


Hunting active Brownian particles: Learning optimal behaviorMarcel Gerhard, Ashreya Jayaram, Andreas Fischer, and Thomas Speck *Institut für Physik, Johannes Gutenberg-Universität Mainz, Staudingerweg 7-9, 55128 Mainz, Germany*

(Received 9 July 2021; accepted 16 November 2021; published 30 November 2021)

We numerically study active Brownian particles that can respond to environmental cues through a small set of actions (switching their motility and turning left or right with respect to some direction) which are motivated by recent experiments with colloidal self-propelled Janus particles. We employ reinforcement learning to find optimal mappings between the state of particles and these actions. Specifically, we first consider a predator-prey situation in which prey particles try to avoid a predator. Using as reward the squared distance from the predator, we discuss the merits of three state-action sets and show that turning away from the predator is the most successful strategy. We then remove the predator and employ as collective reward the local concentration of signaling molecules exuded by all particles and show that aligning with the concentration gradient leads to chemotactic collapse into a single cluster. Our results illustrate a promising route to obtain local interaction rules and design collective states in active matter.

DOI: [10.1103/PhysRevE.104.054614](https://doi.org/10.1103/PhysRevE.104.054614)**I. INTRODUCTION**

The defining feature of microscopic active particles is their persistent motion: in contrast to passive diffusion they carry an orientation along which displacements are more likely. Such persistent motion is exhibited by bacteria propelled by flagella [1] and synthetic colloidal particles that exploit phoretic mechanisms to propel themselves through self-generated gradients [2]. The interplay of this self-propulsion with excluded volume and particle shape leads to a range of collective dynamic states [3] such as dense clusters, swarms, and flocks [4]. Given the wealth of possible collective states emerging already from minimal models, data-driven modeling has attracted interest recently [5,6].

What sets living active matter apart from synthetic active matter is the *intrinsic* ability to respond to external stimuli beyond the direct physical forces generated by gradients. Bacteria and other microorganisms can sense their environment and, *inter alia*, adapt their motility [7]. For example, some bacteria organize into biofilms [8] through quorum sensing [9]; i.e., they respond to the concentration of certain signaling molecules exuded by all members of the population. Transferring similar capabilities to synthetic active matter opens the route to program novel collective behavior with the outlook to perform useful tasks [10–12]. Ultimately, this requires internal degrees of freedom computing a response [13,14]. As an intermediate step to elucidate the basic principles, this computation could be performed by an external agent which then acts back on the system.

Recent advances of feedback techniques have demonstrated exquisite control over self-propelled colloidal particles allowing us to implement interaction rules that go beyond steric volume exclusion and phoretic and hydrodynamic coupling. Janus particles propelled through the local phase separation of a binary solvent can be addressed individually to implement motility switching [15–17] and perception-based interactions [18]. Active dimers can adapt the propulsion

speed depending on their orientation [19]. Local heating of a colloidal particle can be exploited for steering [20,21], pattern formation [22], motility control [23], and recently has been combined with reinforcement learning to guide a single particle toward a target side [24]. Reinforcement learning together with computer simulations has been employed to navigate flows [25–27], to control shape deformations of microswimmers [28,29], to induce flocking of active particles [30,31], and to steer a single active particle through an external potential [32,33].

In this work, we further explore reinforcement learning [34] to determine optimal single-particle actions in response to “states” representing the information that is available to the particle. The policy which maximizes the reward is identified as the optimal behavior. We consider a small set of actions that have been implemented experimentally for colloidal Janus particles: turning motility on or off [16], and exerting a torque that makes the particle turn left or right [35]. The particle dynamics is that of active Brownian particles with propulsion speed and torque depending on the chosen action. As a concrete illustration, we study a simple predator-prey system [36,37] with an absorbing boundary (the “predator”), which induces a particle current. Avoiding a predator is complementary to optimal search (and foraging) strategies, which have received intensive theoretical scrutiny [38–42]. Using as reward the distance to the predator, we evaluate three state-action sets leading to different currents. Finally, we apply reinforcement learning to an interacting suspension of active Brownian particles, demonstrating that it leads to chemotactic collapse for sufficiently large torques overcoming rotational diffusion [43,44].

II. PREDATOR-PREY SYSTEM

We first consider an ideal gas of N noninteracting particles moving in two dimensions in the presence of a single “predator.” To conserve density, whenever a “prey” particle comes

within distance a of the predator it is removed and placed at a random position within the system. This corresponds to an absorbing circular boundary at $r = a$ inducing a total current J leaving the system. The accessible area is $A = L^2 - \pi a^2$ excluding the disk around the predator with global density $\bar{\rho} = N/A$. In the simulations, we employ a square box with edge length L employing periodic boundaries, whereas in the analytical calculations we place the predator at the origin. All numerical results are reported employing the length parameter σ as unit of length and σ^2/D_0 as unit of time with bare diffusion coefficient D_0 . For noninteracting particles, σ still determines the effective particle size. Throughout, the radius of the absorbing circular boundary is $a = 1.66\sigma$.

