

RESEARCH ARTICLE

Development of swarm behavior in artificial learning agents that adapt to different foraging environments

Andrea López-Incera^{1*}, Katja Ried¹, Thomas Müller², Hans J. Briegel^{1,2}¹ Institute for Theoretical Physics, University of Innsbruck, Innsbruck, Austria, ² Fachbereich Philosophie, Universität Konstanz, Konstanz, Germany* andrea.lopez-incera@uibk.ac.at

Abstract

Collective behavior, and swarm formation in particular, has been studied from several perspectives within a large variety of fields, ranging from biology to physics. In this work, we apply Projective Simulation to model each individual as an artificial learning agent that interacts with its neighbors and surroundings in order to make decisions and learn from them. Within a reinforcement learning framework, we discuss one-dimensional learning scenarios where agents need to get to food resources to be rewarded. We observe how different types of collective motion emerge depending on the distance the agents need to travel to reach the resources. For instance, strongly aligned swarms emerge when the food source is placed far away from the region where agents are situated initially. In addition, we study the properties of the individual trajectories that occur within the different types of emergent collective dynamics. Agents trained to find distant resources exhibit individual trajectories that are in most cases best fit by composite correlated random walks with features that resemble Lévy walks. This composite motion emerges from the collective behavior developed under the specific foraging selection pressures. On the other hand, agents trained to reach nearby resources predominantly exhibit Brownian trajectories.

OPEN ACCESS

Citation: López-Incera A, Ried K, Müller T, Briegel HJ (2020) Development of swarm behavior in artificial learning agents that adapt to different foraging environments. PLoS ONE 15(12): e0243628. <https://doi.org/10.1371/journal.pone.0243628>

Editor: Thilo Gross, University Of Bristol, UNITED KINGDOM

Received: July 1, 2020

Accepted: November 24, 2020

Published: December 18, 2020

Copyright: © 2020 López-Incera et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The data underlying the results presented in the study were obtained from the code LearnToForage_public, available from: https://github.com/qic-ibk/LearnToForage_public.

Funding: AL, KR and HJB were supported by the Austrian Science Fund (FWF) through grant SFB BeyondC F7102. HJB was supported by the Ministerium für Wissenschaft, Forschung, und Kunst Baden-Württemberg (AZ:33-7533-30-10/41/1). HJB and TM were supported by Volkswagen

1 Introduction

Collective behavior is a common but intriguing phenomenon in nature. Species as diverse as locusts, and some families of fish or birds exhibit different types of collective motion in very different environments and situations. Although the general properties of swarms, schools and flocks have been widely studied (see e.g. [1] for a review), the emergence of global, coordinated motion from the individual actions is still a subject of study. Different approaches, ranging from statistical physics to agent-based models, have led to new insights and descriptions of the phenomenon. Statistical physics models are very successful at describing macroscopic properties such as phase transitions and metastable states [2–4], but in order to apply the powerful tools of statistical mechanics, these models normally simplify the individuals to particles that interact according to certain rules dictated by the physical model adopted, as for instance the

Foundation (Az:97721) and a BlueSky project of the University of Konstanz.

Competing interests: The authors have declared that no competing interests exist.

Ising-type interaction of the spins in a lattice. A different type of models are the so-called self-propelled particle (SPP) models [5–8], which enable higher complexity in descriptions at the individual level but still allow one to employ the tools of statistical physics. They describe individuals as particles that move with a constant velocity and interact with other individuals via fixed sets of rules that are externally imposed. In SPP models, the description of the interactions is not restricted to physically accepted first principles, but can include ad hoc rules based on specific experimental observations.

In this work, we follow a different approach and model the individuals as artificial learning agents. In particular, we apply Projective Simulation (PS) [9], which is a model of agency that can incorporate learning processes via a reinforcement learning mechanism. The individuals are thus described as PS agents that interact with their surroundings, make decisions accordingly and learn from them based on rewards provided by the environment. This framework allows for a more detailed, realistic description in terms of the perceptual apparatus of the agent. One of the main differences with respect to previous models is that the interaction rules between agents are not imposed or fixed in advance, but they emerge as the result of learning in a given task environment. This type of agent-based models that employ artificial intelligence to model behavior are gaining popularity in the last few years. Artificial neural networks (ANN) have been used, for instance, in the context of navigation behaviors [10, 11] and reinforcement learning (RL) algorithms have been applied to model collective behavior in different scenarios, such as pedestrian movement [12] or flocking [13, 14].

In contrast to other learning models such as neural networks, PS provides a transparent, explicit structure that can be analyzed and interpreted. This feature is particularly useful in modeling collective behavior, since we can study the individual decision making processes, what the agents learn and why they learn it. This way, we can directly address the questions of how and why particular individual interactions arise that in turn lead to collective behaviors. Initial work by Ried et al. [15], where the authors use PS to model the density-dependent swarm behavior of locusts, laid the foundations of the present work.

Since the interaction rules are developed by the agents themselves, the challenge is to design the environment and learning task that will give rise to the individual and, consequently, collective behavior. In previous works, the agents are directly rewarded for aligning themselves with the surrounding agents [15] or for not losing neighbours [14]. Instead of rewarding a specific behavior, in this work we set a survival task that the agents need to fulfill in order to get the reward, and then analyze the emergent behavioral dynamics.

As a starting hypothesis, we consider the need to forage as an evolutionary pressure and design a learning task that consists in finding a remote food source. Due to this particular survival task, our work relates to the investigation of foraging theories and optimal searching behavior.

There is a vast number of studies devoted to the analysis of foraging strategies in different types of environments e.g., [16–19]. In the particular case of environments with sparsely distributed resources (e.g. patchy landscapes), there are two main candidates for the optimal search model: Lévy walks [20–22] and composite correlated random walks (CCRW) [23, 24]. The former are described by a single distribution of step lengths that is characterized by a power-law $p(\ell) \sim \ell^{-\mu}$ with exponent $1 < \mu \leq 3$, whereas the latter consider that the movement is composed of two different modes, characterized by two exponential distributions with different decay rates. Although the mathematical models behind them are fundamentally different, they have some common features that make the movement patterns hard to distinguish [24–28]. In broad terms, both models can produce trajectories that are a combination of short steps (with large turning angles in 2D), which are useful for exploring the patch area, and long, straight steps, which are efficient to travel the inter-patch distances. Even though both models

have theoretical [22, 23] and experimental (e.g. [29, 30]) support, it is not yet clear if animal foraging patterns can be described and explained by such models or if they are too complex to admit such simplifications.

Furthermore, regarding the Lévy walks, there is an ongoing debate on the question whether they emerge under certain animal foraging strategies. Currently there exist two main hypotheses, referred to as the evolutionary and the emergentist. The evolutionary hypothesis (also called Lévy flight foraging (LFF) hypothesis) states that certain species have evolved according to natural selection to develop an optimal foraging strategy consisting of Lévy walk movement patterns (see e.g. [31] and references therein). On the other side, the emergentist hypothesis argues that the LFF hypothesis is not sufficient to account for the complexity of animal behavior since it does not explain certain anomalies observed experimentally (see [32] and references therein). It argues that Lévy walks can emerge spontaneously as a consequence of the features of the environment, which lead to certain responses from the foraging organism. Thus, these responses are not part of an evolved strategy developed over the course of generations, but can arise from innate behaviors and lead to Lévy patterns spontaneously when the animal is confronted with certain environmental conditions.

Due to the fact that our learning task is directly related to foraging strategies, we link the present work to the aforementioned studies by analyzing the individual trajectories the agents produce as a consequence of the behavior developed in the different learning contexts.

The paper is organized as follows: an introduction to Projective Simulation and a detailed description of the model and the learning setup are given in Sec. 2. In Sec. 3, we present different learning tasks and analyze the resulting learned behaviors. In Sec. 4, we study the emergent group dynamics and individual trajectories within the framework of search models to determine if they can be described as Lévy walks or composite correlated random walks. Finally, we summarize the results and conclude in Sec. 5.

2 Methods and model

A wide range of models and techniques have been applied to the study of collective behavior. In this work, we apply Projective Simulation, a model for artificial agency [9, 33–37]. Each individual is an artificial agent that can perceive its surroundings, make decisions and perform actions. Within the PS model, the agent's decision making is integrated into a framework for reinforcement learning (RL) that allows one to design concrete scenarios and tasks that the individuals should solve and then study the resulting strategies developed by the agents. We remark that the notion of *strategy* employed throughout this work does not imply that the agents are able to *plan*. We use the word “strategy” to refer to the behavior the agents develop given a certain learning task. In addition, each agent's motor and sensory abilities can be modeled in a detailed, realistic way.

In our model of collective behavior, the interaction rules with other individuals are not fixed in advance; instead the agents develop them based on their previous experience and learning. The most natural interpretation of this approach is that it describes how a group of given individuals change their behavior over the course of their interactions, for example human children at play. However, our artificial learning agents can also be used to model simpler entities that do not exhibit learning in the sense of noticeable modifications of their responses over the course of a single individual's lifetime, but only change their behavior over the course of several generations. In this case, a single simulated agent does not correspond to one particular individual, in one particular generation, but rather stands as an avatar for a generic individual throughout the entire evolution of the species. The environmental pressures driving behavioural changes over this time-scale can be easily encoded in a RL scenario, since

the reward scheme can be designed in such a way that only the behaviors that happen to be beneficial under these pressures are rewarded. This allows us to directly test whether the environmental pressures are a possible *causal* explanation for the observed behavior or not. Our approach interprets the reinforcement of certain responses from an evolutionary perspective. It differs from genetic algorithms, extended classifier systems [38], and similar advanced machine learning methodology in that it does not model evolution in an explicit manner. Such machinery, e.g., the encoding of genes, mutations, and crossover, usually comes at the cost of a larger model complexity (number of free parameters; see [33]) and additional computational overhead. Alternatively, neural network models might be employed, but these are typically difficult to interpret and thus not useful in our context. Unlike genetic algorithms, Projective Simulation provides a model of agency that describes a stochastic decision-making process of each individual, which can be used beyond mere optimization by focusing on the resulting causal explanations.

Although other reinforcement learning algorithms may be used to model a learning agent, Projective Simulation is particularly suitable for the purpose of modeling collective behavior, since it provides a clear and transparent structure that gives direct access to the internal state of the agent, so that the deliberation process can be analyzed in an explicit way and can be related to the agent's behavior. This analysis can help us gain new insight into how and why the individual interactions that lead to collective behaviors emerge.

2.1 Projective simulation

Projective Simulation (PS) is a model for artificial agency that is based on the notion of episodic memory [9]. The agent interacts with its surroundings and receives some inputs called percepts, which trigger a deliberation process that leads to the agent performing an action on the environment.

In the PS model, the agent processes the percepts by means of an internal structure called episodic and compositional memory (ECM), whose basic units are called clips and represent an episode of the agent's experience. Mathematically, the ECM can be represented as a directed, weighted graph, where each node corresponds to a clip and each edge corresponds to a transition between two clips. All the edge weights are stored in the adjacency matrix of the graph, termed h matrix. For the purpose of this work, the most basic two-layered structure is sufficient to model simple agents. Percept-clips are situated in the first layer and are connected to the action-clips, which constitute the second layer (see Fig 1). Let us define these components of the ECM more formally.

- The *percepts* are mathematically defined as N -tuples $s = (s_1, s_2, \dots, s_N) \in \mathcal{S}$, where \mathcal{S} is the Cartesian product $\mathcal{S} \equiv \mathcal{S}_1 \times \mathcal{S}_2 \times \dots \times \mathcal{S}_N$. As it can be seen from this mathematical definition, the percept s has several categories, represented by \mathcal{S}_i . Each component of the tuple is denoted by $s_i \in \{1, \dots, |\mathcal{S}_i|\}$, where $|\mathcal{S}_i|$ is the number of possible states of \mathcal{S}_i . The total number of percepts is thus given by $|\mathcal{S}_1| \cdot \dots \cdot |\mathcal{S}_N|$.
- Analogously, the *actions* are defined as $a = (a_1, a_2, \dots, a_N) \in \mathcal{A}$, where $\mathcal{A} \equiv \mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_N$ and $a_i \in \{1, \dots, |\mathcal{A}_i|\}$, where $|\mathcal{A}_i|$ is the number of possible states of \mathcal{A}_i . The total number of actions is given by $|\mathcal{A}_1| \cdot \dots \cdot |\mathcal{A}_N|$.

As an example, consider an agent that perceives both its internal state, denoted by \mathcal{S}_1 , with two possible percepts $\mathcal{S}_1 = \{\text{hungry, not hungry}\}$, and some visual input, denoted by \mathcal{S}_2 , with $\mathcal{S}_2 = \{\text{I see food, I do not see food}\}$. Thus, one out of the four possible percepts could

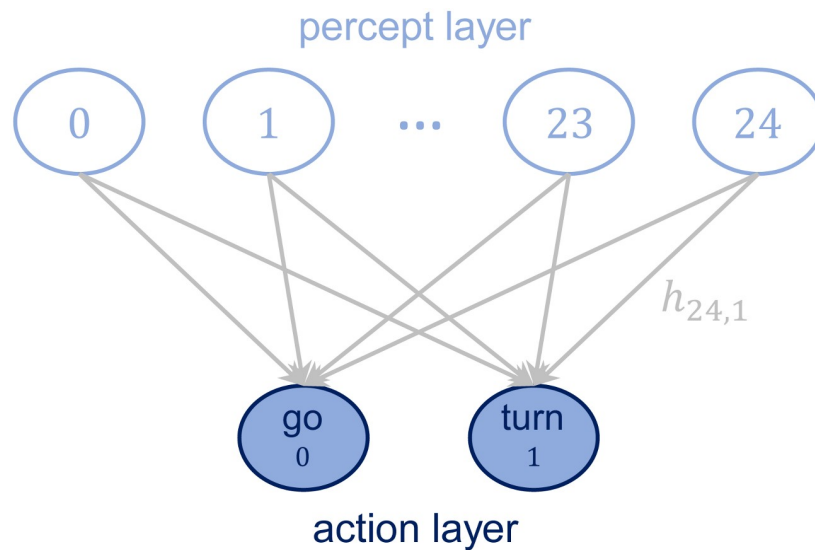


Fig 1. Structure of the ECM. The ECM consists of two layers, one for the percepts and one for the actions. Percepts and actions are connected by edges whose weight h_{ij} determines the transition probability from the given percept to each action (see Sec. 2.2 for details on the model).

<https://doi.org/10.1371/journal.pone.0243628.g001>

be $s = (\text{hungry, I see food})$. In this case, the possible actions may be $\mathcal{A} = \{\text{go for food, turn around}\}$.

Fig 1 represents the structure of the ECM in our model, which consists of a total of 25 percepts and 2 actions (see Sec. 2.2 for a detailed description).

Let us introduce how the agent interacts with the environment and makes decisions via the ECM. When the agent receives a percept, the corresponding percept-clip inside the ECM is activated, starting a random walk that only ends when an action-clip is reached, which triggers a real action on the environment. The transition probability $P(j|i)$ from a given percept-clip i to an action-clip j is determined by the corresponding edge weight h_{ij} as,

$$P(j|i) = \frac{h_{ij}}{\sum_k h_{ik}}, \quad (1)$$

where the normalization is done over all possible edges connected to clip i . This process, starting with the presentation of a perceptual input that activates a percept clip and finishing when the agent performs an action on the environment, is termed an (individual) *interaction round*.

The structure of the ECM allows one to easily model learning by just updating the h matrix at the end of each interaction round. The h matrix is initialized with all its elements being 1, so that the probability distribution of the actions is uniform for each percept. Reinforcement learning is implemented by the environment giving a reward to the agent every time that it performs the correct action. The reward increases the h -values, and thus the transition probabilities, of the successful percept-action pair. Hence, whenever the agent perceives again the same percept, it is more likely to reach the correct action. However, in the context of this work, we are setting a learning task in which the agent should perform a sequence of several actions to reach the goal and get the reward. If the reward is given only at the last interaction round, only the last percept-action pair would be rewarded. Thus, some additional mechanism is necessary in order to store a sequence of several percept-action pairs in the agent's memory. This mechanism is called *glow* and the matrix that stores the information about this sequence is denoted by g . The components g_{ji} , corresponding to the percept-action transition $i \rightarrow j$, are

initialized to zero and are updated at the end of *every* interaction round according to:

$$g_{ij}^{(t+1)} = (1 - \eta)g_{ij}^{(t)} + \left\{ \begin{array}{l} 0 \text{ if edge was not traversed} \\ 1 \text{ if edge was traversed,} \end{array} \right\} \quad (2)$$

where $0 \leq \eta \leq 1$ is the glow parameter, which damps the intensity of the given percept-action memory. For η close to one, the actions that are taken at interaction rounds in temporal vicinity to the rewarded action are more intensely remembered than the initial actions. If $\eta = 0$, all actions the agent performed until the rewarded interaction are equally remembered. The g matrix is updated in such a way that the percept-action pairs that are used more often to get to the reward are proportionally more rewarded than the pairs that were rarely used. Note that the agent is not able to distinguish an *ordered* sequence of actions, but this is not necessary for the purpose of this work.

In the context of our learning task, the agent receives a reward from the environment at the end of the interaction round at which it reaches a goal. Then, the learning is implemented by updating the h matrix with the rule,

$$h^{(t+1)} = h^{(t)} + R \cdot g, \quad (3)$$

where $R \geq 0$ is the reward (only non-zero if the agent reached the goal at the given interaction round) and g is the updated glow matrix. Technically, the glow matrix is updated first, and then, if the agent is rewarded, the h matrix is updated.

