Evolution of Collective Behaviors by Minimizing Surprise

Heiko Hamann
Department of Computer Science,
University of Paderborn, Germany
heiko.hamann@uni-paderborn.de

May 20, 2014

Abstract

Similarly to evolving controllers for single robots also controllers for groups of robots can be generated by applying evolutionary algorithms. Usually a fitness function rewards desired behavioral features. Here we investigate an alternative method that generates collective behaviors almost only as a by-product. We roughly follow the idea of Helmholtz that perception is a process based on probabilistic inference and evolve an internal model that is supposed to predict the agent's future perceptions. Separated from this prediction model the agent also evolves a regular controller. Direct selective pressure, however, is only effective on the prediction model by minimizing prediction error (surprise). Our results show that a number of basic collective behaviors emerge by this approach, such as dispersion, aggregation, and flocking. The probability that a certain behavior emerges and also the difficulty of making correct predictions depends on the swarm density. The reported method has potential to be another simple approach to open-ended evolution analogical to the search for novelty.

1 Introduction

The creation of control algorithms for self-organizing, artificial collective systems is challenging because it is difficult to anticipate how a multitude of local interactions will effect the global behavior. Hence, evolutionary algorithms are a good option. A desired behavior is evolved by defining a fitness function that rewards the occurrence of particular behavioral features. Alternatively, an indirect selective pressure can be generated that does not directly influence the evolved behavior while interesting behaviors should still emerge. In this paper we investigate one of such alternatives by evolving an agent, that tries to predict its future perception, while a number of collective behaviors is obtained almost only as a by-product.

1

1.1 Evolution of collective behaviors

The application of swarm intelligence [Bonabeau et al., 1999] to the field of robotics is called swarm robotics [Dorigo and Şahin, 2004, Brambilla et al., 2013]. The problem of designing algorithms for swarm robotics that generate the desired behaviors (swarm engineering, Brambilla et al. [2013]) is challenging because it means to design self-organizing systems. Self-organization relies on feedback processes and a multitude of interactions [Bonabeau et al., 1999] that have effects which are difficult to anticipate. One solution is to develop global models that predict the expected collective behavior and give insights about underlying principles [Hamann, 2010, Martinoli et al., 2004]. Another approach is to apply methods from evolutionary robotics [Nolfi and Floreano, 2000] to swarm robotics, that is, evolutionary swarm robotics [Trianni, 2008]. For example, the evolution of an aggregation behavior in robot simulations [Trianni et al., 2003] and also the evolution of communication in combination with collective motion were reported [Trianni et al., 2004].

In evolutionary robotics the design of fitness functions is a challenging issue. For example, the unfavorable approach of defining elaborated complex fitness functions, that predefine many behavioral features of the expected behavior, is feasible. That way intrinsically complex robot tasks are simplified by exploiting a priori knowledge which foils the key idea of evolutionary robotics that it is a preferential approach whenever there is only little a priori knowledge [Nelson et al., 2009]. Another issue of fitness functions is that the evolutionary algorithm might converge prematurely. To overcome these challenges, new approaches were reported, such as novelty search and related methods [Lehman and Stanley, 2008, Mouret and Doncieux, 2009]. The successful application of novelty search to swarm robotics was recently reported [Gomes et al., 2013]. Artificial life is a related field in this context, there the evolution of collective behaviors within artificial ecologies was reported that either rely on selective pressure generated by ecological features only [Schmickl and Crailsheim, 2006] or in combination with explicit fitness functions [Ward et al., 2001].

1.2 Behavior based on minimization of surprise

Many common approaches in the field of intelligent agents focus on defining condition-action rules or similar methods to generate intelligent behavior. In evolutionary robotics [Nolfi and Floreano, 2000] the common approach is based on reactive control possibly combined with simple internal states which means complex world models are usually not considered. That leaves a considerable gap to mammals but also to insects such as honeybees that have capabilities for learning and memory [Hammer and Menzel, 1995]. An initial step towards the evolution of simple world models would prepare to go beyond the limited capabilities of reactive agents.

There are many theories that try to define a systemic concept of brains which would allow to explain the variety of behaviors we observe in animals. One of these theories goes back to von Helmholtz [1867] who argued that percep-

tion is a process based on probabilistic inference. According to that the main challenge imposed on each animal's brain is to determine coherent causes to its sensory inputs. This abstract concept was picked up, for example, in machine learning to solve the unsupervised learning task of creating a model for a given data set. The so-called Helmholtz machines [Dayan et al., 1995] are trained to serve as a generative model, that is, they learn a probabilistic model of an assumed hidden structure of the input data. Following this concept the brain can be understood as a 'prediction machine' that learns models which describe the causes of its perceptions. Furthermore, a plausible assumption to make is that minimizing the prediction error of this model is advantageous for the considered organism. This line of thought is implemented in a mathematical framework by Friston [2010] which defines an information-theoretic analogon to thermodynamic (Helmholtz) free energy. Basically he defines this free energy as the prediction error. This approach based on a 'free-energy principle' has the potential to be a unified brain theory [Friston, 2010, Friston et al., 2006]. The connection to Darwinian evolution is established by the argument that minimal prediction errors indirectly confine the set of state-action pairs, that are experienced and executed by an agent, which in turn limits the living space and is life prolonging: "By sampling or navigating the environment selectively they restrict their exchange with it within bounds that preserve their physical integrity and allow them to last longer" [Friston et al., 2006]. A minimal prediction error is hence interpreted as an evolutionary advantage. While Friston mostly focuses on simple agent–environment interactions, we apply this concept to agent–agent interactions in a collective system here. In the case of a swarm, the future perceptions of an agent not only depend on its own actions (and a possibly dynamic environment) but also on the actions of other close-by agents. This adds another feedback loop to Friston's framework that indirectly asks the prediction machine to predict other prediction machines. These prediction machines can be interpreted as world models as used in intelligent agents. However, they are not applied as a mere tool to make intelligent decisions but as the main driving force for the emergence of intelligent behavior. Furthermore note that making this connection between the neurosciences and swarm intelligence follows the idea of 'swarm cognition' [Trianni et al., 2010].