A. Passive diffusion

As reference, we first consider passive diffusion of prey particles with diffusion coefficient D_0 . The current is $\mathbf{j} = -D\nabla\rho$, where ρ is the number density of prey particles. Here we have shifted coordinates so that in the following the predator is at the origin and its diffusive motion enters through the (dimensionless) effective diffusion coefficient $D = 1 + D_{\text{pred}}$. In steady state, $\mathbf{j} = j(r)\mathbf{e}_r$ with r the distance from the predator and \mathbf{e}_r the radial unit vector. The total current leaving the system is $J = 2\pi a j(a) < 0$. The diffusion equation becomes $D\nabla^2\rho + s = 0$ with $s = -J/A > 0$ the uniform rate at which removed particles respawn. The solution reads

$$\rho(r) = b \ln(r/a) - \frac{s}{4D}(r^2 - a^2) \quad (1)$$

with boundary condition $\rho(a) = 0$ and integration constant $b \propto \bar{\rho}$ determined by the conservation of the total number of prey particles. The radial current is $j(r) = -D\rho'(r) = -Db/r + sr/2$, whereby the prime denotes the derivative with respect to r . The total current for passive diffusion thus becomes

$$J_0 = -2\pi Db \left(1 + \frac{\pi a^2}{A}\right)^{-1}. \quad (2)$$

Since $s \propto D$ is proportional to the current and thus to the diffusion coefficient, the density profile Eq. (1) is independent of D . In Fig. 1(a), we show the numerically obtained density profile $\rho(r)$ in a square system with periodic boundaries together with Eq. (1). To this end, the equations of motion are integrated with time step $\delta = 1.5 \times 10^{-5}$ employing the Euler-Maruyama scheme. The negative current $-J_0/D > 0$ is plotted in Fig. 1(b) and shows the expected linear increase with the global density.

B. Free active diffusion

In the next step, we consider prey particles undergoing directed motion with propulsion speed v_0 ,

$$\dot{\mathbf{r}}_k = v_0 \mathbf{e}_k + \sqrt{2D_0} \boldsymbol{\xi}_k, \quad \dot{\varphi}_k = \sqrt{2/\tau_r} \eta_k, \quad (3)$$

where the $\boldsymbol{\xi}_k, \eta_k$ are zero mean and unit-variance Gaussian white noise. Each particle has an orientation $\mathbf{e}_k = (\cos \varphi_k, \sin \varphi_k)^T$ described by the angle φ_k with the x axis, which undergoes free rotational diffusion with correlation time τ_r . Again, the effective diffusion coefficient in the frame of reference of the predator is D . The orientational correlation

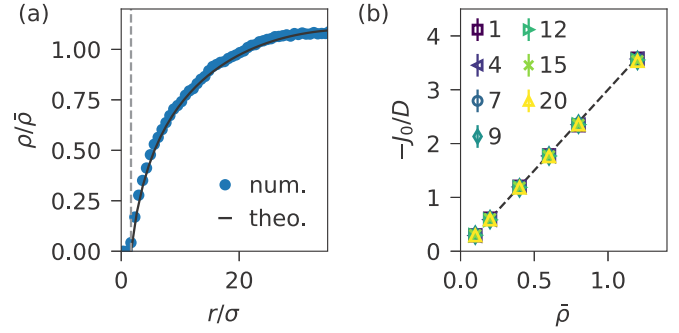


FIG. 1. Passive diffusion ($L = 70$). (a) Density profile $\rho(r)$ away from the predator with absorbing boundary at $r = a$ (dashed line) for effective diffusion coefficient $D = 1$ in a square system. (b) Current $-J_0$ (corresponding to the prey caught per time) as a function of global density $\bar{\rho}$ for several diffusion coefficients D (see legend for values). The dashed line is the prediction Eq. (2) for $-J_0/D$.

time is related through the no-slip boundary condition to the translational diffusion, $1/\tau_r = 3D_0/\sigma_{\text{eff}}^2$, with effective particle diameter $\sigma_{\text{eff}} \simeq 1.10688\sigma$.

In Fig. 2(a) we show numerical density profiles for different speeds. Clearly, increasing the speed leads to a flatter density profile that sharply declines as the absorbing boundary at $r = a$ is approached. The effect of self-propulsion can thus not be captured by an elevated diffusion coefficient alone (as for the mean-square displacement of a free particle [45]). Indeed, Fig. 2(b) shows that the density profile for active particles depends on the value of D in contrast to the passive case.

