Learning of Individual Sensorimotor Mapping to Form Swarm Behavior from Real Fish Data

Hiroyuki Iizuka, Yosuke Nakamoto and Masahito Yamamoto

Graduate school of Information Science and Technology, Hokkaido University
North14, West9, Kita-ku, Sapporo 060-0007, Japan
iizuka@complex.ist.hokudai.ac.jp

Abstract

Swarms of birds and fish produce well-organized behaviors even though each individual only interacts with their neighbors. Previous studies attempted to derive individual interaction rules using heuristic assumptions from data on captured animals. We propose a machine learning method to obtain the sensorimotor mapping mechanism of individuals directly from captured data. Data on swarm behaviors in fish was captured, and individual positions are determined. The sensory inputs and motor outputs are estimated and used as training data. A simple feedforward neural network is trained to learn the sensorimotor mapping of individuals. The trained network is implemented in the simulated environment and resulting swarm behaviors are investigated. As a result, our trained neural network could reproduce the swarm behavior better than the Boids model. The reproduced swarm behaviors are evaluated in terms of three different measures, and the difference from the Boids model is discussed.

Introduction

When many individuals get together, the group shows surprisingly united and coordinated behaviors as if the group behaves like a single unit. The question regarding how the united and coordinated swarm is created from individual local interactions has attracted attention from many scientists for a long time. Recently, real data from individuals forming swarms became available thanks to technological developments and increased computation power, which facilitates collection and tracking of behavior of individuals in a swarm. Those data are very useful for understanding how the swarm behaviors are built and maintained from each individual perspective (Ballerini et al., 2008), to derive individual interaction rules for real animals (Tien et al., 2004; Lukeman et al., 2010; Herbert-Read et al., 2011; Katz et al., 2011) and to build statistical mechanics model to bridge the gap between micro and macro behaviors (Bialek et al., 2012, 2014).

If it is possible to derive complete individual interaction rules for real animals, then we might apply those findings to multiple-robots controllers that can establish coordinated behaviors working like a single unit. And also, we could build hybrid alife system consisting of real swarming animals and artificial agents controlled by the derived interaction rules. The approach becomes popular to understand animal behaviors (Butail et al., 2013; Worm et al., 2017) and the interaction rules make it possible to perform more dynamic experiments than using predefined or playback behaviors. However, building such interaction rules actually used by animals becomes difficult without fully understanding animal cognition. To build the interaction rules, we have no choice but to use heuristic assumptions to decide what they sense, how they recognize other members of the swarm, and how they react. This is similar to good old fashioned artificial intelligence (GOFAI). To develop an intelligent model that can recognize the world, we need to prepare representations as symbolic units of knowledge. With the GOFAI-type approaches, the realized intelligence or swarm behaviors must be limited to pre-defined assumptions.

In contrast to this approach, embodiment and situatedness is a complementary approach, where life-like behaviors including intelligence can be realized in the flow of sensorimotor coordination through interaction with an environment without predefined knowledge (Varela, 1979; Brooks, 1991; Harvey et al., 1997; Pfeifer and Scheier, 2001). Our idea is that the embodiment approach also should be used to describe swarm behaviors. The interaction rules are not defined in advance. Rather, they are obtained as sensorimotor flows of individuals behaviors within a swarm. We use an artificial neural network to determine the interaction rules. The neural network learns the sensorimotor mapping of individuals that participate in real swarms.

Our final goal is to learn and obtain the interaction rules from real swarm behaviors using an artificial neural network, to reproduce swarm behavior where individuals are controlled by the trained neural network, and to understand how swarms are organized and maintained dynamically. As a first attempt, we examine whether the neural network can learn sensorimotor mapping from interactions between individuals and test if swarm behaviors reproduced by the trained network are similar to the real swarm behaviors in this paper. As an evaluation of swarm behavior, we define three swarm characteristics and compare the reproduced swarm with the original swarm and the conventional Boids model (Reynolds, 1987). In the experiment, swarm behaviors in fish are captured and used for learning and evaluation.

Data from fish

12 African lampeyes (*Aplocheilichthys normani*) were used as the target fish swarm. The lampeye is a kind of medaka, which is known as a fish that forms a swarm and act depending on visual cues (Imada et al., 2010).

