Intrinsically motivated collective motion

Henry J. Charlesworth[a] and Matthew S. Turner[a,b,1]

[a]Centre for Complexity Science, University of Warwick, Coventry CV4 7AL, United Kingdom; and [b]Department of Physics, University of Warwick, Coventry CV4 7AL, United Kingdom

Edited by David A. Weitz, Harvard University, Cambridge, MA, and approved June 17, 2019 (received for review January 1, 2019)

Collective motion is found in various animal systems, active suspensions, and robotic or virtual agents. This is often understood by using high-level models that directly encode selected empirical features, such as coalignment and cohesion. Can these features be shown to emerge from an underlying, low-level principle? We find that they emerge naturally under future state maximization (FSM). Here, agents perceive a visual representation of the world around them, such as might be recorded on a simple retina, and then move to maximize the number of different visual environments that they expect to be able to access in the future. Such a control principle may confer evolutionary fitness in an uncertain world by enabling agents to deal with a wide variety of future scenarios. The collective dynamics that spontaneously emerge under FSM resemble animal systems in several qualitative aspects, including cohesion, coalignment, and collision suppression, none of which are explicitly encoded in the model. A multilayered neural network trained on simulated trajectories is shown to represent a heuristic mimicking FSM. Similar levels of reasoning would seem to be accessible under animal cognition, demonstrating a possible route to the emergence of collective motion in social animals directly from the control principle underlying FSM. Such models may also be good candidates for encoding into possible future realizations of artificial "intelligent" matter, able to sense light, process information, and move.

collective motion | intelligent matter | active matter

There have been notable recent advances in our understanding of collective motion motivated by thermodynamics or physical arguments (1–8) and in animal systems (9–13). While generalized hydrodynamic theories (6–8) can be obtained for certain active physical systems, the collective motion of agents capable of information processing can be far more complex. For example, existing generalized hydrodynamic theories do not account for long-ranged interactions, such as those that are likely to arise in higher animals that rely on vision. Agent-based models have been developed that incorporate some of these potential complexities, e.g., distance-dependent attraction, orientation, or repulsions (5, 14, 15) or those relating more directly to vision (16, 17). While these models have had some success in explaining animal data, the starting point is usually an essentially empirical model. This leads to challenges, both in controlling against overfitting and providing low-level explanatory power: "Why and how do agents co-align or remain in cohesive groups?" This question is difficult to answer if the model has coalignment and cohesion hard-wired into it for essentially empirical reasons.

We instead analyze an agent-based system in which each agent senses, and then processes, information in the context of a predictive model of the future. It uses this model to determine its action in the present, recomputing its model of the future from scratch at each discrete time step. Each agent decides how to move, according to a low-level motivational principle that we call *future state maximization* (FSM): It seeks control in the sense that it maximizes the variety of (visual) environments that an agent could access before some time horizon, τ time steps into the future. This is a form of control, as it gives the agent many future options in a potentially uncertain world.

As we report below, FSM spontaneously generates collective motion of a sort that is similar to that observed in animal systems—i.e., moving, cohesive, highly aligned swarms that are stable against small perturbations; Movie S1. While there are even quantitative similarities with the structure and order in flocks of birds (9), the motivation for our work is not to mimic a particular animal system, but rather to analyze a simple, low-level model that may provide a general conceptual basis for collective motion, here based on vision. Crucially, our model does not explicitly include coalignment, cohesion, or any other physical interaction, merely mutual visual perception between agents in infinite (2D) space.

There are several reasons why motivational principles like FSM, that loosely serve to keep options open, may confer fitness, either in artificial intelligence or in nature. FSM increases the control that an agent has over their future. Agents that have many options to reposition themselves relative to their neighbors—e.g., in response to the arrival of a predator—can likely better avoid or confuse that predator.

In general, strategies like FSM that preserve an agent's freedom to reach many different outcomes in an uncertain world are expected to enhance fitness.

Similar strategies are known to be successful in games like chess. Having access to many viable future lines of development is generically preferable, given uncertainty about how the game will actually develop. This confers robustness in defense and strategic maneuverability in attack. Chess players are familiar with the feeling of their options becoming progressively more limited as they lose a game, with the converse being strongly characteristic of winning. One attempt at formalizing this kind of principle is the "empowerment" framework, which does so by using the language of information theory (18–20). Our implementation probably has the most in common with this strand of the literature. FSM is an example of an *intrinsic motivation*

Significance

Our study invokes a low-level principle that we believe might motivate animal behavior in general and collective motion in particular—the principle that an agent seeks to increase the number of states that it is able to encounter in the future. This principle should confer evolutionary fitness for rather general reasons that we discuss. We report how the collective motion that emerges from this principle is similar to that seen in animal systems. In particular, coalignment, cohesion, and collision avoidance all emerge naturally, even though none of these are encoded in the principle itself. In this sense, our work proposes a low-level origin for the emergence of collective motion in animal systems.

Author contributions: H.J.C. and M.S.T. designed research; H.J.C. performed research; H.J.C. and M.S.T. analyzed data; and H.J.C. and M.S.T. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence may be addressed. Email: m.s.turner@warwick.ac.uk.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1822069116/-/DCSupplemental.