C. Learning optimal behavior

We now assume that each prey particle has limited computational capabilities that allow it to determine a state $s_k \in \mathcal{S}$ and to perform an action $a_k \in \mathcal{A}$. Both the possible states \mathcal{S} and the actions \mathcal{A} are discrete sets of a few possibilities, which are related through a Q matrix with entries Q_{sa} . Instead of integrating Eq. (3), particles now evolve according to the following scheme with time step δt :

- (1) determine the state s_k of each particle;
- (2) determine the action a_k that maximizes $Q_{s_k a_k}$;




FIG. 2. Active diffusion ($L = 70$). Density profiles $\rho(r)$ away from the predator for (a) different speeds v_0 at $D = 1$ [the black line shows the passive result Eq. (1)] and (b) different D at speed $v_0 = 10$.

(3) translate all particles positions

$$\mathbf{r}_k(t + \delta t) = \mathbf{r}_k(t) + v(a_k)\delta t \mathbf{e}_k + \sqrt{2D_0\delta t} \boldsymbol{\xi}_k \quad (4)$$

and orientations

$$\varphi_k(t + \delta t) = \varphi_k(t) + \omega(a_k)\delta t + \sqrt{2\delta t/\tau_r}\eta_k, \quad (5)$$

which is repeated. The action a thus determines the propulsion speed $v(a)$ and the torque $\omega(a)$.

To proceed, we need to determine the Q matrix relating state to action. We break the learning into several episodes, and each episode is divided into multiple steps. Each episode represents a simulation, where at the beginning all prey particles are initialized randomly. The predator remains located in the center of the simulation box throughout the learning process. After 10 000 time steps δt , we determine the states s_k and each prey particle receives a reward R_k . Then the prey particles choose new actions, which are applied for the next learning step. The action policy is based upon the current Q matrix according to an ε -greedy exploration scheme,

$$a_k = \begin{cases} \operatorname{argmax} Q_{s_k a}, & \text{with probability } 1 - \varepsilon_n, \\ \text{random action}, & \text{with probability } \varepsilon_n. \end{cases} \quad (6)$$

At the beginning of the learning process, we start with $\varepsilon_0 = 1$ and decrease its value according to $\varepsilon_n = 0.995^n$ with rising experience of the prey, where n enumerates the learning episodes.

The Q matrix is initialized with all entries set to 1. The entry for each state-action pair (s, a) is then updated after each step by the rule

$$Q_{s_k a_k} \leftarrow Q_{s_k a_k} + \alpha_n [R_k + \gamma \max_a Q_{s'_k a} - Q_{s_k a_k}], \quad (7)$$

where s'_k is the new state after advancing the simulation with a_k . The future reward discount factor is set to $\gamma = 0.99$. The learning rate $\alpha_n = 0.8/(0.8 + n)$ depends on the episode n and decreases.

As a first example, we assume that prey particles can somehow estimate their distance to the predator (e.g., through sensing a chemical signal exuded by the predator [36]). The discrete state $s_k = \lfloor (\mathbf{r}_k - \mathbf{r}_p)/\lambda \rfloor$ then measures the distance to the position \mathbf{r}_p of the predator with spacing λ and floor function $\lfloor \cdot \rfloor$. The reward is calculated as

$$R_k = |\mathbf{r}_k - \mathbf{r}_p|^2. \quad (8)$$

The learning process is performed in a square box with edge $L = 50$ and $N = 150$ prey particles over 1000 episodes, each with 20 learning steps. The actions are restricted to switching the motility on/off,

$$v(a) = \begin{cases} 0 & (a = 1), \\ v_0 & (a = 2), \end{cases} \quad (9)$$

where $v_0 = 20$. During the learning process, we measure the success through the average reward at the end of each episode. The reward progress is shown in Fig. 3(a). At the beginning of the learning process, the particles are uniformly distributed over the entire box. In this case, the mean-square distance between prey and predator located in the center is

$$\langle R_k \rangle_{\text{hom}} = \frac{1}{A} \int_{-L/2}^{L/2} dx \int_{-L/2}^{L/2} dy (x^2 + y^2) \approx \frac{L^2}{6}, \quad (10)$$




FIG. 3. Reinforcement learning. (a) Performance of the $N = 150$ prey particles as learning progresses. Plotted is the average reward per particle $\sum_k R_k/N$ at the end of each episode as a function of episodes. Speed of active particles is $v_0 = 20$. The gray dashed line indicates the average reward $\langle R_k \rangle_{\text{hom}} \approx 416$ for a homogeneous system. (b) Representation of the Q matrix, whereby each concentric ring corresponds to one discrete state. Particles in the red rings move actively ($v = v_0$), while particles in the outer green rings diffuse passively ($v = 0$).

which corresponds to a reward of about 416. It reaches a plateau after about 500 episodes at an average value of about 500, which corresponds to an average distance of about 22 to the predator. The resulting policy for five discrete states is shown in Fig. 3(b), where particles in the red area move actively with speed v_0 while particles in the area highlighted in green only undergo diffusive motion. The radius R_* of the active region is a nontrivial function of system size L , discretization, and speed v_0 . Due to the periodic boundaries, to maximize the reward the passive region needs to be sufficiently large otherwise prey particles return to the active region too fast.