1.3 Basic collective behaviors

In general there is a big variety of collective behaviors including rather complex behaviors, such as foraging, sorting, nest building, and cooperative transport [Bonabeau et al., 1999]. At the other end of the scale there are simple collective behaviors that are restricted to agent–agent interactions without explicit agent–environment interactions. These simple behaviors rely exclusively on agent positions and agent motion while not including environmental features. There are four basic collective behaviors that are based on motion (moving or stopped) and relative positions (minimal distances between agents or maximal distances) only. On the basis of these two dimensions they are categorized: dispersion (maximal distances, stopped), aggregation (minimal dis-

(a) agents on ring, arrows indicate clockwise or counterclockwise motion

(b) sensors and their assignment to intervals around the agent (blue circle)
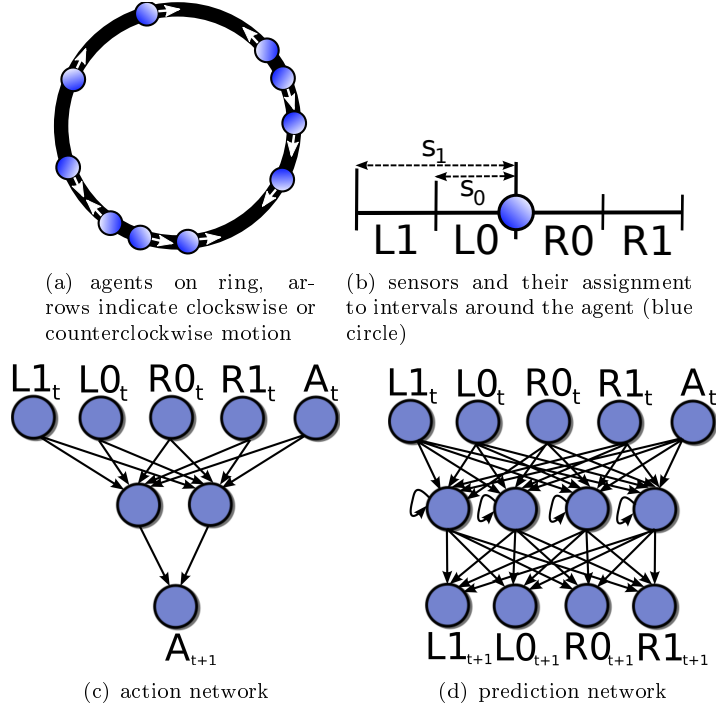
(c) action network

(d) prediction network

Figure 1: Experimental setting, sensor setting, and artificial neural networks for control and prediction.

tances, stopped), random (maximal distances, moving), and flocking (minimal distances, moving).

In the following we focus on these four collective behaviors of low complexity and investigate the necessary conditions that allow them to emerge. Following the motivation of the free-energy principle, our objective is to provoke the emergence of these behaviors without an explicit external force acting on agent behaviors or agent controllers and also without explicit selective pressure due to ecological features such as ecological niches or co-evolution. The investigated hypothesis is that a mere selection for minimal surprise (i.e., minimal prediction error, minimal free energy, Friston [2010]) is sufficient to evolve these four simple collective behaviors. Note that a short overview of this research was published before [Hamann, 2014].

2 Model

Our swarm model is roughly inspired by the desert locust, *Schistocerca gregaria*, that shows collective motion (often called 'marching bands') in the growth stage of a wingless nymph [Buhl et al., 2006]. Buhl et al. reduce the collective

4

parameter	value
swarm size N (num. agents in simulation)	20
agent sensor range s_0	0.5
agent sensor range s_1	1.0
agent speed v	0.1
noise (displacement of agent position)	$[-0.01, 0.01]$
evaluation length in time steps T	500
population size (evolution)	50
number of generations	30
num. of sim. runs per fitness evaluation	10
elitism	1
mutation rate	0.05
ring length L (circumference)	$[5, 50]$

Table 1: Parameter settings.

motion of the locusts to a quasi-one-dimensional system in their experiments by placing the animals in a ring-shaped arena. The observed directional alignment of a majority of locusts is density-dependent and individuals seem to change their direction as a response to neighbors. Here the agents move in 1-d space in the form of a circle, in the following called ring (see Fig. 1a). The ring's circumference is denoted by ring length L. The agents choose from only two available actions: moving clockwise or moving counterclockwise. However, the agents have no global reference frame, that is, whether they are moving clockwise or counterclockwise is unknown to them. Hence, they actually choose from two actions: staying with the current direction or switching the direction. They are not allowed to stop. The agents have four sensors (L1,L0,R0,R1) that cover four intervals of their neighborhood defined by sensor distances s_0 and s_1 (see Fig. 1b and table 1 for the used parameters). The sensors are discrete and output either a '1' if there is at least one neighboring agent within the respective interval or a '0' if there is no agent.