Data collection

Behavioral data were obtained from movements of 12 lampeyes recorded with a camera fixed (Fig. 1 left). The lampeyes were allowed to move around in a styrofoam box. The styrofoam box was filled with shallow water to restrict the motion of lampeyes in two dimensions. Recording was performed for 10 minutes at 15 frames per second (9,000 frames in total), and the resolution of the video was 640 $\times$ 480 pixels. After recording, all video frames were expanded to 960 $\times$ 720 pixels for detection and tracking processes.

The actual area where individuals could move around is W = 615 $\times$ H = 435 pixels in the converted video. Individual lampeyes never moved outside of this area.

The position and velocity of each individual at the t-th frame were obtained as follows:

1. Converting the t-th frame to a binary image. Because the background of the box is white, the lampeyes can be detected as black objects.

2. Calculating the center position of each black area at the t-th frame (Fig. 1 right).

3. Applying the above two operations to the t+1-th frame and calculating the displacement vector between them. This vector represents the t-th velocity of the individual (Fig. 2).

If the number of detected individuals was less than 12 due to missing or overlapping individuals, the frame was omitted. Only clean consecutive data were used. Eventually, the number of effective frames was 6,012, which corresponds to about 400 seconds of data. Because there were 12 individuals in each frame, we obtained 72,144 ($12 \times 6,012$) data points describing individual behavior.

Neural network model

To investigate how individuals react, we propose a neural network be used to determine individual interaction rules and to learn how individuals react to each other.

Structure of neural network

A simple feedforward neural network was trained to obtain sensorimotor mapping of individuals that exhibit swarm behavior. The neural network consists of three layers (i.e., the

Figure 1: Left: Example image of lampeye behavior captured with a camera. Right: Positions of individual lampeyes.

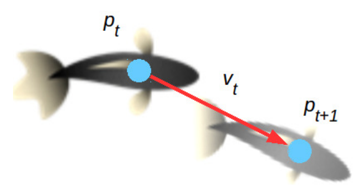

Figure 2: Position and velocity at the t-th frame.

neural network has one hidden layer), and each layer has (l, m, n) = (20, 64, 2) neurons, respectively. All nodes between layers are fully-connected.

The network receives information on the three nearest neighbors, its own velocity, and the distance to the chamber walls. It senses the position and velocity of each neighbor. For the neural network to be sensitive to distance, the positions are given by polar coordinates, and velocities are given by vectors. It should be noted that identity detection of neighbors was not performed. The set of inputs for each individual were sorted in terms of the distances, and the inputs for the three nearest neighbors are given to the network in order. The input nodes for a given individual information changes if the distance order changes. The network can sense its own velocity vectors as well. The wall information is also provided to the network. The network can sense the distance to the nearest wall ($\frac{1}{1+d}$) and the angles ($\frac{1}{1+d}\sin\gamma$ and $\frac{1}{1+d}\cos\gamma$) (Fig. 3).Those inputs become zero if the position is far from wall. In total, the number of input nodes was 20 (= 15 for the three nearest neighbors, 3 for the nearest wall and 2 for its own velocity).

The outputs from the nodes in the hidden and output layers can be calculated as follows:

$$h_j = f(\sum_{i=1}^{l} v_{ji} x_i + a_i) \tag{1}$$

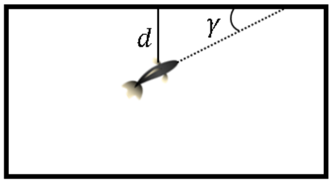

Figure 3: Wall inputs for the network.

$$y_k = g(\sum_{j=1}^{m} w_{kj} h_j + b_j) \quad (2)$$

where x is the set of inputs, y is the set of final outputs from the network, v_{ji} and w_{kj} represent weights, and a_i and b_j represent biases. $f(\cdot)$ is the rectified linear units (ReLU) function and $g(\cdot)$ is the identity function. The network outputs a set of velocity vectors, which cause the position changes for the next time step.

Because the neural network is just a simple feedforward network, it calculates the subsequent movements from the current information for an individuals neighbors. The training data consists of the position and velocity of each neighbor, an individuals own velocity, and the distance from the chamber walls as inputs. The motion over the next time step is estimated from the captured data. We assumed that all individuals share the same neural network, and the behavior of all captured individuals was used as training data for the neural network. After preparing the training data, the neural network was trained using the backpropagation method with the mean square error function. The code for capturing images and training the network were implemented in OpenCV and TensorFlow.