After learning the Q matrix at one speed v_0 and diffusion coefficient D , we perform further simulations with the final Q matrix for different v_0 and D . These simulations integrate the same Eqs. (4) and (5) of motion but the action is determined deterministically through the rule $a_k = \operatorname{argmax} Q_{s_k a}$. In Fig. 4(a), we show numerical density profiles $\rho(r)$ at speed $v_0 = 5$ and global density $\bar{\rho}_0 = 0.2$ for different diffusion coefficients D . For low diffusion coefficients, we see a density discontinuity at the threshold between active and passive particles since the passive particles accumulate in the outer regions of the box, but there is always a certain amount of motile particles in the active region. The higher the diffusion coefficient, and therefore also the rotational diffusion, the narrower the gap becomes, which disappears above about $D = 7$ and the system is dominated by diffusion. In Fig. 4(b), we show for two densities that the current $-J$ through the absorbing boundary increases with v_0 and is always larger than the passive current. On first glance this seems counter-intuitive since the prey particles accumulate away from the predator, but the self-propulsion also leads to an increased probability to encounter the predator. Tentatively replacing the passive diffusion coefficient D in Eq. (2) by an elevated active diffusion leads to the expression

$$-J = -J_0 + \frac{v_0^2}{2D} \frac{1}{c_1 + c_2(v_0/D)^2} \bar{\rho}, \quad (11)$$




FIG. 4. Distance-based strategy. (a) Numerical density profiles $\rho(r)$ away from the predator for $\bar{\rho} = 0.2$ and different diffusion coefficients D employing the Q matrix with speed $v_0 = 5$ (different from the learning speed). The dashed line indicates the transition radius $R_* \approx 21.2$ between passive and active regions; cf. Fig. 3(b). (b) Reduced current $-J/(\bar{\rho}D)$ through the absorbing boundary for two global densities and several diffusion coefficients and speeds. The full line shows Eq. (11).

where we assume that for large speeds the current saturates (trajectories through the active region become basically straight lines with a fixed probability to hit the predator). Figure 4(b) demonstrates that this expression describes the measured current very well with fit parameters $c_1 \simeq 1.75$ and $c_2 \simeq 0.042$. Hence, even though the prey successfully accumulate away from the predator, the goal of not getting caught (on average) is not achieved.

III. NAVIGATING CHEMICAL GRADIENTS

A. External gradient

How can prey particles improve their chances? As already mentioned, one means of communication at the microscale is the release and sensing of signaling molecules, which diffuse quickly and create a gradient. Specifically, let us assume that the predator exudes these signaling molecules with rate γ_c and diffusion coefficient D_c . In principle, predator and prey move much slower than molecules diffuse. In this quasistatic limit, the prey particles effectively move in a concentration field

$$c(\mathbf{r}) = \frac{\gamma_c}{4\pi D_c} \frac{e^{-|\mathbf{r}-\mathbf{r}_p|/\lambda}}{|\mathbf{r}-\mathbf{r}_p|} \quad (12)$$

that is parametrized by the position \mathbf{r}_p of the predator alone. Here we assume that the predator and prey move diffusively in two dimensions close to a substrate, and the concentration profile is that within the semispace above the substrate. A decay of signaling molecules leads to the exponential factor with decay length λ . In the following, we set $\gamma_c/D_c = 1$ and $\lambda = 10$. Importantly, through the concentration gradient prey particles can now respond to the orientation of the predator rather than just distance.

We consider two further state-action sets. The first case is that of motility switching with Eq. (9) but now depending on whether the predator is in front or behind them; see Fig. 5(a). Formally, we distinguish these two states through

$$\mathbf{e}_k \cdot \nabla c|_{\mathbf{r}_k} = \begin{cases} \leq 0, & \text{oriented away from,} \\ > 0, & \text{oriented toward,} \end{cases} \quad (13)$$




FIG. 5. Gradient-based strategies. (a) Motility switching based on the relative orientation with respect to the concentration gradient ∇c (pointing toward the predator). The Q matrix is sketched: to maximize the distance, prey particles become motile when facing away and passive when facing toward the predator. (b) Orientation adaptation. Prey particles are always active and turn away from the predator.

the predator. We perform the learning process over 1000 episodes, 20 steps each, and with learning parameters α and γ as described in Sec. II C using as reward again the squared distance Eq. (8). The outcome policy is that particles which are facing away from the predator should move actively in order to increase the distance to the predator. Particles which are oriented toward the predator should remain passive until either the motion of the predator or rotational diffusive motion leads to a change in state.