In the following experiments the swarm size is fixed to $N = 20$ agents while the ring length L is varied between experiments. This allows to test different swarm densities N/L. Initially the agents are positioned in space by sampling from a random uniform distribution and also their initial direction is chosen with equal probabilities. Time is discrete and each agent's position is updated by adding or subtracting the agent speed of $v = 0.1$. In addition, a small noise is applied to the agent's position by adding a random uniform value sampled from $[-0.01, 0.01]$. This small noise was added to the system based on the heuristic experience that it triggers a bigger diversity of behaviors. The agents are allowed to pass each other without any interference.

Each agent is equipped with two independent artificial neural networks (ANN) that we call 'action network' and 'prediction network'. The action network has 5 input neurons, 2 hidden neurons, and 1 output neuron (see Fig. 1c). The

inputs are the 4 sensor values (L1,L0,R0,R1) and the last action. That way the ANN is a recurrent network although it is implemented as a feedforward network. The output neuron determines the next action (stay with current direction or switch direction) based on a threshold. The prediction network has 5 input neurons, 4 hidden neurons, and 4 output neurons (see Fig. 1d). The input is the same as with the action network. Each hidden neuron has a self-loop which explicitly implements a recurrent network. The output of the prediction network are 4 values that are associated with the 4 intervals of the 4 sensors (L1,L0,R0,R1). Each output neuron determines the predicted value of the respective sensor for the next time step. The agent actually tries to predict where its neighbors will be positioned in the next time step.

We apply a simple evolutionary algorithm. The genomes consist of two sets of weights, one for the action network and one for the prediction network. Note that by applying an evolutionary algorithm we introduce a second concept of a population in this work which is the population of genomes (in difference to the population of agents that are run in the simulations). This population size is 50. Initially a population of random networks is generated (weights are random uniform over $[-0.5, 0.5]$). The genomes are evaluated by applying the above described swarm simulation. For a particular evaluation run all agents of the simulation share the same genetic material, that is, they have the same networks as defined by the genome that is currently evaluated. The fitness function rewards good predictions of the prediction network. It is defined by

$$F = \frac{1}{NT} \sum_{t \in \{0,1,...,T-1\}} \sum_{i \in \{0,1,...,N-1\}} \sum_{j \in \{L1,L0,R0,R1\}} c_{i,j}(t), \qquad (1)$$

whereas N is the swarm size, T is the length of the evaluation in time steps, and $c_{i,j}(t) = 1$ if the prediction of the previous time step for sensor j of agent i matches the current value of sensor j, otherwise $c_{i,j}(t) = 0$. For time $t = 0$ the predictions are set to 0 by definition. The evaluation of a genome is averaged over the results from 10 independent simulation runs. The theoretical best fitness is 4. Note that the behavior of the agents as defined by the output of the action network is not a direct subject to the fitness function because only correct predictions of the prediction network are rewarded.

Each weight of both the action network and the prediction network is mutated with a probability of 0.05. Proportionate selection and elitism of 1 are used. Evolution is run for 30 generations totaling to 1500 evaluations (each based on 10 independent runs of the simulator).

3 Results

All four of the above mentioned basic collective behaviors are found among the best evolved individuals over different swarm densities. Fig. 2 shows example trajectories of all $N = 20$ agents for each of these behaviors (vertical axis gives position x of agents on the ring, horizontal axis gives evolution in time t). The

(a) dispersion ($L = 50$)　　　　(b) aggregation ($L = 15$)

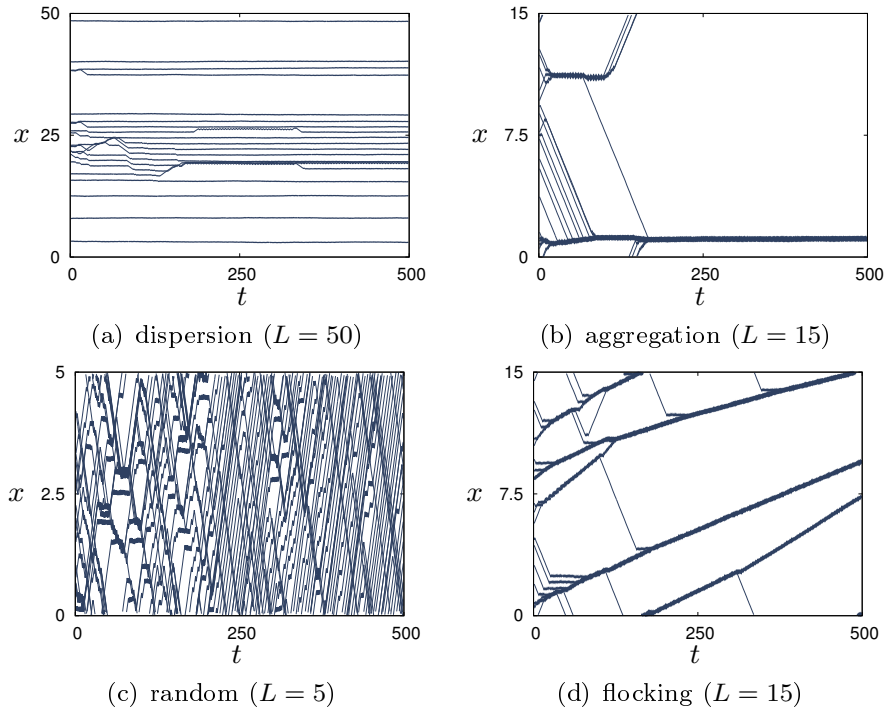(c) random ($L = 5$)　　　　(d) flocking ($L = 15$)

