Evolving Prediction Machines: Collective Behaviors Based on Minimal Surprisal

Heiko Hamann Department of Computer Science, University of Paderborn, Zukunftsmeile 1, 33102 Paderborn, Germany heiko.hamann@uni-paderborn.de

Categories and Subject Descriptors

I.2.2 [Artificial Intelligence]: Automatic Programming— Program synthesis

General Terms

Algorithms

Keywords

swarm intelligence, swarm behavior, evolutionary robotics, free-energy principle

1. INTRODUCTION

Evolutionary algorithms can be applied to evolve controllers for single robots and similarly for groups of robots. Collective behaviors of groups of robots are investigated in the field of swarm robotics [1] which is the application of swarm intelligence to the field of robotics. An option is to to apply methods from evolutionary robotics to swarm robotics, that is, evolutionary swarm robotics [3]. In a standard approach a fitness function is used to reward behavioral features that are desired. In this paper we follow an alternative approach that generates collective behaviors without explicit selection for desired behaviors. We evolve agents, that mainly focus on predictions of their future perceptions, but still observe a number of different collective behaviors as a result. This approach is motivated by the hypothesis that perception is essentially a process of probabilistic inference—an idea that goes back to Helmholtz [4]. Following this concept, the main task of a brain is to figure out appropriate causes to its perceptions. Hence, the brain is interpreted as a 'prediction machine' that learns to model its perceptions. A mathematical framework by Friston [2] defines an informationtheoretic analogon to the thermodynamic (Helmholtz) free energy which is basically the prediction error here. Friston's approach of a 'free-energy principle' might open up opportunities to formulate a unified brain theory [2].

GECCO'14, July 12–16, 2014, Vancouver, BC, Canada. ACM 978-1-4503-2881-4/14/07. http://dx.doi.org/10.1145/2598394.2598507.





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

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

Figure 1: Experimental setting and sensor setting.

Here, we do not learn sophisticated probabilistic models but merely evolve weights of artificial neural networks (ANN). There are two ANNs in each agent: an ANN implementing the prediction machine and an ANN implementing a regular controller. The selective pressure is on the prediction machine by rewarding minimal prediction errors (surprisal) while the actual controller receives no direct selective pressure. Our results show a variety of basic collective behaviors, such as dispersion, aggregation, and flocking. We define 4 basic collective behaviors that are based on two features only: motion (moving or stopped) and relative positions (minimal distances between agents or maximal distances). We categorize the behaviors along these two dimensions: dispersion (maximal distances, stopped), aggregation (minimal distances, stopped), random (maximal distances, moving), and flocking (minimal distances, moving).

2. MODEL

We investigate the collective motion of agents in a onedimensional system in the form of a circle that we call ring (see Fig. 1a). The circumference of this ring is denoted by ring length L. The agents have no global reference frame and cannot discriminate between clockwise or counterclockwise motion. They have two available actions: staying with the current direction or switching the direction while not being allowed to stop. The agents have 4 sensors: L1, L0, R0, and R1. Each sensor covers an interval of the agent's neighborhood as defined by the sensor distances s_0 and s_1 (see Fig. 1b). These are discrete sensors that output 1 if there are agents within the respective interval or 0 if there is no agent.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).



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

In the following the swarm size is fixed to N = 20 and the ring length L is varied (i.e., different swarm densities N/L). The agent positions and directions are initialized to random uniform values. Agents move with speed v = 0.1 with a small noise that is added to the speed and which is also random uniform over [-0.01, 0.01]. Agents can pass each other without any interference.

Each agent has two ANNs that we call 'action network' and 'prediction network'. The action network has 5 input neurons, 2 hidden neurons, and 1 output neuron. The inputs are the 4 sensor values (L1,L0,R0,R1) and the last action. The output neuron determines the next action based on a threshold. The prediction network is a recurrent network due to self-loops of the hidden neurons and has 5 input neurons, 4 hidden neurons, and 4 output neurons. The input is the same as with the action 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 simple evolutionary algorithm (proportionate selection, elitism of 1, population size is 50) operates on genomes consisting of two sets of weights (for the action network and the prediction network). Initially a population of random weights is generated. The genomes are evaluated by applying the swarm simulation. In each particular evaluation run all agents have the same networks. The fitness function rewards good predictions of the prediction network. It is the sum over all correct predictions per sensor averaged over the evaluation period, the swarm, and 10 independent simulation runs. The theoretical best fitness is 4. Each weight of the ANNs is mutated with a probability of 0.05. Evolution is run for 30 generations totaling to 1500 evaluations.

3. RESULTS

The 4 basic collective behaviors discussed above occurred in the experiments for different ring lengths L (i.e., different swarm densities). Examples of these behaviors are shown in Fig. 2 which shows the trajectories of the whole swarm for



Figure 3: Best fitness over generations and comparison of best fitness for different ring lengths L.

each behavior. In Fig. 2a an example for a behavior that can arguably be called dispersion is shown (i.e., agents keep a distance between them). The agents keep switching their direction of motion and hence do not cover any distance. Fig. 2b gives an example of an aggregation behavior. Agents minimize the distances between them and cover no distance. Fig. 2c gives an example of a random behavior. The agents keep moving in the same direction without reacting to each other. A flocking behavior is shown in Fig. 2d. Agents minimize the distances between them and keep moving. Fig. 3a gives the best individual fitness over generations for 200 independent evolutionary runs for ring length L = 25. Initially there is a steep increase up to generation t = 9which is followed by a saturation in the best fitness. The last generation's median is 3.33 which means that the prediction network is predicts an average of 83.3% sensor values cor-40, 45, 50 200 independent evolutionary runs were done. Fig. 3b gives a comparison of the last generation's best fitness for all tested ring lengths L. 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 reached for L = 25(3.33). For L < 30 the medians decrease with increasing ring length and start to increase again for L > 25. It seems that predictions are more difficult for medium densities $(10 \le L \le 35)$. Predictions are simple for high densities as most sensors detect neighbors without too much change. Predictions for low densities are simple as most sensors detect no neighbors. However, sensor input varies for medium densities which complicates making good predictions.

4. CONCLUSION

Motivated by the idea of the free-energy principle [2] we have presented an approach to evolve basic collective behaviors by selecting for good prediction machines. Without a direct selective pressure on behaviors, typical swarm behaviors emerge with dependencies on the swarm density.

5. **REFERENCES**

- M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, 2013.
- [2] K. Friston. The free-energy principle: a unified brain theory? Nature Rev. Neuroscience, 11(2):127–138, 2010.
- [3] V. Trianni. Evolutionary Swarm Robotics, volume 108 of Studies in Comp. Intelligence. Springer, 2008.
- [4] H. von Helmholtz. Handbuch der physiologischen Optik. Ludwig Voss, Leipzig, Germany, 1867.