The second case is an orientation-adaption model, where particles act by either turning themselves to the right or to the left according to

$$\omega(a) = \begin{cases} -\omega_0 & (a = 1), \\ +\omega_0 & (a = 2), \end{cases} \quad (14)$$

with torque (angular speed) $\omega_0 = 13.3$. The speed v_0 in this model is constant. We consider again two possible states, which are sketched in Fig. 5(b). The first one is that the predator is on a particle's left side; the other one is that the predator is on its right side. We can express this with the two-dimensional cross product of the orientation vector and the gradient,

$$\mathbf{e}_k \times \nabla c|_{\mathbf{r}_k} = \begin{cases} \leq 0, & \text{predator right,} \\ > 0, & \text{predator left.} \end{cases} \quad (15)$$

After the learning process, the particles follow the policy derived from the Q matrix shown in Fig. 5(b), which indicates that the prey should always turn away from the predator.

As before, we test the quality of the derived policies by measuring the success of a predator in catching prey particles that follows those policies [46]. In contrast to the previous simulations, now the predator not only passively catches particles that come below a certain distance threshold; it also follows a fixed chasing strategy: it always focuses on chasing the nearest prey particle with constant speed v_p . The prey particles are still randomly set back into the box to maintain a constant density. We use the resulting current induced by the predator to evaluate both policies. We perform simulations with different speeds v_0 of the prey particles while the

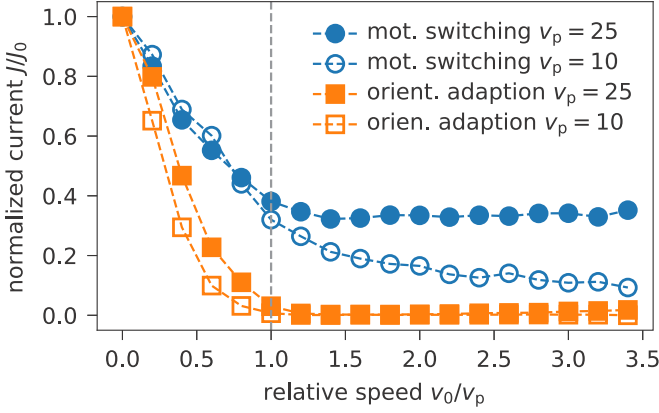


FIG. 6. Normalized current of both gradient-based policies as a function of the relative speed v_0/v_p between predator and prey particles for global density $\bar{\rho} = 0.17$. Shown are results for two predator speeds $v_p = 10$ (open symbols) and $v_p = 25$ (filled symbols). For prey that outruns the predator ($v_0 > v_p$) the current becomes independent of the relative speed and drops to zero for the orientation adaption.

predator moves with speed v_p and rotates with torque $\omega_p = \omega_0 = 13.3$ toward the nearest prey particle.

Figure 6 shows the normalized current J/J_0 of the two policies depending on the relative speed v_0/v_p at global density $\bar{\rho} = 0.17$. The current is normalized by the current J_0 in a system with only passive prey particles. We observe for both models that the current decreases up to a relative speed of about $v_0 \approx v_p$; i.e., prey particles are successful in avoiding the predator. Above that point, the current reaches a nonzero plateau in the motility-switch model for fast predators. In the plateau region, the predator catches only particles that are not actively moving. In the orientation-adaption model the current falls to zero; i.e., all prey particles manage to escape as long as they move faster. The predator is able to catch particles only shortly after initialization, when the prey have to reorient away from the predator.

B. Self-generated gradients

So far, we have considered independent prey particles that react to an external stimulus, here the predator. We now remove the predator and assume that the particles both exude and sense signaling molecules as in quorum sensing [9]. We aim to find a policy so that particles aggregate into clusters at very low densities. Particles now have an excluded volume that we model through the repulsive Weeks-Chandler-Anderson (WCA) pair potential

$$u(r) = \begin{cases} 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^6 \right] + \epsilon, & r/\sigma < 2^{1/6}, \\ 0, & r/\sigma \geq 2^{1/6}, \end{cases} \quad (16)$$

with distance $r = |\mathbf{r}_i - \mathbf{r}_j|$ between two particles at positions \mathbf{r}_i and \mathbf{r}_j . We employ a potential strength $\epsilon = 100k_B T$, which implies hard-disk-like particles with an effective diameter $\sigma_{\text{eff}} = 1.10688\sigma$ [47].

In this model, all particles produce and sense signaling molecules. In regions of high density there is also a high concentration of signaling molecules. We use this concentration




FIG. 7. Concentration field $c(\mathbf{r})$ of signaling molecules produced by the blue particles (red: high c ; blue: low c). The white particle is sensing the concentration gradient with two possible states: the higher density is (a) to the left or (b) to the right.

as a proxy for how far a particle is away from regions with higher density of particles. We define the reward for each particle as the local concentration

$$R_k = c(\mathbf{r}_k) = \frac{\gamma_c}{4\pi D_c} \sum_{i \neq k} \frac{e^{-|\mathbf{r}_k - \mathbf{r}_i|/\lambda}}{|\mathbf{r}_k - \mathbf{r}_i|} \quad (17)$$

that it senses. Again, we define two different states similar to the orientation-adaption model, but instead of responding to an external source of signaling molecules, the particles respond to each other. The two possible states are demonstrated in Fig. 7. We define the states through the sign of the cross product between the orientation of a particle and the local gradient

$$\mathbf{e}_k \times \nabla c|_{\mathbf{r}_k} = \begin{cases} \leq 0, & \text{higher density right,} \\ > 0, & \text{higher density left.} \end{cases} \quad (18)$$

Depending on the state, the particle can choose to turn left or right; see Eq. (14).