Training results

Training data were supplied to the neural network repeatedly. An epoch is defined as a single iteration of training over the entire set of training data. The order of training data were randomized at every epoch, and the network was updated every 100 mini-batches. Figure 4 shows the errors during training. The errors decreased gradually, and the network was successfully trained to produce each subsequent motion from the set of given inputs. However, what is shown with changes in errors is that the neural network can output subsequent with smaller errors at successive time steps. It does not necessarily mean that all individuals are controlled by the trained neural network can show swarm behaviors like a real swarm.

Reproducing swarm behaviors

After training the neural network, we simulated the swarm behavior consisting of individuals controlled by the neural network. The number of individuals N was set to 12. All individuals used the same trained neural network. At the beginning, random initial positions were provided to individuals. The input values for each individual were calculated from local information at every time step. Those inputs were supplied to the network, and the motion of each individual was obtained from the outputs of the network. The resulting behaviors of all individuals were obtained by repetition. The area where individuals can move was set to 615 $\times$ 435, which is nearly the same size as the area where lampeyes were allowed to move in the converted video. When and individual was predicted to move outside this area, the velocity was simply reversed.

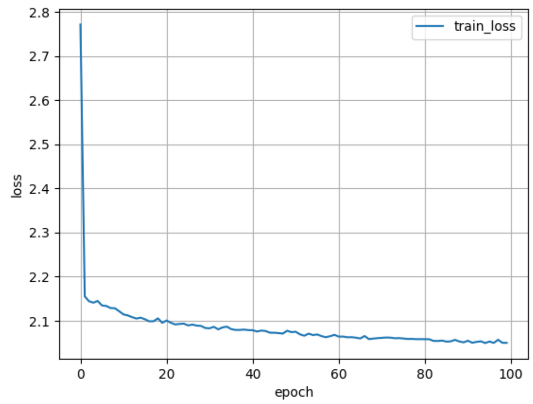

Figure 4: Average mean squared error reduction as a function of the number of epochs

Evaluation of swarm behavior

To investigate whether the swarm controlled by the trained neural network can successfully reproduce the original swarm, it is necessary not only to observe the actual behavior but also to compare the predicted swarm characteristics. There is no common consensus regarding which method should be used to evaluate real swarm behaviors. However, there are some basic methods for characterizing swarm behavior by calculating the position and velocity for each individual (Viscido et al., 2004; Gautrais et al., 2012). In this study, we use the following three similar measures, i.e., D_{com}, D_{min}, and Φ, to evaluate swarm behaviors

D_{com} represents the average distance between each individual and the center of a group, described as:

$$C^t = \frac{1}{N} \sum_{i=1}^{N} p_i^t \quad (3)$$

$$D_{com}^t = \frac{1}{N} \sum_{i=1}^{N} |C^t - p_i^t| \quad (4)$$

where N is the number of individuals in the swarm. D_{com} becomes small when the swarm moves as a single group and becomes large as individuals in the swarm move independently.

D_{min} is the average distance to the nearest individual. if the position of the nearest individual is m_i, then D_{min} can be expressed as:

$$D^t_{min} = \frac{1}{N}\sum_{i=1}^{N} |m^t_i - p^t_i| \tag{5}$$

D_{com} evaluates the cohesiveness as a single large group, while D_{min} evaluates the cohesiveness regardless of the number of groups.

Φ represents the agreement level of the direction of the entire swarm, which is described as:

$$\Phi^t = \frac{1}{N}\left|\sum_{i=1}^{N} \frac{\vec{v^t_i}}{|\vec{v^t_i}|}\right| \tag{6}$$

Φ approaches 1 when all individuals face the same direction, and 0 when all individuals face different directions ($0 \leq \Phi \leq 1$).

Boids model for comparison

Reynolds proposed the Boids model, which simulates the swarm behavior of birds (Reynolds, 1987). Each individual decides how to move based on local information using simple rules. Despite the fact that all individuals are controlled by these rules in a local manner, the results show that the swarm behaves similar to a real swarm. Te swarm behavior reproduced by the trained neural network is evaluated by comparing with results from the Boids model.

In the boids model, each individual moves based on the following three simple rules. (1) Cohesion: each individual moves toward the center of its neighbors, (2) Separation: each individual moves away from its neighbors to avoid collision, and (3) Alignment: each individual moves to align with the average heading of its neighbors. The parameters used in the Boids model are manually adjusted to fit D_{min} observed for real fish.