Figure 2: Trajectories of all agents for the four basic swarm behaviors: dispersion, aggregation, random, and flocking.

agents cannot simply stop but by permanently switching between clockwise and counterclockwise motion (output of action network always 'switch') they cover no distance averaged over several steps as seen in Fig. 2a. Arguably this behavior can be called dispersion although the agents are initially positioned by sampling from a uniform distribution. Agents that are initially close-by actually increase the distance between them. Similarly in some runs the agents also manage to stay together in a cluster of high density after they have aggregated as seen in Fig. 2b. Another behavior is seen in Fig. 2c. By rarely switching the direction (output of action network mostly 'stay') the agents cover maximal distance and logically form two groups (clockwise and counterclockwise) distributed over the whole ring. If the agents aggregate but still cover some distance, then we get a flocking behavior as seen in Fig. 2d. Also note that these flocks manage to travel with different velocities (i.e., different cyclic combinations of clockwise/counterclockwise motion). Fig. 3 shows a few more examples of the obtained collective behaviors for a medium swarm density ($L = 20$) to document the variety of behaviors. Fig. 3a shows an example for flocks moving with different velocities depending on their size. All flocks have positive velocities and small flocks move faster. That way the small flocks catch up and join the bigger flocks. Fig. 3b shows a flock with one agent at each side virtually as 'guards'. Fig. 3c shows a fully aligned swarm that abruptly switches its direction (similarly to what is observed in locusts, see Buhl et al. [2006]. Fig. 3d shows a technique to reduce the size of smaller flocks one by one to form a single flock.

3.1 Fitness of best evolved individuals

Next we investigate a series of experiments for varied ring lengths $L \in \{5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$. Varying the ring length means varying the swarm density because the swarm size is fixed to $N = 20$. For each ring length L, 200 independent evolutionary runs were done. In the following the focus is on the best evolved individuals, that is, the best prediction machines. For the example of ring length $L = 25$ the increase of the best fitness over generations is shown in Fig. 4a (reported before, Hamann [2014]). After a fast increase for generations $t < 15$, a saturation in the best fitness is seen. The median of the last generation ($t = 29$) is 3.33 which means that the prediction network is predicting in average 83.3% of the sensor values correctly.

In Fig. 4b a comparison of the last generation's best individual fitness is given for all settings of the ring length L (reported before, Hamann [2014]). The highest median best fitness is reached for $L = 5$ (3.68) and for $L = 50$ (3.64) whereas the lowest median best fitness is obtained for $L = 25$ (3.33). With increasing ring length (i.e., decreasing swarm density) the medians decrease ($L < 30$) and start to increase again for $L > 25$. From this result we infer that the prediction task is more difficult for medium density ($10 \leq L \leq 35$) without excluding the possibility that better results could be achieved by investing more resources to the optimization process. For very high densities prediction seems easy because most of the time all sensors detect neighbors and hence the perception changes only infrequently. For very low densities prediction seems also easy because
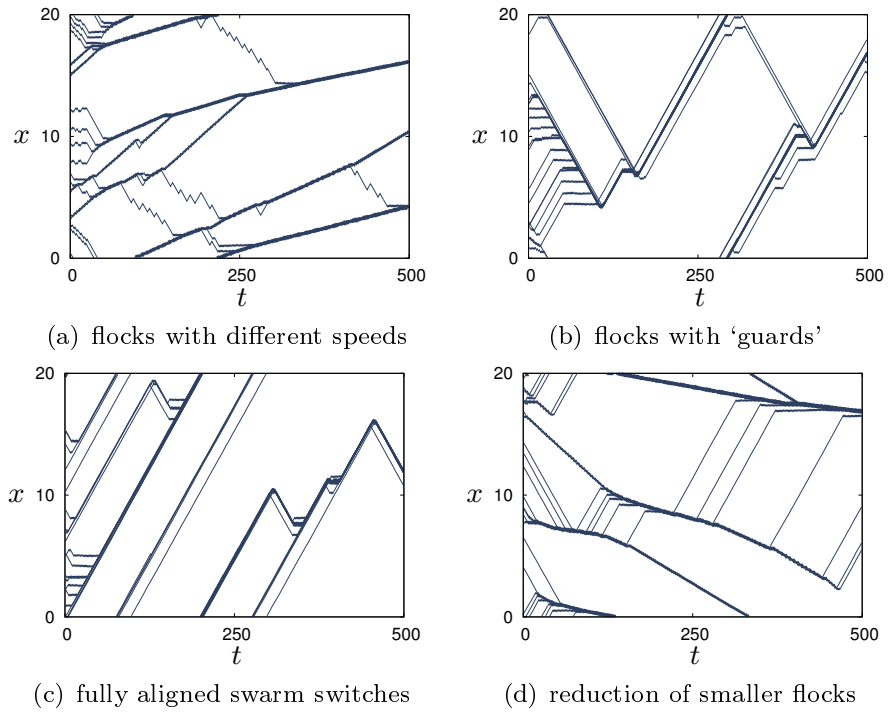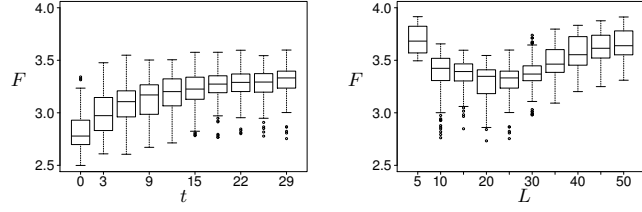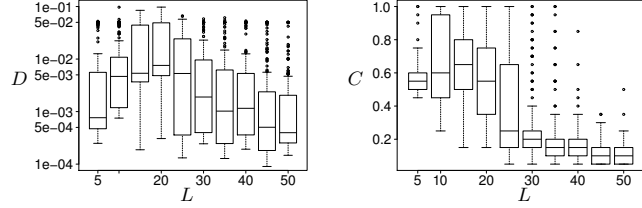
(a) flocks with different speeds