The learning process is performed with $N = 50$ particles in a square system with $L = 40$. All parameters of the reinforcement learning algorithm are the same as in the previous examples. We set the active speed to $v_0 = 20$ and learn over 1000 episodes with 20 steps each. The Q matrix results in a policy in which particles align with the gradient and thus orient toward higher concentrations. We then employ the resulting Q matrix to investigate the clustering process depending on the strength of the reaction of the particles to their self-generated chemical field. While similar to Ref. [44], our simulations differ in the following points: First, we consider a torque that is not depending on the actual value of the gradient but only on the sign of the gradient. Moreover, we do not consider a translational diffusiophoretic motion due to the concentration gradient.

We investigate dilute systems with packing fraction $\phi = N\pi(\sigma_{\text{eff}}/2)^2/(L)^2 \simeq 0.05$ and simulate a total of $N = 255$ particles. We choose an active speed $v_0 = 60$ corresponding to a Péclet number $\text{Pe} = \hat{v}_0/\sqrt{2D_0/\tau_r} \simeq 27$ (with dimensionful speed \hat{v}_0). Figure 8 shows that for small reorientation torque ω_0 the system is an active homogeneous gas while for large ω_0 all particles collapse into one single large cluster. We consider a cluster as an assembly of $N_c \geq 2$ particles. A cluster is determined by all particles that are mutually “bonded” (i.e., they are within the cutoff radius of the interaction potential).




FIG. 8. Average fraction N_c/N of particles in a cluster as a function of reorienting torque ω_0 for a system that was initialized either randomly in a homogeneous state or in a single cluster.

In the intermediate regime, we study the system from two initial configurations: either we start the simulation with all particle positions initialized randomly or we start with a single cluster. Figure 8 shows that there is considerable hysteresis for our small system and that the formation of the cluster from the homogeneous state occurs through nucleation with small clusters decaying. In contrast, once a large cluster has formed it is stable down to small values of ω_0 .

IV. CONCLUSIONS

To summarize, we have extended the model of active Brownian particles to speeds and torques that depend on some discrete action $a_k \in A$; see Eqs. (4) and (5). This action is

chosen through a Q matrix, the maximal entry of which for a given state $s_k \in S$ determines a_k . This Q matrix can be determined through well-established algorithms known as reinforcement learning, which require as further input a reward function that evaluates the “utility” of the current state to the system. We have verified this approach for two toy models. First, we have studied prey particles reacting to a predator, where the reward is given by the distance of the prey from the predator. Particles do not interact directly but only through the predator, aggregating into (in a periodic system) domains away from the predator. Still, we found strong differences between state-action pairs with respect to the success avoiding the predator (measured as a current through an absorbing boundary). These different state-action pairs represent the information that might be available and possible actions. Here we have focused on the local concentration of some signaling molecules, but other cues like light [48], gravity [26,49], viscosity [50,51], etc., leading to the different x taxis might be used. Alternatively, run-and-tumble bacteria like *E. coli* might adopt their tumbling rate. It will also be interesting to consider the influence of unavoidable concentration fluctuations of the signaling molecules on optimal search strategies [42]. Second, we have considered as reward the local concentration gradient. To achieve aggregation, the reorientation torque needs to overcome the rotational diffusion. Our framework can easily be extended to learn the interactions underlying more complex collective behavior.

ACKNOWLEDGMENTS

We acknowledge financial support through the Emergent AI Center funded by the Carl-Zeiss-Stiftung. We thank Michael Wand for useful discussions.

