



HAL
open science

Morphological computation and decentralized learning in a swarm of sterically interacting robots

Matan Yah Ben Zion, Jeremy Fersula, Nicolas Bredeche, Olivier Dauchot

► **To cite this version:**

Matan Yah Ben Zion, Jeremy Fersula, Nicolas Bredeche, Olivier Dauchot. Morphological computation and decentralized learning in a swarm of sterically interacting robots. *Science Robotics*, 2023, 8 (75), 10.1126/scirobotics.abo6140 . hal-03840764

HAL Id: hal-03840764

<https://hal.science/hal-03840764v1>

Submitted on 12 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Morphological computation and decentralized learning in a swarm of sterically interacting robots*

Matan Yah Ben Zion^{1,2,3,†}, Jeremy Fersula^{1,2}, Nicolas Bredeche², and Olivier Dauchot¹

¹*Gulliver UMR CNRS 7083*

ESPCI, PSL Research University, 75005 Paris, France

²*Institut des Systèmes Intelligents et de Robotique*

Sorbonne Université, CNRS, ISIR, F-75005 Paris, France and

³*School of Physics and Astronomy, and the Center for Physics and Chemistry of Living Systems
Tel Aviv University, Tel Aviv 6997801, Israel*

Whereas naturally occurring swarms thrive when crowded, physical interactions in robotic swarms are either avoided or carefully controlled, thus limiting their operational density. Here we present a mechanical design rule that allows robots to act in a collision-dominated environment. We introduce the Morphobots — a robotic swarm platform developed to implement embodied computation through a morpho-functional design. By engineering a 3D-printed exoskeleton we encode a re-orientation response to an external body force (such as gravity) or a surface force (such as a collision). We show that the force-orientation response is generic, and can augment existing swarm-robotic platforms (e.g Kilobots) as well as custom robots even 10 times larger. At the individual level, the exoskeleton improves the motility and stability, and also allows to encode two contrasting dynamical behaviors in response to an external force or a collision (including collision with a wall or a movable obstacle, and on a dynamically tilting plane). This force-orientation response adds a mechanical layer to the robot’s sense-act cycle at the swarm level, leveraging steric interactions for collective phototaxis when crowded. Enabling collisions also promotes information flow, facilitating online distributed learning. Each robot runs an embedded algorithm that ultimately optimizes collective performance. We identify an effective parameter that controls the force-orientation response and explore its implications in swarms that transition from dilute to crowded. Experimenting with both physical swarms (of up to 64 robots), and simulated swarms (of up to 8192 agents) show that the effect of morphological-computation increases with growing swarm size.

INTRODUCTION

A robotic swarm is a synergetic ensemble capable of a greater task than its constituents [1]. In a swarm, self-organization is an emergent property, making the swarm resilient to the malfunction of an individual. Swarm engineering poses an interdisciplinary challenge, combining electrical and mechanical engineering, computer science, and non-equilibrium statistical physics [2, 3]. Discovering robust design rules for a swarm offers an opportunity to simplify the specifications of the individual robotic units, reducing manufacturing complexity and making production at-scale accessible [4–7]. Recent development in robotic swarms demonstrated potential applications in collective construction [8], coordinated motion for flying UAVs [9–12], patrol in open water [13] or underwater exploration [14, 15]. These applications are characterized

by swarms where robots explicitly avoid physical contact [16]. Another approach for swarm design is to keep a cohesive swarm, where robots are constantly touching, keeping a physical interaction that allows continuous transmission of information [17–22]. This approach showed success in self-assembly and morphogenesis [23–25] as well as coordination of a multi-cellular robotic body [26–29]. To date, artificial swarms are designed to operate exclusively in a dilute, collision-avoided setting, or a cohesive dense population. Resorting to *in silico* experiments to guide swarm design in a collision-dominated swarm proves limited by the ability to effectively account for the full mechanical interaction. There is a need for physical design rules and robust algorithms where robots maintain a high degree of autonomy in solitude, but also retain functionality in a crowded environment.

Unlike artificial robotic swarms, naturally occurring swarms show excellent flexibility in both dilute and highly dense environments throughout the animal kingdom. At the cellular scale, an individual bacterium can explore space by performing a simple run-and-tumble [30], but can also form a dense, highly active bacterial colony that displays internal turbulent flow [31] thereby mixing in nutrients [32, 33]. On the millimeter scale, an individual larva can search for food through chemotaxis but self-organize into a living fountain when a large localized food source is present [34]. Macroscopic

* This manuscript has been accepted for publication in Science Robotics. This version has not undergone final editing. Please refer to the complete version of record at <https://www.science.org/doi/10.1126/scirobotics.abo6140>. The manuscript may not be reproduced or used in any manner that does not fall within the fair use provisions of the Copyright Act without the prior, written permission of AAAS.

† To whom correspondence should be addressed.
E-mail: matanbz@gmail.com

animals such as fish and birds can individually nest and mate but can come together into giant schools and flocks to evade predation [35, 36]. Even pedestrians flow can emerge from simple body collisions [37]. Swarm cooperation is found on a range of length scales, suggesting that only simple primitive behaviors are required and physical interactions can be used for the swarm’s advantage [38]. Interestingly, self-organization in dense assemblies is not a prerogative of living beings. It has been shown that collective motion can emerge from independent mindless robots physically interacting with one another, considering a stripped-down flavor of swarm robotics as a particular kind of active matter with no explicit decision computed at the level of the robots [39–44].

Our motivation is to expand the roboticists’ toolbox by taking advantage of a generic mechanical response: the tendency to re-orient in response to an external force. We then demonstrate the role of this morphological response in different robotic tasks, both for an individual robot, and a robotic swarm. We introduce two designs to show that this force-orientation response can come in different flavors, aligning with, or against an external force. We show that this force-orientation response can also be used in a robotic swarm platform, giving control over the outcome of a robot-robot collision and the resulting collective behavior.

To achieve this we augment robots with a flexible exoskeleton and show that the morphological difference is encoded in a single effective parameter, labeled κ , that controls the outcome of collisions. We show two different designs that have similar dynamics for free individual robots (i.e a single robot moving on a smooth, flat surface, without experiencing any external force, nor touching any obstacles or robots). But when experiencing an external force, or collision, the two designs have an opposite response. We show that the passive dynamics of the design allow tuning the robot’s response to align or oppose an external force. Upon collision, a robot can push against, or slide along an obstacle or another robot, based solely on its mechanical design. The emerging collective dynamics add a morphology layer to a robot’s computation, facilitating a combined physical and logical swarm architecture [45, 46]. We quantitatively analyze the response to an external force and find empirically that it can be described by a single effective parameter, κ . We show experimentally, theoretically, and numerically that a change in the sign of κ inverts the response to an external force, making a robot go down or up a force field. We find κ to be a key design parameter in the response of the robots when encountering an obstacle (be it a wall or another robot). Using positive and negative κ makes two primitives that qualitatively alter the swarm’s dynamics in a collision-dominated, crowded environment. We also show that κ can be used to guide the design of collective transport.

We then add a computational layer, where in parallel




FIG. 1. **Article at a glance.** Swarm robotics design space spans both morphology and computation. Experiments with 3D-printed flexible exoskeletons show they can direct the outcome of collisions between robots. The mechanical design of the exoskeleton, frontier versus aligner, controls whether a robot ascends or descends a hill, which translates to pushing against or sliding along a wall. The exoskeletons’ morphological computation extends the ability of the swarm to execute collective tasks, including transport, phototaxis, and decentralized learning.

to the implicit morphological computation, robots execute an explicit sense-act cycle. We experimentally study phototaxis in a partially lit arena, setting a swarm of $N = 64$ robots to search for the lit region. We set up the illumination field to be unoriented (coming directly from above), and without local light gradients, thus ensuring robots cannot use local information to guide their way to the lit region. Moreover, the lit region is designed to be too small to host the whole swarm. This restriction forces the robots to physically interact, turning the phototaxis task into an exercise of collective aggregation with steric interactions, going from dilute (in the dark) to crowded (in the light). Studying swarms with positive and negative κ reveals a qualitative and a quantitative different collective phototaxis, directly attributed to the nature of this effective parameter.

Finally, we show that the swarm can collectively learn to execute the phototactic task when operating in an environment where individual robots transition between dilute and crowded regions. We implement a decentralized reinforcement learning algorithm inspired by social learning [47–50], and find that when collisions are allowed the swarm fluidizes, and effectively converges to a successful phototactic strategy despite the sparse and intermittent communication network. We further support our experimental observations with agent-based model simulations, finding that morphology is central to collective behavior and becomes increasingly important with growing swarm

size. The interplay of morphology and computation is graphically illustrated in Fig. 1, summarizing the findings of our work.

INDIVIDUAL DYNAMICS

The Morphobots

We make Morphobots by augmenting Kilobots [23] with a 3D-printed exoskeleton. The Kilobot is a 3.4 cm tall, 3.3 cm diameter unit, standing on 3 rigid legs made of thin metallic rods. The Kilobots are capable of differential drive locomotion using two vibrators. When either is activated, a Kilobot turns at about 45 deg per second, and when both are activated, the Kilobot moves forward. The power source is a rechargeable lithium-ion battery that can keep the Kilobot running for a few hours. Kilobots are equipped with an infrared transmitter and receiver so that they can communicate with each other. The transmitter of a robot sends light toward the surface which reflects up to the receiver of another nearby robot. Obtaining straight motion over long distances is difficult and requires lengthy calibration of the individual vibrators [51], moreover, the calibration itself drifts over time.

To both improve the robot's nominal speed, and eliminate the need for individual calibration, we encapsulate the Kilobots within exoskeletons that change how collision and locomotion are performed. We made two tripod-based designs (a pair of flexible legs and an opposing stiff leg), with both designs having a circular chassis (diameter $d = 4.8$ cm). The two designs, which we shall call *fronter* versus *aligner*, differ in the fore-versus-aft positioning of their flexible legs (see Fig. 2-A,B,C and Methods). The natural vibration of the semi-flexible legs is tuned near resonance with the vibration motors, to maximize the coupling with their drive. Using an anisotropic flexible leg design (1mm thick, 6 mm wide, and 20 mm long, see Supplementary Materials), vibrations are coupled predominately to the forward stick-slip motion of the robots. The exoskeleton does not interfere with the ability of the robots to communicate, and has multiple roles including energy storage of vibrations, near-resonance coupling to motors, anisotropic design to translate motor-vibrations into a forward motion. The exoskeleton allows a Morphobot to operate at an order of magnitude higher speed than the Kilobots (see Fig. 2-D and Supporting Movie 1). More importantly, it allows control of the response to external forces and collisions.