(21, 22), where an incentive for behavior is provided, even in the absence of any specific tasks to complete or immediate external rewards to be gained. Intrinsically motivated behavior has been studied extensively in the psychology literature (23, 24), as well as more recently in the field of reinforcement learning (25–27), where it is used to aid exploration in environments where rewards are sparse. The key principle is that such behavior should offer a generic and universal benefit to the agent, not because it is useful for solving any one particular problem, but because it is beneficial for a wide range of scenarios that the agent may encounter in the future. A similar idea arises in the analysis of (hypothetical) causal entropic forces (28, 29). These forces generate motion that increases an entropy-like measure of all paths into the future and can lead to behavior with features usually thought to be characteristic of intelligence, including evidence for the spontaneous emergence of tool use and social cooperation.

Other work on decision making has some similarities with FSM (30). In that study, a formalism similar to ref. 28 was used to model agents making a group-level decision: Agents reach a consensus on a (single) decision, made in the same sensory context for all agents, without those agents perceiving states (more than a single step) in the future. No explicit dynamical model was defined or analyzed in ref. 30. In the present work, FSM is applied to a group of agents that can move, perceive their own distinct environments, and build independent models of future states that are accessible to them, guiding their decision making. This leads to the emergence of rich collective dynamics of a kind not previously realized.

Our work can also be seen as motivating the development of artificial particles that can sense, compute, and move; so-called "intelligent matter." This is a natural direction in which to develop existing active systems—e.g., phoretic colloids (31), swimming cells (32), or active biological suspensions (33) that have limited, and rigid, information processing capabilities. Having candidate algorithms to encode into this intelligent matter will help motivate its development. Heuristics that mimic FSM, as discussed below, may represent a particularly powerful choice for such algorithms.

FSM Applied to Collective Motion

Methods. We use deterministic computer simulation to study the motion of agents executing FSM. These agents are unit radius, phantom (i.e., able to overlap without repulsion), circular disks that are free to move on an infinite 2D plane. Their speed is the distance moved in each unit time step, with all lengths measured in disk radius units. Fig 1*A* shows the movement options available to each agent at each time step. These options are taken relative to its direction of motion in the previous time step. They are, in order: continue in the same direction with a choice of three different speeds, v_0 (nominal), $v_0 - \Delta v$ (slow), or $v_0 + \Delta v$ (fast). Alternatively, they are able to turn left or right by a small angle $\Delta\theta$, with speed v_0. Unless specified otherwise, in what follows, the nominal speed $v_0 = 10$, the speed variation $\Delta v = 2$, and the angular rotation $\Delta\theta = 15°$. At each time step, the agent must choose one of these five actions $z \in \{z_1, z_2, z_3, z_4, z_5\}$ and does so by executing a form of FSM, as described below.

Fig 1*B* shows how each visual state is constructed by using the positions of the agents.

This visual state is constructed for each agent by geometrically projecting all $N-1$ other disks down onto its center. This involves constructing pairs of lines that each pass through the focal disk's center and are tangent to both sides of the other disks. Each of these lines can be specified by an angle, measured relative to the agent's direction of motion. This allows us to define angular regions in which a line of sight will intersect with one or more other disks, shown as solid blue regions on the perimeter of Fig. 1*B*. We construct n_s discrete visual sensors that each relate to an angular region of size $2\pi/n_s$. The radial dotted lines in Fig. 1*B* denote the angular sensors (not the tangent lines). Each sensor registers 1 if more than half occupied by angles along which a line of sight will intersect other disk(s), i.e., the fraction of solid blue; 0 otherwise. Unless stated otherwise, $n_s = 40$ in all simulations. Fig. 1*C* shows how each agent constructs its future decision tree, given a model for the motion of all other agents, here simply that they continue on their previous trajectory at nominal speed v_o, as illustrated by the blue agent. In this way, the agent can compare each of the five moves available to it based on the absolute number of different visual states on all nodes accessible from that move. It chooses the move that maximizes this measure.