- [1] J. Elgeti, R. G. Winkler, and G. Gompper, Physics of microswimmers—single particle motion and collective behavior: A review, *Rep. Prog. Phys.* **78**, 056601 (2015).
- [2] C. Bechinger, R. Di Leonardo, H. Löwen, C. Reichhardt, G. Volpe, and G. Volpe, Active particles in complex and crowded environments, *Rev. Mod. Phys.* **88**, 045006 (2016).
- [3] G. Gompper, R. G. Winkler, T. Speck, A. Solon, C. Nardini, F. Peruani, H. Loewen, R. Golestanian, U. B. Kaupp, L. Alvarez, T. Kioerboe, E. Lauga, W. Poon, A. D. Simone, F. Cichos, A. Fischer, S. M. Landin, N. Soeker, R. Kapral, P. Gaspard *et al.*, The 2020 motile active matter roadmap, *J. Phys.: Condens. Matter* **32**, 193001 (2020).
- [4] T. Vicsek and A. Zafeiris, Collective motion, *Phys. Rep.* **517**, 71 (2012).
- [5] F. Cichos, K. Gustavsson, B. Mehlig, and G. Volpe, Machine learning for active matter, *Nat. Mach. Intell.* **2**, 94 (2020).
- [6] A. R. Dulaney and J. F. Brady, Machine learning for phase behavior in active matter systems, *Soft Matter* **17**, 6808 (2021).
- [7] A. Wilde and C. W. Mullineaux, Light-controlled motility in prokaryotes and the problem of directional light perception, *FEMS Microbiol. Rev.* **41**, 900 (2017).
- [8] M. R. Parsek and E. Greenberg, Sociomicrobiology: The connections between quorum sensing and biofilms, *Trends Microbiol.* **13**, 27 (2005).
- [9] S. P. Brown and R. A. Johnstone, Cooperation in the dark: Signalling and collective action in quorum-sensing bacteria, *Proc. Biol. Sci.* **268**, 961 (2001).
- [10] N. Koumakis, A. Lepore, C. Maggi, and R. D. Leonardo, Targeted delivery of colloids by swimming bacteria, *Nat. Commun.* **4**, 2588 (2013).
- [11] M. Rubenstein, A. Cornejo, and R. Nagpal, Programmable self-assembly in a thousand-robot swarm, *Science* **345**, 795 (2014).
- [12] M. Sitti, H. Ceylan, W. Hu, J. Giltinan, M. Turan, S. Yim, and E. Diller, Biomedical applications of untethered mobile milli/microrobots, *Proc. IEEE* **103**, 205 (2015).
- [13] C. Horsman, S. Stepney, R. C. Wagner, and V. Kendon, When does a physical system compute? *Proc. R. Soc. A* **470**, 20140182 (2014).
- [14] L. G. Nava, R. Großmann, and F. Peruani, Markovian robots: Minimal navigation strategies for active particles, *Phys. Rev. E* **97**, 042604 (2018).
- [15] D. F. B. Haeufle, T. Bäuerle, J. Steiner, L. Bremicker, S. Schmitt, and C. Bechinger, External control strategies for self-propelled particles: Optimizing navigational efficiency in the presence of limited resources, *Phys. Rev. E* **94**, 012617 (2016).
- [16] T. Bäuerle, A. Fischer, T. Speck, and C. Bechinger, Self-organization of active particles by quorum sensing rules, *Nat. Commun.* **9**, 3232 (2018).