(b) flocks with 'guards'

(c) fully aligned swarm switches

(d) reduction of smaller flocks

Figure 3: Trajectories of all agents for interesting collective behaviors ($L = 20$).

(a) best fitness over generations t, for ring length $L = 25$ (Hamann, 2014)

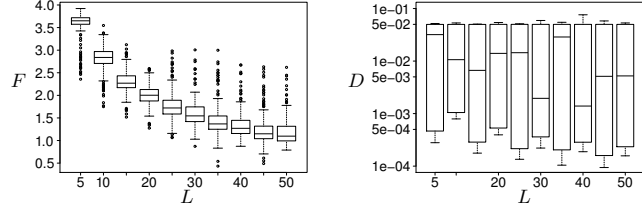(b) best fitness of last generation over ring length L (Hamann, 2014)

(c) covered distance over L, note logarithmic scale

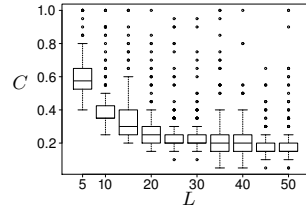(d) ratio of swarm in largest cluster over ring length L
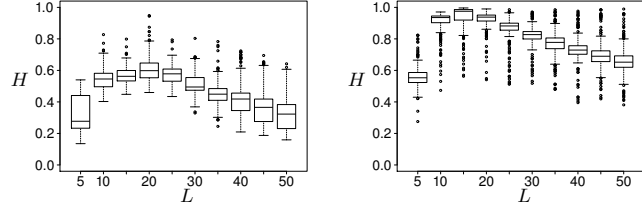
(e) entropy H of sensor input over L

(f) control experiments, entropy H
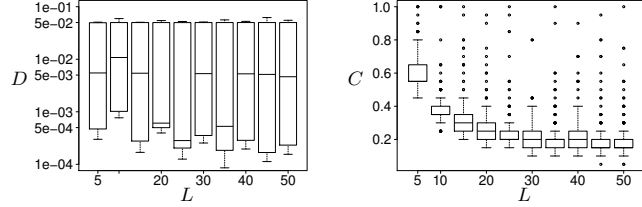
(g) control experiments, covered distance, note logarithmic scale

(h) control experiments, ratio of swarm in largest cluster

(i) inhomogeneous swarm, best fitness of last generation over ring length L

(j) inhomogeneous swarm, covered distance over L, note logarithmic scale

(k) inhomogeneous swarm, ratio of swarm in largest cluster over L

Figure 4: Results – fitness, cov. dist., largest cluster, entropy.

10

most of the time all sensors detect no neighbors. For medium densities, however, the detection of neighbors might change often hence always predicting the detection of neighbors or always predicting that no neighbors will be detected is insufficient.

3.2 Covered distance

In the following we investigate the evolved behaviors. This is done along the above mentioned two dimensions of basic collective behaviors, namely covered distance and distance between agents. The covered distance is measured over a period of time $\tau = \frac{L}{2v}$ at the end of the evaluation. It is the sum over the covered distances of all agents. It is normalized by the length of the time period and the swarm size N. We get

$$D = \frac{1}{N\tau} \sum_{i \in \{0,1,\ldots,N-1\}} |x_i(T - \tau) - x_i(T)|, \tag{2}$$

whereas $x_i(t)$ is the position of agent i at time t. The covered distance has a maximum of $D_{\mathrm{max}} = 0.1 = v$. In Fig. 4c the results averaged over the population of the last generation of the 200 independent evolutionary runs for varied ring length are given (note logarithmic scale). With increasing ring length (i.e., decreasing swarm density) the median covered distance increases ($L < 25$) and starts to decrease again for $L > 20$. Hence, it is very likely to see a static swarm for ring length $L = 50$ and rather unlikely for $L = 20$.