Since we model collective behavior, we consider a group of several agents, each of which has its own and independent ECM to process the surrounding information. Details on the specific learning task and the features of the agents are given in the following section.

2.2 Details of the model

We consider an ensemble of N individuals that we model as PS learning agents, which possess the internal structure (ECM) and the learning capabilities described in section 2.1. This description of the agents can be seen as a simplified model for species with low cognitive capacities and simple deliberation mechanisms, or just as a theoretical approach to study the optimal behavior that emerges under certain conditions.

With respect to the learning, we set up a concrete task and study the strategy agents develop to fulfill it. In particular, we consider a one-dimensional circular world with sparse resources, which mimics patchy landscapes such as deserts, where organisms need to travel long distances to find food. Inspired by this type of environments, we model a task where agents need to reach a remote food source to get rewarded. The strategy the agents learn via the reinforcement learning mechanism does not necessarily imply that the individual organisms should be able to *learn* to develop it, but can also be interpreted as the optimal behavior that a species would exhibit under the given environmental pressures.

Let us proceed to detail the agents' motor and sensory abilities. The positions that the agents can occupy in the world are discretized $\{0, 1, 2, \dots, W\}$, where W is the world size (total number of positions). Several agents can occupy the same position. At each interaction round, the agent can decide between two actions: either it continues moving in the same direction or it turns around and moves in the opposite direction. The agents move at a fixed speed of 1 position per interaction round. For the remainder of this work, we consider the distance between two consecutive positions of the world to be our basic unit of length. Therefore, unless stated otherwise, all distances given in the following are measured in terms of this unit. We remark that, in contrast to other approaches where the actions are defined with respect to other

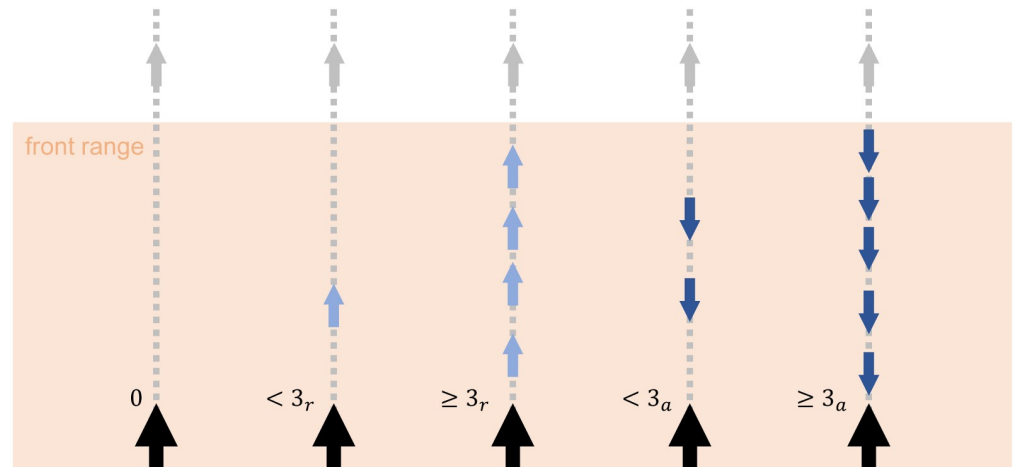


Fig 2. Graphical representation of the percepts' meaning. Only the front visual range (colored region) is considered, which corresponds to the values that category s_f can take. The focal agent is represented with a larger arrow than the frontal neighbors. The agent can only see its neighbors inside the visual range and it can distinguish if the majority are receding (light blue) or approaching (dark blue) and if they are less or more than three.

<https://doi.org/10.1371/journal.pone.0243628.g002>

individuals [39], the actions our agents can perform are purely motor and only depend on the previous orientation of the agent.

Perception is structured as follows: a given agent, termed the focal agent, perceives the relative positions and orientations of other agents inside its visual range (radius with center at the agent's position) V_R , termed its neighbors. The percept space S (see Sec. 2.1) is structured in the Cartesian product form $S \in S_f \times S_b$, where S_f is the region in front of the focal agent and S_b the region at the back. More precisely, each percept $s = (s_f, s_b)$ contains the information of the orientation of the neighbors in each region with respect to the focal agent and if the density of individuals in this region is high or low (see Fig 2). Each category of percepts can take the values $s_f, s_b \in \{0, < 3_r, \geq 3_r, < 3_a, \geq 3_a\}$ (25 percepts in total), which mean:

- 0. No agents
- $< 3_r$. There are less than 3 neighbors in this region and the majority of them are receding from the focal agent.
- $\geq 3_r$. There are 3 or more neighbors in this region and the majority of them are receding from the focal agent.
- $< 3_a$. There are less than 3 neighbors in this region and the majority of them are approaching the focal agent.
- $\geq 3_a$. There are 3 or more neighbors in this region and the majority of them are approaching the focal agent.

In the following discussions, we refer to the situation where the focal agent has the same orientation as the neighbors as a percept of *positive flow* (majority of neighbors are receding at the front and approaching at the back). If the focal agent is oriented against its neighbors (these are approaching at the front and receding at the back), we denote it as a percept of *negative flow*. Note that the agents can only perceive information about the neighboring agents inside their visual range, but they are not able to see any resource or landmark present in the surroundings. This situation can be found in realistic, natural environments where the

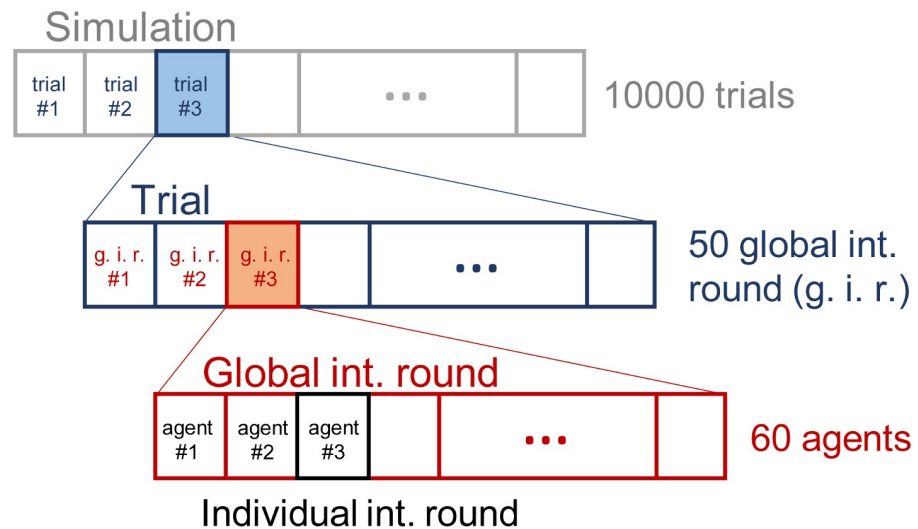


Fig 3. Structure of the simulation. Each ensemble of agents is trained for 10^4 trials, where each trial consists of 50 global interaction rounds (g.i.r.). At each g.i.r., the agents interact sequentially (see text for details).

<https://doi.org/10.1371/journal.pone.0243628.g003>

distance between resources is large and the searcher has no additional input while moving from one patch to another. Furthermore, the important issue of body orientation is thereby taken into account in our model [32].

The interactions between agents are assumed to be sequential, in the sense that one agent at a time receives a percept, deliberates and then takes its action before another agent is given its percept. Technically, agents are given a label at the beginning of the simulation to keep track of the interaction sequence but we remark that they are placed at random positions in the world. There are two reasons for assuming a sequential interaction. For one, in a group of real animals (or other entities), different individuals typically take action at slightly different times, with perfect synchronization being a remarkable and costly exception. The second argument in favor of sequential updating is that it ensures that a given agent's circumstances do not change from the time it receives its percept until the time when it acts. If the actions of all agents were applied simultaneously, a given focal agent would not be able to react to the actions of the other agents in the same round. Such a simplification would not allow us to take into account any sequential, time-resolved interactions between different agents of a group. In the real situation, while one focal agent is deliberating, other agents' actions may change its perceptual input. Therefore, an action that may have been appropriate at the beginning of the round, would no longer be appropriate at this agent's turn.

The complete simulation has the structure displayed in Fig 3, where:

- With each ensemble of $N = 60$ agents, we perform a simulation of 10^4 trials during which the agents develop new behaviors to get the reward (RL mechanism). This process is denoted as *learning process* or *training* from this point on.
- Each trial consists of $n = 50$ global interaction rounds. At the beginning of each trial, all agents of the ensemble are placed in random positions within the initial region (see Fig 4).
- We define a global interaction round to be the sequential interaction of the ensemble, where agents take turns to perform their individual interaction round (perception-deliberation-action). Note that each agent perceives, decides and moves only once per global interaction round.

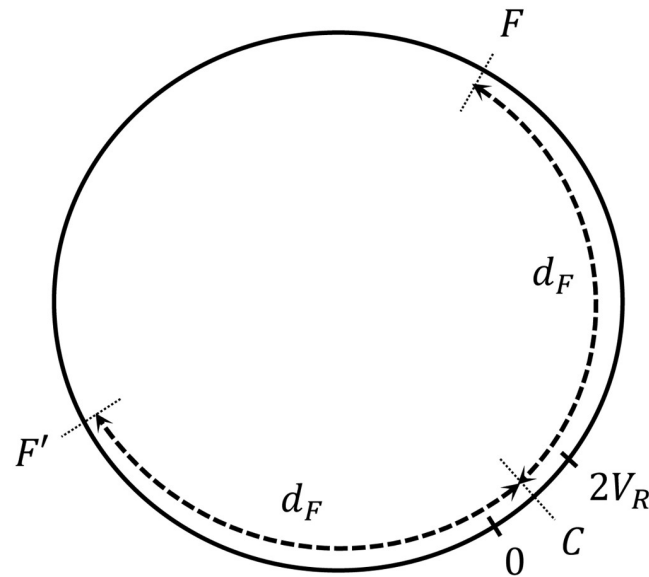


Fig 4. 1D environment (world). Agents are initialized randomly within the first $2V_R$ positions. Food is located at positions F and F' . d_F is the distance from the center of the initial region C to the food positions.

<https://doi.org/10.1371/journal.pone.0243628.g004>

The learning task is defined as follows: at the beginning of each trial, all the agents are placed at random positions within the first $2V_R$ positions of the world, with orientations also randomized. Each agent has a fixed number n of interaction rounds over the course of a trial to get to a food source, located at positions F and F' (Fig 4). At each interaction round, the agent first evaluates its surroundings and gets the corresponding percept. Given the percept, it decides to perform one out of the two actions (“go” or “turn”). After a decision is made, it moves one position. If the final position of the agent at the end of an interaction round is a food point, the agent is rewarded ($R = 1$) and its ECM is updated according to the rules specified in Sec. 2.1. Each agent can only be rewarded once per trial. Note that the h matrices of the agents are only updated following Eq (3), and we do not consider any other transformation to explicitly model evolution as it could be done, in principle, using genetic algorithms to explicitly represent evolutionary mechanisms of mutation, crossover etc. (see also the discussion at the beginning of Sec. 2).

We consider different learning scenarios by changing the distances d_F at which food is positioned. However, note that a circular one-dimensional world admits a trivial strategy for reaching the food without any interactions, namely going straight in one direction until food is reached. Thus, in order to emulate the complexity that a more realistic two-dimensional scenario has in terms of degrees of freedom of the movement, we introduce a noise element that randomizes the orientation of each agent every s_r steps (it changes orientation with probability $1/2$). Not all agents randomize the orientations at the same interaction round, which would lead to random global behavior. This randomization can be also interpreted biologically as a fidgeting behavior or even as a built-in behavior to escape predators [40]. Protean movement has been observed in several species [41–44] and there exist empirical studies that show that unpredictable turns [45] and complex movement patterns [46] decrease the risk of predation. In addition, if the memory of the organism is not very powerful, we can also consider that, at these randomization points, it forgets its previous trajectory and needs to rely on the

neighbors' orientations in order to stabilize its trajectory. The agent can do so, since the randomization takes place right before the agent starts the interaction round.

Under these conditions, we study how the agents get to the food when the only input information available to them is the orientation of the agents around them.

3 Results I: Learned behavior in different scenarios

We consider different learning scenarios characterized by the distance d_F (see Fig 4). We study how the dynamics that the agents develop in order to reach the food source change as the distance d_F increases. In particular, we focus on two extreme scenarios: one where the resource is within the initial region ($d_F < V_R$)—agents are initialized within the first $2V_R$ positions of the world—, and the other one where the resource is at a much larger distance. As a scale for this distance, we consider how far an agent can travel on average with a random walk, which is $d_{RW} = \sqrt{n}$ providing that it moves one position per interaction round. Hence, the other extreme scenario is such that $d_F \gg d_{RW}$. Note that the scale of d_F for this regime depends on the total number n of interaction rounds that the agents perform in one trial. The maximum value of d_F that we can choose thus depends on the maximum distance the agents can travel within n rounds following an unbiased random walk (for $n = 50$, this threshold is approximately at $d_F = 21$).

The situation where $d_F < V_R$ mimics an environment with densely distributed resources, whereas the regime with $d_F \gg d_{RW}$ resembles a resource-scarce environment where a random walk is no longer a valid strategy for reaching food sources.

The parameters of the model that are used in all the learning processes are given in Table 1. Providing that $d_{RW} = \sqrt{50} \simeq 7$, we consider values of d_F ranging from 2 to 21 and focus on the cases with $d_F = 4, 21$ as the representative examples of resource-dense and resource-scarce environments, respectively. All agents start the learning process with a newly initialized h matrix, so they perform each action (“go” or “turn”) with equal probability. Fig 5 shows the learning curves for three different scenarios, where the food is placed at $d_F = 4, 10, 21$. The learning processes are independent from each other, that is, the distance d_F does not change within one complete simulation of 10^4 trials. In this way, we can analyze the learned behaviors separately for each d_F . The learning curve displays the percentage of agents that reach the food source and obtain a reward at each trial. As a baseline for comparison, we also set the same learning task with $d_F = 21$ for non-interacting (n.i.) agents (we set $V_R = 0$, so they cannot see the neighbors). The n.i. agents learn to go straight almost deterministically—the probability for the action “go” at the end of the learning process is almost 1 for percept (0, 0)—. Therefore, these agents perform a random walk with $n/s_r = 50/5 = 10$ steps of length $s_r = 5$, which allows it to cover a distance of $5\sqrt{10} \simeq 16$ positions. The rest of percepts are never encountered, so the initial h values remain the same. Due to the periodic randomization of the agents' orientation, it can be seen that they do not reach the efficiency rate of the interacting agents (see Fig 5) and only one out of three agents reaches the reward at each trial. Fig 5 shows that, for $d_F = 4$, the

Table 1. Description of the parameters used in the learning simulations with PS.

Agent		Environment	
Description	Value	Description	Value
Visual range (V_R)	6	Number of agents (N)	60
Reorient. freq. (s_r)	5	World size (W)	500
Glow (η)	0.2	Int. rounds per trial (n)	50
Reward (R)	1	Number of trials	10^4

<https://doi.org/10.1371/journal.pone.0243628.t001>

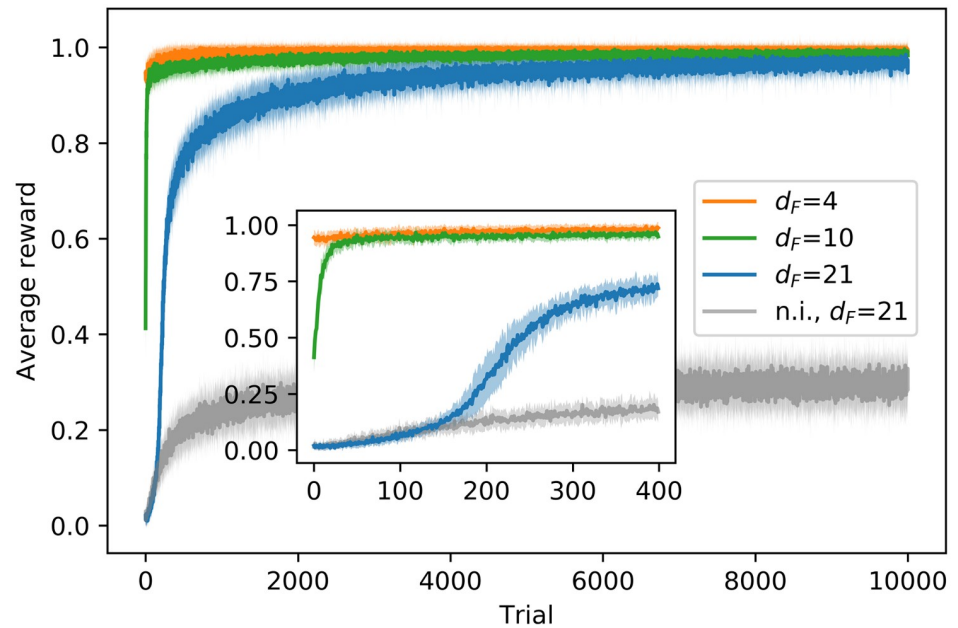


Fig 5. Learning curves for $d_F = 4, 10, 21$ and $d_F = 21$ for non-interacting (n.i.) agents. The curve shows the percentage of agents that reach the food source and obtain a reward of $R = 1$ at each trial. For each task, the average is taken over 20 (independent) ensembles of 60 agents each and the shaded area indicates the standard deviation. Zooming into the initial phase of the learning process, the inset figure shows a faster learning in the task with $d_F = 10$ than in the task with $d_F = 21$. In the case of $d_F = 21$, no agent is able to reach the food source in the first trial, and it takes the interacting agents approx. 200 trials to outperform the n.i. agents.