Results

Figure 5 (a), (b), and (c) show example trajectories for real fish (i.e., lampeyes), individuals controlled by our trained neural network, and Boid rules, respectively. Real fish exhibited jerky motion, but individuals formed a swarm rather than moving on their own. The example shows results when a single swarm was divided into two subgroups. In the behaviors reproduced by our trained network, the results also showed an aggregation tendency, and the swarm moved around as a single unit. However, the swarm was sometimes divided into two groups. The example trajectory shows the division. The swarm was divided into two groups at the left

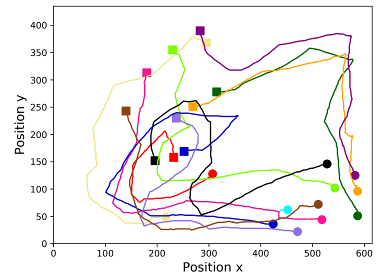

(a)

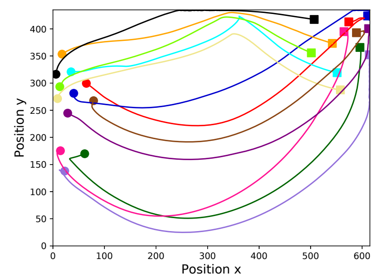

(b)

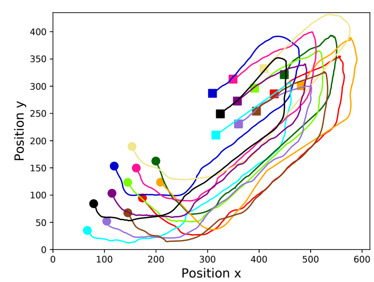

(c)

Figure 5: Example trajectories of (a) lampeyes, (b) individuals controlled by our trained neural network, and (c) Boids models (circle $\rightarrow$ square)

side. Both groups moved to the right and became one large group at the top-right corner. These two examples of division are arbitrarily selected to show that the division could happen without including perturbations. On the other hand, it is quite difficult to observe such division of groups in the Boids model without any external perturbations like obstacles. The Boids rules show a strong aggregation tendency, which stabilizes the swarm as a single group. The trajectories of Boids became jittery like real fish. This could happen because the parameters were adjusted to fit D_{min} observed for real fish and the separation powers, which was applied when the neighbors came into a visual range, became stronger than usual . Our trained neural network resulted in aggregation, and separation into subgroups appeared as an instability in a real fish swarm. It should be noted that instability of the Boids model might change in response to the

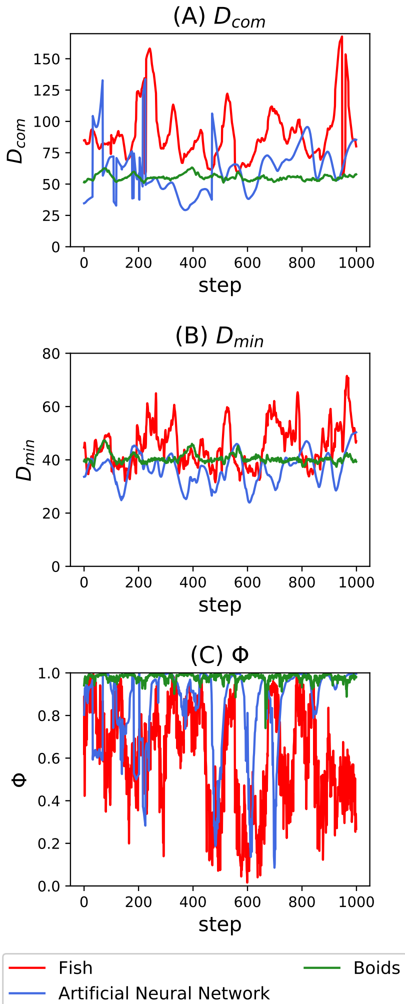

Figure 6: Time evolution of D_{com}, D_{min}, Φ in the lampeye swarm, our reproduced swarm, and Boids.

