Interaction Rules Supporting Effective Flocking Behavior

Nicola Milano*
National Research Council
Institute of Cognitive Science
and Technologies
nicola.milano@unina.it

Stefano Nolfi
University of Naples Federico II
Natural and Artificial Cognition
Laboratory "Orazio Miglino"

Abstract Several simulation models have demonstrated how flocking behavior emerges from the interaction among individuals that react to the relative orientation of their neighbors based on simple rules. However, the precise nature of these rules and the relationship between the characteristics of the rules and the efficacy of the resulting collective behavior are unknown. In this article, we analyze the effect of the strength with which individuals react to the orientation of neighbors located in different sectors of their visual fields and the benefit that could be obtained by using control rules that are more elaborate than those normally used. Our results demonstrate that considering only neighbors located on the frontal side of the visual field permits an increase in the aggregation level of the swarm. Using more complex rules and/or additional sensory information does not lead to better performance.

Keywords
Behavioral rules, flocking, self-organization, autonomous agents, swarm behavior

1 Introduction

Collective behaviors exhibited by bird flocks, fish schools, and mammal herds are paradigmatic examples of self-organizing phenomena in which the global pattern emerges from the local interaction between individuals with no need for external stimuli. The collective behavior exhibited by these species is also robust to environmental perturbations and predator attacks. For these reasons, the study of these phenomena is relevant not only for ethology and evolutionary biology but also for social sciences and robotics.

The models proposed so far demonstrate that flocking behavior emerges from the interaction among multiple agents that simply steer toward the relative direction of their neighbors (Ginelli & Chaté, 2010; Gregoire et al., 2003; Reynolds, 1987; Vicsek et al., 1995). However, the precise nature of these steering rules and the relation between the detailed characteristics of the rules and the efficacy of the resulting collective behavior have not yet been analyzed in detail.

In this study, we analyze the effect of the strength with which individuals react to neighbors located in different sectors of their visual fields. Moreover, we analyze the potential benefit of more complex control rules that include nonlinearities, elaborate sensory information into internal states, and regulate steering based on the relative distance of neighbors. This is realized by evolving the parameters that determine the behavior of agents through an evolutionary algorithm and by carrying out multiple experiments in which we systematically vary the information available in the observation vector of agents and the architecture of the agents' controller.

Our results demonstrate that the utility of the information provided by neighbors located in the frontal sectors is greater than that of the neighbors located in the lateral sectors. Moreover, our results demonstrate that a more compact control strategy that relies on the neighbors located in the

* Corresponding author.

frontal sectors only, and eventually on a single neighbor located frontally, outperforms the standard model that also relies on the information provided by neighbors located on the lateral sectors. Finally, the analysis of experiments performed with neural controllers that are computationally more powerful than the standard model indicates that the additional capabilities do not permit better performance and are thus unnecessary.

After reviewing related research in the following section, we describe our experimental setting in section 3. We present our results in section 4 and discuss the implications of our findings in section 5.

2 Relation to the State of the Art

Several models of collective motion have been investigated in previous works. In this article, we focus on self-propelled particle methods that were successfully used to model detailed aspects of natural collective behavior (Ginelli & Chaté, 2010; Gregoire et al., 2003; Vicsek et al., 1995). In these methods, the agents are represented by point-like particles that move synchronously, and collisions are not simulated. The agents modify their direction of motion to align with nearby agents and eventually to escape predators. For examples of models in which the behavior of the agents is also influenced by separation and cohesion rules, see Reynolds (1987) and Gregoire et al. (2003).

In their seminal paper, Couzin et al. (2002) demonstrated that a model of this kind can generate four classes of collective behaviors observed in natural species: (a) swarm behavior, characterized by good aggregation combined with a low level of polarization (parallel alignment); (b) torus behavior, in which the agents perpetually rotate around an empty core; (c) dynamic parallel group behavior, in which the group is much more mobile compared to the previous case; and (d) highly parallel group behavior, in which the group self-organizes into a highly aligned arrangement. The model also accounts for rapid transitions between these four types of behavior and the instability of intermediate behavioral types. Indeed, it demonstrates that the transition between behaviors of different types can be triggered by minor variations of the parameters that regulate the interaction among individuals. Couzin et al. (2002) also demonstrated that differences in speeds and turning rates of individuals, possibly caused by different motivation strengths, influence the positions occupied by the agents in the group. This produces an implicit sorting of individuals in which agents with high, medium, and low speed and/or turning rates are located at the front, center, and rear of the swarm, respectively. In other words, the individuals assume different relative positions even if they are unable to detect their positions in the swarm.

