119**, e2122269119 (2022).

- [50] B. Nakayama, H. Nagase, H. Takahashi, Y. Saito, S. Hatayama, K. Makino, E. Yamamoto, and T. Saiki, Tunable pheromone interactions among microswimmers, *Proc. Natl. Acad. Sci. USA* **120**, e2213713120 (2023).
- [51] B. Liebchen and H. Löwen, Synthetic chemotaxis and collective behavior in active matter, *Acc. Chem. Res.* **51**, 2982 (2018).
- [52] J. Grauer, H. Löwen, A. Be'er, and B. Liebchen, Swarm hunting and cluster ejections in chemically communicating active mixtures, *Sci. Rep.* **10**, 5594 (2020).
- [53] E. Moreno, R. Großmann, C. Beta, and S. Alonso, From single to collective motion of social amoebae: A computational study of interacting cells, *Front. Phys.* **9**, 750187 (2022).
- [54] A. V. Ivlev, J. Bartnick, M. Heinen, C.-R. Du, V. Nosenko, and H. Löwen, Statistical mechanics where Newton's third law is broken, *Phys. Rev. X* **5**, 011035 (2015).
- [55] Z. You, A. Baskaran, and M. C. Marchetti, Nonreciprocity as a generic route to traveling states, *Proc. Natl. Acad. Sci. USA* **117**, 19767 (2020).
- [56] S. Saha, J. Agudo-Canalejo, and R. Golestanian, Scalar active mixtures: The nonreciprocal Cahn-Hilliard model, *Phys. Rev. X* **10**, 041009 (2020).
- [57] M. Fruchart, R. Hanai, P. B. Littlewood, and V. Vitelli, Non-reciprocal phase transitions, *Nature (London)* **592**, 363 (2021).
- [58] J. P. Banerjee, R. Mandal, D. S. Banerjee, S. Thutupalli, and M. Rao, Unjamming and emergent nonreciprocity in active ploughing through a compressible viscoelastic fluid, *Nat. Commun.* **13**, 4533 (2022).
- [59] K. L. Kreienkamp and S. H. L. Klapp, Nonreciprocal alignment induces asymmetric clustering in active mixtures, *Phys. Rev. Lett.* **133**, 258303 (2024).
- [60] M. Mangeat, S. Chatterjee, J. D. Noh, and H. Rieger, Emergent complex phases in a discrete flocking model with reciprocal and non-reciprocal interactions, *Commun. Phys.* **8**, 186 (2025).
- [61] J. Choi, J. D. Noh, and H. Rieger, Flocking with random nonreciprocal interactions, *Phys. Rev. Res.* **7**, L042059 (2025).
- [62] A. T. Mohite and H. Rieger, Stochastic thermodynamics of nonreciprocally interacting particles and fields, *Phys. Rev. E* **113**, 064136 (2026).
- [63] C.-U. Woo, H. Rieger, and J. D. Noh, Collective behavior in the nonreciprocal multi-species Vicsek model, *arXiv:2512.18756*.
- [64] C.-U. Woo, J. D. Noh, and H. Rieger, Extensive spatio-temporal chaos in non-reciprocal flocking, *arXiv:2604.06462*.
- [65] J. Taktikos, V. Zaburdaev, and H. Stark, Collective dynamics of model microorganisms with chemotactic signaling, *Phys. Rev. E* **85**, 051901 (2012).
- [66] D. J. Amit, G. Parisi, and L. Peliti, Asymptotic behavior of the "true" self-avoiding walk, *Phys. Rev. B* **27**, 1635 (1983).
- [67] P. Grassberger, Self-trapping self-repelling random walks, *Phys. Rev. Lett.* **119**, 140601 (2017).
- [68] S. Ro and Y. W. Kim, Parallel random target searches in a confined space, *Phys. Rev. E* **96**, 012143 (2017).
- [69] S. Ro and Y. W. Kim, Optimal searcher distribution for parallel random target searches, *Phys. Rev. E* **106**, 024101 (2022).
- [70] R. Monchaux, M. Bourgoïn, and A. Cartellier, Preferential concentration of heavy particles: A Voronoi analysis, *Phys. Fluids* **22**, 103304 (2010).
- [71] J.-S. Ferenc and Z. Nédá, On the size distribution of Poisson Voronoi cells, *Physica A* **385**, 518 (2007).
- [72] G. A. Arteca and S. Zhang, Quantitative measure of folding in two-dimensional polymers, *Phys. Rev. E* **59**, 4209 (1999).
- [73] L. Tweedy, D. A. Knecht, G. M. Mackay, and R. H. Insall, Self-generated chemoattractant gradients: Attractant depletion extends the range and robustness of chemotaxis, *PLoS Biol.* **14**, e1002404 (2016).
- [74] L. Tweedy, P. A. Thomason, P. I. Paschke, K. Martin, L. M. Machesky, M. Zagnoni, and R. H. Insall, Seeing around corners: Cells solve mazes and respond at a distance using attractant breakdown, *Science* **369**, eaay9792 (2020).
- [75] D. R. Zusman, A. E. Scott, Z. Yang, and J. R. Kirby, Chemosensory pathways, motility and development in *Myxococcus xanthus*, *Nat. Rev. Microbiol.* **5**, 862 (2007).
- [76] M. R. D'Orsogna, M. A. Suchard, and T. Chou, Interplay of chemotaxis and chemokinesis mechanisms in bacterial dynamics, *Phys. Rev. E* **68**, 021925 (2003).
- [77] H. Meyer and H. Rieger, Optimal number of agents in a collective search and when to launch them, *Phys. Rev. E* **111**, 064112 (2025).
- [78] F. Cichos, K. Gustavsson, B. Mehlig, and G. Volpe, Machine learning for active matter, *Nat. Mach. Intell.* **2**, 94 (2020).
- [79] C. Kaspar, B. J. Ravoo, W. G. van der Wiel, S. V. Wegner, and W. H. P. Pernice, The rise of intelligent matter, *Nature (London)* **594**, 345 (2021).
- [80] L. Piro, B. Mahault, and R. Golestanian, Optimal navigation of microswimmers in complex and noisy environments, *New J. Phys.* **24**, 093037 (2022).
- [81] M. Nasiri, H. Löwen, and B. Liebchen, Optimal active particle navigation meets machine learning, *Europhys. Lett.* **142**, 17001 (2023).