<https://doi.org/10.1371/journal.pone.0243628.g005>

food source is so close (inside the initial region) that the agents get the reward in all the trials from the beginning. On the other hand, the tasks with $d_F = 10, 21$ show a learning process that takes more trials for the agents to come up with a behavior that allows them to get to the reward. In particular, only 40% of the agents are able to reach the goal with the initial behavior (Brownian motion) in the scenario with $d_F = 10$ and this percentage drops to almost 0% in the case with $d_F = 21$. Note that it takes more trials for the agents to learn how to get to the furthest point ($d_F = 21$) than it takes for $d_F = 10$ (see inset in Fig 5). The interacting agents start outperforming the n.i. agents in the task with $d_F = 21$ at trial 200, where they start to form aligned swarms, as one can also see from the increase in the alignment parameter at the same trial (see Sec. 3.2.1 for details).

3.1 Individual responses

The behavior the agents have learned at the *end* of the training can be studied by analyzing the final state of the agents' ECM, from where one obtains the final probabilities for each action depending on the percept the agents get from the environment (see Eq (1)). These final probabilities are given in Fig 6 for the learning tasks with $d_F = 4, 21$.

Tables of Fig 6 show the probability of taking the action “go” for each of the 25 percepts. We focus on the learning tasks with $d_F = 4, 21$, which represent the two most distinctive behaviors that we observe.

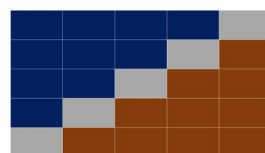
Let us start with the case of $d_F = 21$ (Fig 6(a)), which corresponds to a task where the food is located much further away than the distance reachable with a random walk. In this case, highly aligned swarms emerge as the optimal collective strategy for reaching the food (see also Sec. 3.2 and figures therein), since the orientations of the surrounding neighbors allow the focal

BACK FRONT	Approaching (≥ 3)	Approaching (< 3)	No agent	Receding (< 3)	Receding (≥ 3)
Receding (≥ 3)	0.99±0.01	0.95±0.01	0.81±0.03	0.86±0.02	0.86±0.01
Receding (< 3)	0.90±0.02	0.90±0.02	0.65±0.02	0.77±0.03	0.46±0.03
No agent	0.73±0.03	0.78±0.03	0.77±0.02	0.60±0.03	0.35±0.03
Approaching (< 3)	0.75±0.02	0.72±0.03	0.42±0.04	0.43±0.03	0.16±0.02
Approaching (≥ 3)	0.63±0.03	0.44±0.03	0.36±0.03	0.20±0.03	0.03±0.01

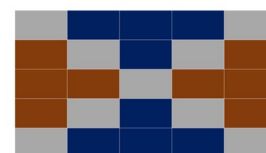
(a) $d_F=21$

BACK FRONT	Approaching (≥ 3)	Approaching (< 3)	No agent	Receding (< 3)	Receding (≥ 3)
Receding (≥ 3)	0.90±0.01	0.84±0.02	0.76±0.02	0.92±0.01	0.87±0.01
Receding (< 3)	0.03±0.00	0.50±0.00	0.50±0.00	0.50±0.00	0.05±0.01
No agent	0.08±0.01	0.50±0.00	0.50±0.00	0.50±0.00	0.13±0.01
Approaching (< 3)	0.17±0.01	0.50±0.00	0.50±0.00	0.50±0.00	0.18±0.02
Approaching (≥ 3)	0.81±0.02	0.82±0.01	0.79±0.02	0.94±0.01	0.80±0.02

(b) $d_F=4$



(c) Ideal probabilities for alignment



(d) Ideal probabilities for cohesion

Fig 6. Learned behavior at the end of the training process. The final probabilities in the agents’ ECM for the action “go” are shown for each of the 25 percepts (5x5 table). (a) and (b) Final probabilities learned in the scenarios with $d_F = 21$ and $d_F = 4$ respectively. The average is taken over 20 ensembles (each learning task) of 60 agents each. Background colors are given to easily identify the learned behavior, where blue denotes that the preferred action for that percept is “go” and orange denotes that it is “turn”. More specifically, the darker the color is, the higher the probability for that action, ranging from grey ($p \simeq 0.5$), light ($0.5 < p < 0.7$) and normal ($0.7 \leq p < 0.9$) to dark ($p \geq 0.9$). Figures (c) and (d) show what the tables would look like if the behavior is purely based on alignment (agent aligns to its neighbors with probability 1) or cohesion (agent goes towards the region with higher density of neighbors with probability 1), respectively. See text for details.

<https://doi.org/10.1371/journal.pone.0243628.g006>

agent to stabilize its orientation against the periodic randomization. The individual responses that lead to such collective behavior can be studied by looking at table (a): the diagonal corresponds to percepts with a clear reaction leading to alignment, i.e. to keep going when there is a positive flow of neighbors and to turn if there is a negative flow. More specifically, one can see that when the agent is in the middle of a swarm and aligned with it, the probability that it

keeps going is 0.99 for dense swarms [percept ($\geq 3_r, \geq 3_a$)] and 0.90 for sparse swarms [percept ($< 3_r, < 3_a$)]. In the same situations, the agent that is not aligned turns around with probability 0.97 for dense swarms [percept ($\geq 3_a, \geq 3_r$)] and 0.57 for sparse swarms [percept ($< 3_a, < 3_r$)]. Outside the diagonal, one observes that the probability of turning is high when a high density of agents are approaching the focal individual from the front (last row) and the agents in the back are not approaching. We can also analyze the learned behavior at the back edge of the swarm, which is important to keep the cohesion of the swarm. When an agent is at the back of a dense swarm and aligned with it [percept ($\geq 3_r, 0$)], the probability of keeping the orientation is 0.81. If instead, the agent is oriented against the swarm [percept ($0, \geq 3_r$)] the probability of turning around to follow the swarm is 0.65. This behavior is less pronounced when the swarm is not so dense [percepts ($< 3_r, 0$), ($0, < 3_r$)], in fact, when a low density of neighbors at the back are receding from the focal agent [percept ($0, < 3_r$)], the focal agent turns around to rejoin the swarm with probability 0.4, which results in this agent leaving the swarm with higher probability. If the agent is alone [percept ($0, 0$)], it keeps going with probability 0.77.

A very different table is observed for $d_F = 4$ (Fig 6(b)). In this task, the food source is located inside the initial region where the agents are placed at the beginning of the trials, so the agents perceive, in general, high density of neighbors around them. For this reason, they rarely encounter the nine percepts encoding low density—that correspond to the ones at the center of the table, with grey background (Table (b) in Fig 6)—throughout the interaction rounds they perform until they get the reward. The corresponding probabilities are the initialized ones, i.e. 1/2 for each action. For the remaining percepts, we observe that the agents have learned to go to the region with higher density of neighbors, which leads to very cohesive swarms (see also Sec. 3.2.2). Since the food source is placed inside the initialization region in this case—which is also within the region agents can cover with a random walk—, there is a high probability that there are several agents already at the food source when an agent arrives there, so they learn to go to the regions with higher density of agents. This behavior can be observed, for instance, for percepts in the first column (high density at the back) and second, third and fourth row (low/no density at the front), where the agents turn around with high probability. In addition, we observe that there is a general bias towards continuing in the same direction, which can be seen for example in percepts with the same density in both regions (e.g. percepts at the corners of the table). The tendency to keep walking is always beneficial in one-dimensional environments to get to the food source (non-interacting agents learn to do so deterministically, as argued for Fig 5). In general, we observe that, in order to find the resource point at $d_F = 4$, agents do not need to align with their neighbors because the food is close enough that they can reach it by performing a Brownian walk.

Fig 6(c) and 6(d) show what the tables would look like if the agents had deterministically (with probability 1) learned just to align with the neighbors (c) or just to go to the region inside the visual range with higher density of neighbors (d). In these figures, percepts for which there is no pronounced optimal behavior have grey background.

In Fig 7, we select four representative percepts that show the main differences in the individual behaviors and plot the average probability of taking the action “go” at the end of a wide range of different learning scenarios where the distance to the food source is increasingly large. We observe that there are two clear regimes with a transition that starts at $d_F = 6$. This is the end of the initial region (see Fig 4, with $V_R = 6$ in our simulations) where the agents are positioned at the beginning of each trial (see S1 Appendix for details on why this transition occurs at $d_F = 6$). The main difference between regimes is that, when the food is placed near the initial positions of the agents, they learn to “join the crowd”, whereas, if the food is placed farther away, they learn to align themselves and “go with the flow”. More specifically, for $d_F < 6$, the orientations of the surrounding neighbors do not play a role, but the agents learn to go

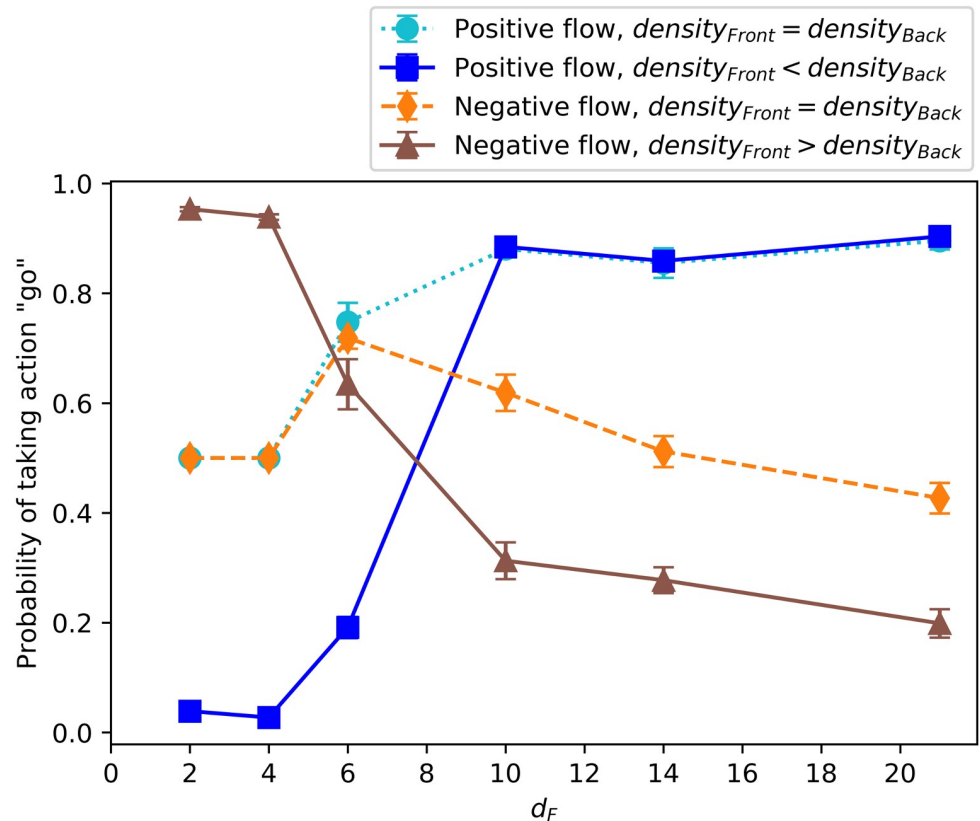


Fig 7. Final probability of taking the action “go” depending on the learning task (increasing distance to food source d_F) for four significant percepts. The percepts are $(< 3_r, < 3_a)$, $(< 3_r, \geq 3_a)$, $(\geq 3_a, < 3_r)$, $(\geq 3_a, \geq 3_r)$, respectively (see legend). The average is taken over the agents’ ECM of 20 independently trained ensembles (1200 agents) at the end of the learning process. Each ensemble performs one task per simulation (d_F does not change during the learning process).

<https://doi.org/10.1371/journal.pone.0243628.g007>

to the region (front/back) with higher number of neighbors, which leads to unaligned swarms with high cohesion. On the contrary, for the tasks with $d_F > 6$, the agents tend to align with their neighbors. This difference in behavior can be observed, for instance, in the dark blue (squares) curve of Fig 7, which corresponds to the percept “positive flow and higher density at the back”. We observe that for $d_F = 2, 4$, the preferred action is “turn” (the probability of taking action “go” is low), since there are more neighbors at the back. However, for $d_F = 10, 14, 21$, the agents tend to continue in the same direction, since there is a positive flow (neighbors have the same orientation as the focal agent). Analogously, the brown curve (triangles) shows the case where there is a negative flow and higher density at the front, so agents trained to find nearby food ($d_F = 2, 4$) have high probability of going, whereas agents trained to find distant food ($d_F = 10, 14, 21$) have high probability of turning.

We remark that, even though the learning task is defined in terms of the distances d_F , the results from this section and S1 Appendix show that the main features of these two types of dynamics do not only depend on the choice of the absolute value of d_F , but also on the location where the food is placed relative to the initial distribution of the agents. Therefore, we observe that the two regimes correspond to the situations when (i) the food is placed within the initial region where the other agents are also situated at the beginning and (ii) when the food is placed far away from the initial region. In case (i) agents learn to rely on the information about the positions of the neighbors and in (ii) they learn to rely on the information about their

orientations. Note that for instance agents trained with $d_F = 4$ have low probability of taking the action “go” (which leads to alignment) if the initial region is $2V_R$ (Fig 7), whereas they take the action “go” with higher probability if the initial region is reduced to V_R (see Fig 1 in S1 Appendix).

In general, we observe that agents with the same motor and sensory abilities can develop very different behaviors in response to different reward schemes. Agents start with the same initial ECM in all the learning scenarios, but depending on the environmental circumstances, in our case the distance to food, some responses to sensory input happen to be more beneficial than others in the sense that they eventually lead the agent to get a reward. For instance, agents that happen to align with their neighbors are the ones that reach the reward when the food is far away, so this response is enhanced in that particular scenario, but not in the one with nearby food.

3.2 Collective dynamics

In this section, we study the properties of the collective motion that emerges from the learned individual responses analyzed in the previous section. We focus on two main properties of the swarms, namely alignment and cohesion. Figs 8 and 9 show the trajectories of the agents of one ensemble before (Fig 8) any learning process and at the end of the learning processes with $d_F = 4, 21$ (Fig 9). One can see that the collective motion developed in the two scenarios differs greatly in terms of alignment and cohesion. Thus, we quantify and analyze these differences in the following.

3.2.1 Alignment. The emergence of aligned swarms as a strategy for reaching distant resources is studied by analyzing the order parameter, defined as

$$\phi = \frac{1}{N} \left| \sum_{i=1..N} v_i \right|, \quad (4)$$

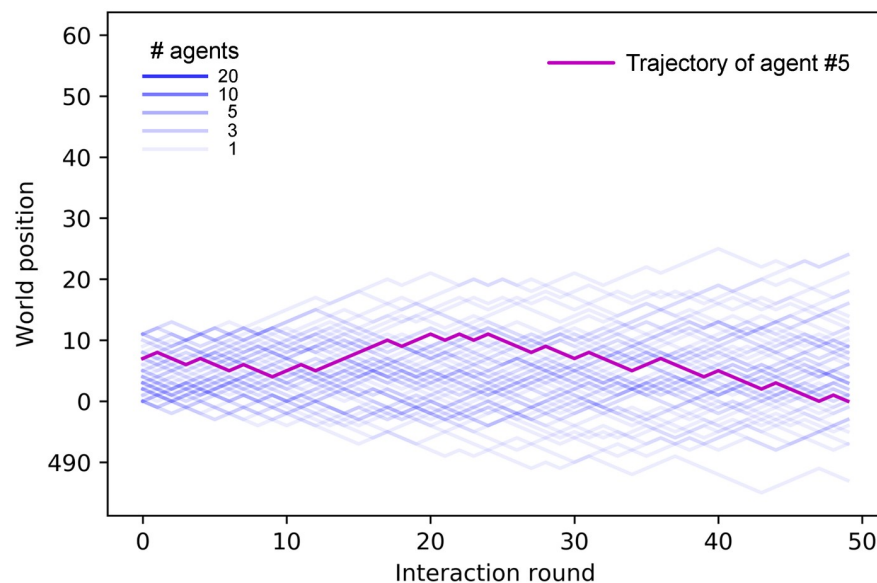
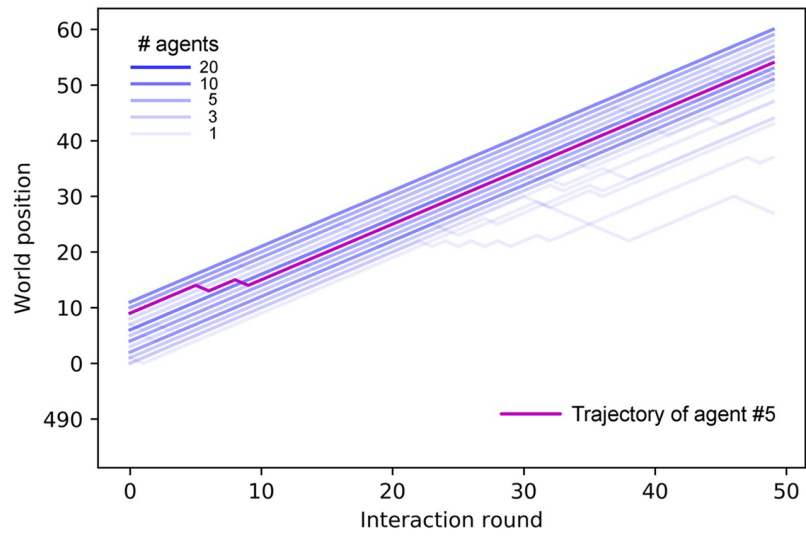
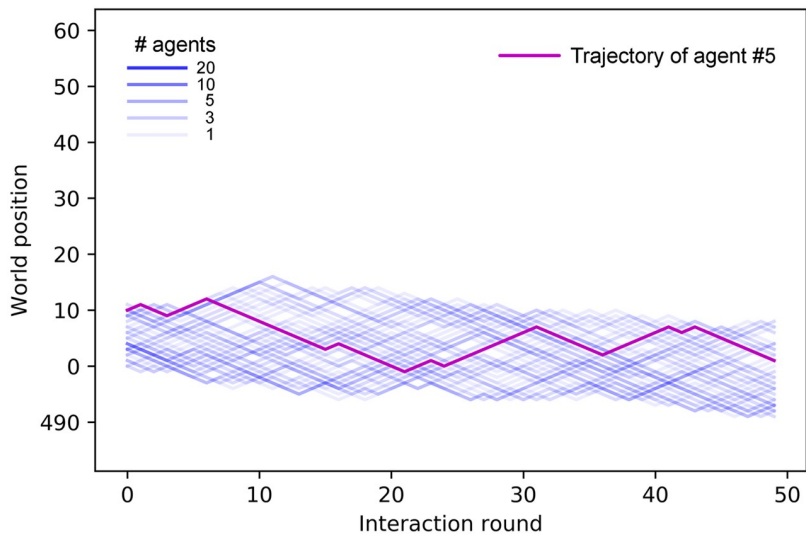


Fig 8. Trajectories (position vs. time) of an ensemble of 60 agents in one trial prior to any learning process. The vertical axis displays the position of the agent in the world and the horizontal axis the interaction round (note that the trial consists of $n = 50$ rounds). Each line corresponds to the trajectory of one agent. However, some agents' trajectories overlap, which is indicated by the color intensity. The trajectory of one particular agent is highlighted for clarity.