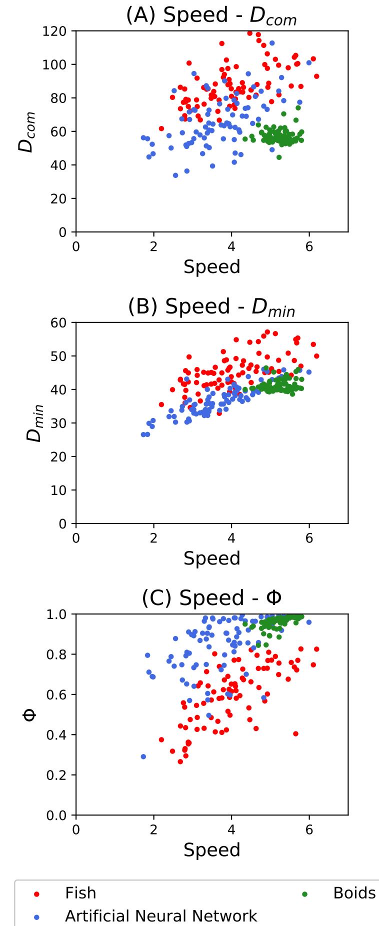

Figure 7: Correlation between group speed and three basic measures, D_{com}, D_{min}, and Φ. Red, blue and green points show the results for a real fish swarm, the swarm behavior reproduced by our trained neural network, and Boids results, respectively.

number of individuals (Ikegami et al., 2017).

The time evolution of three characteristics described in the above are shown in Fig. 6 to clarify the difference of between swarm behaviors. In the lampeye swarm, D_{com} oscillated up and down greatly while D_{min} remained small with smaller variances. Those characteristics are also observed in our reproduced behaviors. On the other hand, the evaluation measures of the Boids model became much more stable than the remaining measures. The polarities Φ for real fish and our reproduced behaviors often broke down, but they recovered immediately. The Boids model is too stable, and the polarity rarely broke down. These results suggest that our trained neural network could reproduce the swarm behaviors better than the Boids model. The primary difference is the instability of the swarm. Actually, Hartman and Benes (2006) extend the Boids model by introducing a separation power called leadership as a swarm instability in their model.

It is known that there is a correlation between group speed and polarity in real fish swarms (Viscido et al., 2004). To evaluate our trained neural network, correlations between group speed and the three basic measures are investigated. The group speed and the three measures are obtained by averaging over 75 steps, which corresponds to about 5 seconds of data. Figure 7 shows the correlations for a real fish swarm, the swarm behavior reproduced by our trained neural network, and results from the Boids model. Our results indicate that correlations exists between group speed and polarity Φ and between group speed and D_{com} and D_{min} in a real fish swarm. These results also suggest that our trained neural network could successfully reproduce the swarm behaviors.

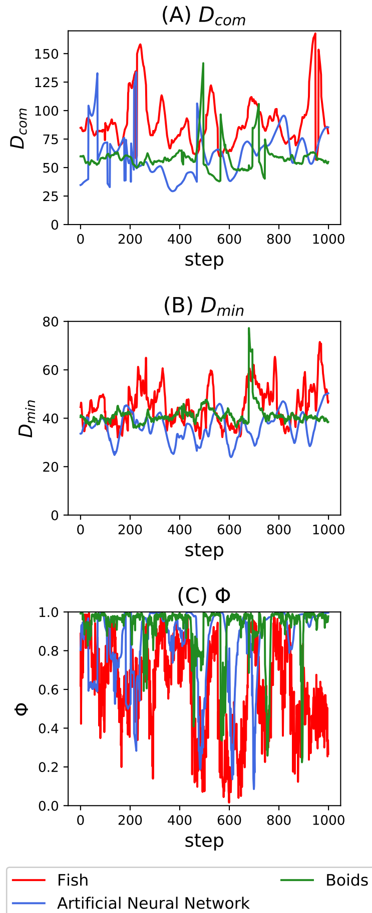

Figure 8: Time evolution of D_{com}, D_{min}, and Φ. These are the same as Fig. 6, except for the Boids results. Noise was added to the individual movements in the Boids model. For comparison, lampeye results and our reproduced results are also shown.

What is the difference in swarm behavior between our model and the conventional Boids model? Because the boids model is too stable without external perturbations as we show, we added noise to the Boids model as an external perturbation. Noise was added to the motion for each individual. Figure 8 shows D_{com} , D_{min}, and Φ when noise was added to the Boids model. The D_{min} results show better agreement with the real fish results, but D_{com} and Φ show less agreement. Thus, the simple instability produced by noise did not improve the Boids model. In other words, our trained neural network could obtain not simple instability of swarm behavior.