By tracking the motion of individual Morphobots, we find that aligners and fronters show near-identical motility. Motility is defined by the magnitude of the displacement of an individual robot over different timescales, and measured by tracking the time evolution of their two-dimensional coordinates, $\vec{r}(t) = (x(t), y(t))$, when run-

ning in the experimental arena [52]. Only a low concentration of robots was used (less than 1% filling fraction) excluding the effect of collisions, and robots were tracked only when away from the walls of the arena. The mean square displacement, $\text{MSD} = \langle \Delta r^2(\Delta t) \rangle_N$, of an ensemble of noisy walkers, presents two dynamical regimes [53]. On short time intervals Δt , the displacement of the robots is aligned with the direction of their individual orientation and displacements are proportional to Δt . The motion is said to be ballistic when robots go roughly on a straight line at a constant speed: $\langle \Delta r^2(\Delta t) \rangle_N = v_0^2 \Delta t^2$, where v_0 is the instantaneous speed.

On long time intervals, the noise has reoriented the robots several times and displacement grows as $\sqrt{\Delta t}$ as robots explore the space randomly. The motion is now said to be diffusive $\langle \Delta r^2(\Delta t) \rangle_N = 4D_0 \Delta t$, where D_0 is called the diffusion constant. The noisy trajectories of the robots originate from the stochastic nature of their vibrational drive [54], as well as the roughness of the arena and internal electronic drift. The two regimes ($\propto \Delta t$ and $\propto \sqrt{\Delta t}$) are separated by a crossover, which takes place on a time scale called the persistence time, τ_p (see below), to which corresponds a persistence length $\langle \Delta r^2(\tau_p) \rangle_N^{1/2}$. Fig. 2-D displays the mean square displacement measured for a set of Morphobots equipped with the two types of the exoskeleton, together with their ensemble average (the MSD of 8 robots of each design were measured, shown as separate curves in Fig.2D, see also Table I). One sees that the difference in motility between the two designs is no larger than the robot-to-robot variation within each design. The observed ballistic nature of the mean square displacement ($\propto v_0^2 \Delta t^2$) indicates that the persistence length of the motion is larger than the size of the circular arena (diameter $D = 150$ cm). Fitting the ensemble mean square displacement in the ballistic regime, we find that the mean nominal speeds are 5.2 ± 0.9 cm/s and 4.8 ± 0.1 cm/s for the aligners and fronters respectively. The measured effective parameters are summarized in Table I.

Altogether, the anisotropic leg design allows us to use higher motor drive while keeping the robot's trajectories straight, alleviating the need for individual robot calibration, and resulting in a nominal speed of $v_0 \simeq 5$ cm/s, a ten-fold increase, relative to the bare Kilobot design (see Supporting Movie 1). The notion of using passive-dynamic machines [55] allowed to achieve a similar feat in bristle bots, using rows of soft legs [56], and in the commercially available *HEXBUG* [28, 29, 57, 58]. Using the differential drive, robots can also pivot at 45 deg per second. Programming alternated straight and pivoting motion of random duration and random direction, we can set the Morphobots into a run-and-tumble motion [59], where, on average, the robots run straight during a run time $\tau_{run} = 2$ s, with a velocity to be specified, and tumble during a mean tumble duration $\tau_{tumble} = 4$




FIG. 2. Augmenting Kilobots with exoskeletons makes Morphobots. **A** a Kilobot with a light sensor, IR communication, and two differential drive vibration motors, is fitted into one of two 3D printed exoskeleton designs - (i) fronter, and (ii) aligner. Both fronter **B** and aligner **C** exoskeletons are tripods, but the restitution vector, \hat{r} of the aligners (fronters) is parallel (anti-parallel) to the orientation vector \hat{n} . **D** mean square displacement versus time shows the ballistic response for both designs ($\propto t^2$), with a typical speed of 5.2 cm/s (4.8 cm/s) for the aligners (fronters). Also shown, the mean square displacement of a bare Kilobot which after calibration can also move ballistically, but at a speed one order of magnitude slower (0.53 cm/s). Inset: 12-second long trajectories of fronters (green) and aligners (cyan), and a 36-second long trajectory of bare Kilobots (gray) inside the 150 cm diameter arena.

s. The total duration of a run-and-tumble sequence sets $\tau_{RT} = \tau_{run} + \tau_{tumble}$, with the resulting persistent time being $\tau_P \approx 18$ s (See SM).

Fronters vs. Aligners response to an external force

We systematically measure the response of the two Morphobots designs to an external body force by placing the robots on a plane tilted by a small angle, φ . The two designs have a polar different response: aligners go downhill, while fronters climb against gravity (Fig. 3A,B and Supporting Movies 1-6). The effect is consistent, regardless of the initial orientation of the robot being along the positive y direction (downhill to its right), or negative y direction (downhill to its left), excluding the effect of an internal left/right bias (see Movies 3-6). We systematically increase the slope (see Fig. 3C), and find that with increasing φ , the trajectories turn increasingly curved, until the robot orients downhill (aligner) or uphill (fronter).

To quantify the different responses of the robots to an external body force, we modify a model introduced in the context of active matter to describe self-aligning particles [57]. A Morphobot has two translational $\vec{r} = (x, y)$ and one orientational $\hat{n} = (\cos(\theta), \sin(\theta))$ degree of freedom. In the overdamped limit, the deterministic dynamics of the velocity \vec{v} and orientation \hat{n} of a Morphobot with a nominal speed v_0 subjected to a body force \vec{f} (see

Fig. 3D) obey the following equations :

$$\frac{d\vec{r}}{dt} = \vec{v} = v_0\hat{n} + \mu\vec{f} \quad (1)$$

$$\frac{d\hat{n}}{dt} = \kappa(\hat{n} \times \vec{v}) \times \hat{n}, \quad (2)$$

where μ is the particle mobility and κ is the alignment parameter — an effective parameter describing the reorientation by force. The first equation simply states that the velocity is the additive combination of the propulsion velocity $v_0\hat{n}$ and the speed induced by the external force, $\mu\vec{f}$. The second equation contains the key ingredient of the model. It describes the reorientation of the Morphobot along the direction of its motion. The coupling between an external force and the orientation is captured through an effective alignment parameter, κ . This parameter measures to what extent an external force, \vec{f} , affects the orientation of a robot, \hat{n} . In the most simple case, $\kappa = 0$, and an external force does not affect the orientation of a robot (Eq. 2 becomes redundant). However, generically force and orientation do couple, and κ is finite ($\kappa \neq 0$). Moreover, κ is signed as it can be positive or negative. This means that when subjected to an external force the robot does not only translate, but it also rotates. The larger the magnitude of the alignment parameter, $\|\kappa\|$, the quicker a robot's orientation will respond to an external force. There are two important types: when positive ($\kappa > 0$) a robot aligns parallel with an external force. Here we call robots with positive κ aligners. When negative ($\kappa < 0$) a robot aligns anti-parallel to an external force. Here we call robots with




FIG. 3. The orientation response of aligners and fronters to an external force is equal and opposite. **A** Sketch of the experimental setup for measuring the dynamical response on an inclined plane (see Movies 1, 3-6). **B** Force diagram of a robot placed on an inclined plane. For small slopes ($\varphi \leq 6^\circ$), robots experience an external body force. **C** Trajectories of aligners (cyan) and fronters (green) moving on different slopes (φ). Aligners and fronters have an opposite response to the external force - aligners' orientation turns downwards as they quickly descend, while fronters turn upwards, and slowly ascend. **D** Throughout its motion, the orientation of the robot, \hat{n} , forms an angle θ with the external force, \vec{f} , together setting the net velocity \vec{v} (see Eq. 1). **E** Tracking the orientations shows that aligners turn downhill ($\theta \rightarrow 0$), aligning with the external force, while fronters turn uphill ($\theta \rightarrow \pi$), fronting the external force (see also Movies 3-6). Solid curves are fits to Eq. 4, from which κ can be extracted (see SM). **F** Average fitted κ for each incline angle and robot-design (error bars are the standard error). For finite tilt angle, the response of the two designs is approximately equal and opposite $\kappa_{aligner} = -\kappa_{fronter} \approx 0.3/d$. At a slope near zero, the response is dominated by the robot's internal bias (left/right curving and internal noise [54]), but averages to zero. Scale bar is 20 cm.

negative κ fronters. κ is an effective parameter and results from the design details of the robot (material properties, mass distribution, drive mechanism, and so on). And yet κ is empirically well defined as it can be experimentally measured and quantitatively compared between different robotic platforms. For a constant force in the \hat{x} direction $\vec{f} = f\hat{x}$, with $f = mgs\sin\phi$, where m is the mass of the Morphobot and g is the acceleration of gravity, the dynamical evolution of the orientation, $\theta = \theta(t)$ relative to the x axis obeys a closed form equation :

$$\frac{d\theta}{dt} = -\kappa\mu f \sin\theta. \quad (3)$$

Eq. 3 is the equation for an overdamped simple pendulum, which can be solved directly to give $\tan\frac{\theta}{2} = e^{-\kappa\mu f t}$ (for initial condition $\theta(t=0) = 90^\circ$). In the limit where $v_0 \ll \mu f$, when the orientation is expressed as a function of the arclength, $\theta(s)$, it is found to be independent of the external force and mobility (see Supplementary Materials),

$$\theta(s) = 2\text{atan}(e^{-\kappa s}). \quad (4)$$

To find the alignment parameter, κ , we tracked the ori-

entation of individual robots as they move across a fixed inclined plane (see Fig. 3E, Movies 3, 4, and Methods). In the absence of an external force (zero slope, $\phi = 0^\circ$), both fronters and aligners move on roughly straight lines. For steeper inclines ($\phi = 3^\circ$ and $\phi = 6^\circ$), both designs show a similar rate for the convergence of the orientation. Fitting the evolution of the orientation to Eq. 4, we extract the alignment parameter, κ . The magnitude of κ for both designs is similar, but with opposite signs: $-\kappa_{fronter} \approx \kappa_{aligner} \approx 0.06 \text{ cm}^{-1}$. The measured effective parameters are summarized in Table I.

We find that the design rules controlling the alignment parameter apply also in a dynamically changing environment as well as to robotic platforms of different sizes (see Fig. 4). Placing a Morphobot of either type on a seesaw gives rise to a unique behavior (Fig. 4): An aligner turns quickly downhill and falls off the seesaw, while a fronter turns upwards and climbs. Once passing the center of the seesaw, the mass of the robot suffices to tip the plane of the swinging table. The fronter is now facing downhill, yet, again, as it moves, it turns around and recommences climbing, until passing the center of the seesaw once more, where these dynamics repeat, effec-




FIG. 4. Fronters and aligners respond differently to a dynamic environment. **A** A Morphobot on a swinging table setup for testing morphological response to a dynamically changing environment. **B** A fronter continually turns uphill, moving in a circular motion while staying on the swinging table (see also supporting Movie 2), whereas an aligner **(C)** quickly turns downhill and falls off. Top shows schematics, bottom shows measured trajectories. **(D)** The design rules apply also to custom-made robots (not Kilobot based), producing the expected force alignment response with Morphobots as much as 10 times larger than a Kilobot (see supporting Movies 5,6). Scale bars are 10 cm.

tively tracing loops (see Fig. 4B and supporting Movie 2). Moreover, the design rules for the Morphobot generalizes to a custom robotic platform independent of the Kilobots (Fig. 4D) and can be made 10 times larger (see Movies 5 and 6).

f

	v_0 (d/s)	κ (1/d)
Aligners	$1.1 \pm 17\%$	$0.3 \pm 10\%$
Fronters	$1.0 \pm 2\%$	$-0.3 \pm 20\%$

TABLE I. Measured nominal speed and alignment parameter of aligner and fronter designs.