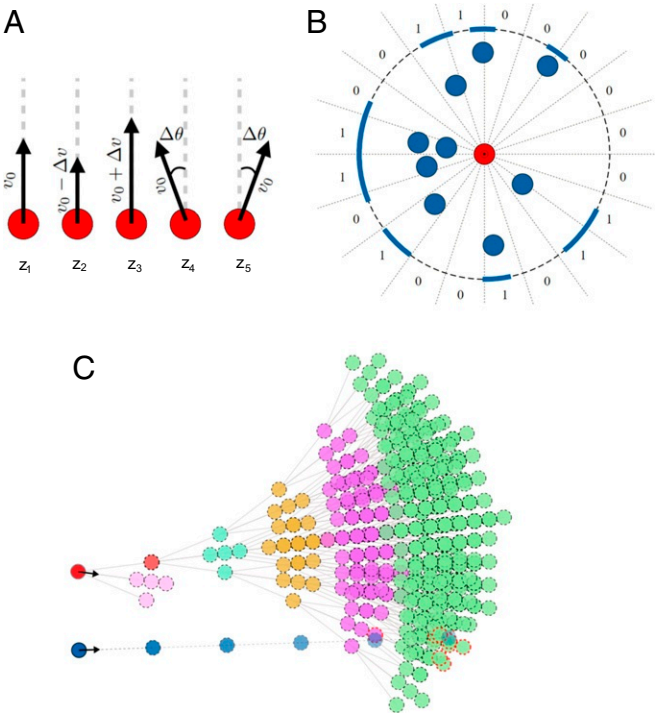

Fig. 1. Sketch showing an agent's movement options, a representation of the visual state of the world around it, and its future decision tree. (*A*) The five actions available to each agent at every time step, given that its previous move was in the direction of the dashed line, continue in the same direction at a nominal/slow/fast speed or turn left/right, respectively. (*B*) A representative agent (red) sees the other agents (blue) geometrically projected onto a retina-like sensor array. Each sensor registers 1 if a line of sight through more than half of its angular region intersects other disk(s), corresponding to the solid blue regions on the perimeter; 0 otherwise. This n_s-dimensional Boolean vector is the agent's sensory input and represents its "state." Here, we show $n_s = 20$, for clarity. (*C*) The spatial positions that an agent, shown as red, can access in the future form nodes on a fan-like tree, color-coded by the time into the future: pink/red (one step), cyan (two steps), orange (three steps), magenta (four steps), and green (five steps); in this cartoon, the maximum future time horizon is therefore $\tau = 5$. The branch of this tree that the agent explores is contingent on its next move (here shown as a turn to the left, in red). A similar branch exists for the four other possible moves, but these are omitted for clarity. The red agent computes the future sensory states accessible to it at each future node, as described in *B*, choosing the move in the next time step that leads to the branch with the largest number of distinct visual states. The nodes that are highlighted in dotted red correspond to positions that the agent anticipates will overlap ("collide") with other agents. Here, a single other colliding agent is shown in blue, for clarity. When computing the number of distinct visual states, we exclude those from nodes that correspond to, or follow after, such a collision.

In more mathematical language, we define the visual state $\mathbf{f_i} \in \{0,1\}^{n_s}$ of an agent on the i^{th} node of its tree of potential future states, as discussed above. Each of the five available moves in the next time step leads to branch α of the tree of potential future states. For each of these five branches, we then construct a set S_α consisting of all of the unique visual states $\mathbf{f}_i^{(\alpha)}$ within that branch. The future time horizon (tree depth) is $\tau = 4$ in our simulations, unless stated otherwise. Each branch is then given a weight $W_\alpha = |S_\alpha|$, and the agent then chooses the current action z_{α^*}, such that $\alpha^* = \text{argmax}_\alpha |S_\alpha|$.

Consider a toy example of this process, in which there are only $n_s = 4$ sensors and two possible actions. Imagine that the branch $\alpha = 1$, following action z_1, leads to three nodes with visual states of $\{1,0,1,0\}$, $\{1,0,0,0\}$, $\{1,0,0,0\}$, and $\{1,0,1,0\}$, while the branch $\alpha = 2$, following action z_2, leads to four nodes with visual states of $\{1,0,1,0\}$, $\{1,0,0,0\}$, $\{1,1,0,0\}$, and $\{1,0,1,1\}$. In this example, branch $\alpha = 2$, and hence action z_2, would be chosen because it leads to a future with four distinct Boolean vectors (states), whereas the branch $\alpha = 1$ contains only two distinct states; the vectors $\{1,0,1,0\}$ and $\{1,0,0,0\}$ being repeated.

Some nodes on the decision tree correspond to collisions and are highlighted in Fig 1*B* with a dotted red outline. An agent considers any branch of its decision tree to terminate on collision—i.e., this and any subsequent nodes are deemed inaccessible. In this way, the agent tends to avoid collisions because they contribute no states to its future. We find a strong reduction of collisions in the FSM trajectories that result, typically two to three orders of magnitude below a control collision rate (*SI Appendix*, Fig. S2). This is despite the fact that there is no explicit suppression of collisions, e.g., via physical interactions.

In *SI Appendix*, we discuss how to generalize this to a continuous measure of the degeneracy of future visual states.

In ref. 28, a Gibbs measure of the accessible state space, rather than a count of the number of distinct states, is used to quantify the future freedom. Our work could be extended in this direction in the future.

Results

Swarms similar to those shown in Fig. 2*A* arise from these FSM dynamics across a broad range of parameter values; see *SI Appendix* for a comparison. However, there are some restrictions—e.g., the number of sensors can neither be too large (so that all states become unique) nor too small (sensory resolution is lost), and the time horizon must be sufficiently long. For time horizons that are too short ($\tau < 4$ for $N = 50$), the swarm becomes less stable, with agents separating from the main swarm increasingly frequently. In general, the initial conditions must be chosen to be roughly commensurate with the steady state. If the system is prepared in an initial configuration from which the agents' decision trees cannot perceive the steady state within τ time steps, then the swarm fragments, typically into cohesive subgroups; see Movie S2 for an example of this phenomenon with $N = 500$. Such initial conditions correspond to widely separated and/or orientationally disordered agents. Robustness to variation of the initial conditions improves with increasing τ.