3.3 Largest cluster

As an indicator for the distances between agents we determine the largest cluster. For each agent we determine the total number of agents within a distance of s_1 (i.e., we count all agents that are in any of the four intervals of sensors L1, L0, R0, R1). The maximum over these numbers divided by swarm size N determines the size C of the largest cluster. This measure is more meaningful than just calculating the sum over all agent-agent distances because it allows, for example, to distinguish situations when the swarm is separated in two flocks from situations with uniform distributions. The results averaged over the population of the last generation of the 200 independent evolutionary runs are given in Fig. 4d. For $L = 5$ the median of the largest cluster size is $C = 0.55$ (11 agents) and hence relatively big as expected because an agent covers $2s_1$ with its sensors which is 40% of the ring for $L = 5$. With increasing ring length the median of the largest cluster size increases ($L < 20$) and starts to decrease again for $L > 15$ until it reaches the low value of $C = 0.15$ (3 agents) for $L = 50$. Hence, it is very likely to see an aggregated swarm, for example, for ring length $L = 15$ and rather unlikely for $L = 50$. To make a statement concerning the expected frequency of basic collective behaviors, both indicators (covered distance and largest cluster size) have to be considered in combination. For example, flocking is likely to occur for $L = 15$ because fast and well aggregated swarms are frequent for that setting. However, from these results one cannot tell whether both
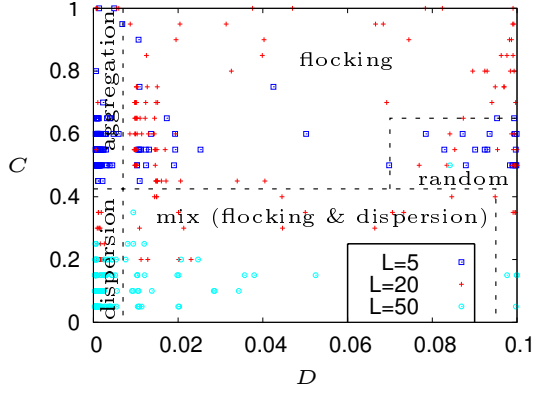
11

Figure 5: Plot of largest cluster size over covered distance for $L \in \{5, 20, 50\}$.

features coexist. In Fig. 5 the largest cluster size is plotted over covered distance for $L \in \{5, 20, 50\}$ and hence allowing to classify the behaviors at least roughly. Aggregation and random behavior is mostly observed for $L = 5$. Dispersion is mostly observed for $L = 50$. Flocking is mostly observed for $L = 20$, whereas it is possible to distinguish between slow moving flocks and fast moving flocks (cf. Fig. 2, flocks can move with different velocities). The behaviors seen in the region $0.07 \leq D \leq 0.095$, $C < 0.42$ are inconclusive as they are mixes between flocking and dispersion (i.e., a fraction of the swarm moves as flock and another fraction is dispersed).

3.4 Entropy of sensor input

The difficulty of predicting future sensor inputs and minimizing surprise is conceptually connected to information theory by the concept of Shannon entropy. We are interested in the Shannon entropy of the sensor inputs. The entropy is measured for each agent over the whole period of the simulation. For each sensor $X \in \{L1, L0, R0, R1\}$ we measure the probability p_X that the sensor gives a 1 as output and calculate the entropy

$$H_X = -p_X \log_2(p_X) - (1 - p_X) \log_2(1 - p_X). \tag{3}$$

The overall entropy for the considered agent is the average over all sensors $H = \frac{1}{4}(H_{\mathrm{L1}} + H_{\mathrm{L0}} + H_{\mathrm{R0}} + H_{\mathrm{R1}})$. Finally, we average over all individual entropies H of the swarm. The results averaged over the population of the last generation of the 200 independent evolutionary runs are given in Fig. 4e. Note that high entropy is measured for $L \in \{10, 15, 20, 25, 30\}$ which corresponds to relatively low fitness for the same ring lengths as seen in Fig. 4b. The inverse accordance between entropy and fitness indicates that diverse sensor input is inherent for intermediate swarm densities.

3.5 Control experiments – random genomes

A series of control experiments was done to test whether the results reported above are different from randomly generated behaviors and whether especially the random changes to the action networks by mutations are effective and relevant. A set of 200 random genomes was generated for each ring length $L \in \{5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$ and a swarm controlled by the respective action networks was simulated for each of them. The performance of these random prediction networks is low for all ring lengths. The median fitness is about 2.0, that is, in average the prediction is correct for two out of the four sensors. This would be expected for randomly generated predictions and equally distributed sensor values which is, however, not necessarily the case as it depends on swarm density and the actual agent behaviors.

The entropy of the sensor input for these random behaviors is shown in Fig. 4f. For all ring lengths the entropy of random behaviors is significantly higher than for the entropy of evolved behaviors as shown in Fig. 4e. We also note that the difference between Figs. 4e and f is mostly a mere displacement by $\Delta H \approx 0.4$, that is, the overall shape of the curve is preserved. Hence, the entropy also depends inherently on the swarm density. Given that a randomly generated population is also the starting point of the evolutionary approach, it is justified to say that the evolution of prediction machines minimizes the entropy of the sensor input, too.

The covered distance D is shown in Fig. 4g. It is also invariant to the ring length (no significant differences) and it is in average shorter compared to the evolved behaviors. This is expected because by random control the swarm is unlikely to align and form a flock.

The size of the largest clusters C is shown in Fig. 4h. Large clusters are observed for $L = 5$ as expected because an agent covers already 40% of the ring with its sensors. Still, the system cannot be understood as a mere point process because the agents might react to each other based on their random controllers. For an assumed uniform distribution of agents we would expect largest clusters of size $C = 0.4$ (8 agents), however, the median is 0.55 (11 agents). With increasing ring length L the size of the largest clusters decreases fast as expected. For ring lengths $L \in \{10, 15, 20, 25\}$ the largest clusters are significantly different and the evolved behaviors discussed above achieve bigger clusters, that is, better organized swarm behaviors. Interestingly, in the case of random behaviors there are a few extreme outliers even for large ring lengths up to $L = 50$ reaching up to $C = 1$ (all agents clustered) while the evolved behaviors achieve only $C \leq 0.5$ for $L > 40$.