Interaction with a wall

Using the response of the Morphobots to an external force we can predict their steric interaction with stationary walls, movable objects, and with one another. Morphological interaction with a boundary (stationary or movable) can be seen as a series of collisions, at each of which the robot experiences an external force (see Fig. 5-A). With each collision, a Morphobot slightly re-orient. An aligner orientation approaches the normal force from the wall (but leaves when parallel to the wall). A fronter turns against the normal force (and towards the wall). This is confirmed by the direct observation of Morphobots sent to move towards a wall of stationary robots (see Fig. 5-B). Arranging turned off robots into three, tight rows forms effective wall with which a moving robot mechanically interacts. On average, both designs spend similar time near the wall ($\sim 3\tau_p \approx 1$ min) but aligners travel a path along the wall which is twice longer than that of the fronters (see Fig. 5-C and Movie 1). Using a multi-agent Brownian dynamics simulation of active soft discs, we were able to predict the outcome of collisions based on the alignment strength parameter κ . During a collision, an agent with $\kappa > 0$ (aligner), turns towards the normal force, and away from the obstacle, whereas an agent with $\kappa < 0$ (fronter) turns against the normal force and into the obstacle thus pushing it. Also, simulating agents with $\kappa < 0$ near a passive particle shows that the agents collectively move the passive object (see Fig. 5-D). This has been confirmed experimentally by using a group of fronters that work together to push a circular object (see Fig. 5-E and Supporting Movie 1). Previous work showed that collective transport can be implemented using global control [60]. Here we extend this notion and show that a single, effective alignment parameter, κ , captures embodied dynamics of sterically interacting agents, reducing the required computing resources for simulating a large swarm in a crowded environment.

COLLECTIVE DYNAMICS

We turn to study the ability of a swarm of Morphobots to execute a collective task, focusing on phototaxis, where robots explicitly use their sensory and computational power in addition to the morphological interaction. Phototaxis is a common behavioral trait found in nature where an organism changes its motion in response to locally measured light. Phototaxis pervades the animal kingdom, from individual bacterium to groups of vertebrates. Phototaxis is classified as either positive, where the organism seeks to be present in the light (such as moths [61]) or negative, where the organism prefers the dark (such as golden shiners [62]). Here we test the ability of a robotic swarm to perform this widespread natural behavior, with an emphasis on the collective abil-




FIG. 5. Tuning interaction with a static and a moving obstacle using morphological computation. The morphological interaction of the Morphobots with walls can be explained using their response to an external force. When a Morphobot collides with a wall, each impingement can be seen as a momentary normal force. **A** Aligners turn to align their orientation (black arrow) with the normal force (red). Once aligned along the wall, there is no further rotation as there are no more collisions, and the robot slides along. **B** When a fronter collides with a wall, its orientation (black) turns anti-parallel with the normal force (red), and the robot continues to push into the wall. When programmed to run-and-tumble, aligners and fronters interact differently with a wall of stationary robots where aligners (**A**) cover a much larger distance along the wall relative to fronters (**B**). **C** Multiple collision experiments show that aligners and fronters spend similar time near a wall ($3\tau_p \approx 1$ min) yet aligners move twice as fast. **D** Brownian dynamics simulations predict that active particles with a negative alignment parameter, $\kappa < 0$ (fronters), push a movable object. **E** With real robots: placing a round movable obstacle shows that a swarm of fronters is capable of transporting an object. Scale bar is 10 cm.

ity to maximize the time spent in the light (positive phototaxis). We program the robots to change their speed in response to the locally measured light intensity (also known as photokinesis [63]) and show that the robots' morphology can be used to improve their collective phototaxis when crowded. Stimulus-induced locomotory movements (or *taxis*) constitute an important building block in swarm robotics [2, 3, 19, 64–66], and several bio-inspired algorithmic implementations have been proposed for phototaxis using ambient light sensors [26, 67–70].

Collective phototaxis in the presence of steric interactions

The experimental setup for collective phototaxis was designed to prevent an individual robot from accessing global information that could guide its sense-act cycle. For this, we set a swarm ($N = 64$ bots) of either aligner ($\kappa > 0$) or fronter ($\kappa < 0$) to explore a circular arena (diameter $D = 150$ cm) with a lit spot that covers only $\sigma = 6\%$ of the area of the arena (diameter $b = 36$ cm, see Fig. 6-A and Section in SM). The light is unoriented (coming from above), and there are no local light intensity gradients to guide a robot to the lit region (the

light field is essentially flat). This excludes the use of traditional phototactic algorithms such as steepest descent search [25, 71] or reorientation [23]. The lit region can not accommodate more than 80% of the robots (about 50 Morphobots), challenging the swarm to transition between sparse and crowded.

As a first step, we consider ad hoc calibrated swarms, where the Morphobots' policies are defined as follows. Each Morphobot continuously monitors local light intensity: when in the dark (below a predefined intensity threshold, P_{th}) the robot performs a run-and-tumble motion, with a running velocity V_0 . When entering a bright region (light intensity above P_{th}), the robot stops, $V_1 = 0$. At early times, when there are only a few robots in the lit region, the two designs show similar performance as measured by the fraction of the swarm that is in the light

$$\mathcal{F} \equiv \frac{N_1}{N}, \quad (5)$$

where N_1 is the number of bots taxied in the light. At the beginning of every experiment, robots are placed in the arena with random positions and orientations. At early times, \mathcal{F} is quantitatively captured as a diffusion-limited process (see Fig. 6-F), with the kinetic reaction constant found using Smoluchowski first passage computation for diffusive particles with an effective diffusion constant, $D_{eff} \approx 4.2^2$ cm/s (See Supplementary Materials) [72]. After approximately 10 minutes, both swarms deposit the first layer of robots at the perimeter of the lit region (Fig. 6B, D). From this point onward, the phototaxis rate slows down and deviates from the simple diffusion-limited reaction, as crowding of taxiing robots at the perimeter of the light spot physically prevents further robots from entering the lit region (Fig. 6C, E). This steric effect is quantitatively more pronounced for the aligner design, as can be seen for the overall performance given by the fraction of the swarm that has entered the lit region over long times (Fig. 6F).

The reason is that the collective success of each swarm depends on the morphological interactions of a robot impinging from the dark region on the first layer of robots at the perimeter of the lit region. Individual Morphobots arriving at the wall of stationary robots show qualitatively different dynamics for either design: an aligner ($\kappa > 0$) that arrives at the wall tends to reorient along the wall until it finally leaves, whilst a fronter Morphobot ($\kappa < 0$) tends to “dig in”, pushing already taxied robots into the lit region until taxied itself (see Fig. 6B, D). These individual events add up: fronting Morphobots progressively push the perimeter inwards, constantly filling the bulk of the lit region with robots. By contrast, the aligning Morphobots keep reorienting away from the robots at the perimeter, leaving a hollow in the lit region (see Fig. 6C, E and Supporting Movie 1). This discrepancy exacerbates as the swarm size increases. By simulating up to




FIG. 6. The collective ability of a robotic swarm to phototaxis is constrained by steric interactions. **A** A swarm of phototactic robots is placed in an arena patterned with an un-oriented light field. As there are no local gradients to guide a robot to the light source, robots perform a simple run-and-tumble in the dark and stand in the light. **B** When a fronter arrives at the light spot, it collides with robots already taxied at the perimeter, but progressively pushes its way in. Repeating this process results in the complete coverage of the light spot (**C**). When an aligner arrives at the light spot (**D**) it slides along the wall of robots and leaves. This leads to only partial filling of the light spot (**E**), leaving an empty void in the center. **F** The resulting fraction of swarm in the light as a function of time shows that at early times ($\tau \approx 30\tau_p \approx 10$ mins), both design show equal performance, quantitatively consistent with a diffusion-limited reaction (dashed line). At later times, fronters outperform aligners (average of 4 realizations with error bars being standard deviation). **G** Simulating swarms with over two orders of magnitude more agents show a sharp slowdown in the phototaxis of aligners with swarm size. **H** The typical time for half of the swarm to enter the light $T_{1/2}$ grows rapidly for an increasingly larger swarm of aligners, while having only a small effect on a swarm of fronters of similar size. Scale bars are 20 cm.

8192 phototactic agents that follow Eqs. 1,2 we find that a swarm of aligners ($\kappa > 0$), becomes considerably slower to phototaxis relative to an equivalent swarm of fronters ($\kappa < 0$). The typical time, $T_{1/2}$, to have half of the swarm in the light ($\mathcal{F} = 1/2$) grows much faster for aligners than for fronters (see Fig. 6G, H). The above findings demonstrate that the collective ability of the swarm to perform a phototactic task at scale hinges on a subtle morpho-functional difference.

Distributed learning of a phototactic strategy

The swarms considered thus far had the environmental information hard-coded into individual robots before deployment — the sense-act cycle of every robot had a pre-programmed, ad-hoc knowledge of the correct light threshold for phototaxis (which is an unrealistic assumption when robots are to be deployed in an unknown environment). Concurrent challenges however require deployment of robotic systems in unknown terrains [2, 3, 64, 65],

including extra-terrestrial exploration, underwater surveying, and search-and-rescue missions [73–75]. Under such circumstances, robots will lack an a priori model of the environment that can be pre-programmed. This requires an added layer to the robots’ controller for on-line tuning of the sense-act cycle based on continually acquired knowledge. In such conditions, a distributed robotic system offers robustness with scale, improving both exploration and survival. A distributed system however also poses a challenge as the swarm has to learn collectively.

In this final section we test whether a swarm can collectively learn even when robots collide. We show that in a swarm of Morphobots, both execution and learning can be made fully distributed (see Fig. 7). We then proceed to test the effect of the two designs on the collective performance and find numerically that over long enough time, a swarm of frontiers out-performs a swarm of aligners. Since the robots are allowed to collide, the swarm is fluidized [76]. An individual collision event may be stochastic, but the mean data flow within the swarm becomes predictable. Unlike a solid or a jammed state (where individuals do not move), here the swarm acts as a fluid, meaning that it satisfies good mixing, as robots frequently exchange neighbors and visit various parts of the arena, sustaining a flow of information (see Figs. S4, S5 in supporting information). Each robot evaluates its performance and compares it to other encountered robots. The selection algorithm can be made distributed by implementing a decentralized online reinforcement learning scheme, to make the swarm adaptive.