CONCLUSIONS

In this paper, we proposed a machine learning method that can be applied to describe the behavior of real fish, and we show that the trained neural network could learn sensorimotor mapping of individuals from data on real fish. In our model, we assume that all individuals share the same feedforward neural network, but individuals must have different identities to demonstrate a tendency for some individuals to act as leaders or followers. Such a variety causes different information to flow in the swarm (Reebs, 2000). If we can track the behavior of real individuals at all times, it is possible to train each network to obtain the sensorimotor mapping of each individual or to train a recurrent network to have memories. Our machine learning method can be a powerful tool to analyze and understand the behavior of more complicated swarms.

References

Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, E., Giardina, I., Lecomte, V., Orlandi, A., Parisi, G., Procaccini, A., Viale, M., and Zdravkovic, V. (2008). Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the national academy of sciences*, 105(4):1232–1237.

Bialek, W., Cavagna, A., Giardina, I., Mora, T., Pohl, O., Silvestri, E., Viale, M., and Walczak, A. M. (2014). Social interactions dominate speed control in poising natural flocks near criticality. *Proceedings of the National Academy of Sciences*, 111(20):7212–7217.

Bialek, W., Cavagna, A., Giardina, I., Mora, T., Silvestri, E., Viale, M., and Walczak, A. M. (2012). Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences*, 109(13):4786–4791.

Brooks, R. A. (1991). Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159.

Butail, S., Bartolini, T., and Porfiri, M. (2013). Collective response of zebrafish shoals to a free-swimming robotic fish. *PLoS One*, 8(10):e76123.

Gautrais, J., Ginelli, F., Fournier, R., Blanco, S., Soria, M., Chaté, H., and Theraulaz, G. (2012). Deciphering interactions in moving animal groups. *Plos computational biology*, 8(9):e1002678.

Hartman, C. and Benes, B. (2006). Autonomous boids. *Computer Animation and Virtual Worlds*, 17(3-4):199–206.

Harvey, I., Husbands, P., Cliff, D., Thompson, A., and Jakobi, N. (1997). Evolutionary robotics: the sussex approach. *Robotics and autonomous systems*, 20(2-4):205–224.

Herbert-Read, J. E., Perna, A., Mann, R. P., Schaerf, T. M., Sumpter, D. J., and Ward, A. J. (2011). Inferring the rules of interaction of shoaling fish. *Proceedings of the National Academy of Sciences*, 108(46):18726–18731.

Ikegami, T., Mototake, Y., Kobori, S., Oka, M., and Hashimoto, Y. (2017). Life as an emergent phenomenon: studies from a large-scale boid simulation and web data. *Philosophical Transactions A*, 375(2109).

Imada, H., Hoki, M., Suehiro, Y., Okuyama, T., Kurabayashi, D., Shimada, A., Naruse, K., Takeda, H., Kubo, T., and Takeuchi, H. (2010). Coordinated and cohesive movement of two small conspecific fish induced by eliciting a simultaneous optomotor response. *PLoS One*, 5(6):e11248.

Katz, Y., Tunstrøm, K., Ioannou, C. C., Huepe, C., and Couzin, I. D. (2011). Inferring the structure and dynamics of interactions in schooling fish. *Proceedings of the National Academy of Sciences*, 108(46):18720–18725.

Lukeman, R., Li, Y.-X., and Edelstein-Keshet, L. (2010). Inferring individual rules from collective behavior. *Proceedings of the National Academy of Sciences*, 107(28):12576–12580.

Pfeifer, R. and Scheier, C. (2001). *Understanding intelligence*. MIT press.

Reebs, S. G. (2000). Can a minority of informed leaders determine the foraging movements of a fish shoal? *Animal behaviour*, 59(2):403–409.

Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. *ACM SIGGRAPH computer graphics*, 21(4):25–34.

Tien, J. H., Levin, S. A., and Rubenstein, D. I. (2004). Dynamics of fish shoals: identifying key decision rules. *Evolutionary Ecology Research*, 6(4):555–565.

Varela, F. (1979). Principles of biological autonomy.

Viscido, S. V., Parrish, J. K., and Grünbaum, D. (2004). Individual behavior and emergent properties of fish schools: a comparison of observation and theory. *Marine Ecology Progress Series*, 273:239–250.

Worm, M., Kirschbaum, F., and von der Emde, G. (2017). Social interactions between live and artificial weakly electric fish: Electrocommunication and locomotor behavior of mormyrus rume proboscirostris towards a mobile dummy fish. *PLoS One*, 12(9):e0184622.