3.6 Control experiments – inhomogeneous swarm

Another series of experiments was done to test whether collective behaviors can also be evolved in an inhomogeneous swarm. Before all agents of a particular simulation run (simulated agents on ring) were clones because they shared the same genetic material. Now each agent of a simulation run has its own

genome which most likely differs from all other genomes in the swarm. Hence, we merge the two populations that we had before, that is, the population of the evolutionary algorithm is identical to the population in the simulation runs. An advantage of this approach is that the need for computational resources is decreased significantly. Before we were doing 10 repetitions of simulation runs for each genome which sums to 500 simulation runs per generation. Now we do only 10 repetitions of simulation runs per generation because we evaluate all genomes within the same simulation run. However, this approach is ineffective as seen in Fig. 4i. While for ring length $L = 5$ the approach is competitive to the homogeneous approach (cf. Fig. 4b, note different scales), it obtains poor fitness for bigger ring lengths. Also the results for covered distance (Fig. 4j) and size of the largest cluster (Fig. 4k) indicate the poor performance as they are similar to the results of the random genomes (Figs. 4g and h). The evolved inhomogeneous swarms always show random behavior (data not shown).

4   Discussion and Conclusion

Following the idea of the free-energy principle [Friston, 2010, Friston et al., 2006], that is, minimizing surprise (prediction errors), we have shown that a variety of basic collective behaviors can be evolved by selecting for the ability of making good predictions only. Without generating any direct selective pressure on the actual evolved behavior these typical swarm behaviors emerge depending on the swarm density. Our findings are supported by control experiments. The comparison to random behaviors as generated by random genomes are significantly different. The comparison to evolving behaviors in an inhomogeneous swarm shows that homogeneity in the genomes and hence in the behaviors of a swarm is crucial. However, we do not rule out the possibility to evolve interesting collective behaviors in an inhomogeneous swarm if the inhomogeneity is limited (e.g., initializing to a homogeneous population, minimizing the effect of mutations, greedy selection).

The comparison of this approach to the concept of novelty search [Lehman and Stanley, 2008, Mouret and Doncieux, 2009] is interesting. Novelty search generates selective pressure towards unpopulated regions of behavior space, that is, it tries to trigger the evolution of novel behaviors that were not observed before. Our approach of minimizing prediction error seems to implement the opposite, that is, it seems to reward repetitive and hence easy to predict behaviors (on the time scale of a genome's evaluation while novelty search operates on the time scale of generations). Still, the approach is effective and generates a variety of different behaviors. In contrast to novelty search no measure of behavioral distance needs to be defined. In order to apply the proposed approach to the goal-oriented evolution of collective behaviors, additional research is necessary to determine options to guide the evolution.

The similarity between the trajectory plots of the evolved behaviors especially for swarms showing flocking (see Fig. 2d) and space-time diagrams of cellular automata (e.g., see Crutchfield and Mitchell [1995, Fig. 2]) is probably more

than mere coincidence. Also the concept of particles in cellular automata seems comparable to the concept of flocks here (see Crutchfield and Mitchell [1995, Table 2]). The reported swarm model is spatially 'almost discrete' (agents are displaced by a constant distance) and the agents move in continuous space only due to the continuous initialization of positions and the additive noise. Furthermore, the sensors also discretize an agent's neighborhood spatially (see Fig. 1a), hence the action network might implement simple update rules as seen in cellular automata.

There is also a low similarity to cellular automata in the context of computation 'at the edge of chaos' as discussed in [Langton, 1990, Crutchfield and Young, 1990]. The results concerning the best fitness, the covered distance, and the size of the largest cluster (see Fig. 4b-d) indicate that for ring lengths of $L \in [10, 25]$ the evolutionary optimization problem is possibly more difficult and the required behaviors are possibly more complex. For high ($L < 10$) and for low swarm densities ($L > 30$) the sensory input is rather monotonous (see results of entrop measurements, Fig. 4e, a majority of sensors give frequently a '1' for high densities or a majortiy of sensors give frequently a '0' for low densities). For medium densities the sensory input is more likely to change over time and consequently more complex behaviors seem to emerge (e.g., flocking with flocks having different velocities depending on the flock size).

From the perspective of defining optimal conditions for collective behavior, the existence of an optimal density for good swarm performance is known [Hamann, 2013, 2012]. The universality of these findings seems to be confirmed by the results concerning the best fitness, the covered distance, and the size of the largest cluster (see Fig. 4b-d).

Future work includes the investigation of how this approach can be utilized also as an approach to open-ended evolution to evolve desired behaviors similar to the concept of novelty search. Another subject will be to investigate whether minimizing prediction error triggers the evolution of positive and negative feedback loops within swarms.

References

Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford Univ. Press, New York, NY, 1999.

Manuele Brambilla, Eliseo Ferrante, Mauro Birattari, and Marco Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, 2013. ISSN 1935-3812. doi: 10.1007/s11721-012-0075-2. URL http://dx.doi.org/10.1007/s11721-012-0075-2.

Jerome Buhl, David J. T. Sumpter, Iain D. Couzin, Joe J. Hale, Emma Despland, E. R. Miller, and Steve J. Simpson. From disorder to order in marching locusts. *Science*, 312 (5778):1402–1406, 2006. doi: 10.1126/science.1125142.

James P. Crutchfield and Melanie Mitchell. The evolution of emergent computation. *Proc. Natl. Acad. Sci. USA*, 92:10742–10746, November 1995.