The state shown in Fig 2*A* and Movie S1 has further qualitative similarities with animal systems and, in particular, large flocks of starlings: Its alignment order parameter is within 1% of a typical value for starling flocks (9), and it is in a state of *marginal opacity*, in which the fraction of sensors in state 0 to state 1 is order unity (16) (see *SI Appendix*, Fig. 2 for more details). Finally, the correlation length scales with the system size, as shown in Fig 2*B*. This is indicative of scale-free correlations, another feature of starling flocks (9), and systems close to criticality more generally (34). Fig 2*C* shows snapshots of a larger swarm ($N = 500$, $\tau = 5$), sequentially in time, with motion determined by FSM on the continuous measure of visual state degeneracy described in *SI Appendix*. While for smaller swarms, the two approaches gave virtually identical results (compare Movies S1 and S3), for larger swarms, the continuous measure had more variety in its steady-state collective dynamics and was more robust to fragmentation (contrast Movies S2 and S4).

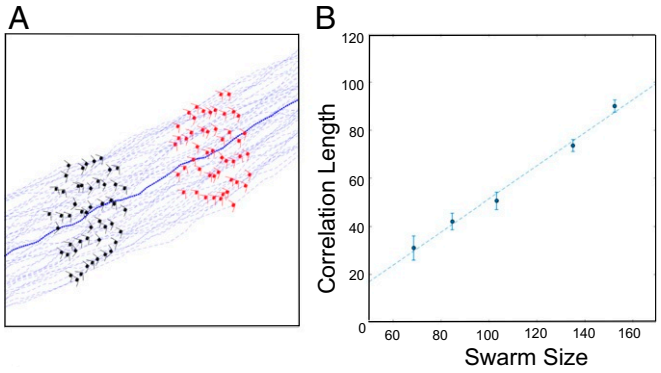

Fig. 2. (*A*) Structure of collective swarms that emerge under FSM dynamics, as described in Fig. 1*A*. Shown are snapshots of a typical realization at two different times showing the trajectories of the agents (light dashed lines) and center of mass (dark dotted line), with $N = 50$, $n_s = 40$, $v_0 = 10$, $\Delta v = 2$, and $\Delta\theta = 15^\circ$ and a time-horizon of $\tau = 4$. Wedges show agents' direction of motion; Movie S1. (*B*) The center-of-mass frame velocity correlation function for agents is computed for systems with the same parameter values, except that the data points correspond to $N = 50, 75, 100, 150, 200$ agents. Shown is the correlation length thereby obtained, here defined as the distance at which this correlation function crosses zero (nearby agents are positively correlated; distant ones are negatively correlated). This correlation length is compared against the corresponding swarm size, with the square root of the area of a convex hull containing all agents. See *SI Appendix* for details.

It is perhaps somewhat counterintuitive that such a highly ordered state emerges, given that FSM can be interpreted as preferring highly varied (roughly, high entropy) distributions of states. This is because FSM is insensitive to the variety/disorder of the swarm in the present. It is from such a highly ordered state that the swarm can access the greatest variety of states in the future. Thus, it targets this state and remains there. The state is cohesive because nearby agents then have the most freedom to rearrange their relative positions, Marginal opacity is selected because most configurations have sensor states roughly evenly split between 0 and 1.

Changing the Heuristic Used to Model Hypothetical Future Trajectories. A key ingredient of the FSM model is an assumption for how the other agents will move in the future. Without such an assumption, their future positions remain undetermined, and the corresponding visual projections cannot therefore be computed. Fig 1*C* shows the simplest of four different assumptions, or *heuristics*, that we report in this work: All other agents [only a single (blue) one is shown] are assumed to continue on ballistic trajectories, without turning, at speed v_o. The structure of the cohesive, coaligned swarms that spontaneously emerge under this assumption are shown in Fig 2*A* (see also Movie S1). The ballistic-motion assumption is an approximate model for the motion of the other agents and is not strictly self-consistent, insofar as all agents are identical and actually move according to FSM. Hence, the (other) agents won't move in exactly such a ballistic fashion, as can be seen from the individual trajectories in Fig 2*A*. Nonetheless, this assumption is quite good for the highly ordered (coaligned) swarms that do emerge from FSM. All agents would indeed continue moving in exactly the same direction under perfect coalignment. The alignment order

parameter is here defined as $\phi = \langle \frac{1}{N} \sum_i^N \hat{\mathrm{v}}_i(t) \rangle$ with the average performed over time and $\hat{\mathrm{v}}_i(t)$ a unit vector in the direction of motion of the i^{th} agent at time step t. The swarm in Fig 2*A* has order $\phi \simeq 0.98$.

Other heuristics can be made self-consistent with FSM. Examples include: 1) Agents are assumed to collectively target a particular value of order. At each time step, every agent, in random order, turns in either direction if this brings the collective order closer to the target order ϕ_A, otherwise continuing at speed v_o. 2) Agents are assumed to move at speed v_o, according to a topological version of the Vicsek model (5), in which coaligning neighbors are those that share edges under a Delaunay triangulation. As usual, this model involves a variable noise η, with a one-to-one relationship between this and the average order parameter at that noise $\phi_B(\eta)$. Fig 4 shows that both of these heuristics can be made self-consistent with FSM at the level of the order realized: The FSM trajectories that are generated, by using these heuristics as a model for the motion of all (other) agents, then have the same order as is generated by the bare heuristic, a value that was not known a priori. Any evolutionary pressure to adopt FSM-like dynamics should, presumably, also favor the ability to self-consistently predict the behavior of other members of the group in this way.