Moreover, coordinated behavior in which individuals align with other nearby individuals permits minimizing collisions, facilitating the movement of the group, and transferring information among individuals implicitly. Indeed, the tendency to align permits transferring the information sensed by an individual that detects an obstacle to nearby individuals, which react accordingly (Couzin et al., 2002).

By analyzing the 3-D positions of individual birds during flocking behavior, Ballerini et al. (2008) demonstrated that interactions among birds depend not on the metric distance but rather on the topological distance. In other words, the birds vary their behavior based not on the position and direction of neighbors located within a maximum distance but rather on the six to seven neighbors that are closest within different sectors of their visual fields. Such a topological strategy is compatible with the lack of information caused by visual occlusion and is more effective than the metric strategy. Indeed, simulation studies comparing metric and topological strategies demonstrate that the latter lead to behaviors that are more robust to perturbations and scale better to variations in the density of individuals (Ballerini et al., 2008; Camperi et al., 2012).

The fact that birds consider only six to seven conspecifics seems to be due not to cognitive limitations of agents (Ballerini et al., 2008) but rather to functional reasons. Indeed, simulation studies show that the optimal number of conspecifics is three to five for 2-D settings and six to seven for 3-D settings (Inada & Kawachi, 2002; Tegeder & Krause, 1995).

On the other hand, the detailed characteristics of the rules that regulate the interactions among individuals are unknown. Current models implicitly assume that the relation between the steering response of individuals and the relative orientation of neighbors is linear, that the impact of the relative orientation of neighbors should be the same regardless of their relative positions, and that the orientation of neighbors is the only information that matters. In this work, we investigate the validity of these implicit assumptions by varying these characteristics and analyzing the control strategies optimized through an evolutionary algorithm. Our results indicate that focusing only on the closest neighbors located in a subset of visual sectors leads to significantly better performance and that information provided by frontal sectors is more useful than information provided by lateral sectors. Finally, our results indicate that availability of nonlinear rules and distance information does not provide an advantage; that is, the relative orientation of neighbors is the only information that matters.

Several works used artificial evolution to study selective pressure promoting the evolution of aggregating behaviors (Biswas et al., 2014; Demšar et al., 2015; Demšar et al., 2016; Kunz et al., 2006; Olson et al., 2013; Olson et al., 2015; Olson et al., 2016; Tosh, 2011; Wood & Ackland, 2007). Other works carried out evolutionary experiments to analyze the impact of the number of neighbors considered (Morrell et al., 2015), to study the evolution of laterally expanded visual perception (Wood & Ackland, 2007), and to model collective perception (Van Diggelen et al., 2022). As mentioned, instead, we use artificial evolution to study the role of the rules determining how agents react to their neighbors.

3 Method

To investigate the role of control rules that regulate interactions among individuals, we carried out several sets of experiments in which we compared the performance and behavior of individuals provided with manually designed and evolved controllers. The swarms considered are composed of 512 homogeneous individuals located in a 2-D toroidal arena of 1,000 × 1,000 m that contains a predator. The individuals' goal is to survive the predator.

At the beginning of each episode, we initialize the initial position and orientation of individuals and the predator randomly with a uniform distribution within the 1,000 × 1,000 m area and within [−pi, pi], respectively. The velocity of individuals and the predator is set to 5 m/s and 6 m/s, respectively. The visual field of individuals has a total amplitude of 270° and is divided into seven sectors of 38.57° each. The topological neighbors influencing the behavior of each agent correspond to the closest individuals located within each sector, if any. Consequently, each individual determines its action based on the relative orientation of up to seven neighbors. To account for the fact that the perception view of individuals is limited, we exclude individuals located at a distance greater than 100 m (see Figure 1). We will refer to the individuals selected with the term *topological* neighbors* to clarify that they exclude individuals that are outside the view range and to stress that the way in which we select them differs slightly from the way in which they are identified in Ballerini et al. (2008) and Camperi et al. (2012).