<https://doi.org/10.1371/journal.pone.0243628.g008>



(a) Trajectories after training with $d_F=21$



(b) Trajectories after training with $d_F=4$

Fig 9. Trajectories of all agents of an ensemble in the last trial of the learning process for (a) $d_F = 21$ and (b) $d_F = 4$. Ensembles of agents trained to find distant food form aligned swarms (a), whereas agents trained to find nearby food form cohesive, unaligned swarms (b). With the same number of interaction rounds, aligned swarms (a) cover larger distances than cohesive swarms (b). In addition, observe that trajectories in panel (b) spread less than in Fig 8.

<https://doi.org/10.1371/journal.pone.0243628.g009>

where N is the total number of agents and $v_i \in \{1, -1\}$ the orientation of each agent (clockwise or counterclockwise). The order parameter or *global alignment parameter* goes from 0 to 1, where 0 means that the orientations of the agents average out and 1 means that all of them are aligned. In addition, we also evaluate the *local alignment parameter*, since the visual perception of the agent only depends on its local surroundings, and so does the action it takes. In this case, the order parameter ϕ_i is computed for each agent i , considering only the orientation of its neighbors.

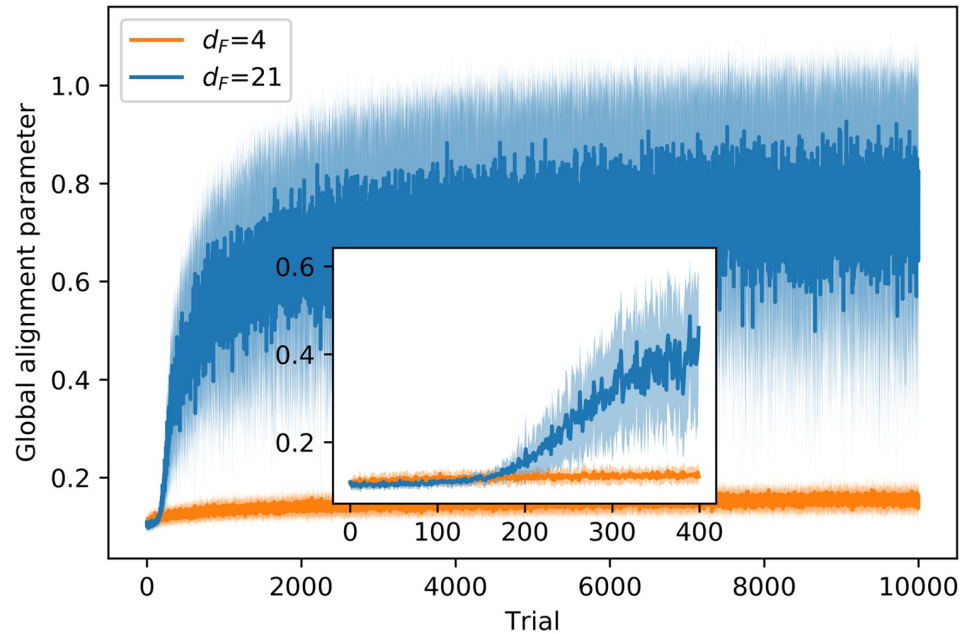


Fig 10. Evolution of the global alignment parameter through the learning processes with $d_F = 4, 21$. At each trial, there is one data point that displays the average of the order parameter, first over all the (global) interaction rounds of the trial and then over 20 different ensembles of agents, where each ensemble learns the task independently. Shaded areas represent one standard deviation.

<https://doi.org/10.1371/journal.pone.0243628.g010>

Fig 10 shows how agents that need to find nearby food do not align, whereas those whose task is to find distant resources learn to form strongly aligned swarms as a strategy for getting the reward, as can be seen from the increase in the order parameter over the course of the training. The inset in Fig 10 shows that agents with the reward at $d_F = 21$ start to align with the neighbors from trial 200, which leads to the conclusion that increasing the alignment is the behavior that allows them to get to the reward (note that the agents start to be rewarded also from trial 200, as can be seen in the inset of Fig 5). The large standard deviation in the $d_F = 21$ case is due to the fact that, in some trials, agents split in two strongly aligned groups that move in opposite directions (see Fig 3 (a) in S1 Appendix for details).

3.2.2 Cohesion. In this section, we study the cohesion and stability of the different types of swarms. In particular, we quantify the cohesion by means of the average number of neighbors (agents within visual range of the focal agent),

$$M = \frac{1}{N} \sum_{i=1}^N m_i, \quad (5)$$

where m_i is the number of neighbors of the i th agent.

Fig 11 shows the evolution of the average number of neighbors through the learning processes with $d_F = 4, 21$. In the training with $d_F = 21$, we observe a decay in M in the first 200 trials, due to the fact that agents start to learn to align locally (see S1 Appendix and Fig 4 therein for details), but the global alignment is not high enough to entail an increase in the average number of neighbors. Therefore, as agents begin to move in straight lines for longer intervals (instead of the initial Brownian motion), they tend to leave the regions with a higher density of agents and M drops. From trial 200 onwards, agents start to form aligned swarms—global alignment parameter increases (see inset of Fig 10)—to get to the food, which leads to an

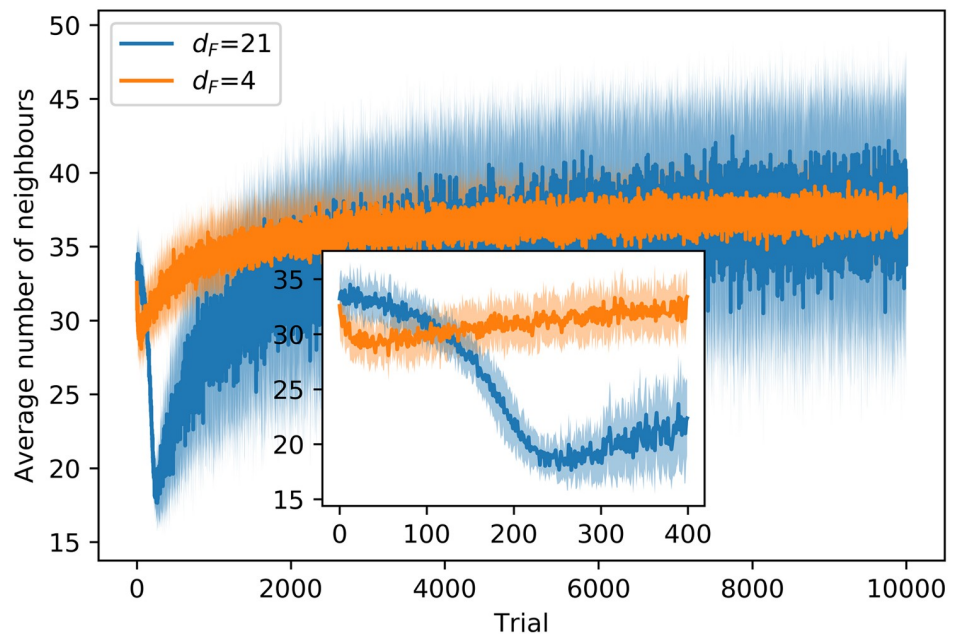


Fig 11. Evolution of the average number of neighbors around each agent through the learning processes with $d_F = 4, 21$. At each trial, there is one data point that displays the average of M , first over all the (global) interaction rounds of the trial and then over 20 different ensembles of agents, where each ensemble learns the task independently. Shaded areas represent one standard deviation.

<https://doi.org/10.1371/journal.pone.0243628.g011>

increase in M (see inset of Fig 11). In the training with $d_F = 4$, agents learn quickly (first 50 trials) to form cohesive swarms, so M increases until a stable value of 36 neighbors is attained.

Up to this point, all the analyses have been done with trials of 50 interaction rounds. However, this is insufficient for assessing the stability of the swarm. For this purpose, we take the already trained ensembles and let them walk for longer trials—the agents do not learn anything new in these simulations, i.e. their ECMs remain unchanged—so that we can analyze how the cohesion of the different swarms evolves with time. We place the agents (one ensemble of 60 agents per simulation) in a world that is big enough so that they cannot complete one cycle within one trial. This resembles infinite environments insofar as agents that leave the swarm have no possibility of rejoining it. This allows us to study the stability of the swarm cohesion and the conditions under which it disperses.

Fig 12 shows the trajectories of ensembles of agents trained with different distances d_F . In the case with $d_F = 21$ (Fig 12(a)), there is a continuous drop of agents from the swarm until the swarm completely dissolves. On the other hand, agents trained with $d_F = 4$ (Fig 12(b)) present higher cohesion and no alignment (see inset of Fig 12(b)). Note that this strong cohesion makes individual trajectories spread less than the Brownian motion exhibited by agents prior to the training (see Fig 8). The evolution of the average number of neighbors throughout the simulation is given in Fig 13, where we compare the cohesion of ensembles of agents trained with $d_F = 2, 4, 21$. In the latter case, the agents leave the swarm continuously, so the average number of neighbors decreases slowly until the swarm is completely dissolved. For $d_F = 2$ ($d_F = 4$) the individual responses are such that the average number of neighbors increases (decreases) in the first 30 rounds until the swarm stabilizes and from then on M stays at a stable value of 57 (35) neighbors. The average number of neighbors is correlated to the swarm size, which we measure by the difference between the maximum and minimum world positions occupied by the agents (modulo world size $W = 500$). As one can see in Fig 12(b), all

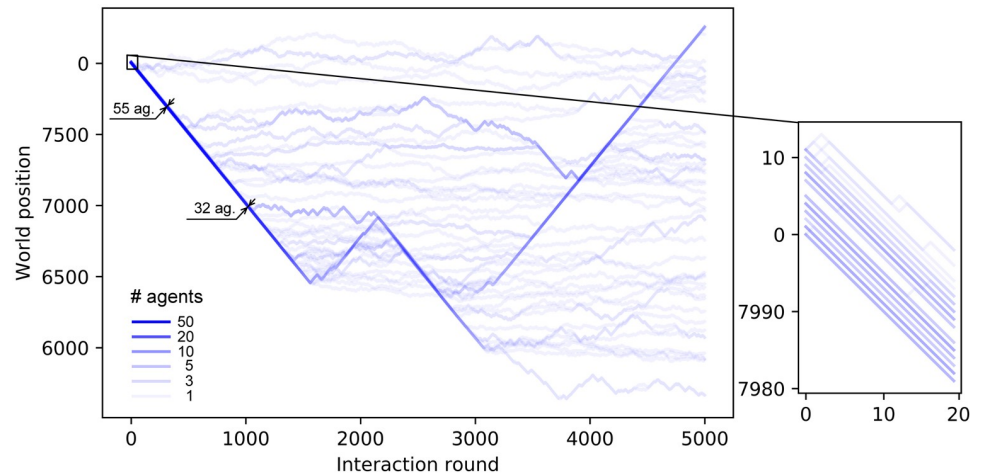
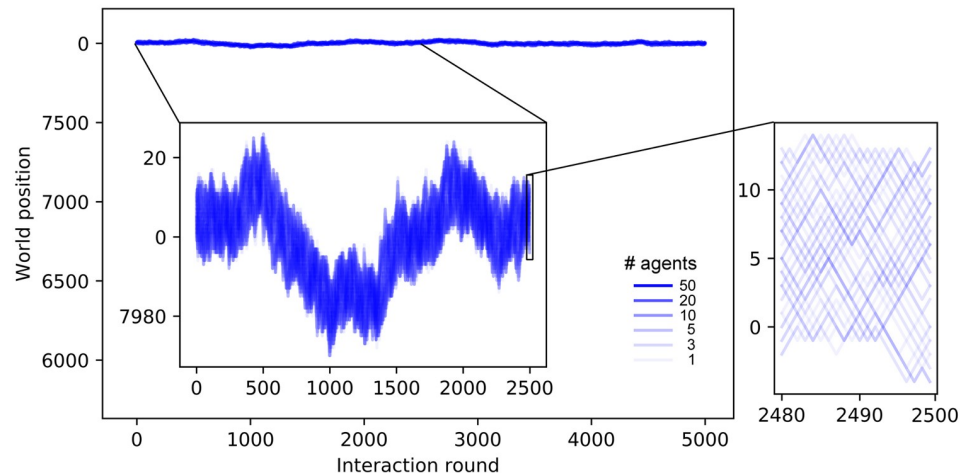
(a) Ensemble trained with $d_F=21$ (b) Ensemble trained with $d_F=4$

Fig 12. Trajectories of an ensemble of 60 agents, in a world of size $W = 8000$, shown over 5000 interaction rounds. (a) Agents trained with $d_F = 21$ form a swarm that continuously loses members until it dissolves completely. (b) Agents trained with $d_F = 4$ form a highly cohesive swarm for the entire trial. The centered inset of this plot shows the first 2500 rounds, with a re-scaled vertical axis to observe the movement of the swarm. Insets on the right zoom in to 20 interaction rounds so as to resolve individual trajectories.

<https://doi.org/10.1371/journal.pone.0243628.g012>

agents remain within the swarm. If the swarm size increases, the average number of neighbors decreases, since the agents are distributed over a wider range of positions. The swarm stabilizes at a given size depending on the individual responses learned during the different trainings. For instance, swarms formed by agents trained with $d_F = 2$ stabilize at swarm sizes of approximately 9 positions, whereas those trained with $d_F = 4$ stabilize at larger swarm sizes (around 17 positions, see e.g. inset of Fig 12(b)), which explains the lower value of M observed in Fig 13 for $d_F = 4$.

3.2.3 Comparison between learning scenarios. Finally, we compare how the alignment and cohesion of the swarms change as a function of the distance at which the resource is placed

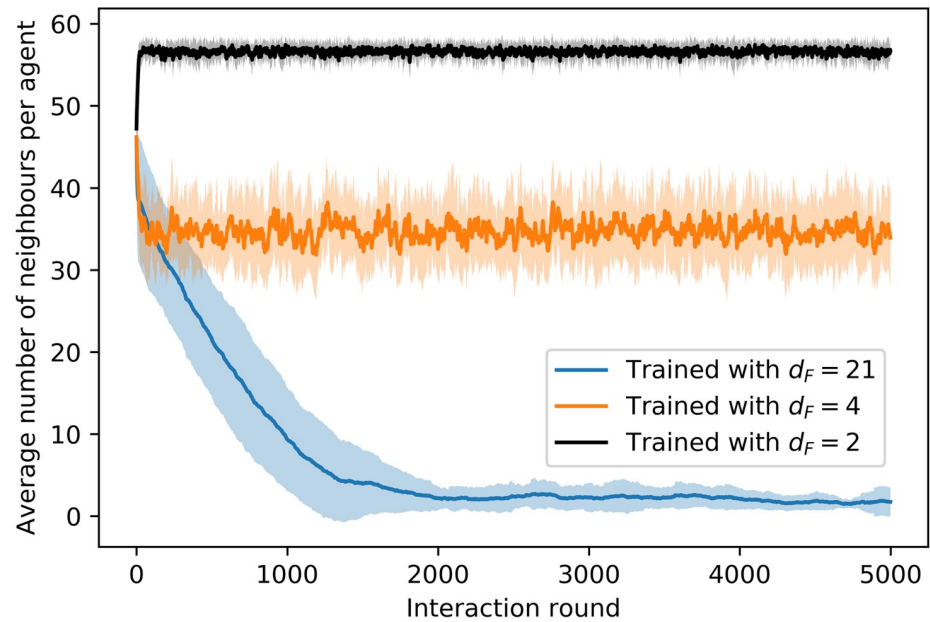


Fig 13. Evolution of the average number of neighbors throughout the trial of 5000 interaction rounds. Average is taken over 20 ensembles of 60 agents each, where for each ensemble the simulation is performed independently. Shaded areas indicates one standard deviation.