Previous work showed that a swarm can learn a phototactic strategy with a policy defined by the weights of an Artificial Neural Network as a controller, $\pi = \{w_i\}$ [77–79]. There, the learning was implemented using an evolutionary algorithm to act as a direct policy search, using a reinforcement learning method, incrementally refining behavioral policies by directly tuning the policies’ parameters [80, 81]. The swarm is randomly initialized with a pool of sense-act policies, and in an iterative selection process, only the best-performing policies are carried over to the next generation. In these algorithms, the execution is indeed distributed over the swarm, but it is a central computer performing selection and optimization (see Fig. 7D). Such schemes are impossible to implement outside the lab if the swarm has no access to an external operator for guidance, nor a global communication channel to compare performances. In what follows we show that despite having a sparse and intermittent network, communication in a swarm of Morphobots becomes effectively connected across time thanks to the swarm’s fluidity, facilitating decentralized learning (see Fig. 7F).

For a successful reinforcement learning algorithm to converge, a suitable reward function, ρ , must be defined [82]. As an input, the reward function receives the robot’s sensory information $\rho = \rho(\{i_i\})$ and out-

puts a scalar score (see Fig. 7A and Algorithm 1 in SM). The combination of an individual robot sense-act policy, $\pi = \pi(\{w_i\})$, and the environment it sampled, results in a robot’s reward at a given time. Robots can also communicate when within communication range and exchange their policies and rewards (Fig. 7 C, F1-F4). Following such an exchange, a robot can choose to adopt the received policy based on its associated reward, operating under the assumption that a higher reward is the result of a superior policy. Unlike classical evolutionary algorithms that use a centralized computer to compare performance, the scheme presented here allows embedding both the evaluation and the selection within the robot. In other words, robots collectively perform a kind of social learning, where successful individual innovations diffuse over the population [47, 48, 50]. Note that the assumption that a higher reward is a result of a superior policy may not be always true. From time to time robots with an incidental favorable history can have a disproportional higher reward. Such bots will propagate a poor policy globally to the swarm, effectively spreading ‘fake news’. This local-to-global gap challenges the chosen reward function to be robust to such fluctuations. The fluid nature of the swarm of Morphobots allows us to choose ρ that overcomes this culprit. Most swarm robotic systems employ a collision avoidance algorithm that slows down their collective dynamics, jamming the information flow, and creating deadlocks [8]. By contrast, the Morphobots are allowed to collide, sustaining good mixing, and offering a quantitative link between global and local information (see Figs. S4, S5 in Supplementary Materials).

To formulate a proper reward function we note that the collective phototaxis, \mathcal{F} , as defined in Eq. 5 is an ensemble average, $\langle \rangle_N$, counting the fraction of the swarm that is in the light at a given time. \mathcal{F} can not be directly optimized by an individual robot as it stores global information. However in a system with good mixing, a time average can approximate an ensemble average $\phi \equiv \langle \rangle_T \approx \langle \rangle_N$. A time average equivalent to the ensemble average given in Eq. 5 is the mean time an individual robot spent in the light, $\Phi \equiv T_{\text{in light}}/T \approx \mathcal{F}$. In a distributed robotic system with good mixing, a reward function that estimates the mean time the robot spent in the light can serve as a proxy for the global performance of a swarm of its kind, $\rho \propto \Phi \approx \mathcal{F}$. Robot i can evaluate the mean time it spent in the light through a temporal average of its light sensor signal:

$$\rho^i = \frac{1}{M} \sum_{k=0}^{M-1} P_k^i, \quad (6)$$

where M is the number of averaged measurements, and P_k^i is the the k ’th measurement taken at constant time intervals. Since the time a robot spends in the light is inversely proportional to its speed in the light, $\Phi \propto \rho \propto$




FIG. 7. Distributed online reinforcement learning algorithm and deployment of a swarm of Morphobots. **A** The embedded algorithm for updating the phototactic policy of each robot, is composed of two main moduli: a sense-act cycle and a learning cycle. **B** The sense-act cycle is encoded into a Perceptron: external stimuli are weighted and summed, evaluated through a sigmoid function, the output of which is binarized to select an action. The weights, w , define the policy π of the robot. **C** The learning cycle is based on the conditional diffusion of the policies across the population, biased by a reward function ρ evaluated by each robot: when robot 2 receives a message containing the policy control parameters $\{w_i\}_1$ and the reward ρ_1 of robot 1, robot 2 will inherit the new policy, if it is associated with a higher reward $\rho_1 > \rho_2$. **D** Illustration of a centralized communication network where all agents (gray circles) are connected to a central operator (red circle). **E** An illustration of a sparsely connected network where agents communicate only when in close proximity. **F1-4** Illustration of a sparse and intermittent communication where a robot (black curve) already in the light source (**F1**) departs (**F2**) and broadcasts within communication range (yellow circle) its reward and policy to a robot in the dark (red curve) (**F3**), which inherits the new policy (**F4**). Note that the broadcasting robot may have a wrong policy, effectively spreading false information, with neither robot knowing what is the best policy for phototaxis. **G** A swarm of Morphobots with randomly initialized light thresholds, collectively learns its environment, increasing the swarm fraction in the light. Results are the average of 4 runs with error bars being the standard deviation. Insets show that fronters (top) and aligners (bottom) develop a different spatial arrangement in the highlighted region. **H** Evolution of the swarm fraction in the light for multi-agent simulations of swarms of different sizes learning to phototaxis using the distributed on-line, reinforcement learning. At earlier times, swarms of fronters and aligners develop the spatial arrangement found in the experiments. At later times (inaccessible experimentally) a swarm of fronters compactly arrange inside the light spot a larger fraction of the swarm in the light compared to the swarm of aligners (see also insets).

$1/V_1$ the above choice of reward function is analytically linked to a given policy. This link allows us to show that when formulated as a dynamical process, collective learning will indeed converge on average (see SM for detailed derivation).

To test the validity of the proposed decentralized learning scheme, we deployed a swarm of Morphobots in the same light-patterned arena, but this time with *randomly initialized policies* (See Fig. 7, Supporting Movie 1). A

“correct” policy corresponds to a set of weights $\{w_i\}$ of the perception, where the robot slows down in the light. A “wrong” policy corresponds to a set of weights, where the robot does not slow down in the light. When randomly distributed in the large arena used (150 cm), robots are on average over 16 cm apart (see for example 01:41 in Movie 1). This is more than twice the maximum communication distance supported by the Kilobot platform (7 cm). At such a separation, many of the robots

are out of communication range from one another making the network sparse. Not only the network is sparse, but the network is also intermittent. Even when two robots do come within communication range, they quickly move away. On average, the subgroup of the swarm that interacts at a given time makes up only a small fraction of the swarm, and this subgroup is constantly changing. Altogether making a sparse and intermittent communication network. We also program the robots to start with random initial policies by drawing their weights at equal probability from values ranging between 0-255 (see Methods Section and Supplementary Materials for details). At very early times only a few robots slow down in the light — most robots either stand in the dark or move quickly through the lit region. And yet the swarm fraction in the light progressively increases with time. We tested swarms of both aligners and fronters, and found a consistent increase in the fraction of the swarm in the light reaching approximately $\mathcal{F} \approx 0.4$. At these swarm fractions in the light, the difference between fronters and aligners is not expected to be significant as is evident here (compare with Fig. 6F at early times). The learning swarms had 8 times more robots in the light than the case of random placement that is, the fraction of the lit region, $\sigma = 0.06$. Given that Kilobots are ten times slower, and diffusion scales as the speed square ($\propto v_0^2$, see SI), a swarm of bare Kilobots would have taken 100 times longer to converge, making collective learning impossible. The fact that the adaptive swarm performed almost, but not quite as well as, the pre-calibrated swarm is not surprising and may be attributed to the fact that dynamics of about half of the randomly initialized population are very slow (as they stand in the dark), and given enough time, they will also eventually learn. This may be resolved by expanding the learning space of the sense-act cycle, to have the robots also learn the speed in the light, V_1 .

We find that the difference between learning swarms of aligners and fronters increases with time. By extending the multi-agent numerical engine, we tested the learning process over a longer duration (inaccessible experimentally) and for larger swarms (up to 8 times larger, see Fig. 7H and Fig. S3 in supporting information). Implementing Algorithm 1 numerically, we find that after approximately $\tau \approx 300 - 400$ from their initialization, swarms of either aligners or fronters reach the same fraction of robots in the light ($\mathcal{F} \approx 0.4$, Fig. 7G), consistent with Morphobots studied in the laboratory. In both experiments and simulations, we find that at these intermediate times, the swarm of fronters tends to occupy the lit region while aligners remain on the perimeter (see Fig. 7G, H). Simulations show that at later times ($\tau > 400$) this qualitative difference grows: the fraction of aligners in the light remains the same, whereas the swarm of fronters continues to crowd the lit region, reaching $\mathcal{F} \approx 0.55$. This effect is consistent with growing swarm size, showing the significance of morphological

interaction at scale in adaptive swarms.

CONCLUSION

In this research article, we introduced a robotic swarm that can operate in both dilute and dense settings, where collisions between robots occur frequently. We show that the robots' physical morphology plays an important role in terms of collective behaviors, where a seemingly small change in the shape leads to very different outcomes. Applying tools from statistical physics of active-matter to swarm robotics, we show that behavior found in natural swarms (such as ant colonies [83]) including collective self-aggregation and transport can be attained with limited use of software-controlled behaviors, by capitalizing on the morphological properties of robots, demonstrating morphological computation at work within the swarm.

We also showed that distributed online evolutionary reinforcement learning can be implemented in a dense robot swarm, to exploit or mitigate physical contingencies after deployment in the open. We find that this kind of artificial social learning, when robots locally exchange information about their behavioral strategies, can be modeled as a conditional diffusion process.

Future collective learning implementations in swarm robotics can go beyond the diffusion of behavioral strategies. An extreme case is of sparse swarms, where individual robots or groups may be temporarily disconnected from the whole [84]. While dense and sparse robot swarms may be considered separately, a robot swarm in the open will likely have to switch from one configuration to another, requiring versatile social learning capabilities. A natural extension can be motivated by cultural evolution found in nature [85], which is built on the reformulation and combining previously learned behaviors to continuously discover behaviors of growing complexity.

Future physical models of a learning swarm should go beyond the mean-field approximation used here and account for the spatial density fluctuations (as is done for describing the motility-induced phase separation observed in active fluids [59, 86]). From that point of view, our work shows that building upon the recent theoretical progress made in describing the population of active particles, one can envision a rapid development of a statistical mechanics description of smart active matter, for future swarm engineering.