- [17] A. Fischer, F. Schmid, and T. Speck, Quorum-sensing active particles with discontinuous motility, *Phys. Rev. E* **101**, 012601 (2020).
- [18] F. A. Lavergne, H. Wendehenne, T. Bäuerle, and C. Bechinger, Group formation and cohesion of active particles with visual perception-dependent motility, *Science* **364**, 70 (2019).
- [19] A. R. Sprenger, M. A. Fernandez-Rodriguez, L. Alvarez, L. Isa, R. Wittkowski, and H. Löwen, Active Brownian motion with orientation-dependent motility: Theory and experiments, *Langmuir* **36**, 7066 (2020).
- [20] B. Qian, D. Montiel, A. Bregulla, F. Cichos, and H. Yang, Harnessing thermal fluctuations for purposeful activities: The manipulation of single micro-swimmers by adaptive photon nudging, *Chem. Sci.* **4**, 1420 (2013).
- [21] A. P. Bregulla, H. Yang, and F. Cichos, Stochastic localization of microswimmers by photon nudging, *ACS Nano* **8**, 6542 (2014).
- [22] U. Khadka, V. Holubec, H. Yang, and F. Cichos, Active particles bound by information flows, *Nat. Commun.* **9**, 3864 (2018).
- [23] N. A. Söker, S. Auschra, V. Holubec, K. Kroy, and F. Cichos, How Activity Landscapes Polarize Microswimmers without Alignment Forces, *Phys. Rev. Lett.* **126**, 228001 (2021).
- [24] S. Muñios-Landin, A. Fischer, V. Holubec, and F. Cichos, Reinforcement learning with artificial microswimmers, *Sci. Robot.* **6**, eabd9285 (2021).
- [25] M. Gazzola, A. Tchieu, D. Alexeev, A. de Brauer, and P. Koumoutsakos, Learning to school in the presence of hydrodynamic interactions, *J. Fluid Mech.* **789**, 726 (2016).
- [26] S. Colabrese, K. Gustavsson, A. Celani, and L. Biferale, Flow Navigation by Smart Microswimmers via Reinforcement Learning, *Phys. Rev. Lett.* **118**, 158004 (2017).
- [27] S. Verma, G. Novati, and P. Koumoutsakos, Efficient collective swimming by harnessing vortices through deep reinforcement learning, *Proc. Natl. Acad. Sci. U.S.A.* **115**, 5849 (2018).
- [28] A. C. H. Tsang, P. W. Tong, S. Nallan, and O. S. Pak, Self-learning how to swim at low Reynolds number, *Phys. Rev. Fluids* **5**, 074101 (2020).
- [29] B. Hartl, M. Hübl, G. Kahl, and A. Zöttl, Microswimmers learning chemotaxis with genetic algorithms, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2019683118 (2021).
- [30] M. Durve, F. Peruani, and A. Celani, Learning to flock through reinforcement, *Phys. Rev. E* **102**, 012601 (2020).
- [31] M. J. Falk, V. Alizadehyazdi, H. Jaeger, and A. Murugan, Learning to control active matter, *Phys. Rev. Res.* **3**, 033291 (2021).
- [32] E. Schneider and H. Stark, Optimal steering of a smart active particle, *Europhys. Lett.* **127**, 64003 (2019).
- [33] B. Liebchen and H. Löwen, Optimal navigation strategies for active particles, *Europhys. Lett.* **127**, 34003 (2019).
- [34] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA, 2018).
- [35] C. Lozano, B. ten Hagen, H. Löwen, and C. Bechinger, Phototaxis of synthetic microswimmers in optical landscapes, *Nat. Commun.* **7**, 12828 (2016).
- [36] A. Sengupta, T. Kruppa, and H. Löwen, Chemotactic predator-prey dynamics, *Phys. Rev. E* **83**, 031914 (2011).
- [37] M. Schwarzl, A. Godec, G. Oshanin, and R. Metzler, A single predator charging a herd of prey: Effects of self volume and predator-prey decision-making, *J. Phys. A* **49**, 225601 (2016).
- [38] O. Bénichou, M. Coppey, M. Moreau, P.-H. Suet, and R. Voituriez, Optimal Search Strategies for Hidden Targets, *Phys. Rev. Lett.* **94**, 198101 (2005).
- [39] O. Bénichou, C. Loverdo, M. Moreau, and R. Voituriez, Intermittent search strategies, *Rev. Mod. Phys.* **83**, 81 (2011).
- [40] F. Bartumeus, E. P. Raposo, G. M. Viswanathan, and M. G. E. da Luz, Stochastic optimal foraging: Tuning intensive and extensive dynamics in random searches, *PLoS ONE* **9**, e106373 (2014).
- [41] K. Schwarz, Y. Schröder, and H. Rieger, Numerical analysis of homogeneous and inhomogeneous intermittent search strategies, *Phys. Rev. E* **94**, 042133 (2016).
- [42] J. A. Kromer, N. de la Cruz, and B. M. Friedrich, Chemokinetic Scattering, Trapping, and Avoidance of Active Brownian Particles, *Phys. Rev. Lett.* **124**, 118101 (2020).
- [43] I. Theurkauff, C. Cottin-Bizonne, J. Palacci, C. Ybert, and L. Bocquet, Dynamic Clustering in Active Colloidal Suspensions with Chemical Signaling, *Phys. Rev. Lett.* **108**, 268303 (2012).
- [44] O. Pohl and H. Stark, Dynamic Clustering and Chemotactic Collapse of Self-Phoretic Active Particles, *Phys. Rev. Lett.* **112**, 238303 (2014).
- [45] J. R. Howse, R. A. L. Jones, A. J. Ryan, T. Gough, R. Vafabakhsh, and R. Golestanian, Self-Motile Colloidal Particles: From Directed Propulsion to Random Walk, *Phys. Rev. Lett.* **99**, 048102 (2007).
- [46] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevE.104.054614> for two videos demonstrating the gradient-based strategies.
- [47] J. A. Barker and D. Henderson, Perturbation theory and equation of state for fluids. II. A successful theory of liquids, *J. Chem. Phys.* **47**, 4714 (1967).
- [48] G. Jékely, J. Colombelli, H. Hausen, K. Guy, E. Stelzer, F. Nédélec, and D. Arendt, Mechanism of phototaxis in marine zooplankton, *Nature (London)* **456**, 395 (2008).
- [49] B. ten Hagen, F. Kümmel, R. Wittkowski, D. Takagi, H. Löwen, and C. Bechinger, Gravitaxis of asymmetric self-propelled colloidal particles, *Nat. Commun.* **5**, 4829 (2014).
- [50] B. Liebchen, P. Monderkamp, B. ten Hagen, and H. Löwen, Viscotaxis: Microswimmer Navigation in Viscosity Gradients, *Phys. Rev. Lett.* **120**, 208002 (2018).
- [51] C. Datt and G. J. Elfring, Active Particles in Viscosity Gradients, *Phys. Rev. Lett.* **123**, 158006 (2019).