<https://doi.org/10.1371/journal.pone.0243628.g013>

in the training. Fig 14 shows the average local and global alignment parameters, together with the average number of neighbors (at the end of the training) as a function of the distance d_F with which the ensembles were trained. We observe that the farther away the resource is placed, the more strongly the agents align with their neighbors (local alignment) in order to reach it. This is directly related to the individual responses analyzed in Fig 7, where one can see that for $d_F \geq 6$ the agents react to positive and negative flow by aligning themselves with their neighbors. Specifically, the observed collective dynamics can be explained in terms of individual responses as follows. The probability of turning around when there is a negative flow and there are not a lot of neighbors (orange-diamonds curve in Fig 7) becomes higher as the d_F increases, from $\simeq 0.3$ at $d_F = 6$ to 0.6 at $d_F = 21$. The change in the other individual alignment responses (in particular, the other curves in Fig 7) is not so large in the region where $d_F > 6$, which suggests that the increase in the local alignment and cohesion we observe for $d_F > 6$ is mostly due to the strength of the tendency the agents have to turn around when there is a negative flow, even when there are not a lot of neighbors. In addition, the lower values of the global alignment parameter observed in the grey (circles) curve in Fig 14 for $d_F \geq 6$ correspond to the behavior analyzed in Sec. 3.2.1, where it is shown that strongly aligned swarms split into two groups in some of the trials (see S1 Appendix for details). With respect to the average number of neighbors, we observe that almost all the agents are within each other's visual range when $d_F = 2$. As d_F increases, swarms become initially less cohesive, but once $d_F > 6$, they become strongly aligned and consequently once again more cohesive (see discussion in Sec. 3.2.2 and S1 Appendix for details).

3.3 Foraging efficiency

In this section, we study how efficient each type of collective motion is for the purpose of foraging. First, we perform a test where we evaluate how the trained ensembles explore the

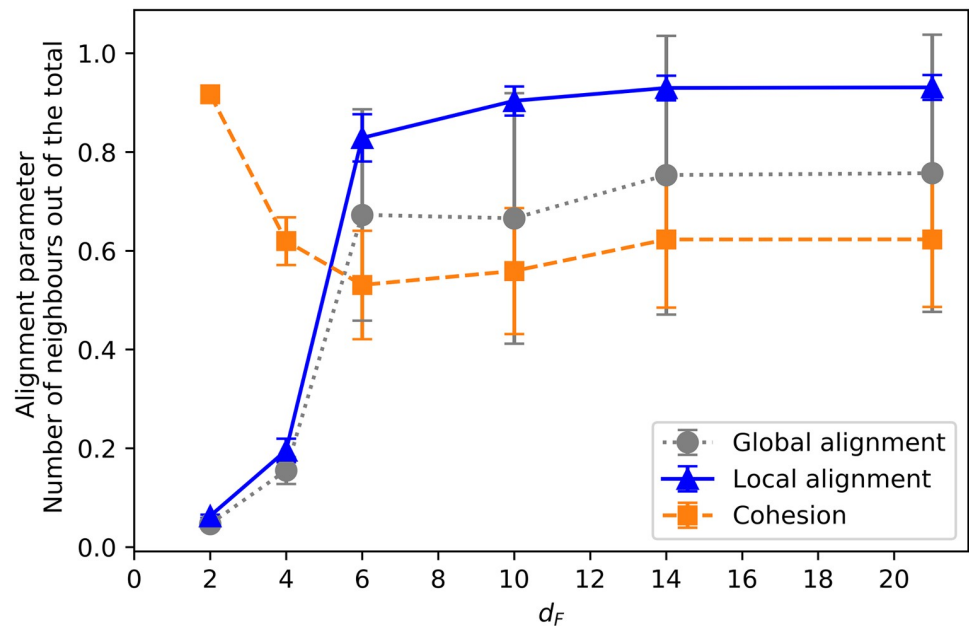


Fig 14. Average number of neighbors (in percentage), global and local alignment parameter as a function of the distance d_F . Note that d_F is the distance to the point where food is placed during the training. Each point is the average of the corresponding parameter over all interaction rounds (50) of one trial, and over 100 trials. 20 already trained ensembles are considered.

<https://doi.org/10.1371/journal.pone.0243628.g014>

different world positions. For this test, we analyze which positions in the world are visited by which fraction of agents. The results are given in Fig 15. We observe that, for positions within the initial region, agents trained with $d_F = 4$ perform better than the others, since they do a random walk that allows them to explore all these positions exhaustively (as evidenced by high percentages of agents that explored positions before the edge of the initial region in Fig 15). On the other hand, agents trained with $d_F = 6$, 21 perform worse when exploring nearby regions, since they form aligned swarms and move straight in one direction. This behavior prevents the agents that are initialized close to the edge of the initial region from exploring the positions inside it. The closer the position is to the edge of the initial region, the more agents visit it because they pass through it when traveling within the swarm. Thus, we conclude that the motion of these swarms is not the optimal to exploit a small region of resources that are located close to each other (a patch).

Non-interacting (n.i.) agents trained with $d_F = 21$ perform slightly better at the intermediate distances than agents trained with $d_F = 4$, since they typically travel five steps in a straight line before being randomly reoriented, thereby covering an expected total of 16 positions in one trial (see Sec. 3). Both curves (grey diamonds and orange squares) show a faster decay in this region than the other two cases ($d_F = 6, 21$), which is due to the fact that agents do not walk straight for long distances in these two types of dynamics, since they do not stabilize themselves by aligning.

Agents trained with $d_F = 21$ reach the best performance for longer distances. In particular, their performance is always better than the performance of agents trained with $d_F = 6$, showing that the strategy developed by agents trained with $d_F = 21$, namely strong alignment, is the most efficient one for traveling long distances (distance from patch to patch). Agents trained with $d_F = 6$ do not align as strongly (see local alignment curve in Fig 14) and there are more agents that leave the swarm before reaching the furthest positions (see also Fig 3 in

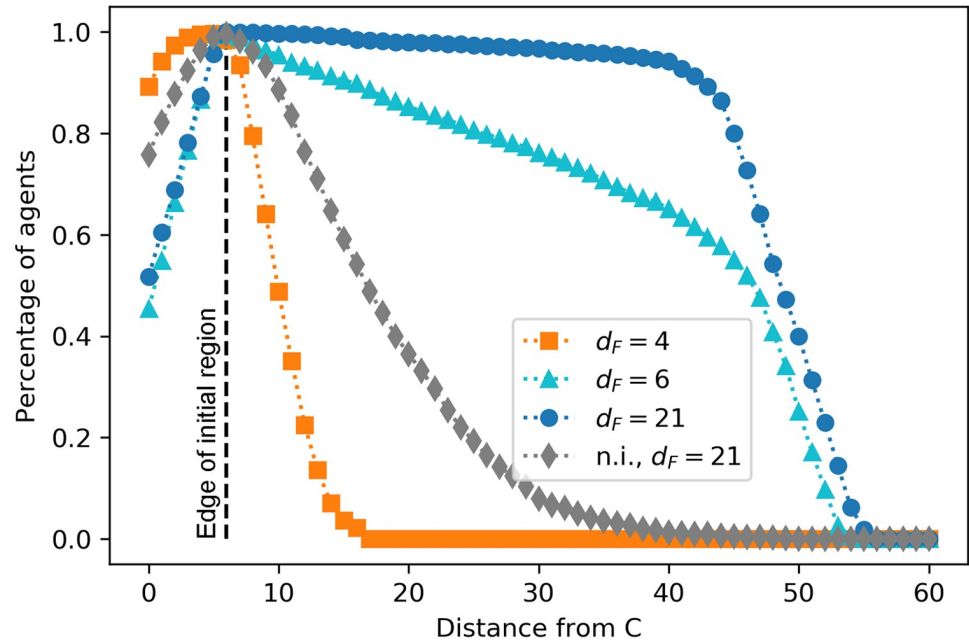


Fig 15. Percentage of agents that visit the positions situated at a distance from C given on the horizontal axis. Since C is located at world position 6 (see Fig 4), a distance of e.g. 10 on the horizontal axis refers to the world positions 16 and 496. The already trained ensembles walk for one trial of 50 interaction rounds. For each of the four trainings (see legend), the performance of 20 ensembles is considered.

<https://doi.org/10.1371/journal.pone.0243628.g015>

S1 Appendix), which explains the lower performance at intermediate/long distances (the light blue curve (triangles) has a linear decrease that is stronger in this region than the dark blue curve (circles)). Note that the maximum distance reached by agents is 56; this is simply due to the fact that each trial lasts 50 rounds and the initial positions are within $C \pm 6$ (see Fig 4).

In addition, we study the swarm velocity for the different types of collective motions. To do so, we compute the average net distance traveled per round. Considering that the swarm walks for a fixed number of rounds (50), we define the normalized swarm velocity as,

$$\langle \xi \rangle = \frac{1}{N} \sum_{i=1}^N \frac{s_i}{50}, \tag{6}$$

where N is the number of agents and s_i is the net distance traveled by the i th agent from the initial position $(x_{i,(r=1)})$ to the final position after 50 interaction rounds $(x_{i,(r=50)})$, that is,

$$s_i = \min((x_{i,(r=50)} - x_{i,(r=1)}) \bmod W, (x_{i,(r=1)} - x_{i,(r=50)}) \bmod W), \tag{7}$$

where r stands for interaction round and W is the world size. Note that the maximum distance agents can travel in 50 rounds is 50 because they move at a fixed speed of 1 position per round.

Fig 16 displays the swarm velocity as a function of the distance d_F at which food was placed during the learning process. Agents trained to find distant resources (e.g. $d_F = 14, 21$) are able to cover a distance almost as large as the number of rounds for which they move. However, while the ensembles trained to find nearby resources (e.g. $d_F = 2, 4$) form very cohesive swarms, they are less efficient in terms of net distance traveled per interaction round. We observe that the transition between the two regimes happens at $d_F = 6$ —corresponding to the

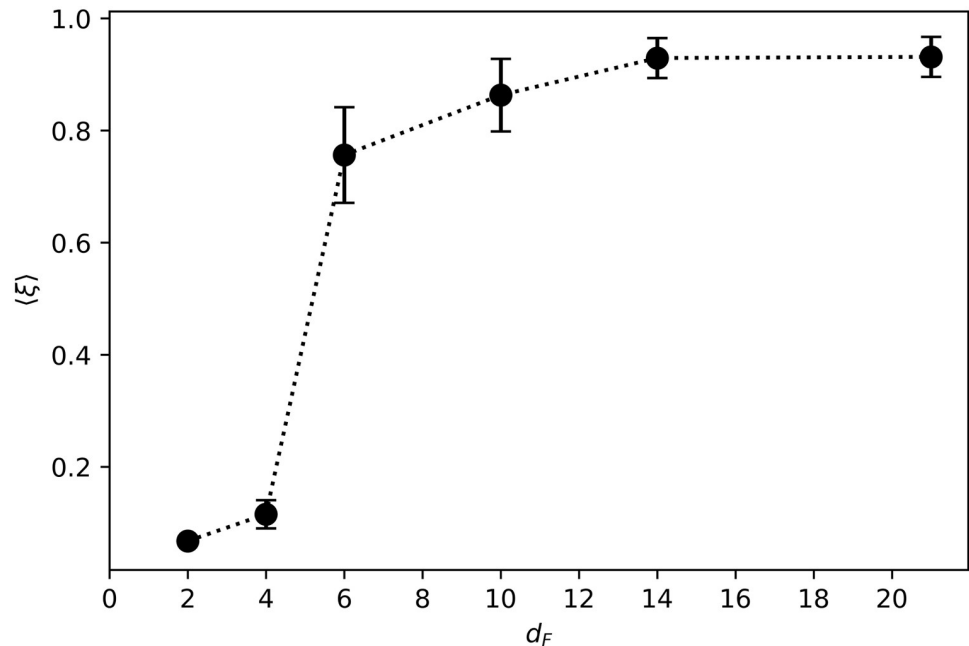


Fig 16. Swarm velocity $\langle \xi \rangle$ as a function of the training distance d_F . Each point is the average over the agents of 20 independently trained ensembles that have performed 50 independent trials each.

<https://doi.org/10.1371/journal.pone.0243628.g016>

end of the initialization region—, which is consistent with the transitions observed in Figs 7 and 14 (see discussion in S1 Appendix for more details).

4 Results II: Analysis of the trajectories

In this section, we analyze the individual trajectories that result from the different types of swarm dynamics. In order to gather enough statistics, we consider ensembles of agents that have been trained under various conditions, as described above, and let them walk for longer trials so that the individual trajectories are long enough to obtain reliable results. During this process, the agents do not learn anything new anymore; that is, the agents' ECMs remain as they are at the end of the training. Thus, we study the trajectories that emerge from the behavior at the end of the learning process, which can be interpreted as the behavior developed on the level of a population in order to adapt to given environmental pressures. The individuals' capacity for learning does not play a role in this analysis.

We focus on the two most representative types of swarms we have observed, i.e. the swarms that emerge from the training with close resources (e.g. $d_F = 4$), characterized by strong cohesion; and the swarms that result from the training with distant resources (e.g. $d_F = 21$), characterized by strong alignment. For easier readability, in the following we will refer to the swarms formed by agents trained with $d_F = 4$ as *cohesive swarms*, and to the swarms formed by agents trained with $d_F = 21$ as *aligned swarms*.

In the simulations for this analysis, we let each ensemble of agents perform 10^5 interaction rounds—where each agent moves one position per interaction round—in a world of size $W = 500$ and analyze the individual trajectories. An example of such individual trajectories for the case of agents trained with $d_F = 21$ is given in Fig 17. We observe that some agents leave the swarm at certain points; however, due to the 'closed' nature of our world model, they have the possibility of rejoining the swarm once it completes a cycle and starts a new turn around the world. Due to these environmental circumstances, the agents exhibit two movement

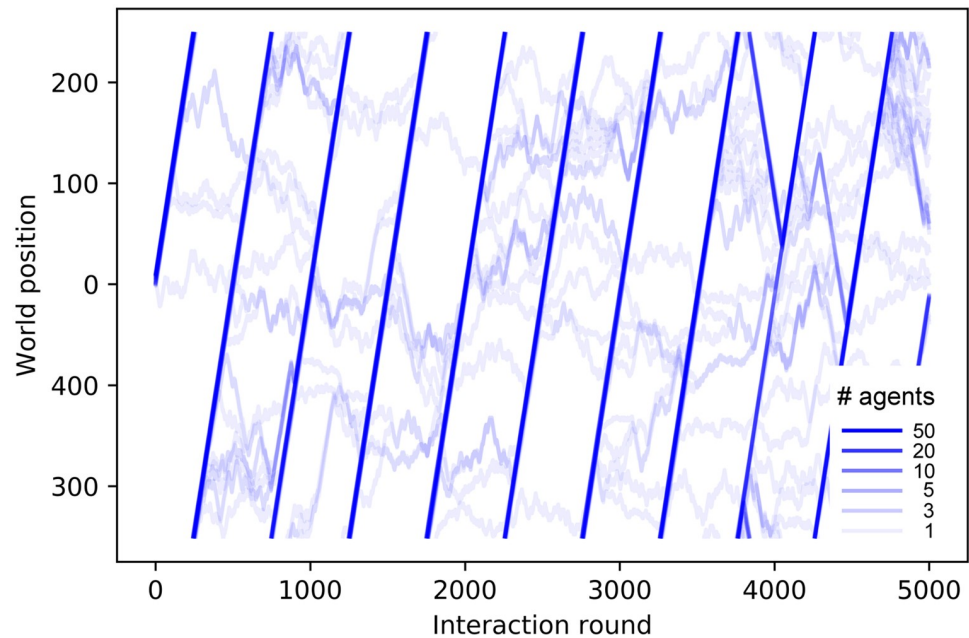


Fig 17. Trajectories of one ensemble of 60 agents that were trained with $d_p = 21$. The world size is $W = 500$. Color intensity indicates the number of agents following the same trajectory, i.e. moving within the swarm. Some agents leave the swarm and then rejoin it when the swarm completes the cycle and starts a new turn. Only the first 5000 interaction rounds (of a total of 10^5) are shown.

<https://doi.org/10.1371/journal.pone.0243628.g017>

modes: when they are alone and when they are inside the swarm. By looking at Fig 17, one can see how agents exhibit directional persistence when they move within the swarm, since they have learnt to align themselves with their neighbours as a strategy for stabilizing their orientations. However, trajectories become more tortuous as agents leave the swarm and walk on their own. Note that it is only possible for individuals to leave the swarm [47] because of the weaker cohesion exhibited by aligned swarms (see Sec. 3.2.2). This bimodal behavior can occur in nature (see e.g. collective motion and phase polyphenism in locusts [48, 49]), where individuals may benefit from collective alignment, for instance, to travel long distances in an efficient way, but they move independently to better explore nearby resources (see Sec. 3.3 for details on exploration efficiency of the different collective dynamics).

In the following sections, we characterize the trajectories and assess how well the agents' movement patterns fit to well-known foraging models such as Lévy walks or composite correlated random walks.

4.1 Theoretical foraging models

This work is directly related to foraging theory, since the task we set for the learning process is to find food in different environmental conditions. For this reason, we will analyze our data to determine whether the movement patterns that emerge from this learning process support any of the most prominent search models. For environments with scarce resources (e.g. patchy landscapes), these models are the Lévy walks [22] and the composite correlated random walks (CCRW) [23].