In standard models, each agent i steers in each step to align with the average direction of topological* neighbors and eventually steers away from the predator based on the following equations:

$$\varphi = \theta_i - (\rho - \pi), \ \alpha_j = W_j \left(\theta_i - \theta_{j \in S_j}\right) \tag{1}$$

$$\theta_i = W \cdot \varphi + (1.0 - W) \cdot \sum_j \alpha_j + \xi \tag{2}$$

where θ_i is the steering angle in radians, j are the topological* neighbors, α is the relative orientation of the neighbors with respect to the individual in radians, ρ is the angle between the agent

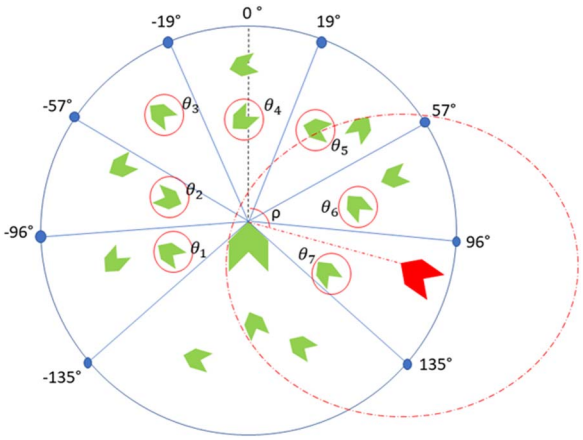

Figure 1. Schematic of the perceptual information of the individual located at the center of the figure. The individuals and the predator are shown in green and red, respectively. The blue circle represents the maximum perceptual range. The blue lines represent the limits of the seven visual sectors spanning over 270°. The red circles indicate topological* neighbors, that is, the closest individuals located in each visual sector. The red line represents the relative angle of the predator. The dashed red circle represents the hunting area of the predator.

and the predator, and φ is the opposite of the relative angle of the predator in radians. Finally, W is the relative weight that the orientation of topological* neighbors and the relative direction of the predator have on the steering response, and $\xi \in [-10^\circ, 10^\circ]$ is a random value extracted with a uniform distribution that accounts for noise.

The predator moves by steering toward the nearest individual based on the following equations:

$$\rho = a\tan 2(y, x)$$

$$y = y_{\text{prey}} - y_{\text{predator}} \text{ and } x = x_{\text{prey}} - x_{\text{predator}} \tag{3}$$

Because the probability of being captured by a predator generally correlates with the distance from the predator and the isolation of the individual (Olson et al., 2016), the probability of individuals being captured in each time step is modeled based on the following equation:

$$P = \left(1 - \frac{d}{r}\right) - \frac{1}{n}\sum_{j=1}^{n} e_j \tag{4}$$

where d is distance of individual from predator, r is the hunting range of the predator, n is the number of visual sectors, and e_j is a binary value that is 1 when the sector contains at least one neighbor, and otherwise is 0.

The *standard model* can be implemented in a linear two-layer neural network with eight sensory neurons and one motor neuron (Figure 2), in which (a) the sensory neurons encode the relative orientation of the seven topological* neighbors and the opposite direction of the predator, normalized in the range [−1.0, 1.0]; (b) the motor neuron encodes the steering angle of the individual normalized in the range $[-\pi/2, \pi/2]$; (c) the connection weights are set to [1/7, 1/7, 1/7, 1/7, 1/7, 1/7, 1/7, 1]; and (d) the activation of the motor neuron is updated based on a linear function. Notice that using these connection weights implies that W corresponds to 0.5 and that the contribution of the information detected by different visual sectors is equal.