ACKNOWLEDGEMENTS

We thank Y. Lahini and N. Oppenheimer, and S. Elul. **Funding:** This work was supported by the MSR project funded by the Agence Nationale pour la Recherche under Grant No ANR-18-CE33-0006. M.Y.B.Z. gratefully acknowledges support from the Ministry of Aliyah and

Immigrant Absorption of Israel. **Author contributions:** O.D, N.B, and M.Y.B.Z conceived the research. M.Y.B.Z Designed the experiments, theoretical models, and analysis. M.Y.B.Z and J.F performed the experiments and the data analysis. M.Y.B.Z, O.D, and N.B wrote the manuscript. **Competing interests:** The authors declare no competing interests. **Data and materials availability:** All (other) data needed to evaluate the conclusions in the article are presented in the article or the Supplementary Materials.

DATA AVAILABILITY

All data and code required for reproducing the results in the manuscript are found in online repositories [87].

SUPPLEMENTARY MATERIALS

Exoskeletons CAD files
Code Phototaxis
Movies S1 and S2
Materials and Methods
Supplementary Text
Figs. S1, S2, S3
Algorithm 1

-
- [1] G. Beni and J. Wang, in *Robots and biological systems: towards a new bionics?* (Springer, 1993) pp. 703–712.
- [2] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, *Swarm Intelligence* **7**, 1 (2013).
- [3] H. Hamann, *Swarm Robotics: A Formal Approach* (Springer International Publishing, Cham, 2018) p. 215.
- [4] R. A. Brooks and A. M. Flynn, *Fast, cheap and out of control*, Tech. Rep. (Massachusetts Inst of Tech Cambridge Artificial Intelligence Lab, 1989).
- [5] M. Dorigo, G. Theraulaz, and V. Trianni, *Science Robotics* **5** (2020).
- [6] D. Floreano and H. Lipson, From individual robots to robot societies (2021).
- [7] M. Dorigo, G. Theraulaz, and V. Trianni, *Proceedings of the IEEE* **109**, 1152 (2021).
- [8] J. Werfel, K. Petersen, and R. Nagpal, *Science* **343**, 754 (2014).
- [9] S. Hauert, J.-C. Zufferey, and D. Floreano, *Autonomous Robots* **26**, 21 (2008).
- [10] C. Virágh, G. Vásárhelyi, N. Tarcai, T. Szörényi, G. Somorjai, T. Nepusz, and T. Vicsek, *Bioinspiration & biomimetics* **9**, 025012 (2014).
- [11] G. Vásárhelyi, C. Virágh, G. Somorjai, T. Nepusz, A. E. Eiben, and T. Vicsek, *Science Robotics* **3** (2018).
- [12] K. McGuire, C. De Wagter, K. Tuyls, H. Kappen, and G. C. de Croon, *Science Robotics* **4** (2019).
- [13] J. Gomes, M. Duarte, P. Mariano, and A. L. Christensen, in *International Conference on Parallel Problem Solving from Nature* (Springer, 2016) pp. 591–601.
- [14] P. Zahadat and T. Schmickl, *Adaptive Behavior* **24**, 87 (2016).
- [15] F. Berlinger, M. Gauci, and R. Nagpal, *Science Robotics* **6** (2021).
- [16] M. Schranz, M. Umlauf, M. Sende, and W. Elmenreich, *Frontiers in Robotics and AI* **7**, 36 (2020).
- [17] D. Saldana, B. Gabrich, M. Whitzer, A. Prorok, M. F. Campos, M. Yim, and V. Kumar, *IEEE International Conference on Intelligent Robots and Systems* **2017-September**, 2736 (2017).
- [18] S. Li, R. Batra, D. Brown, H. D. Chang, N. Ranganathan, C. Hoberman, D. Rus, and H. Lipson, *Nature* **567**, 361 (2019).
- [19] W. Savoie, T. A. Berrueta, Z. Jackson, A. Pervan, R. Warkentin, S. Li, T. D. Murphey, K. Wiesenfeld, and D. I. Goldman, *Science robotics* **4**, 10.1126/scirobotics.aax4316 (2019).
- [20] X. Du, H. Cui, T. Xu, C. Huang, Y. Wang, Q. Zhao, Y. Xu, and X. Wu, *Advanced Functional Materials* **30**, 1 (2020).
- [21] S. Li, B. Dutta, S. Cannon, J. J. Daymude, R. Avinery, E. Aydin, A. W. Richa, D. I. Goldman, and D. Randall, *Science Advances* **7**, 1 (2021), arXiv:2009.05710.
- [22] Y. Ozkan-Aydin, D. I. Goldman, and M. Saad Bhamla, *Proceedings of the National Academy of Sciences of the United States of America* **118**, 10.1073/pnas.2010542118 (2021).
- [23] M. Rubenstein, A. Cornejo, and R. Nagpal, *Science* **345**, 795 (2014).
- [24] I. Slavkov, D. Carrillo-Zapata, N. Carranza, X. Diego, F. Jansson, J. Kaandorp, S. Hauert, and J. Sharpe, *Science Robotics* **3**, eaau9178 (2018).
- [25] G. Wang, T. V. Phan, S. Li, M. Wombacher, J. Qu, Y. Peng, G. Chen, D. I. Goldman, S. A. Levin, R. H. Austin, and L. Liu, *Physical Review Letters* **126**, 108002 (2021).
- [26] S. Li, R. Batra, D. Brown, H.-D. Chang, N. Ranganathan, C. Hoberman, D. Rus, and H. Lipson, *Nature* **567**, 361 (2019).
- [27] G. Oliveri, L. C. van Laake, C. Carissimo, C. Miette, and J. T. Overvelde, *Proceedings of the National Academy of Sciences* **118** (2021).
- [28] A. Deblais, T. Barois, T. Guerin, P. H. Delville, R. Vaudaine, J. S. Lintuvuori, J. F. Boudet, J. C. Baret, and H. Kellay, *Physical Review Letters* **120**, 188002 (2018).
- [29] J. F. Boudet, J. Lintuvuori, C. Lacouture, T. Barois, A. Deblais, K. Xie, S. Cassagnere, B. Tregon, D. B. Brückner, J. C. Baret, and H. Kellay, *Science Robotics* **6**, eabd0272 (2021).
- [30] H. C. Berg, *Physics Today* **53**, 24 (2000).
- [31] H. H. Wensink, J. Dunkel, S. Heidenreich, K. Drescher, R. E. Goldstein, H. Löwen, and J. M. Yeomans, *Proceedings of the National Academy of Sciences* **109**, 14308 (2012), arXiv:abs/1302.5277 [arxiv.org].
- [32] H. Xu, J. Dauparas, D. Das, E. Lauga, and Y. Wu, *Nature Communications* **10**, 1 (2019).
- [33] J. W. Costerton, P. S. Stewart, and E. P. Greenberg, *Science* **284**, 1318 (1999).
- [34] O. Shishkov, M. Hu, C. Johnson, and D. L. Hu, *Journal of The Royal Society Interface* **16**, 20180735 (2019).
- [35] A. Procaccini, A. Orlandi, A. Cavagna, I. Giardina, F. Zoratto, D. Santucci, F. Chiarotti, C. K. Hemelrijk, E. Alleva, G. Parisi, and C. Carere, *Animal Behaviour* **82**, 759 (2011).