In order to analyze the trajectories and determine which type of walk fits them best, the distribution of step lengths is studied, where a *step length* is defined as the distance between two consecutive locations of an organism. Intuitively, the optimal strategy for navigating a patchy

landscape allows for both an exhaustive exploration inside patches and an efficient displacement between patches, employing some combination of short and long steps. Lévy walks have a distribution of step lengths in which short steps have higher probability of occurrence but arbitrarily long steps can also occur due to its power-law (PL) tail. In two- and three-dimensional scenarios, the direction of motion is taken from a uniform distribution from 0 to 2π , which implies that Lévy walks do not consider directionality in the sense of correlation in direction between consecutive steps [32]. On the other hand, CCRW and the simpler version thereof, composite random walks (CRW), consist of two modes, one intensive and one extensive, which are mathematically described by two different exponential distributions of the step lengths. The intensive mode is characterized by short steps (with large turning angles in 2D) to exploit the patch, whereas the extensive mode—whose distribution has a lower decay rate—is responsible for the inter-patch, straight, fast displacement. CCRW in addition allow for correlations between the directions of successive steps.

Even though the models are conceptually different, the resulting trajectories may be difficult to distinguish [24, 50, 51], even more if the data is incomplete or comes from experiments where animals are difficult to track. In the past years, many works have been published that try to provide techniques to uniquely identify Lévy walks [52–54] and to differentiate between the two main models [24, 55, 56]. For instance, some of the experiments that initially supported the hypothesis that animals perform Lévy walks [25, 26, 57] were later reanalyzed to support the conclusion that more sophisticated statistical techniques are, in general, needed [27, 28, 55, 58]. Apart from that, there exist several studies that relate different models of collective dynamics to the formation of Lévy walk patterns under certain conditions [59, 60]. For instance, it has been shown [61] that Lévy walk movement patterns can arise as a consequence of the interaction between effective leaders and a small group of followers, where none of them has information about the resource.

In our study, we consider the three models we have already mentioned (PL, CCRW and CRW), together with Brownian motion (BW) as a baseline for comparison. Since our model is one-dimensional, a distribution of the step lengths is sufficient to model the trajectories we observe, and no additional distributions, such as the turning angle distributions, are needed. In addition, the steps are unambiguously identified: a step has length ℓ if the agent has moved in the same direction for ℓ consecutive interaction rounds. Finally, since space in our model is discretized, we consider the discrete version of each model's probability density function (PDF). More specifically, the PDFs we consider are,

1. Brownian motion (BW):

$$p(\ell) = (1 - e^{-\lambda}) e^{-\lambda(\ell-1)}, \quad \ell \geq 1, \quad (8)$$

where λ is the decay rate and the minimum value a step length can have is, in our case, known to be 1, since agents move at a constant speed of one position per interaction round.

2. Composite random walk (CRW):

$$p(\ell) = p(1 - e^{-\beta_I}) e^{-\beta_I(\ell-1)} + (1 - p)(1 - e^{-\beta_E}) e^{-\beta_E(\ell-1)}, \quad \ell \geq 1, \quad (9)$$

where p is the probability of taking the intensive mode, β_I is its decay rate and β_E is the decay rate of the extensive mode. In this case, again, the minimum step length is 1.

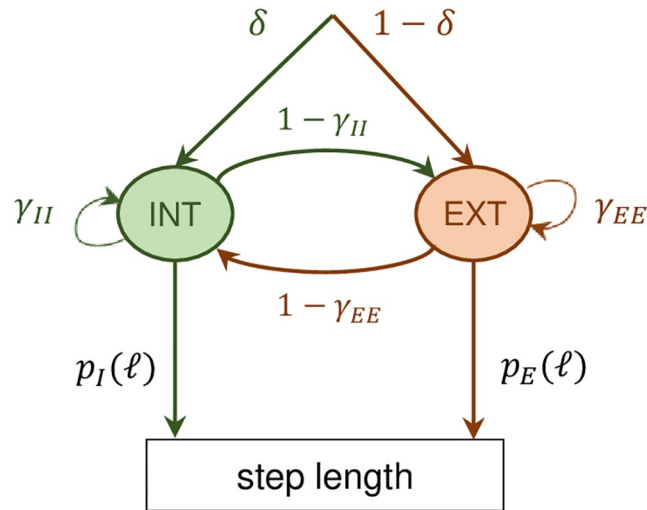


Fig 18. Hidden Markov model for the CCRW. There are two modes, the intensive and the extensive, with probability distributions given by p_I and p_E (see text for details). The probability of transition from the intensive (extensive) to the extensive (intensive) mode is given by $1 - \gamma_{II}$ ($1 - \gamma_{EE}$), where γ_{II} and γ_{EE} are the probabilities of remaining in the intensive and extensive mode respectively. δ is the probability of starting in the intensive mode.

<https://doi.org/10.1371/journal.pone.0243628.g018>

3. Composite correlated random walk (CCRW):

$$p_I(\ell|I) = (1 - e^{-\lambda_I}) e^{-\lambda_I(\ell-1)}, \quad \ell \geq 1, \tag{10}$$

$$p_E(\ell|E) = (1 - e^{-\lambda_E}) e^{-\lambda_E(\ell-1)}, \quad \ell \geq 1, \tag{11}$$

$$p(m' = E|m = I) = 1 - \gamma_{II}, \tag{12}$$

$$p(m' = I|m = E) = 1 - \gamma_{EE}, \tag{13}$$

where $p_I(\ell|I)$ and $p_E(\ell|E)$ are the PDFs of the step lengths ℓ corresponding to the intensive and extensive mode respectively. Denoting the mode in which the agent is as m and the mode to which the agent transitions as m' , $p(m' = E|m = I)$ is the transition probability from the intensive to the extensive mode and $p(m' = I|m = E)$, from the extensive to the intensive mode. λ_I , λ_E , γ_{II} and γ_{EE} are parameters of the model. The main difference between the CRW and the CCRW models is that, in the latter, the step lengths are correlated, i.e. the order of the sequence of step lengths, and thus the order in which the movement modes alternate, matters. The CCRW is modeled as a hidden Markov model (HMM) (see [56, 62]) with two modes, the intensive and the extensive. Fig 18 shows the details of the model and the notation for the transition probabilities between modes.

4. Power-law (PL):

$$p(\ell) = \frac{\ell^{-\mu}}{\zeta(\mu, 1)}, \quad \ell \geq 1, \tag{14}$$

where the normalization factor $\zeta(\mu, 1) = \sum_{a=0}^{\infty} (a + 1)^{-\mu}$ is the Hurwitz zeta function [63]. The parameter μ gives rise to different regimes of motion: Lévy walks are characterized by a heavy-tailed distribution, with exponents $1 < \mu \leq 3$, which produces superdiffusive trajectories, whereas $\mu > 3$ corresponds to normal diffusion, as exhibited by Brownian walks. We

note that the above distribution starts at $\ell = 1$, which is the shortest possible distance that our agents move in a straight line. The scale of this minimum step length is determined by the embodied structure of the organism and is typically considered to be one body length [32]. Some other works (e.g. [53, 63]) consider a variant of the above distribution that only follows the PL form for steps longer than some threshold ℓ_0 , for example when analysing experimental data that become increasingly noisy at short step-lengths. However, since the step lengths resulting from our simulations are natively discrete, the unbounded PL distribution given in Eq (14) seems appropriate. Moreover, if one were to introduce a lower bound $\ell_0 > 1$, one would need to add more parameters in the model to account for the probabilities $p(\ell)$ for all $1 \leq \ell < \ell_0$, which we consider an unnecessary complication. This is particularly relevant when it comes to comparing PL to BW, CRW or CCRW as models for fitting our data: since none of the other models include lower bounds, we achieve a more consistent comparison by a parsimonious approach that includes all step lengths $\ell \geq 1$ in the PL model and thereby abstains from additional free parameters.

4.2 Visual analysis

In this section, we study the general characteristics of the trajectories of both types of swarm dynamics. We start by analyzing how diffusive the individual trajectories are depending on whether the agents belong to an ensemble trained with $d_F = 21$ (dynamics of aligned swarms) or $d_F = 4$ (dynamics of cohesive swarms). More specifically, we analyze the mean squared displacement (MSD), defined as,

$$\langle \delta r^2 \rangle = \langle |x(t) - x_0|^2 \rangle, \quad (15)$$

where x_0 is the reference (initial) position and $x(t)$ is the position after time t elapsed. In general, the MSD increases with the time elapsed as $\langle \delta r^2 \rangle \sim t^\alpha$. Depending on the exponent α , the diffusion is classified as normal diffusion ($\alpha = 1$), subdiffusion ($\alpha < 1$) or superdiffusion ($\alpha > 1$), which is called ballistic diffusion when $\alpha = 2$. For instance, a Brownian particle undergoes normal diffusion, since its MSD grows linearly with time.

Fig 19 shows that the dynamics of aligned swarms leads to superdiffusive individual trajectories (ballistic, with $\alpha = 2$), whereas the trajectories of agents that belong to cohesive swarms exhibit close-to-normal diffusion. The anomalous diffusion (superdiffusion) exhibited by the agents trained with $d_F = 21$ (curve with blue circles in Fig 19) favors the hypothesis that the swarm behavior may induce Lévy-like movement patterns, since Lévy walks are one of the most prominent models describing superdiffusive processes. However, CCRW can also produce superdiffusive trajectories [23, 24]. In contrast, agents trained with $d_F = 4$ do not align with each other and the normal diffusion shown in Fig 19 is indicative of Brownian motion.

The analysis presented above already shows a major difference between the two types of swarm dynamics but it is in general not sufficient to determine which theoretical model (Lévy walks or CCRW) best fits the data from aligned swarms. According to [24], one possible way to distinguish between composite random walks and Lévy walks is to look at their survival distributions, which is the complement of the cumulative distribution function, giving the fraction of steps longer than a given threshold. Lévy walks would exhibit a linear log-log relationship when this type of distribution is plotted, whereas CCRW exhibit a non-linear relationship. Fig 20 compares the survival distributions of two trajectories, one from each type of swarm, to those predicted by the best-fitting models of each of the four classes. The maximum length observed in the $d_F = 4$ -trajectory is of the order of 10, whereas in the case of the $d_F = 21$ -trajectory, it is one order of magnitude larger. The most prominent features one infers

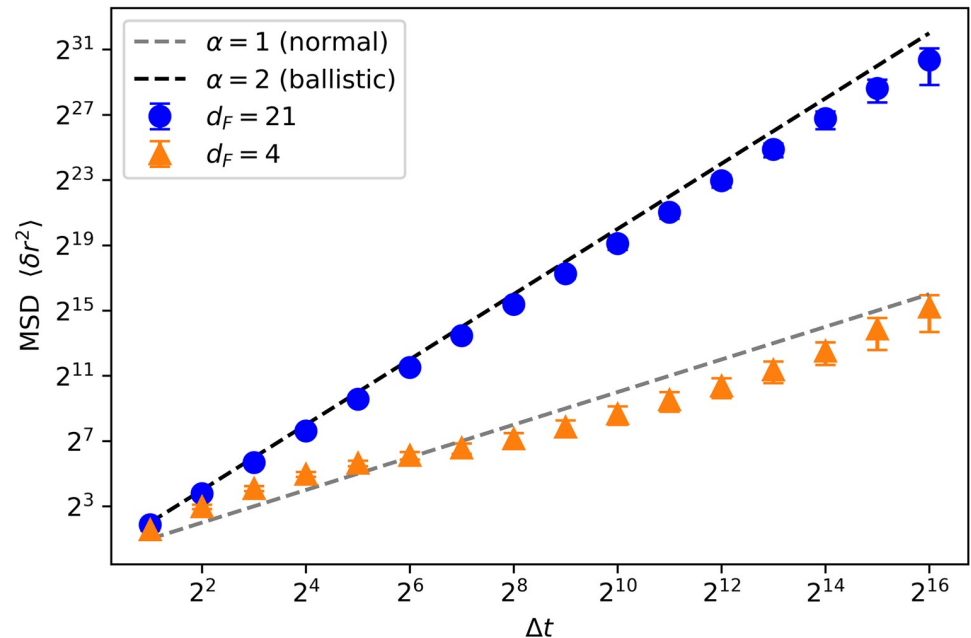


Fig 19. Mean squared displacement (MSD). Log-log (base 2) plot of the MSD as a function of the time interval for two types of trajectories: trajectories performed by agents trained with $d_F = 21$ (blue curve, circles) and by agents trained with $d_F = 4$ (orange curve, triangles). We observe that the former present ballistic diffusion, whereas the latter exhibit close-to-normal diffusion. 600 individual trajectories (10 ensembles of 60 agents) are considered for each case.

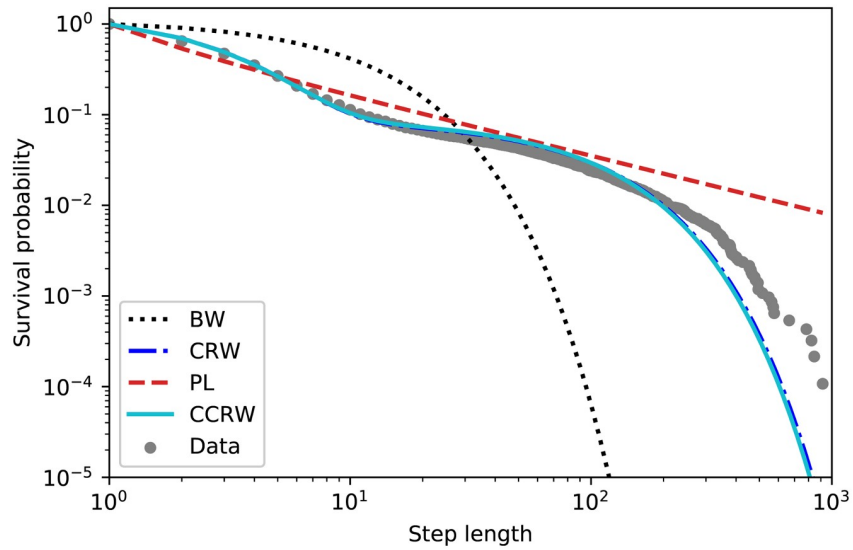
<https://doi.org/10.1371/journal.pone.0243628.g019>

from these figures are that all models except PL seem to fit the data of the $d_F = 4$ -trajectory, and that Brownian motion is clearly not a good model to describe the $d_F = 21$ -trajectory. In addition, Fig 20(a) is curved and seems to be better fit by the CCRW. However, when other trajectories of agents trained with $d_F = 21$ are plotted in the same way, we see that data seems to better follow the straight line of the PL rather than the CCRW (see for example Fig 4 in S2 Appendix).

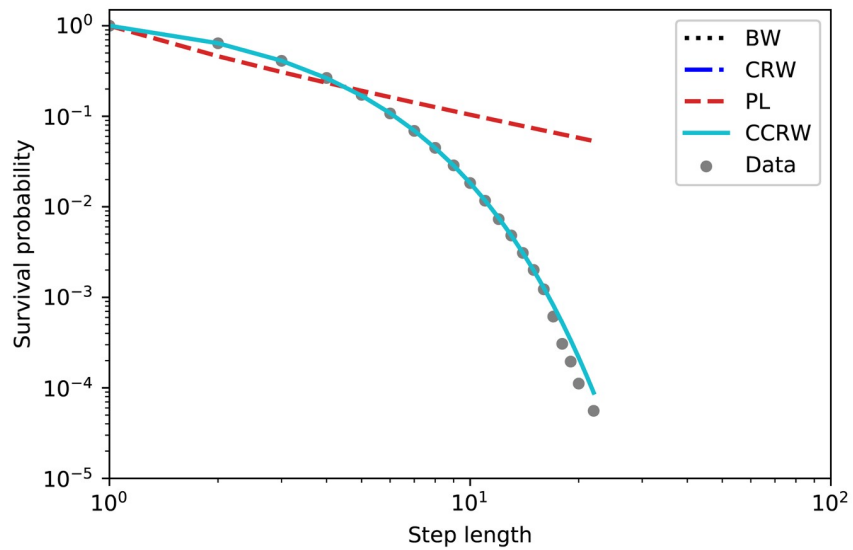
While visual inspection may be an intuitive way of assessing model fit, and one that is easy to apply at small scales, it would be preferable to use a method that yields quantitative and objective, repeatable assessments of how well various models fit a given data set. Moreover, we generated 600 individual trajectories per type of swarm, in order to support statistically meaningful conclusions, and at this scale visual inspection quickly becomes infeasible. For this reason, we now turn to a more rigorous statistical analysis of the individual trajectories.

4.3 Statistical analysis

In order to determine which of the mentioned models best fits our data, we perform the following three-step statistical analysis for each individual trajectory: (i) first, we optimize each family of models to get the PDF that most likely fits our data via a maximum likelihood estimation (MLE) of the model parameters. (ii) Then, we compare the four different candidate models among them by means of the Akaike information criterion (AIC) [64] and (iii) finally, we apply an absolute fit test for the best model. We repeat this analysis for agents trained with $d_F = 4$ and $d_F = 21$, yielding a total of 600 individual trajectories per type of training (10 ensembles of aligned swarms and 10 of cohesive swarms, where each ensemble has 60 agents). The simulation of 10^5 interaction rounds is performed for each ensemble independently. In order to do



(a) Trajectory of an agent that belongs to an aligned swarm.



(b) Trajectory of an agent that belongs to a cohesive swarm.

Fig 20. Survival probability as a function of the step length. The survival probability is the percentage of step lengths larger than the corresponding value on the horizontal axis. Each panel depicts the data from the trajectory of one agent picked from (a) aligned swarms and (b) cohesive swarms, so that this figure represents the most frequently observed trajectory for each type of dynamics. The survival distributions of the four candidate models are also plotted. The distributions for each model are obtained considering the maximum likelihood estimation of the corresponding parameters (see Sec. 4.3 for details). The curve for the CCRW model is obtained by an analytic approximation of the probabilities of each step length, given the maximum likelihood estimation of its parameters. Since the order of the sequence of step lengths is not relevant for this plot, we estimate the probabilities of each step length as

$$p(\ell) = p'(1 - e^{-\hat{\lambda}_I})e^{-\hat{\lambda}_I(\ell-1)} + (1 - p')(1 - e^{-\hat{\lambda}_E})e^{-\hat{\lambda}_E(\ell-1)} \text{ (see Eq (9)) with } p' \simeq \frac{1}{1-\gamma_H} \left(\frac{1}{1-\gamma_H} + \frac{1}{1-\gamma_{EE}} \right)^{-1}.$$

<https://doi.org/10.1371/journal.pone.0243628.g020>

the statistical analysis, each individual trajectory is divided into steps, which are defined in our case as the distance the agent travels without turning. We obtain sample sizes that range from 4000 to 17000 steps for trajectories of agents trained with $d_F = 21$ and from 20000 to 40000 steps in the case with $d_F = 4$.