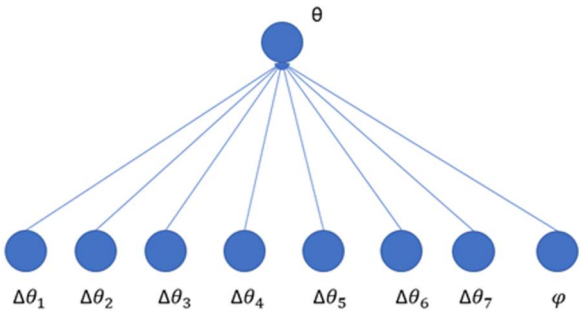

Figure 2. Neural architecture of the standard and linear adapted model. The bottom disks represent the sensory neurons that encode the relative orientations of the seven corresponding topological* neighbors and the inverse of the relative orientation of the predator. The top disk represents the motor neuron encoding the steering angle of the agent.

The performance of this model and following models is evaluated by calculating the average number of individuals captured by the predator during 16 evaluation episodes lasting 500 evaluation steps.

In a second series of experiments, we used a *linear adapted model* that has the same architecture shown in Figure 2 but in which the connection weights have been adapted through the OpenAI-ES (Salimans et al., 2017) evolutionary algorithm. This was realized by (a) setting the initial value of parameters to [1/7, 1/7, 1/7, 1/7, 1/7, 1/7, 1/7, 1], (b) using the hyperparameters described in the appendix, and (c) continuing the evolutionary process until a total of 100 million steps were performed. The initial population of evolving individuals was initialized with connection weights used for the standard model. This model permits discovering solutions in which the absolute and relative weight of information encoded in different sensory neurons differs from the standard model. Moreover, it permits discovering solutions in which the relative weight of orientation of neighbors and the relative position of the predator differs from the standard model.

In a third series of experiments, we used a *nonlinear adapted model* in which the neural network controller of the individuals includes an internal layer with 32 neurons whose activation is updated based on the tanh function (Figure 3). These additional neurons and associated additional connection

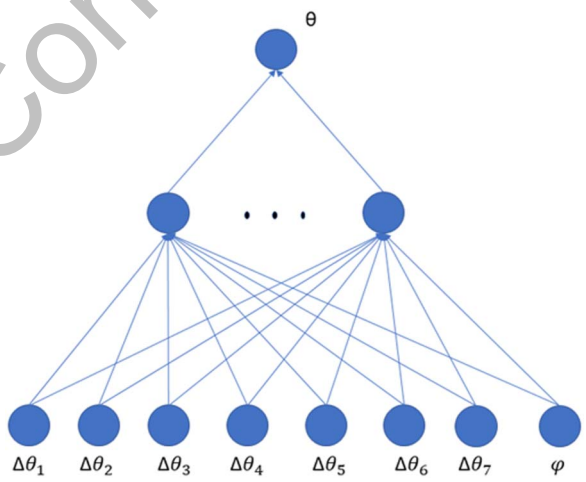

Figure 3. Neural architecture of the nonlinear adapted model. The bottom disks represent the sensory neurons that encode the relative orientations of the seven corresponding topological* neighbors and the inverse of the relative orientation of the predator. The intermediate disks represent the internal neurons. The top disk represents the motor neuron encoding the steering angle of the agent.

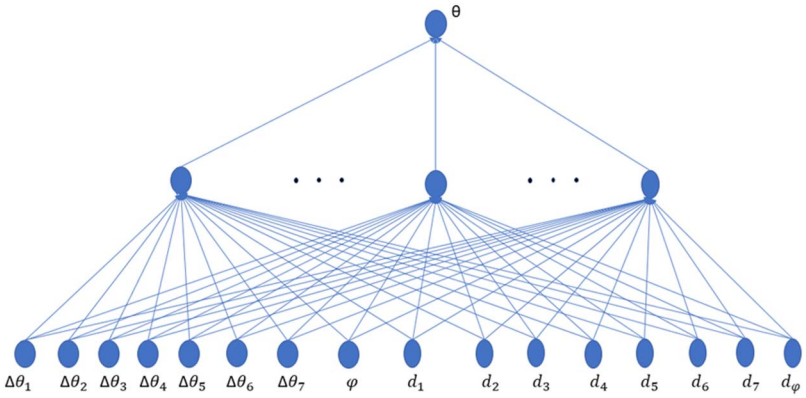

Figure 4. Architecture of the nonlinear distance adapted model. The bottom disks represent the sensory neurons that encode the relative orientations of the seven corresponding topological* neighbors, the inverse of the relative orientation of the predator, and the distances of neighbors and predator. The intermediate disks represent internal neurons. The top disk represents the motor neuron encoding the steering angle of the agent.