Training a Neural Network to Mimic the FSM Algorithm. While the full FSM algorithm is computationally demanding, an artificial neural network could serve as an example of a heuristic that can closely mimic FSM and fitness benefits arising therefrom. Crucially, once trained, it is computationally simple and fast. Similar levels of reasoning could be expected to operate under

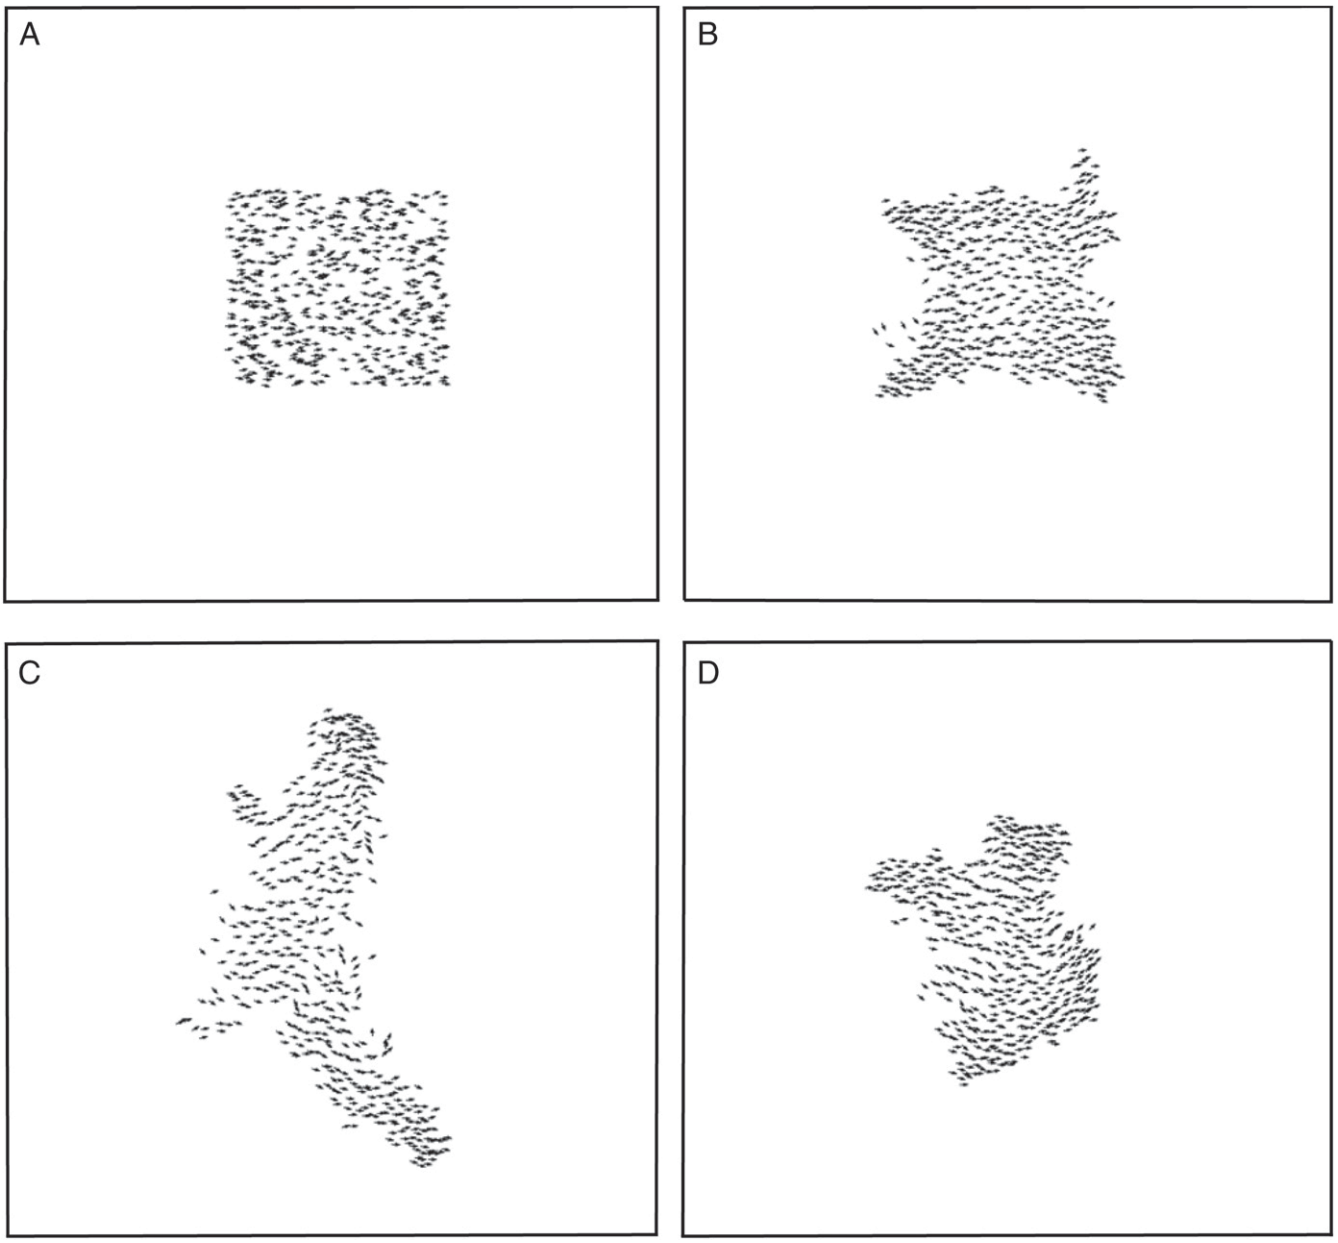

Fig. 3. Snapshots of a swarm made up of $N = 500$ agents with $\tau = 5$, shown at different times in a frame comoving with the swarm's center of mass. *A* shows the initial state of the swarm, and then *B–D* show snapshots of its subsequent evolution (in chronological order). In this example, we use a continuous measure of visual degeneracy (see *SI Appendix* for details). The full simulation is shown in Movie S4.

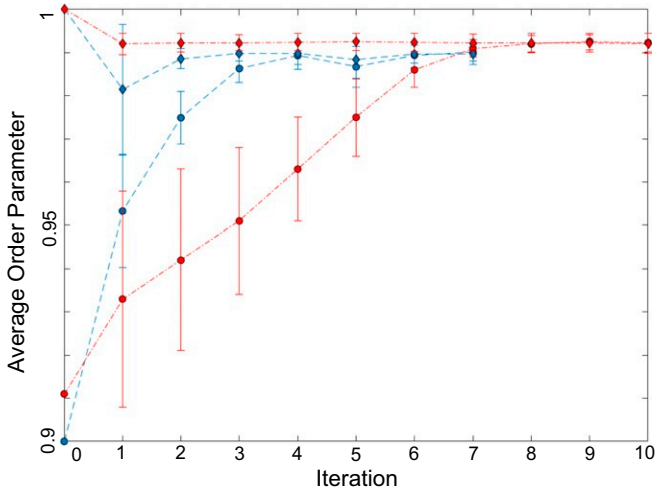

Fig. 4. Convergence of heuristics A (order targeting; blue) and B [topological Vicsek (5); red] to a value of the order parameter that is self-consistent with the value realized by FSM in each case. An initial (iteration 0) order parameter for the heuristic (ϕ_A and ϕ_B, respectively) is chosen. This parameterizes the model of all (other) agents to be used when constructing their trajectories into the future to apply FSM on each agent's predicted future visual states. The average order realized by the FSM simulation then serves as the order parameter for the heuristic in the next iteration, and the process is repeated. The order converges, both from above and below, to an average order parameter that is the same, both for the heuristic and the motion generated by FSM using that heuristic to model the behavior of other agents. FSM under heuristic A is unstable for values of $\phi_A \lesssim 0.9$, leading to flock fragmentation into (ordered) subgroups. Parameter values are as given in Fig. 2. See also Movie S5.

animal cognition. We do not claim that an artificial neural network would be a direct model for (wet) neural networks, even though the former field was indeed motivated by the latter. We only argue that reasoning with this level of complexity could be encoded in an animal brain. This heuristic, like the others described above, could also be used as a model for the behavior of other agents during FSM.