James P. Crutchfield and Karl Young. Computation at the onset of chaos. In Wojciech Zurek, editor, *Complexity, Entropy, and Physics of Information*. Addison-Wesley, Reading, MA, 1990.

Peter Dayan, Geoffrey E. Hinton, Radford M. Neal, and Richard S. Zemel. The Helmholtz machine. *Neural Computation*, 7(5):889–904, 1995.

Marco Dorigo and Erol Şahin. Guest editorial: Swarm robotics. *Autonomous Robots*, 17(2-3): 111–113, 2004.

Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.

Karl Friston, James Kilner, and Lee Harrison. A free energy principle for the brain. *Journal of Physiology Paris*, 100(1):70–87, 2006.

Jorge Gomes, Paulo Urbano, and Anders Lyhne Christensen. Evolution of swarm robotics systems with novelty search. *Swarm Intelligence*, 7(2-3):115–144, 2013.

Heiko Hamann. *Space-Time Continuous Models of Swarm Robotics Systems: Supporting Global-to-Local Programming*. Springer, Berlin, Germany, 2010.

Heiko Hamann. Towards swarm calculus: Universal properties of swarm performance and collective decisions. In Marco Dorigo, Mauro Birattari, Christian Blum, Anders Lyhne Christensen, Andries Petrus Engelbrecht, Roderich Groß, and Thomas Stützle, editors, *Swarm Intelligence: 8th International Conference, ANTS 2012*, volume 7461 of *LNCS*, pages 168–179, Berlin, Germany, 2012. Springer. URL http://dx.doi.org/10.1007/978-3-642-32650-9_15.

Heiko Hamann. Towards swarm calculus: Urn models of collective decisions and universal properties of swarm performance. *Swarm Intelligence*, 7(2-3):145–172, 2013. URL http://dx.doi.org/10.1007/s11721-013-0080-0.

Heiko Hamann. Evolving prediction machines: Collective behaviors based on minimal surprisal. In *Int. Conf. on Genetic and Evolutionary Computation (GECCO 2014)*. ACM, 2014. [extended abstract], in press.

Martin Hammer and Randolf Menzel. Learning and memory in the honeybee. *The Journal of Neuroscience*, 15(3):1617–1630, 1995.

Christopher G. Langton. Computation at the edge of chaos: phase transitions and emergent computation. In *Proceedings of the ninth annual international conference of the Center for Nonlinear Studies on Self-organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks*, pages 12–37, 1990.

Joel Lehman and Kenneth O. Stanley. Exploiting open-endedness to solve problems through the search for novelty. In S. Bullock, J. Noble, R. Watson, and M. A. Bedau, editors, *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, pages 329–336. MIT Press, 2008.

Alcherio Martinoli, Kjerstin Easton, and William Agassounon. Modeling swarm robotic systems: A case study in collaborative distributed manipulation. *Int. Journal of Robotics Research*, 23(4):415–436, 2004.

Jean-Baptiste Mouret and Stéphane Doncieux. Using behavioral exploration objectives to solve deceptive problems in neuro-evolution. In *Proceedings of the 11th annual conference on Genetic and Evolutionary Computation (GECCO'09)*, pages 627–634. ACM, 2009.

Andrew L. Nelson, Gregory J. Barlow, and Lefteris Doitsidis. Fitness functions in evolutionary robotics: A survey and analysis. *Robotics and Autonomous Systems*, 57:345–370, 2009.

Stefano Nolfi and Dario Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, 2000.

Thomas Schmickl and Karl Crailsheim. Bubbleworld.Evo: Artificial evolution of behavioral decisions in a simulated predator-prey ecosystem. In *From Animals to Animats 9*, volume 4095 of *LNCS*, pages 594–605. Springer, 2006.

Vito Trianni. *Evolutionary Swarm Robotics - Evolving Self-Organising Behaviours in Groups of Autonomous Robots*, volume 108 of *Studies in Computational Intelligence*. Springer, Berlin, Germany, 2008.

Vito Trianni, Roderich Groß, T. H. Labella, Erol Şahin, and Marco Dorigo. Evolving aggregation behaviors in a swarm of robots. In Wolfgang Banzhaf, Jens Ziegler, Thomas Christaller, Peter Dittrich, and Jan T. Kim, editors, *Advances in Artificial Life (ECAL 2003)*, volume 2801 of *Lecture Notes in Artificial Intelligence*, pages 865–874. Springer, 2003.

Vito Trianni, Thomas H. Labella, and Marco Dorigo. Evolution of direct communication for a swarm-bot performing hole avoidance. In Marco Dorigo, Mauro Birattari, Christian Blum, Luca Maria Gambardella, Francesco Mondada, and Thomas Stützle, editors, *Ant Colony Optimization and Swarm Intelligence (ANTS 2004)*, volume 3172 of *LNCS*, pages 130–141. Springer, 2004.

Vito Trianni, Elio Tuci, Kevin M. Passino, and James A.R. Marshall. Swarm cognition: an interdisciplinary approach to the study of self-organising biological collectives. *Swarm Intelligence*, 5(1):3–18, 2010.

Hermann von Helmholtz. *Handbuch der physiologischen Optik*. Ludwig Voss, Leipzig, Germany, 1867.

Christopher R. Ward, Fernand Gobet, and Graham Kendall. Evolving collective behavior in an artificial ecology. *Artificial Life*, 7(2):191–209, 2001.