weights enable evolving individuals to use nonlinear control rules and eventually determine steering strength based on internal states that combine the activation of multiple sensory neurons. The connection weights of the initial evolving population are initialized randomly.

Finally, in a fourth series of experiments, we used a *nonlinear distance adapted model* in which the neural network controller of individuals includes eight additional sensory neurons encoding the distance of the seven topological* neighbors and of the predator (Figure 4). These additional sensors enable evolving swarms to discover solutions in which the steering response to relative orientations of topological* neighbors and the relative position of the predator is modulated by their proximity. Connection weights of the initial evolving population are initialized randomly.

For some models, we also replicated experiments by varying the size of swarms and consequently the average density of individuals.

The experiments were performed using the EvoJAX tool (Tang et al., 2022). They can be replicated by using the source code available from the GitHub repository indicated in the appendix.

4 Results

The results of the experiments indicate that adapted solutions outperform the standard model (Wilcoxon nonparametric test p-value < 0.001; Figure 5) and that three adaptive models produce equally good performance (Wilcoxon nonparametric test p-value > 0.05; Figure 5). All models converged after 400 generations (Figure 6), as demonstrated by the fact that the performance at Generation 400 and Generation 500 does not differ in all cases (Wilcoxon nonparametric p-value > 0.05). The initial performance and the converge speed are much higher in the linear condition in which the parameters are initially set on the basis of the standard models. In the other experimental conditions, the parameters are initialized randomly. Because the only difference between the adaptive linear model and the standard model concerns the value of parameters, the fact that the former outperforms the latter implies that the parameters of the standard model are suboptimal. We analyze how adapted parameters differ in what follows.

The fact that the performance of linear and nonlinear adapted models does not differ implies that using nonlinear control rules and determining steering actions based on internal states that combine activation of multiple sensory neurons does not provide an advantage.

Finally, the fact that the performance of nonlinear and nonlinear distance adapted models does not differ implies that modulating steering action based on the distance of topological* neighbors and predator does not provide an advantage. These conclusions are also supported by additional analysis reported later.

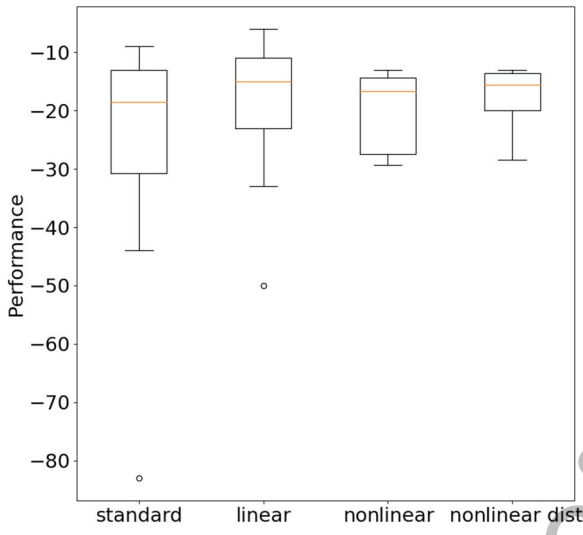

Figure 5. Performance obtained with the standard model and with linear, nonlinear, and nonlinear distance adapted models. Performance represents fitness obtained in 50 evaluation episodes. In the case of adapted models, 50 episodes are realized by performing 5 episodes with each of 10 solutions obtained in 10 corresponding replications of experiments. Boxes represent interquartile ranges of data, and horizontal lines inside boxes mark median values. Whiskers extend to most extreme data points within 1.5 times interquartile range from the box. Circles indicate outliers.