We trained a multilayered neural network on the simulated trajectories that arise under FSM over 200,000 time steps, as sketched in Fig 5. We gather training data by running the full FSM algorithm using nominal parameters ($N = 50$, $\tau = 4$, $n_s = 40$, $\Delta\theta = 15^\circ$, $v_0 = 10$, and $\Delta v = 2$). Future visual states are computed under the assumption that other agents will move ballistically in their future trajectories, i.e., at speed v_0 in their current direction of motion. We generated 800 separate simulations, each with agents initially placed randomly in a square region with dimensions that vary between 80 and 160 disk radii. Each agent's initial orientation was chosen randomly from a Gaussian distribution with mean orientation along the nominal x-direction and a SD of $2\Delta\theta$. We chose these different initial conditions to provide representative examples of trajectories that recover from perturbations. This allowed the trained network to make decisions that mimic FSM in situations that vary from the steady state, improving its robustness. In each of the simulations, we recorded the current speed and the current and previous visual state of every agent at every time step, along with the actual decision made by the FSM algorithm in that situation (represented as an integer between 1 and 5). Note that including memory, via the previous visual state, is found to be crucial to train a network which qualitatively reproduces the behavior of the full FSM algorithm. The training process is a supervised learning problem in which we have 10 million labeled example decisions, each corresponding to a vector input of dimension 81 ($2n_s + 1$, for the speed) with each output an integer between 1 and 5.

The neural network architecture we used consisted of a hidden layer of 200 fully connected neurons connected to the input with three further fully connected hidden layers of sizes 100; 50; and 25 respectively. The last of these was connected to a softmax classifier which output an integer between 1 and 5, corresponding to the selected action. All of the hidden layers used the RelU activation function. We trained the neural network on all of the data for 500 epochs using the ADAM optimizer under Keras with an initial learning rate of 0.0001.