- [36] BBC, Feeding frenzy, from the hunt: Hunger at sea (oceans) (2015).
- [37] M. Moussaïd, D. Helbing, and G. Theraulaz, Proceedings of the National Academy of Sciences **108**, 6884 (2011).
- [38] C. Bechinger, R. Di Leonardo, H. Löwen, C. Reichhardt, G. Volpe, and G. Volpe, Reviews of Modern Physics **88**, 045006 (2016), arXiv:1602.00081.
- [39] F. Hara and R. Pfeifer, eds., *Morpho-functional Machines: The New Species* (Springer Japan, 2003) pp. 88–100.
- [40] J. Deseigne, O. Dauchot, and H. Chaté, Physical review letters **105**, 098001 (2010).
- [41] T. Vicsek and A. Zafeiris, Physics Reports **517**, 71 (2012).
- [42] K.-D. N. T. Lam, M. Schindler, and O. Dauchot, New Journal of Physics **17**, 113056 (2015).
- [43] A. Bricard, J.-b. Caussin, N. Desreumaux, O. Dauchot, and D. Bartolo, Nature **503**, 95 (2013).
- [44] J. Boudet, J. Lintuvuori, C. Lacouture, T. Barois, A. Deblais, K. Xie, S. Cassagnere, B. Tregon, D. Brückner, J. Baret, *et al.*, Science Robotics **6** (2021).
- [45] R. Pfeifer and J. Bongard, *How the body shapes the way we think: a new view of intelligence* (MIT press, 2006).
- [46] R. Pfeifer and G. Gómez, in *Creating brain-like intelligence* (Springer, 2009) pp. 66–83.
- [47] R. A. Watson, S. G. Ficici, and J. B. Pollack, Robotics and Autonomous Systems **39**, 1 (2002).
- [48] N. Bredeche, E. Haasdijk, and A. Prieto, Frontiers in Robotics and AI **5**, 12 (2018).
- [49] N. Fontbonne, O. Dauchot, and N. Bredeche, in *2020 IEEE Congress on Evolutionary Computation (CEC)* (2020) pp. 1–8.
- [50] N. Bredeche and N. Fontbonne, Philosophical Transactions of the Royal Society B **377**, 20200309 (2022).
- [51] M. Rubenstein, C. Ahler, N. Hoff, A. Cabrera, and R. Nagpal, Robotics and Autonomous Systems 10.1016/j.robot.2013.08.006 (2013).
- [52] D. Allan, C. van der Wel, N. Keim, T. A. Caswell, D. Wieker, R. Verweij, C. Reid, Thierry, L. Grueter, K. Ramos, Apiszcz, Zoeith, R. W. Perry, F. Boulogne, P. Sinha, Pfigliozzi, N. Bruot, L. Uieda, J. Katins, H. Mary, and A. Ahmadi, soft-matter/trackpy: Trackpy v0.4.2 (2019).
- [53] É. Fodor and M. Cristina Marchetti, Physica A: Statistical Mechanics and its Applications **504**, 106 (2018), arXiv:1708.08652.
- [54] C. Scholz, S. D’Silva, and T. Pöschel, New Journal of Physics **18**, 10.1088/1367-2630/18/12/123001 (2016).
- [55] S. Collins, A. Ruina, R. Tedrake, and M. Wisse, Science **307**, 1082 (2005).
- [56] L. Gioni, N. Hawley-Weld, and L. Mahadevan, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **469**, 20120637 (2013), arXiv:1302.5952.
- [57] O. Dauchot and V. Démery, Physical Review Letters **122**, 1 (2019), arXiv:1810.13303.
- [58] P. Baconnier, D. Shohat, C. H. López, C. Coulais, V. Démery, G. Düring, and O. Dauchot, Nature Physics 10.1038/s41567-022-01704-x (2022), arXiv:2110.01516.
- [59] J. Tailleur and M. E. Cates, Physical Review Letters **100**, 218103 (2008), arXiv:abs/0803.1069 [arxiv.org].
- [60] S. Shahrokhi and A. T. Becker, IEEE International Conference on Automation Science and Engineering **2016-November**, 561 (2016).
- [61] D. Lees and A. Zilli, *Moths: A Complete Guide to Biology and Behavior* (Smithsonian Books, 2019) p. 208.
- [62] A. Berdahl, C. J. Torney, C. C. Ioannou, J. J. Faria, and I. D. Couzin, Science **339**, 574 (2013).
- [63] W. NULTSCH, Photochemistry and Photobiology **14**, 705 (1971).
- [64] E. Şahin, in *International workshop on swarm robotics* (Springer, 2004) pp. 10–20.
- [65] L. Bayindir, Neurocomputing **172**, 292 (2016).
- [66] W. Savoie, S. Cannon, J. J. Daymude, R. Warkentin, S. Li, A. W. Richa, D. Randall, and D. I. Goldman, Artificial Life and Robotics **23**, 459 (2018), arXiv:1711.01327.
- [67] H. Hamann, H. Worn, K. Crailsheim, and T. Schmickl, in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IEEE, 2008) pp. 1415–1420.
- [68] T. Schmickl, R. Thenius, C. Moeslinger, G. Radspieler, S. Kernbach, M. Szymanski, and K. Crailsheim, Autonomous Agents and Multi-Agent Systems **18**, 133 (2009).
- [69] S. Kernbach, R. Thenius, O. Kernbach, and T. Schmickl, Adaptive Behavior **17**, 237 (2009).
- [70] F. Arvin, K. Samsudin, A. R. Ramli, and M. Bekravi, International Journal of Computational Intelligence Systems **4**, 739 (2011).
- [71] T. Schmickl and H. Hamann, Bio-inspired computing and communication networks , 95 (2011).
- [72] M. V. Smoluchowski, Physik Zeitcher, **17**, 557 (1916).
- [73] K. Uckert, A. Parness, N. Chanover, E. J. Eshelman, N. Abcouwer, J. Nash, R. Detry, C. Fuller, D. Voelz, R. Hull, D. Flannery, R. Bhartia, K. S. Manatt, W. J. Abbey, and P. Boston, Astrobiology **20**, 1427 (2020).
- [74] I. Fiorello, E. Del Dottore, F. Tramacere, and B. Mazzolai, Bioinspiration and Biomimetics **15**, 10.1088/1748-3190/ab7416 (2020).
- [75] F. Berlinger, M. Gauci, and R. Nagpal, Science Robotics **6**, 1 (2021).
- [76] R. Solé, M. Moses, and S. Forrest, Philosophical Transactions of the Royal Society B: Biological Sciences **374**, 10.1098/rstb.2019.0040 (2019).
- [77] V. Trianni, R. Groß, T. H. Labella, E. Şahin, and M. Dorigo, in *European Conference on Artificial Life* (Springer, 2003) pp. 865–874.
- [78] M. Dorigo, V. Trianni, E. Şahin, R. Groß, T. H. Labella, G. Baldassarre, S. Nolfi, J.-L. Deneubourg, F. Mondada, D. Floreano, *et al.*, Autonomous Robots **17**, 223 (2004).
- [79] O. Soysal, E. Bahçeci, and E. Şahin, Turkish Journal of Electrical Engineering & Computer Sciences **15**, 199 (2007).
- [80] S. Nolfi and D. Floreano, *Evolutionary Robotics: the Biology, Intelligence, and Technology* (MIT Press, Cambridge MA, USA, 2000).
- [81] S. Doncieux, N. Bredeche, J.-B. Mouret, and A. Eiben, Frontiers in Robotics and AI **2**, 1 (2015).
- [82] L. P. Kaelbling, M. L. Littman, and A. W. Moore, Journal of artificial intelligence research **4**, 237 (1996).
- [83] O. Feinerman, I. Pinkoviezky, A. Gelblum, E. Fonio, and N. S. Gov, Nature Physics **14**, 683 (2018).
- [84] D. Tarapore, R. Groß, and K.-P. Zauner, Frontiers in Robotics and AI **7**, 83 (2020).
- [85] A. Whiten, Science **372** (2021).
- [86] M. Y. Ben Zion, Y. Caba, A. Modin, and P. M. Chaikin, Nature Communications **13**, 184 (2022), arXiv:2012.15087.

- [87] Matan Yah Ben Zion, Morphological computation and decentralized learning in a swarm of sterically interacting robots (2023).
- [88] A. D. Edelstein, M. A. Tsuchida, N. Amodaj, H. Pinkard, R. D. Vale, and N. Stuurman, *Journal of Biological Methods* **1**, e10 (2014).
- [89] C. T. Rueden, J. Schindelin, M. C. Hiner, B. E. DeZonia, A. E. Walter, E. T. Arena, and K. W. Eliceiri, *BMC Bioinformatics* **18**, 529 (2017).
- [90] N. Oppenheimer, D. B. Stein, and M. J. Shelley, *Physical Review Letters* **123**, 148101 (2019), arXiv:1903.00940.
- [91] N. Oppenheimer, D. B. Stein, M. Y. Ben Zion, and M. J. Shelley, *Nature Communications* **13**, 804 (2022), arXiv:2103.00296.
- [92] M. Y. Ben Zion, A. Modin, and P. M. Chaikin, *Arxiv cond-mat.soft* doi.org/10.48550/arXiv.2203.11051 (2022), arXiv:2203.11051.
- [93] M. Y. Ben Zion, X. He, C. C. Maass, R. Sha, N. C. Seeman, and P. M. Chaikin, *Science* **358**, 633 (2017).
- [94] G. Zhu, M. Hannel, R. Sha, F. Zhou, M. Y. Ben Zion, Y. Zhang, K. Bishop, D. Grier, N. Seeman, and P. Chaikin, *Proceedings of the National Academy of Sciences* **118**, e2023508118 (2021).
- [95] J. R. Howse, R. A. L. Jones, A. J. Ryan, T. Gough, R. Vafabakhsh, and R. Golestanian, *Physical Review Letters* **99**, 048102 (2007), arXiv:0706.4406.
- [96] M. Y. Ben Zion, Y. Caba, R. Sha, N. C. Seeman, and P. M. Chaikin, *Soft Matter* **16**, 4358 (2020).
- [97] M. Yoonis, M. Y. Ben Zion, O. Dauchot, and N. Bredeche, in *GECCO 2022 - Proceedings of the 2022 Genetic and Evolutionary Computation Conference* (Association for Computing Machinery, 2022) pp. 104–112.
- [98] M. Y. Ben Zion, N. Bredeche, and O. Dauchot, *Arxiv cond-mat.soft*, 1 (2021), arXiv:2111.06953.
- [99] M. J. Schnitzer, *Physical Review E* **48**, 2553 (1993).

MATERIALS AND METHODS

Exoskeleton design and manufacturing

Exoskeleton Manufacturing

Exoskeletons were 3D printed using either one of the following 3D printers and materials to give similar results: *startAsys Objet350 Connex3* printer, using *vero-Clear* (modulus of elasticity $E = 1 - 2 \text{ GPa}$, density $\rho = 1.2 \text{ g/cm}^3$), *3D systems projet 2500plus* printer using *VisiJet M2R-CL (MJP)* (density $\rho = 1.14 \text{ g/cm}^3$, Elastic modulus $E \approx 1 \text{ GPa}$), *Prusa i3* printer using *PLA* (density $\rho = 1.2 \text{ g/cm}^3$, elastic modulus $E = 5 \text{ GPa}$).

Data acquisition and experimental setups

Phototaxis experiments

Phototaxis experiments were performed on a 5 mm translucent plexiglass sheet placed on the floor (PMMA 27000 Diffusant Laitoux 40%, WEBER METAUX). A

circular region 6% of the arena was lit using a projector (EPSON EB-1795F), at RGB = (10,255,255), at approximately the spectral sensitivity peak (570 nm) of the Kilo-bot’s light sensor (TEPT5700, Vishay Semiconductors). To reduce interference of ambient light with the light sensor, the arena was lit using four desk lamps directed at the surrounding walls (for a homogeneous diffusive illumination) and covered with red cellophane (where the robot’s light sensor’s sensitivity is at a minimum). Red cellophane also covered the camera used for acquisition (PIXELINK.PL-D734MU), to reduce the saturation from the lit region, and imaging of the robots. Image acquisition was carried out using *micromanager* [88], at 2 frames per second. Robots are counted when fully in the light and results shown in Figs. 6F, 7G are the average of four realizations with error bars being the standard deviation.

Code used

Robots were programmed to phototaxis using variations of the following two codes: 1. **phototaxisAdhoc.c** — robots stand when their photoperceptor is activated given a predefined threshold. 2. **phototaxisLearning.c** — Robot stand when photoperceptor is activated but light threshold are randomly initialized and evolve through the learning algorithm.

Inclined plane experiments

Coupling between external force and robot’s orientation was measured by letting robots run on a plane made of 5 mm thick plexiglass placed at an incline, and imaged using SonyAlphaS7 camera, with video-rate image acquisition at 30 frames per second. See supporting Movies 3-6.

Data analysis

Raw images were preprocess using ImageJ [89], or FFMpeg, followed by particle locating using trackpy package [52], or custom code, and trajectories were linked using trackpy [52].

Simulating Morphobots

Computer simulations were carried out by adding a stochastic term to the Eqs. 1, 2 and using 5th order Runge-Kutta for the time propagation of Brownian dynamics of active soft-discs by adopting code from previous publications [90–92].

Custom Morphobots

Design files and bill of materials for custom Morphobots (Fig. 4 in the main text) are found in the extended data. At their core, custom Morphobots are made from a 3D-printed flat chassis onto which a battery (LIR 2477), an electric switch, and two vibration motors (Best-Tong 14kRPM 3 Volts) are mounted. The chassis has two holes for mounting flexible curved legs (which can be 3D printed separately).