Table 2. Average values of the MLE parameters for the different models.

Model		$d_F = 21$	$d_F = 4$
BW	$\hat{\lambda}$	0.083 ± 0.032	0.305 ± 0.052
CRW	$\hat{\beta}_I$	0.37 ± 0.15	0.6 ± 1.1
	$\hat{\beta}_E$	0.0126 ± 0.0069	0.302 ± 0.048
	\hat{p}	0.879 ± 0.040	0.196 ± 0.059
PL	$\hat{\mu}$	1.59 ± 0.10	1.657 ± 0.066
CCRW	$\hat{\delta}$	0.23 ± 0.31	0.12 ± 0.13
	$\hat{\lambda}_I$	0.37 ± 0.15	0.36 ± 0.21
	$\hat{\lambda}_E$	0.0134 ± 0.0066	0.300 ± 0.047
	$\hat{\gamma}_{II}$	0.839 ± 0.078	0.749 ± 0.060
	$\hat{\gamma}_{EE}$	0.013 ± 0.020	0.09 ± 0.26

600 trajectories are analyzed for each type of swarm.

<https://doi.org/10.1371/journal.pone.0243628.t002>

The following provides more detail on the analysis, starting with the MLE method, which consists in maximizing the likelihood of each model candidate with respect to its parameters. The likelihood function is generally defined as,

$$\mathcal{L}(\theta|\ell_{i=1..S}) = \prod_{i=1}^S p(\ell_i, \theta), \tag{16}$$

where S is the sample size and $p(\ell_i, \theta)$ is the PDF of the given model —that depends on the model parameters θ — evaluated at the data point ℓ_i . Details on the maximization process and the computation of the likelihood function in the case of CCRW, which is more complicated since consecutive step lengths ℓ_i are not sampled independently, are given in [S2 Appendix](#). In the following, we denote the values of the parameters that maximize the likelihood and the value of the maximum likelihood with hatted symbols.

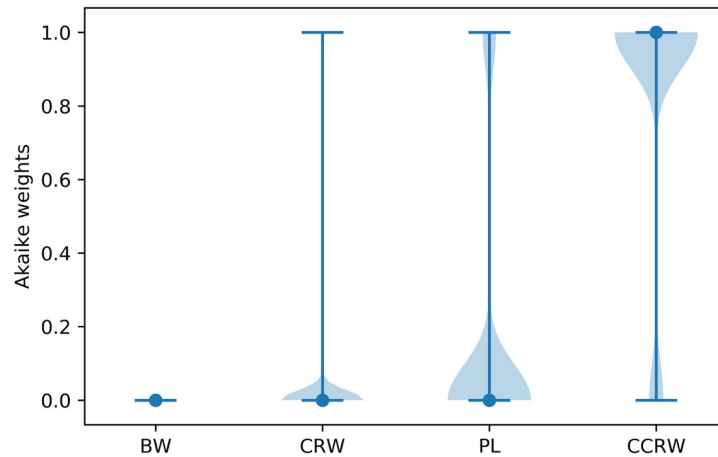
[Table 2](#) shows the MLE parameters we have obtained for each model and for each swarm type. We observe that, in the $d_F = 21$ case, the decay rates of the exponential distributions ($\hat{\lambda}, \hat{\beta}_E, \hat{\lambda}_E$) are very small (approx. of the order of 0.01) compared to the decay rates in the $d_F = 4$ case (approx. of the order of 0.3), which implies that the former allows for longer steps to occur with higher probability. The decay rates of the intensive modes ($\hat{\beta}_I, \hat{\lambda}_I$) are comparable to the BW decay rate of $d_F = 4$ because they account for the shorter, more frequent steps, which occur in both types of dynamics —in the $d_F = 21$ case, agents perform shorter steps when they leave the swarm and move on their own—. Also note that the power-law coefficient $\mu \simeq 1.6$ in the $d_F = 21$ case implies that the PL model is that of a Lévy walk.

Once the value of the maximum likelihood $\hat{\mathcal{L}}$ is obtained for each model, it is straightforward to compute its Akaike value,

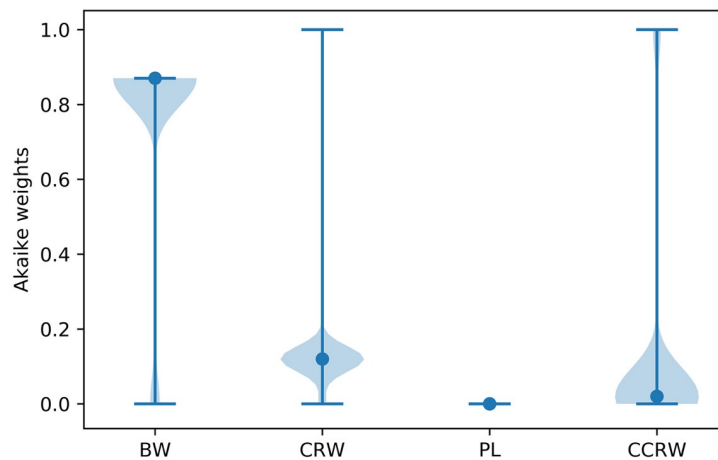
$$AIC = 2k - 2 \ln(\hat{\mathcal{L}}), \tag{17}$$

where k is the number of parameters of the model. The model with the lowest AIC (AIC_{\min}) is the best model (out of the ones that are compared) to fit the data [64]. In order to compare the models in a normalized way, the Akaike weights are obtained from the Akaike values as,

$$w_i = \frac{e^{-\frac{1}{2}\Delta_i(AIC)}}{\sum_{k=1}^K e^{-\frac{1}{2}\Delta_k(AIC)}} \tag{18}$$



(a) Akaike weights of the 600 trajectories of agents trained with $d_F=21$.



(b) Akaike weights of the 600 trajectories of agents trained with $d_F=4$.

Fig 21. Violin plots that represent the Akaike weights obtained for each model. (a) Akaike weights of trajectories of agents trained with $d_F = 21$ (aligned swarms). (b) Akaike weights of trajectories of agents trained with $d_F = 4$ (cohesive swarms). 600 individual trajectories —per type of swarm— were analyzed for each plot. The ‘•’ symbol represents the median and the vertical lines indicate the range of values in the data sample (e.g. PL model in figure (a) has extreme values of 0 and 1). Shaded regions form a smoothed histogram of the data (e.g. the majority of Akaike weights of the CCRW model in figure (a) have value 1, and there are no values between 0.2 and 0.8). See text for more details.

<https://doi.org/10.1371/journal.pone.0243628.g021>

where w_i is the Akaike weight of the i th model, and $\Delta_i(AIC) = AIC_i - AIC_{\min}$, with AIC_i the Akaike value of the i th model. The interpretation of w_i is not straightforward but, as it was argued in [65], “Akaike weights can be considered as analogous to the probability that a given model is the best approximating model”. K is the total number of models under comparison, so that the Akaike weights are normalized as $\sum_{i=1}^K w_i = 1$. In S2 Appendix, we present detailed tables with the results of this statistical analysis for three trajectories, two for training with $d_F = 21$ and one with $d_F = 4$.

Fig 21 shows the results of the Akaike weights obtained for each of the 600 trajectories analyzed for each type of swarm. In the case of the aligned swarms (Fig 21(a)), we observe that the BW model is discarded in comparison to the other models, since its Akaike weight is zero for

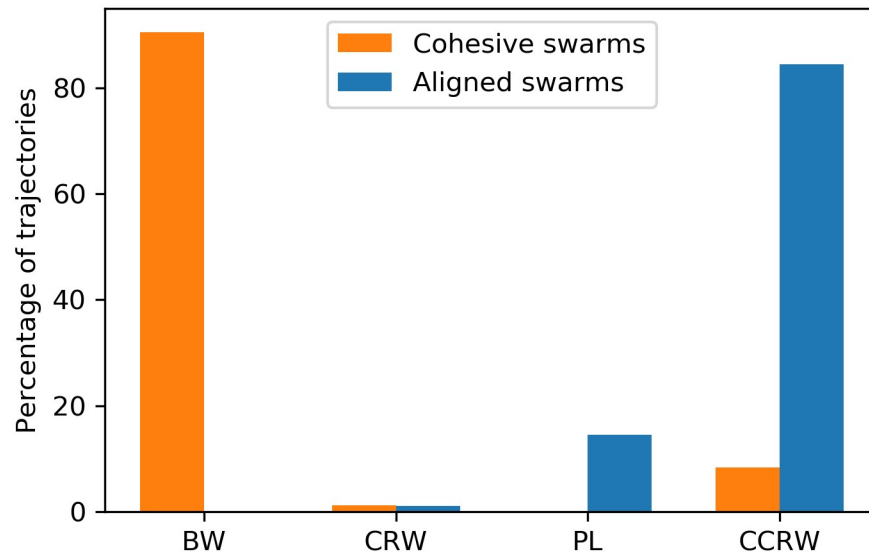


Fig 22. Percentage of trajectories that are best fit by each model according to the BIC criterion. A model is considered to best fit the data of a given trajectory if it has the lowest BIC value and its difference with respect to the rest of the models is larger than 10. 600 individual trajectories —per type of swarm— were analyzed for each histogram.

<https://doi.org/10.1371/journal.pone.0243628.g022>

all trajectories. 85% of trajectories have Akaike weight of 1 for the CCRW model [66] and 0 for the rest of the models, whereas 14% of trajectories have Akaike weight of 1 for the PL model and 0 for the rest.

On the other hand, the cohesive swarms (Fig 21(b)) show high Akaike weights for all models except the PL. 92% of trajectories have Akaike weights 0.87, 0.12 and 0.01 for the BW, CRW and CCRW models, respectively. The remaining 8% of trajectories have $w_{CCRW} = 1$.

In order to exclude the worry that the CCRW and CRW are chosen even though they have more parameters than the power-law or the BW, we also consider the Bayesian information criterion (BIC), which penalizes more strongly the number of parameters of the model. The BIC value is given by,

$$BIC = k \ln(S) - 2 \ln(\hat{\mathcal{L}}), \quad (19)$$

where S is the sample size, k the number of parameters of the model and $\hat{\mathcal{L}}$ the maximum likelihood. Fig 22 shows the results of this analysis, where we consider that a model best fits the data of a given trajectory if it has the lowest BIC value (BIC_{\min}) and its difference with respect to the rest of the models is larger than 10 ($\Delta_i BIC > 10$, where $\Delta_i BIC = BIC_i - BIC_{\min}$) [64].

Fig 22 leads to the same conclusions regarding the aligned swarms, i.e. 85% of the trajectories are best fit by CCRW and 14% are best fit by PL. However, for the cohesive swarms, the BW is the best model to fit the data of 90% of the trajectories according to the BIC criterion. Therefore, we conclude that the dynamics shown by the cohesive swarms lead to individual Brownian motion. Regarding the aligned swarms, our results are in agreement with previous works that claim that “selection pressures give CCRW Lévy walk characteristics” [67]. The majority of individual trajectories are best fit by CCRW with two exponential distributions whose means are $\hat{\lambda}_I^{-1} \simeq 2.7$ and $\hat{\lambda}_E^{-1} \simeq 75$, which give the movement patterns Lévy-walk features. In addition, a considerable percentage of trajectories are indeed best fit by a power-law distribution with exponent $\mu = 1.6$, that is, a Lévy walk.

Finally, we study the goodness-of-fit (GOF) of the different models. For models that deal with i.i.d. variables (BW, CRW, PL), it is enough to perform a likelihood ratio test, whose p -value indicates how well the data is fit by the model. Within our framework, a low p -value, namely $p < 0.05$, means that the model can be rejected as a description for the observed data with a confidence of 95%. The closer p is to 1, the better the model fits the data. In the case of CCRW, a more involved method is needed due to the correlation in the data. Specifically, we first compute the uniform pseudo-residuals (see [62]) and then we perform a Kolmogorov-Smirnov (KS) test to check for uniformity of the mid-pseudo-residuals. Details on the methods used in both GOF tests are given in S2 Appendix. Even though a visual inspection of Fig 20 suggests that the CRW, PL, CCRW models fit the data reasonably well, a quantitative analysis gives p -values of $p < 0.01$ for most of the trajectories fitted by the BW and PL models. Some trajectories fitted by the CRW model give better fittings, e.g. the best p -values are $p = 0.97$ and $p = 0.36$ for a trajectory of an agent trained with $d_F = 4$ and $d_F = 21$, respectively. In the CCRW case, we give the average value of the KS distance obtained in the 600 trajectories, which is $D_{KS} = 0.134 \pm 0.016$ and $D_{KS} = 0.189 \pm 0.046$ for $d_F = 4$ and $d_F = 21$ -trajectories respectively (note that a perfect fit gives $D_{KS} = 0$ and the worst possible fitting gives $D_{KS} = 1$). More details on the GOF tests and their results are given in S2 Appendix.

A closer inspection reveals that this relatively poor fit is mostly due to irregularities in the tails of the observed distributions. However, more importantly, we note that the trajectories were in fact not drawn from a theoretical distribution chosen for its mathematical simplicity, but result from individual interactions of agents that have learned certain behaviors. In this regard, with respect to the sometimes low goodness-of-fit values, our simulations lead to similar challenges as the analysis of experimental data from real animals (see e.g. [56]). Nonetheless, the above analysis does provide a more robust account of key features of the collective dynamics.

5 Discussion

We have studied the collective behavior of artificial learning agents, more precisely PS agents, that arises as they attempt to survive in foraging environments. More specifically, we design different foraging scenarios in one-dimensional worlds in which the resources are either near or far from the region where agents are initialized.

This ansatz differs from existing work in that PS agents allow for a complex, realistic description of the sensory (percepts) and motor (actions) abilities of each individual. In particular, agents can distinguish how other agents within visual range are oriented and if the density of agents is high or low in the front and at the back of their visual area. Based on this information, agents can decide whether to continue moving in their current direction or to turn around and move in the opposite direction. Crucially, there are no fixed interaction rules, which is the main difference that sets our work apart from previous approaches, like the self-propelled particle (SPP) models or other models from statistical physics. Instead, the interactions emerge as a result of the learning process agents perform within a framework of reinforcement learning. The rewards given as part of this learning process play a role analogous to environmental pressures in nature, by enhancing the behaviors that led the agent to be rewarded. Therefore, by varying the task and reward scheme and studying the resulting behaviors, our approach allows us to test different *causal explanations* for specific observed behaviors, in the sense of environmental pressures proposed to have led to these behaviors.

In this work, we have considered scenarios where the food is situated inside or far from the region where agents are initialized and we have observed that the initially identical agents develop very different individual responses—leading to different collective dynamics—

depending on the distance they need to cover to reach the reward (food source). Agents learn to form strongly aligned swarms to get to distant food sources, whereas they learn to form cohesive (but weakly aligned) swarms when the distance to the food source is short.

Since we model each individual as an artificial learning agent, we are able not only to study the collective properties that arise from the given tasks, but also to analyze the individual responses that agents learn and that, in turn, lead to the swarm formation. Thus, we observe for instance that the tendency to align with the neighbors in the $d_F = 21$ case increases with the density of neighbors surrounding the agent. In the case of a training with $d_F = 4$, we observe that the individuals tend to move to the region with higher number of neighbors, which leads to high cohesion at the collective level.

We note that the task faced by our artificial agents, of reaching a food source, is closely related to the behaviors studied in the context of foraging theory. For this reason, we compare the individual trajectories that result from the learning process to the principal theoretical models in that field. We show that most of the individual trajectories resulting from the training with distant resources—which leads to strongly aligned swarms—are best fitted by composite correlated random walks consisting of two modes, one intensive and one extensive, whose mean step lengths are $\hat{\lambda}_I^{-1} \simeq 2.7$ and $\hat{\lambda}_E^{-1} \simeq 75$, respectively. A smaller fraction of these trajectories is best fitted by power-law distributions with exponents $\hat{\mu} \simeq 1.6$, that is, Lévy walks. The exponent of the power-law distribution we obtain is close to 2, which is the optimal Lévy walk for maximizing the rate of target encounters in environments with sparsely distributed, renewable resources [22, 31, 68]. Moreover, our results are in agreement with the study of Reynolds [67] that shows that animals can approximate Lévy walks by adopting a composite correlated random walk.

In contrast, agents that were trained to find nearby resources and follow the dynamics of cohesive swarms present normal-diffusive, Brownian-like trajectories that do not exhibit two movement modes but just one.

One crucial point of this analysis is that our simulated agents move in a multi-agent context and their movement patterns are therefore determined by the swarm dynamics they have developed through the learning process. In particular, we provide a new perspective and additional insight on the studies mentioned above regarding Lévy walks and CCRW, since the individual trajectories that are best fit by these two models arise from a collective motion with very specific features such as strong alignment and decaying cohesion. This, together with the fact that the individual responses emerge as a result of the learning process, provides an example of how trajectories with features that resemble Lévy walks can emerge from individual mechanisms that are not generated by a Lévy walk process. In this sense, our work provides an unusual example to consider within the emergentist versus evolutionary debate on Lévy walks (see e.g. [31, 32]). In particular, our work supports the former view point insofar as the agents do not have built-in interaction rules that come from a certain mathematical model such as that of the Lévy walk, but nevertheless exhibit trajectories that resemble Lévy walks as a result of the swarm dynamics that emerged from certain foraging environmental pressures.