The output from our artificial neural network was seen to closely mimic the FSM trajectories; Movie S6.

In summary, we propose a form of intrinsically motivated collective motion based on FSM. This involves a minimal

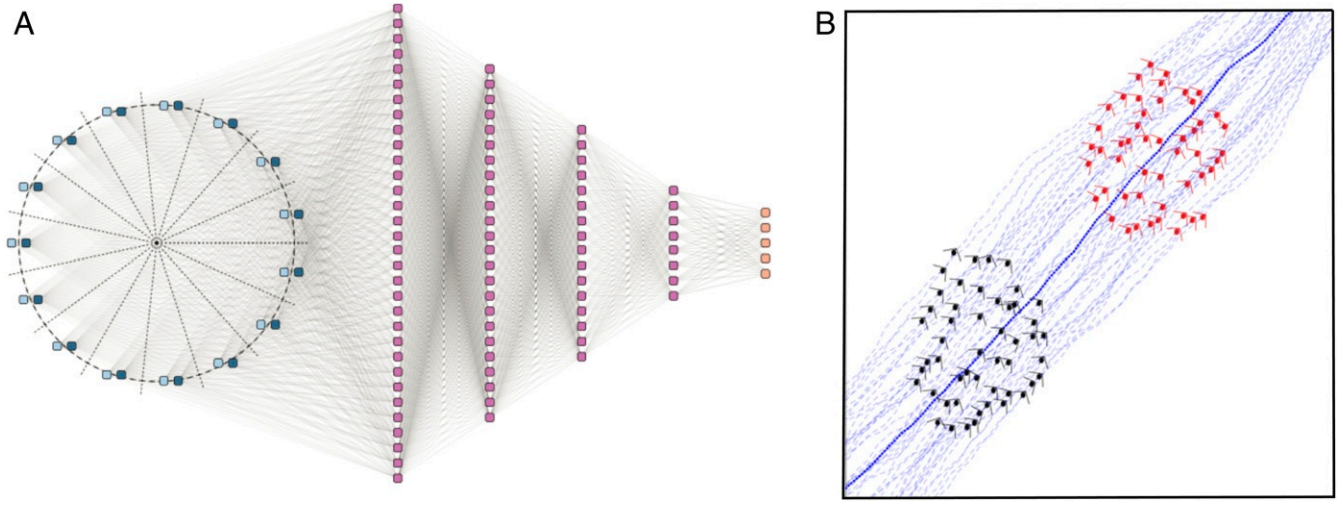

Fig. 5. Training a neural network as a heuristic approximating FSM. (*A*) Sketch of the network architecture. The network takes as its input the agent's current speed and the state of each sensor in both the current and previous time steps, represented as light and dark blue squares on each sensor (left-hand side). This is then passed through four hidden layers of neurons of sizes 200; 100; 50; and 25, which have RelU activation functions. These are attached to a softmax classifier which outputs an integer between 1 and 5, identifying the next move (final output; right-hand side). The network was trained to mimic FSM trajectories using 10 million examples (data from 200,000 simulation time steps). (*B*) The output dynamics from this network is seen to closely mimic the FSM trajectories shown in Fig. 2; Movie S6.

representation of vision in which agents seek to increase their control of the visual world around them. Specifically, they target being able to reach the greatest variety of future environments. The potential fitness benefits of this lie in the fact that it gives the agent freedom to access different outcomes in an uncertain world. Cohesive, ordered swarms that resemble natural animal systems spontaneously emerge under FSM. This behavior can be encoded in heuristics, mimicking full FSM. A neural network is an example of the kind of heuristic that could mimic FSM under animal cognition, providing a possible route for the evolutionary selection of this behavior. Such heuristics could also lie within the processing power of future realizations of "intelligent" materials that may incorporate sensors, as well as the ability to move.

ACKNOWLEDGMENTS. This work was partially supported by the UK Engineering and Physical Sciences Research Council though the Mathematics for Real-World Systems Centre for Doctoral Training Grant EP/L015374/1 (to H.J.C.). Computing facilities were provided by the Scientific Computing Research Technology Platform of the University of Warwick. We thank George Rowlands (Warwick) and Hugues Chaté (Saclay) for stimulating discussions.