Morphological computation and decentralized learning in a swarm of sterically interacting robots — Supplementary Materials*

Matan Yah Ben Zion^{1,2,3,†}, Jeremy Fersula^{1,2}, Nicolas Bredeche², and Olivier Dauchot¹

¹*Gulliver UMR CNRS 7083*

ESPCI, PSL Research University, 75005 Paris, France

²*Institut des Systèmes Intelligents et de Robotique*

Sorbonne Université, CNRS, ISIR, F-75005 Paris, France and

³*School of Physics and Astronomy, and the Center for Physics and Chemistry of Living Systems
Tel Aviv University, Tel Aviv 6997801, Israel*

Exoskeleton design

Exoskeletons were designed as a round chassis (diameter $d = 4.8$ cm) with an off-center cavity where the Kilobot is pressed in. The exoskeleton stands on three, $L = 2$ cm tall legs — one leg is stiff and round, and two opposing legs that are flexible and with a rectangular cross-section (thickness $T = 1$ mm, and width $W = 0.6$ cm, see Fig. 2 in the main text as well as previous work on elastic beams [93, 94]). The flexible legs were designed to have their first natural vibration near resonance with the vibration motors frequency (Pololu 10X2.0 mm vibration motors, $f_{motor} \approx 250$ Hz). This was done by estimating the legs to be elastic beams which are described dynamically using the Euler–Bernoulli beam theory for a homogeneous thin beam

$$EI \frac{\partial^4 h}{\partial x^4} = -\mu \frac{\partial^2 h}{\partial t^2}, \quad (7)$$

where h is the deflection, E is the elastic modulus, μ is the linear mass density (given by $\mu = TW\rho$), and I is the beam moment of inertia (given by $I = \frac{W \times T^3}{12}$). In Fourier space, Eq. 7 becomes

$$\omega = \sqrt{\frac{EI}{\mu}} k^2, \quad (8)$$

where ω is the angular frequency, and k the wave number. The wavelength of the first harmonic λ_1 of a beam of length L fixed at one end is $\lambda_1 = 4L$, along with Eq. 8 allows us to find that the natural first harmonic of the leg is to leading order

$$f_1 = \frac{\pi}{16\sqrt{6}} \sqrt{\frac{E}{\rho}} \frac{T}{L^2}. \quad (9)$$

* This manuscript has been accepted for publication in Science Robotics. This version has not undergone final editing. Please refer to the complete version of record at <https://www.science.org/doi/10.1126/scirobotics.abo6140>. The manuscript may not be reproduced or used in any manner that does not fall within the fair use provisions of the Copyright Act without the prior, written permission of AAAS.

† To whom correspondence should be addressed.
E-mail: matanbz@gmail.com

Note that Eq. 9 is independent of the width. Given the materials and geometry used (see the previous section), $f_1 \approx 250$ Hz $\approx f_{motor}$.

Individual robot dynamics

Speed of an individual persistent particle

The effective diffusion constant of a run-and-tumble particle, that covers a distance l_p , during the time interval τ_{RT} is $D_{eff} \approx \frac{1}{4} l_p^2 / \tau_{RT}$ [95]. In the run and tumble mode, each Morphobot is programmed to run for $\tau_{run} = 2$ s at its nominal speed of $v_0 \approx 5$ cm/s, then perform a tumble where it pivots on average for $\tau_{pivot} = 4$ s. The step size is thus $l_p = v \cdot \tau_{run} \approx 10$ cm, and the total duration is $\tau_{RT} = 6$ s. The resulting effective diffusion constant is $D_{eff} \approx \frac{1}{4} (10 \text{ cm})^2 / 6 \text{ s} = 4 \text{ cm}^2/\text{s}$.

Particle orientation as a function of trajectory length under constant force

The dynamics of the orientation, θ , relative to the positive \hat{x} direction, for an active particle in the over-damped limit in the absence of noise is found to be that of a simple overdamped pendulum (Eq. 3 in the main text) [96]:

$$\dot{\theta} = -\kappa \mu f \sin \theta. \quad (10)$$

It is helpful to formulate Eq. 10 as a function of s , the trajectory traveled. In 2D, s is defined locally by the Euclidean distance:

$$ds^2 = dx^2 + dy^2. \quad (11)$$

Dividing Eq. 11 by dt^2 reminds us that

$$\left(\frac{ds}{dt}\right)^2 = \left(\frac{dx}{dt}\right)^2 + \left(\frac{dy}{dt}\right)^2 = v^2 \quad (12)$$

We can find v^2 by multiplying Eq. 1 (main text) by itself:

$$v^2 = v_0^2 + (\mu f)^2 + v_0 \mu f \cos \theta. \quad (13)$$

Using the chain rule, $\frac{d\theta}{dt} = \frac{d\theta}{ds} \frac{ds}{dt}$, and plugging into Eqs. 12 and 13, we get an ODE describing θ as a function of s , the arclength:

$$\frac{d\theta}{ds} = \frac{\kappa f \cos\theta}{\sqrt{v_0^2 + f^2 + v_0 f \cos\theta}}. \quad (14)$$

Equation 14 can be solved exactly, however, it is more instructive to examine the limit of $\frac{v_0}{\mu f} \ll 1$ which then simplifies to

$$\frac{d\theta}{ds} = \kappa \cos\theta. \quad (15)$$

For $\theta(s=0) = \frac{\pi}{2}$, Eq. 15 gives Eq. 4 in the main text:

$$\theta(s) = 2\text{atan}(e^{-\kappa s}). \quad (16)$$

Phototaxis without learning

Phototaxis Arena




Fig. S1. Phototaxis arena.

Photoperceptron

Each robot has a photoperceptron encoded to measure environmental cues, and take a specific action given whether the perceptron fired or not. The minimal architecture for the photoperceptron implemented has a neural network of a single neuron (perceptron), with two inputs, bias $b = x_0$, and the locally measure light power $p = x_1$, and their associated weights w_0 and w_1 , which together effectively define the threshold for firing. The perceptron output is p , a logistic function of the form

$$p = \frac{1}{1 + e^{-h}}. \quad (17)$$

where $h = \sum_{i=0}^1 w_i x_i$

Sense-act cycle

The sense-act cycle uses the robot's local light intensity measurement, where the robot chooses between two meta-behaviors: *run* and tumble or *walk* and tumble, given the photoperceptron is below threshold $P < 0.5$, or above threshold $P \geq 0.5$, respectively. In both cases, the motion is a sequence of a moving phase (run/walk) and a tumble phase. In the tumble phase, the robot randomly chooses a direction (clockwise or counterclockwise) and a random duration with a mean of 6 seconds. Randomizing the duration compensates for the robots' internal dynamical variability and environmental variability. For example, robots with stronger/weaker motors will take less/more time to tumble and complete the same reorientation. Similarly, a robot near a wall (or another robot) or in solitude. When in *run* mode, the robot sets both motor values to high, approximately moving at a straight line (persistence length much larger than the robot's size), at its nominal speed $V_0 = v_0$. When walking, the robot moves for a duration τ_{walk} , and then stands for a duration τ_{stand} . The ratio between the two controls the mean speed of the robot when walking $V_1 = \frac{\tau_{\text{walk}}}{\tau_{\text{walk}} + \tau_{\text{stand}}} V_0$. When $\tau_{\text{walk}} = 0$ the robot stands when the photoperceptron fires.

Phototaxis rate for low concentration noninteracting robots estimated using first arrival time

The rate of robot arrival to the light spot at early times (where the lit region is mostly empty), can be found from the reaction rate constant, k , using Smoluchowski first arrival given a number concentration c . In 2D the mean first arrival time, $\tau_{FA} = 1/kc$, of a Brownian particles with diffusion constant D , arriving at a target of size b in an arena of diameter d is given by

$$k = \frac{2\pi D}{\log \frac{d}{b}}. \quad (18)$$

The rate at which robots are subtracted from the dark region $\frac{dN_{\text{bulk}}}{dt}$, is given from a first-order reaction rate

$$\frac{dN_{\text{bulk}}}{dt} = -kN_{\text{bulk}}, \quad (19)$$

which solves readily to $N_{\text{bulk}} = A\exp\{-kt\}$. Given an initial number of robots, N_0 , the number of robots in the light, N , is given from the robot conservation $N_0 = N + N_{\text{bulk}}$,

$$N(t) = N_0(1 - e^{-kt}), \quad (20)$$

which is plotted in Fig. 6.

In-silico results of phototactic swarms

The spatial distribution of a phototactic swarm of aligners and frontiers becomes increasingly distinct with increasing swarm size. Figure S2 shows snapshots from simulations of the ad-hoc calibrated phototactic swarm. Showing that frontiers fill the lit region, while aligners develop a ring at the perimeter of the light spot, and that the effect grows with swarm size. The effect of morphology becomes more pronounced with growing swarm size also for learning swarms as can be seen in Fig. S3.




Fig. S2. Snapshots from the phototaxis process of ad-hoc calibrated agents in swarms of size 64, 128, 256, 512 (left to right) of frontiers (top) and aligners (bottom). Aligners tend to stop at the perimeter and leave a hollow while swarms of frontiers display a more cooperative phototaxis where robots push already taxied robots deeper into the lit region




Fig. S3. Snapshots from the learning process of simulated phototactic swarms of size 64, 128, 256, 512 (left to right) of frontier (top) and aligners (bottom). Swarms of aligners are arranged more loosely in the lit region relative to the tightly packed swarms of frontiers.

Collective learning in a sterically interacting swarm

In this final section, we derive an analytic link between the behavior of individuals in a swarm and the

convergence of a decentralized learning scheme. We first establish empirically that a swarm where robots are allowed to collide has good mixing, a property required to establish individual-to-collective correspondence. We then cast collective learning as a dynamical process where robots in the swarm can exchange their state upon encounter. And finally, we show that given the reward function, ρ , as defined in the main text (Eq. 6), a randomly initialized swarm will converge on a successful phototaxis strategy.

Individual dynamics of pre-calibrated robots

The mean time a robot spends in the light spot, Φ , depends on three quantities: 1. The ratio between the speed of the robot in the light V_1 to its speed outside the light, V_0 ; 2. The area fraction of the light spot, σ ; 3. Whether or not the robot has the correct threshold for phototaxis. Allowing individual, pre-set robots in the light patterned arena to perform phototaxis shows two distinct behaviors: robots with a correct light threshold for phototaxis spend more time in the lit region than robots with a wrong threshold (either too high or too low, see Fig. S4B) for which time in the light is as good as random $\Phi_W \approx \sigma$.

Collective dynamics of a pre-calibrated swarm

To test how the speed ratio V_1/V_0 affects the number of bots in the light, phototaxis experiments were done by increasing the size of the lit region to 25% of the arena. Experiments were performed with either 64 bots or 32 bots (see inset in Fig. S5B). The speed ratio was set by controlling the τ_{stand} as described in Section , changing the speed from standing in light $V_1/V_0 = 0$ to effectively being agnostic to the light $V_1/V_0 = 1$ (see Fig. S5B). The relative number of robots in the light was tracked over up to 2 hours, and the mean and standard deviation of the last 10% of the experiment duration were used to evaluate the steady-state fraction of the swarm in the light, \mathcal{F}_{ss} .

Collective decentralized learning

The collective learning was implemented by deploying a randomly initialized swarm of Morphobots that follow the algorithm detailed below (see Algorithm 1) to give the results presented in Fig.7 in the main text.

Each robot in the population runs Algorithm 1. This Algorithm builds from the HIT algorithm originally presented in [49, 97], with two important modifications. Firstly, the algorithm is designed with phototaxis as an objective (the ρ objective function on line 9 measures the




Fig. S4. Dynamics and performance of individual pre-programmed robots. **A** Mean square displacement of the two meta behaviors: the fast run-and-tumble (pink), or the much slower walk-and-tumble (red), with a fixed speed ratio between run and walk. The run-and-tumble behavior transitions from ballistic ($\propto t^2$) to linear ($\propto t^1$) at approximately one persistent time $\tau \equiv t/\tau_p \approx 1$. Inset shows typical, 20 minutes long trajectories of both run-and-tumble (pink) and walk-and-tumble (red) inside the circular 150 cm wide arena. **B** The time fraction in the light Φ of an individual robot shows that a robot with too high or too low light threshold performs as well as random, while a robot that has a correct policy, is responsive to the light source. The policy of each robot is determined by the weights of the photo-perceptron, setting the light threshold to activate the photo-perceptron (solid curves).