To conclude, we have applied a model of artificial agency (PS) to different foraging scenarios within the framework of collective motion. We have shown that, without any prior hard-wired interaction rules, the same agents develop different individual responses and collective interactions, depending on the distance they need to travel to reach a food source. Agents form strongly aligned swarms to stabilize their trajectories and reach distant resources, whereas they form cohesive, unaligned swarms when the resources are near. In addition, we have shown that Lévy-like trajectories can be obtained from individual responses that do not have a simple theoretical model as the underlying process, but instead are generated and arise

from the interplay of a fine-grained set of learned individual responses and the swarm behavior that emerges from them at a collective level.

This work provides a new framework for the study of collective behavior, which supports more detailed and realistic representations of individuals' sensory and motor abilities and different types of environmental pressures. It would be interesting to apply this approach to the more complex collective behaviors that arise in two- and three-dimensional environments. Furthermore, the PS model allows for a variety of new scenarios to explore in the context of behavioral biology, since different reward schemes can easily be implemented and studied.

Supporting information

S1 Appendix. Additional analysis. In this appendix, we provide additional information about the dynamics presented in Sec. 3 of the main text. In particular, we analyze why there is transition at $d_F = 6$ from the regime where agents form cohesive swarms to the one in which they form aligned swarms to reach the food source. Furthermore, we extend the analysis of alignment and cohesion presented in Sec. 3.2 of the main text.
(PDF)

S2 Appendix. Statistical methods. We describe the statistical methods applied to obtain the results of Sec. 4. First, we explain how to optimize the likelihood function of each of the four models we consider, with a focus on the CCRW model. In addition, we describe the goodness-of-fit tests that we perform, that is, a log-likelihood ratio for the BW, CRW and PL models and a test based on uniform pseudo-residuals for the CCRW model.
(PDF)

Author Contributions

Conceptualization: Andrea López-Incera, Katja Ried, Thomas Müller, Hans J. Briegel.

Formal analysis: Andrea López-Incera, Katja Ried.

Investigation: Andrea López-Incera.

Methodology: Andrea López-Incera, Katja Ried.

Software: Andrea López-Incera.

Supervision: Katja Ried, Hans J. Briegel.

Visualization: Andrea López-Incera, Katja Ried.

Writing – original draft: Andrea López-Incera.

Writing – review & editing: Andrea López-Incera, Katja Ried, Thomas Müller, Hans J. Briegel.

References

1. Zafeiris A, Vicsek T. Collective motion. *Physics Reports*. 2012; 517:71–140. <https://doi.org/10.1016/j.physrep.2012.03.004>
2. Yates CA, Erban R, Escudero C, Couzin ID, Buhl J, Kevrekidis IG, et al. Inherent noise can facilitate coherence in collective swarm motion. *PNAS*. 2009; 106(14):5464–5469. <https://doi.org/10.1073/pnas.0811195106> PMID: 19336580
3. Kolpas A, Moehlis J, Kevrekidis IG. Coarse-grained analysis of stochasticity-induced switching between collective motion states. *PNAS*. 2007; 104(14):5931–5935. <https://doi.org/10.1073/pnas.0608270104>
4. Bode NW, Franks DW, Wood AJ. Making noise: emergent stochasticity in collective motion. *J Theor Biol*. 2010; 267:292–299. <https://doi.org/10.1016/j.jtbi.2010.08.034>

5. Vicsek T, Czirók A, Ben-Jacob E, Cohen I, Shochet O. Novel type of phase transition in a system of self-driven particles. *Phys Rev Lett*. 1995; 75:1226–1229. <https://doi.org/10.1103/PhysRevLett.75.1226>
6. Czirók A, Barabási AL, Vicsek T. Collective motion of self-propelled particles: kinetic phase transition in one dimension. *Phys Rev Lett*. 1999; 82:209–212. <https://doi.org/10.1103/PhysRevLett.82.209>
7. Czirók A, Vicsek M, Vicsek T. Collective motion of organisms in three dimensions. *Physica A*. 1999; 264:299–304. [https://doi.org/10.1016/S0378-4371\(98\)00468-3](https://doi.org/10.1016/S0378-4371(98)00468-3)
8. O’Loan OJ, Evans MR. Alternating steady state in one-dimensional flocking. *J Phys A: Math Gen*. 1999; 32:L99–L105. <https://doi.org/10.1088/0305-4470/32/8/002>
9. Briegel HJ, De las Cuevas G. Projective simulation for artificial intelligence. *Sci Rep*. 2012; 2:400. <https://doi.org/10.1038/srep00400>
10. Morales JM, Fortin D, Frair JL, Merrill EH. Adaptive models for large herbivore movements in heterogeneous landscapes. *Landsc Ecol*. 2005; 20(3):301–316. <https://doi.org/10.1007/s10980-005-0061-9>
11. Mueller T, Fagan WF, Grimm V. Integrating individual search and navigation behaviors in mechanistic movement models. *Theor Ecol*. 2011; 4(3):341–355. <https://doi.org/10.1007/s12080-010-0081-1>
12. Martinez-Gil F, Lozano M, Fernandez F. Emergent behaviors and scalability for multi-agent reinforcement learning-based pedestrian models. *Simul Model Pract Th*. 2017; 74:117–133. <https://doi.org/10.1016/j.simpat.2017.03.003>
13. Shimada K, Bentley P. Learning how to flock: deriving individual behaviour from collective behaviour with multi-agent reinforcement learning and natural evolution strategies. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*; 2018. p. 169–170.
14. Durve M, Peruani F, Celani A. Learning to flock through reinforcement. *arXiv preprint arXiv:191101697*. 2019;.
15. Ried K, Müller T, Briegel HJ. Modelling collective motion based on the principle of agency: General framework and the case of marching locusts. *PLoS ONE*. 2019; 14:e0212044. <https://doi.org/10.1371/journal.pone.0212044>
16. Sinervo B. *Optimal foraging theory: constraints and cognitive processes*. University of California, Santa Cruz; 1997.
17. Stephens DW, Brown JS, Ydenberg RC. *Foraging: Behavior and Ecology*. Chicago: The University of Chicago Press; 2007.
18. Pyke GH. Optimal foraging theory: a critical review. *Annu Rev Ecol Syst*. 1984; 15(1):523–575. <https://doi.org/10.1146/annurev.es.15.110184.002515>
19. Viswanathan GM, da Luz MGE, Raposo EP, Stanley HE. *The physics of foraging: an introduction to random searches and biological encounters*. Cambridge University Press; 2011.
20. Lévy P. *Théorie de l’addition des variables aléatoires*. vol. 1. Gauthier-Villars; 1954.
21. Shlesinger MF, Klafter J. Lévy walks versus Lévy flights. In: Stanley HE, Ostrowski N, editors. *Growth and form*. Amsterdam: Martinus Nijhof Publishers; 1986.
22. Viswanathan GM, Buldyrev SV, Havlin S, Da Luz MGE, Raposo EP, Stanley HE. Optimizing the success of random searches. *Nature*. 1999; 401(6756):911–914. <https://doi.org/10.1038/44831>
23. Benhamou S. Efficiency of area-concentrated searching behaviour in a continuous patchy environment. *J Theor Biol*. 1992; 159(1):67–81. [https://doi.org/10.1016/S0022-5193\(05\)80768-4](https://doi.org/10.1016/S0022-5193(05)80768-4)
24. Benhamou S. How many animals really do the Lévy walk? *Ecology*. 2007; 88(8):1962–1969. <https://doi.org/10.1890/06-1769.1>
25. Viswanathan GM, Afanasyev V, Buldyrev SV, Murphy EJ, Prince PA, Stanley HE. Lévy flight search patterns of wandering albatrosses. *Nature*. 1996; 381(6581):413. <https://doi.org/10.1038/381413a0>
26. Sims DW, Southall EJ, Humphries NE, Hays GC, Bradshaw CJA, Pitchford JW, et al. Scaling laws of marine predator search behaviour. *Nature*. 2008; 451(7182):1098. <https://doi.org/10.1038/nature06518> PMID: 18305542
27. Edwards AM. Overturning conclusions of Lévy flight movement patterns by fishing boats and foraging animals. *Ecology*. 2011; 92(6):1247–1257. <https://doi.org/10.1890/10-1182.1>
28. Edwards AM, Freeman MP, Breed GA, Jonsen ID. Incorrect likelihood methods were used to infer scaling laws of marine predator search behaviour. *PLoS one*. 2012; 7(10):e45174. <https://doi.org/10.1371/journal.pone.0045174>
29. Humphries NE, Weimerskirch H, Queiroz N, Southall EJ, Sims DW. Foraging success of biological Lévy flights recorded in situ. *PNAS*. 2012; 109(19):7169–7174. <https://doi.org/10.1073/pnas.1121201109>
30. Dragon AC, Bar-Hen A, Monestiez P, Guinet C. Comparative analysis of methods for inferring successful foraging areas from Argos and GPS tracking data. *Marine Ecology Progress Series*. 2012; 452:253–267. <https://doi.org/10.3354/meps09618>

31. Wosniack ME, Santos MC, Raposo EP, Viswanathan GM, da Luz MGE. The evolutionary origins of Lévy walk foraging. *PLoS comput biol.* 2017; 13(10):e1005774. <https://doi.org/10.1371/journal.pcbi.1005774>
32. Pyke GH. Understanding movements of organisms: it's time to abandon the Lévy foraging hypothesis. *Methods Ecol Evol.* 2015; 6(1):1–16. <https://doi.org/10.1111/2041-210X.12298>
33. Mautner J, Makmal A, Manzano D, Tiersch M, Briegel HJ. Projective simulation for classical learning agents: a comprehensive investigation. *New Gener Comput.* 2015; 33(1):69–114. <https://doi.org/10.1007/s00354-015-0102-0>
34. Makmal A, Melnikov AA, Dunjko V, Briegel HJ. Meta-learning within projective simulation. *IEEE Access.* 2016; 4:2110–2122. <https://doi.org/10.1109/ACCESS.2016.2556579>
35. Melnikov AA, Makmal A, Dunjko V, Briegel HJ. Projective simulation with generalization. *Sci Rep.* 2017; 7:14430. <https://doi.org/10.1038/s41598-017-14740-y>
36. Melnikov AA, Makmal A, Briegel HJ. Benchmarking projective simulation in navigation problems. *IEEE Access.* 2018; 6:64639–64648. <https://doi.org/10.1109/ACCESS.2018.2876494>
37. Ried K, Eva B, Müller T, Briegel HJ. How a minimal learning agent can infer the existence of unobserved variables in a complex environment. preprint arXiv:191006985v1. 2019;
38. Holland JH. *Adaptation in Natural and Artificial Systems.* University of Michigan Press. 1975.
39. For instance, in the self-propelled particle models [5, 6], the particle changes its orientation at each time step to align itself to the average orientation of the neighboring particles.
40. Humphries DA, Driver PM. Protean defence by prey animals. *Oecologia.* 1970; 5(4):285–302. <https://doi.org/10.1007/BF00815496>
41. Bilecenoğlu M. Observations on the burrowing behaviour of the dwarf blaasop, *Torquigener flavimaculosus* (Osteichthyes: Tetraodontidae) along the coast of Fethiye, Turkey. *Zoology in the Middle East.* 2005; 35(1):29–34. <https://doi.org/10.1080/09397140.2005.10638100>
42. Eifler D, Eifler M. Escape tactics in the lizard *Meroles cuneirostris*. *Amphib-Reptil.* 2014; 35(4):383–389. <https://doi.org/10.1163/15685381-00002963>
43. Yager DD, May ML, Fenton MB. Ultrasound-triggered, flight-gated evasive maneuvers in the praying mantis *Parasphendale agrionina*. I. Free flight. *J Exp Biol.* 1990; 152(1):17–39.
44. Combes SA, Rundle DE, Iwasaki JM, Crall JD. Linking biomechanics and ecology through predator–prey interactions: flight performance of dragonflies and their prey. *J Exp Biol.* 2012; 215(6):903–913. <https://doi.org/10.1242/jeb.059394>
45. Jones KA, Jackson AL, Ruxton GD. Prey jitters; protean behaviour in grouped prey. *Behavioral Ecology.* 2011; 22(4):831–836. <https://doi.org/10.1093/beheco/arr062>
46. Richardson G, Dickinson P, Burman OHP, Pike TW. Unpredictable movement as an anti-predator strategy. *Proc R Soc B.* 2018; 285(1885):20181112. <https://doi.org/10.1098/rspb.2018.1112>
47. The specific probabilities of doing so are given in Fig.6 (a) and analyzed in Sec.3.1.
48. Ariel G, Ayali A. Locust collective motion and its modeling. *PLoS comput biol.* 2015; 11(12). <https://doi.org/10.1371/journal.pcbi.1004522> PMID: 26656851
49. Pener M, Simpson S. Locust phase polyphenism: an update. *Advan Insect Physiol.* 2009; 36:1–272. [https://doi.org/10.1016/S0065-2806\(08\)36001-9](https://doi.org/10.1016/S0065-2806(08)36001-9)
50. Plank MJ, James A. Optimal foraging: Lévy pattern or process? *J R Soc Interface.* 2008; 5(26):1077–1086. <https://doi.org/10.1098/rsif.2008.0006>
51. Plank MJ, Codling EA. Sampling rate and misidentification of Lévy and non-Lévy movement paths. *Ecology.* 2009; 90(12):3546–3553. <https://doi.org/10.1890/09-0079.1>
52. Reynolds AM. Distinguishing between Lévy walks and strong alternative models. *Ecology.* 2012; 93(5):1228–1233. <https://doi.org/10.1890/11-1815.1>
53. Humphries NE, Weimerskirch H, Sims DW. A new approach for objective identification of turns and steps in organism movement data relevant to random walk modelling. *Methods Ecol Evol.* 2013; 4(10):930–938. <https://doi.org/10.1111/2041-210x.12096>
54. Gautestad AO. Brownian motion or Lévy walk? Stepping towards an extended statistical mechanics for animal locomotion. *J R Soc Interface.* 2012; 9(74):2332–2340. <https://doi.org/10.1098/rsif.2012.0059>
55. Jansen VAA, Mashanova A, Petrovskii S. Comment on “Lévy walks evolve through interaction between movement and environmental complexity”. *Science.* 2012; 335(6071):918–918. <https://doi.org/10.1126/science.1215747>
56. Auger-Méthé M, Derocher AE, Plank MJ, Codling EA, Lewis MA. Differentiating the Lévy walk from a composite correlated random walk. *Methods Ecol Evol.* 2015; 6(10):1179–1189. <https://doi.org/10.1111/2041-210X.12412>

57. de Jager M, Weissing FJ, Herman PMJ, Nolet BA, van de Koppel J. Lévy walks evolve through interaction between movement and environmental complexity. *Science*. 2011; 332(6037):1551–1553. <https://doi.org/10.1126/science.1201187>
58. Edwards AM, Phillips RA, Watkins NW, Freeman MP, Murphy EJ, Afanasyev V, et al. Revisiting Lévy flight search patterns of wandering albatrosses, bumblebees and deer. *Nature*. 2007; 449(7165):1044. <https://doi.org/10.1038/nature06199> PMID: 17960243
59. Santos MC, Raposo EP, Viswanathan GM, da Luz MGE. Can collective searches profit from Lévy walk strategies? *J Phys A*. 2009; 42(43):434017. <https://doi.org/10.1088/1751-8113/42/43/434017>
60. Reynolds AM, Ouellette NT. Swarm dynamics may give rise to Lévy flights. *Scientific reports*. 2016; 6:30515. <https://doi.org/10.1038/srep30515>
61. Reynolds AM. Effective leadership in animal groups when no individual has pertinent information about resource locations: How interactions between leaders and followers can result in Lévy walk movement patterns. *EPL*. 2013; 102(1):18001. <https://doi.org/10.1209/0295-5075/102/18001>
62. Zucchini W, MacDonald IL. Hidden Markov models for time series: an introduction using R. vol. 110. Monographs on Statistics and Applied Probability, CRC Press; 2009.
63. Clauset A, Shalizi CR, Newman MEJ. Power-law distributions in empirical data. *SIAM review*. 2009; 51(4):661–703. <https://doi.org/10.1137/070710111>
64. Burnham KP, Anderson DR. Multimodel inference: understanding AIC and BIC in model selection. *Sociol Methods Res*. 2004; 33(2):261–304. <https://doi.org/10.1177/0049124104268644>
65. Symonds MRE, Moussalli A. A brief guide to model selection, multimodel inference and model averaging in behavioural ecology using Akaike's information criterion. *Behav Ecol Sociobiol*. 2011; 65(1):13–21. <https://doi.org/10.1007/s00265-010-1037-6>
66. 0.8% of trajectories have $w_{CRW} = 1$, which is to be expected since the MLE parameters of both CRW and CCRW models are roughly the same.
67. Reynolds AM. Selection pressures give composite correlated random walks Lévy walk characteristics. *J Theor Biol*. 2013; 332:117–122. <https://doi.org/10.1016/j.jtbi.2013.04.035>
68. Raposo EP, Buldyrev SV, da Luz MGE, Viswanathan GM, Stanley HE. Lévy flights and random searches. *J Phys A: Math Theor*. 2009; 42(43):434003. <https://doi.org/10.1088/1751-8113/42/43/434003>