1. L. F. Valadares *et al.*, Catalytic nanomotors: Self-propelled sphere dimers. *Small* **6**, 565–572 (2010).
2. J. Palacci, S. Sacanna, A. P. Steinberg, D. J Pine, P. M. Chaikin, Living crystals of light-activated colloidal surfers. *Science* **339**, 936–940 (2013).
3. S. Ramaswamy, Active matter. *J. Stat. Mech. Theory Exp.* **2017**, 054002 (2017).
4. M. E. Cates, J. Tailleur, When are active Brownian particles and run-and-tumble particles equivalent? Consequences for motility-induced phase separation. *EPL (Europhys. Lett.)* **101** 20010 (2013).
5. H. Chaté, F. Ginelli, Relevance of metric-free interactions in flocking phenomena. *Phys. Rev. Lett.* **105**, 168103 (2010).
6. M. C. Marchetti *et al.*, Hydrodynamics of soft active matter. *Rev. Mod. Phys.* **85**, 1143–1189 (2013).
7. J. Toner Y. Tu, Long-range order in a two-dimensional dynamical XY model: How birds fly together. *Phys. Rev. Lett.* **75**:4326–4329 (1995).
8. E. Bertin, M. Droz, G. Grégoire, Boltzmann and hydrodynamic description for self-propelled particles. *Phys. Rev. E* **74**, 022101 (2006).
9. A Cavagna *et al.*, Scale-free correlations in starling flocks. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 11865–11870 (2010).
10. A. Flack, M. Nagy, W. Fiedler, I. D. Couzin, M. Wikelski, From local collective behavior to global migratory patterns in white storks. *Science* **360**, 911–914 (2018).
11. D. P. Zitterbart, B. Wienecke, J.P. Butler, B. Fabry, Coordinated movements prevent jamming in an emperor penguin huddle. *PLoS ONE*, **6**, e20260 (2011).
12. D. Helbing, P. Molnár, Social force model for pedestrian dynamics. *Phys. Rev. E* **51**, 4282–4286 (1995).
13. M. Moussaïd, D. Helbing, G. Theraulaz, How simple rules determine pedestrian behavior and crowd disasters. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 6884–6888 (2011).
14. M. Ballerini *et al.*, Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 1232–1237 (2008).
15. H. Hildenbrandt, C. Carere, C.K. Hemelrijk. Self-organized aerial displays of thousands of starlings: A model. *Behav. Ecol.* **21**, 1349–1359 (2010).
16. D. J. G. Pearce, A. M. Miller, G. Rowlands, M. S. Turner, Role of projection in the control of bird flocks. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 10422–10426 (2014).
17. A. C. Gallup *et al.*, Visual attention and the acquisition of information in human crowds. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 7245–7250 (2012).
18. A. S. Kluybin, D Polani, C. L. Nehaniv, "Empowerment: A universal agent-centric measure of control" in *The 2005 IEEE Congress on Evolutionary Computation* (IEEE, Piscataway, NJ, 2005), **1**, pp. 128–135.
19. P. Capdepuy, D. Polani, C. L. Nehaniv, "Maximization of potential information flow as a universal utility for collective behaviour" in *IEEE Symposium on Artificial Life* (IEEE, Piscataway, NJ, 2007), pp. 207–213.
20. P. Capdepuy, D. Polani, C. L. Nehaniv, Perception-action loops of multiple agents: Informational aspects and the impact of coordination. *Theory Biosci.* **131**, 149–159 (2012).
21. C. Salge, C. Glackin, D. Polani, Changing the environment based on empowerment as intrinsic motivation. *Entropy* **16**, 2789–2819 (2014).
22. G. Baldassarre, M. Mirolli, *Intrinsically Motivated Learning in Natural and Artificial Systems* (Springer, New York, NY, 2013).
23. R. W. White, Motivation reconsidered: The concept of competence. *Psychol. Rev.* **66**, 297–333 (1959).
24. R. M. Ryan, E. L. Deci, Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemp. Educ. Psychol.* **25**, 54–67 (2000).
25. A. G. Barto, S. Singh, N. Chentanez "Intrinsically motivated learning of hierarchical collections of skills" in *The 3rd International Conference on Development and Learning*, J. Triesch, T. Jebara, Eds. (UCSD Institute for Neural Computation, La Jolla, CA, 2004).
26. S. Mohamed, D. J. Rezende, "Variational information maximisation for intrinsically motivated reinforcement learning" in *Proceedings of the 28th International Conference on Neural Information Processing Systems*, C. Cortes, D. D. Lee, M. Sugiyama, R. Garnett, Eds. (MIT Press, Cambridge, MA, 2015), pp. 2116–2124.
27. B. Eysenbach, A. Gupta, J. Ibarz, S. Levine, Diversity is all you need: Learning diverse skills without a reward function. arXiv:1802.06070 (16 February 2018).
28. A. D. Wissner-Gross, C. E. Freer, Causal entropic forces. *Phys. Rev. Lett.* **110**, 168702 (2013).
29. H. Hornischer, S. Herminghaus, M. G. Mazza, Intelligence of agents produces a structural phase transition in collective behaviour. arXiv:1706.01458 (5 June 2017).
30. R. P. Mann, R. Garnett, The entropic basis of collective behaviour. *J. R. Soc. Interf.* **12**, 20150037 (2015).
31. H.-R. Jiang, N. Yoshinaga, M. Sano, Active motion of a Janus particle by self-thermophoresis in a defocused laser beam. *Phys. Rev. Lett.* **105**, 268302 (2010).
32. V. Kantsler, J. Dunkel, M. Polin, R. E. Goldstein, Ciliary contact interactions dominate surface scattering of swimming eukaryotes. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 1187–1192 (2013).
33. T. Sanchez *et al.*, Spontaneous motion in hierarchically assembled active matter. *Nature* **491**, 431–434 (2012).
34. T. Mora, W. Bialek, Are biological systems poised at criticality? *J. Stat. Phys.* **144**, 268–302 (2011).