Fig. S5. Collective performances of pre-programmed robots. **A** Swarm fraction in the light for a fixed speeds ratio ($V_1/V_0 = 1/15$) different initializations: correctly calibrated swarm (black), low (red) high (pink) and randomly initialized swarm (gray). Insets show typical trajectories showing both low and high are agnostic to the light spot, while the correctly, ad-hoc, initialized robot, dwells longer in the light spot, increasing its reward ρ . **B** The run-to-walk speed ratios determine the collective ability of a swarm to phototaxis, and a steady-state flux balance model (see Eq. 27) quantitatively predicts the percentage of the swarm in the light, with no fitting parameters.

amount of light measured in the last T time steps). Secondly, transfer of **all** control parameters is performed, which implies that whenever a robot receives a parameter set that is deemed to be performing better than its own, this parameter set is fully copied and overwrites existing parameter values. Mutation is not used in this implementation, as the number of robots used (64) and the low dimensionality of the search space (\mathbb{R}^2) allows for a sufficient amount of behavioral diversity in the initial

population of robots (see [49] for a comprehensive study of transfer and mutation rate and operators).

It should also be noted that the performance assessment (ρ^i) is not reset to zero after a parameter update, which implies that the performance assessment of a newly updated policy is under-estimated for T time steps. While this may temporarily lead to theoretically best-performing control parameters being overwritten before T steps elapsed, the elitist selection scheme ensures that

Algorithm 1: The phototactic-HIT algorithm

Data:
 i : the unique identifier of the current robot,
 T : evaluation time,
 π : Policy function,
 \mathbf{w}^i : Random uniform initialisation of policy parameters
 $R[T]$: Empty reward buffer of size T ,
 r : Current light intensity,
 ρ^i : performance self-assessment during T ,
 \mathbf{a} : Null action vector,
 \mathbf{o} : Null observation vector

```

1 begin
2    $t = 0$ 
3   while True do
4      $\mathbf{o}, r = \text{sense}()$ 
5      $R[t \bmod T] = r$ 
6      $\mathbf{a} = \pi(\mathbf{o}|\mathbf{w}^i)$ 
7     act( $\mathbf{a}$ )
8     if  $t > T$  then
9        $\rho^i = \sum_{k=0}^{T-1} R[k]$ 
10      broadcast( $\mathbf{w}^i, \rho^i$ )
11      if new_message then
12         $\mathbf{w}^j, \rho^j = \text{decode\_message}$ 
13        if  $\rho^j > \rho^i$  then
14           $\mathbf{w}^i = \mathbf{w}^j$ 
15        end
16      end
17    end
18     $t = t + 1$ 
19  end
20 end

```

the original best-performing individual is never lost.

Finally, note that in the algorithm the observable o contains information about the current light intensity r . They are separated for semantic reasons as o is used as an input value for the policy π , while r is used to update the performance self-assessment ρ^i .

Collective learning as a dynamical process

The goal is to link policy and environment to find a route for a successful decentralized learning scheme on average. The time evolution of the probability p_C^i , respectively p_W^i , of finding a robot i with a correct, respectively wrong, policy is described by the following master equation :

$$\frac{dp_C^i}{dt} = \omega_{C,W}^i p_W^i - \omega_{W,C}^i p_C^i, \quad (21)$$

where $p_W^i = 1 - p_C^i$ and $\omega_{C,W}^i$, respectively $\omega_{W,C}^i$, are the transition rates at which the robot i with a wrong policy turns into a correct policy and conversely. Defining the concentration of robots with a correct policy as c_C , the concentration of robots with a wrong policy is given by

$c_W = c_0 - c_C$ (where c_0 is the constant overall concentration of the swarm). In a system with good mixing, the encounter rate between robots with correct and wrong policy is proportional to the product of their concentrations, $k c_C c_W$ (to leading order), where k is proportional to mutual diffusion. When encountered, robots compare their rewards and a robot with a wrong policy will adopt the correct policy at a probability f_{CW} . The opposite is also possible (false information scenario) and a robot with a correct policy will adopt a wrong policy at a rate f_{WC} . When combined, the rate equation for the growth in the population with a correct policy becomes

$$\frac{dc_C}{dt} = k (f_{CW} - f_{WC}) c_C c_W. \quad (22)$$

Collective learning is achieved when the concentration of robots with correct policy grows, implying: $dc_C/dt > 0$. Since k , c_W , and c_C are all non-negative, collective learning is guaranteed when $f_{CW} > f_{WC}$. Given the definition of the reward function in Eq. 6 and the selection protocol defined in Fig. 7, robots with the correct threshold will have, on average, a higher reward, $\rho_C > \rho_W$, and the population with the correct policies will prevail (see SM for a full derivation).

We assume space is homogeneous and identify the probability of a given policy for a given robot with the concentration of robots, $c_{C,W}$, having that policy. Since transition between states happens only through local communication, the transition rates are proportional to the robot encounter rate constant, the latter being independent of the policies of the encountering robots. Under the above hypothesis, the transition rates are proportional to the concentrations and the transition elements read $\omega_{C,W} \approx k f_{CW} c_C$, and symmetrically $\omega_{W,C} \approx k f_{WC} c_W$, where k is the encounter rate, and f_{CW} , respectively f_{WC} , is the transfer probability of W , respectively C , becoming C , respectively W , upon encounter. As a result,

$$\frac{dc_C}{dt} = k (f_{CW} - f_{WC}) c_C (c_0 - c_C), \quad (23)$$

where c_0 is the overall concentration of robots in the arena.

To achieve positive mean collective learning, $dc_C/dt > 0$, the factors of the product in Eq. 23 has to be non negative, which amounts to

$$f_{WC} > f_{CW}$$

As learning is decentralized, f is an embedded function and can depend only on information locally accessible to the robot. This includes the robot's policy (here, the weights of its neural network, $\{w_i\}$), its measurement history, and the message received. Here we exclude a direct functional dependence of f upon the policy as it

will artificially bias the learning to a given trait (say turn left) instead of the desired fitness (say measure light).¹ Therefore f will depend on some function of the history of its own sensor input and the history of an encountered robot. In our minimal system, the measurement history contains the M most recent light intensity measurements, $I_{i \in 1..M}$ (stored as a rolling buffer, i.e FIFO).² Since the robots have a finite bandwidth, to allow asynchronous performance, where communication does not slow down dynamics, the message should be kept succinct. This means robots should only exchange a reduced version of their memory that can be contained within a message.³ This will be a function of the memory and will be called the reward function, $\rho = \rho(\{I_i\})$.⁴ In our implementation scheme, robots continuously broadcast a message with their reward and policy; a robot receiving a message compares the Broadcasted reward, ρ_B , to its self-reward, ρ_S , and if $\rho_B > \rho_S$ the policy will be *deterministically* adopted. The question then becomes what functional form does the local ρ take, to satisfy global positive learning *on average*: $\langle \rho_C \rangle > \langle \rho_W \rangle$.

Empirically we found that in an ad-hoc calibrated swarm at steady state, individual robots with a correct policy spend more time in the light $\Phi_C > \Phi_W$ (see Fig.S4), where the mean time in the light of an individual robot is measured by $\Phi = \frac{1}{T} \int_{-T}^0 dt I(t)$. Identifying the reward function as proportional to the mean time in the light, $\rho \propto \Phi$ allows us to bias the swarm towards learning the correct policy. We do this by computing the reward as the average of the light intensity buffer, $\rho = \frac{1}{M} \sum_{i=1}^M I_i$ (Eq. 6 in the main text).

Recalling that the transfer probability function depends only on the received and embedded reward, the condition for learning simply becomes that on average, correct robots have a greater reward:

$$\rho_C > \rho_W \quad (24)$$

Mean reward in a well-mixed swarm

To satisfy the learning condition presented in the previous section (Eq. 24) we recall the good mixing property found empirically in previous sections wherein the locally computed reward can approximate an individual's time average, which is a proxy to the global performance $\rho \propto \Phi \approx \mathcal{F}$ [98]. A swarm with good mixing can be modeled as a fluid that obeys the continuity equation:

³ On the Kilobot platform, the message size is 8 bytes.

⁴ In our case, ρ will be a single byte.

$$\frac{dC}{dt} = -\nabla J \quad (25)$$

where C is concentration and J is the flux. At steady-state, the concentration does not change in time, and so the left-hand side of Eq. 25 vanishes $\frac{dC}{dt} = 0$. In a system made of a collective with spatially varying speed, it has been shown that concentrations are inversely proportional to speeds (see Schnitzer who studied field taxi in bacteria [99]). For the simple case discussed here, space is divided into two regions, lit and dark, where robots have speeds V_1 and V_0 respectively. This results in the following relation of concentration of robots in the light given their speed in the light,

$$C_0 V_0 = C_1 V_1, \quad (26)$$

where C_1 and C_0 are the concentrations of the robots in the light and in the dark respectively. We next derive a link between Eq. 26 and the definition of the swarm fraction in the light, \mathcal{F} as defined in Eq. 27 the main text. The concentration of robots in the light, C_1 is the number of robots in the light over the area of the lit region $C_1 = N_1/\sigma A$. The concentration of robots in the dark is the remaining number of robots in the swarm, $N_0 = N - N_1$ over the remaining area of the arena: $N_0/(1-\sigma)A$. Plugging these into Eq. 26 and rearranging gives,

$$\mathcal{F}_{ss} = \frac{1}{\frac{V_1}{V_0} \left(\frac{1}{\sigma} - 1 \right) + 1}, \quad (27)$$

connecting the collective performance (fraction of the swarm in the light), to individual policies (speed ratio), given the environmental condition (fraction of the arena that is lit). For the case of an agnostic swarm, $V_0 = V_1$, the swarm does not perform better than random $\mathcal{F} \approx \sigma$ (consistent with measured performance of individuals Fig. S4 and collectives Fig. S5). When the robots are responsive to the light, $V_1 < V_0$, they accumulate in the light spot, and both the individual dynamics (Fig. S4B) and the collective dynamics (Fig. S5B), are also consistent with the above derivation with Eq. 27 setting the high bound for the swarm's performance. When robots are responsive, the swarm accumulates in the light, and individual robots spend more time in the light. Since the reward function as defined in Eq. 6 is proportional to the time a robot spends in the light $\rho \propto \Phi$, robots that slow down in the light will satisfy the learning condition (Eq. 24), and their superior policy will spread throughout the swarm.

¹ Note that using the policy directly could also be useful if two robots have similar rewards but one has a superior policy (say more energetically efficient). This can also be directly introduced into the reward.

² In our case I is an array of bytes and $M = 250$.