Visual inspection of the behavior of the evolved swarm shows that the agents aggregate in a few large flocks that fly together in a coordinated manner and move away from the predator when approached, in all experimental conditions (see the video described in the appendix). Formed flocks often split into separate subflocks during predator-avoidance behavior. The difference in performance between the agents obtained in the different experimental conditions is due to small differences in the aggregation rate and in the arbitration between the flocking and predator-avoidance behavior.

To analyze how solutions found by the linear adaptive model differ from those found by the standard model, we plotted parameters that determine the contribution of topological* neighbors

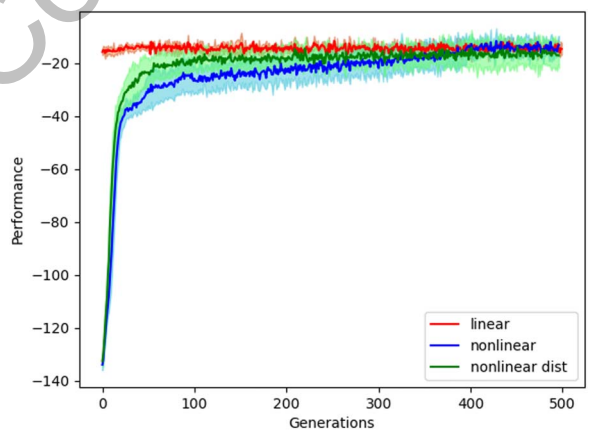

Figure 6. Performance of the best swarm of each generation across 500 generations. The solid lines show the data averaged over 10 replications for each condition for the linear, nonlinear, and nonlinear distance models for the best individual of each generation. The shaded areas represent 90% confidence intervals.

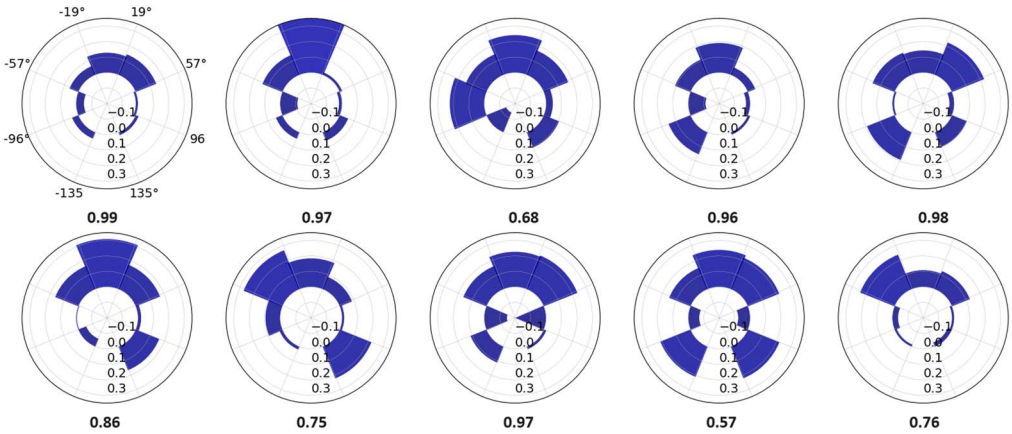

Figure 7. Plots of the parameters that determine the contribution of each of the seven topological* neighbors located in seven corresponding visual sectors to the determination of the steering action. The 10 plots show the value of the parameters of the 10 evolved solutions obtained with the linear adaptive model. The number in bold under each plot indicates the contribution of the predator.

located in different sectors of the visual field to the steering response of individuals. More specifically, we plotted the value of the parameters of 10 evolved solutions obtained in 10 corresponding replications of the experiment (Figure 7) and the average values of the parameters of 10 solutions (Figure 8). As can be seen, in evolved solutions, steering is determined primarily by one to three topological* neighbors. Indeed, only the parameters of a few sectors, generally located in the frontal side, have relatively high values.

To analyze the efficacy of solutions considering a single topological* neighbor and the relative utility of neighbors located in different visual sectors, we conducted seven evaluation tests. In each test, individuals reacted to only one of the seven topological* neighbors. We created a modified version of the control networks evolved with the linear adaptive model in which six of the seven

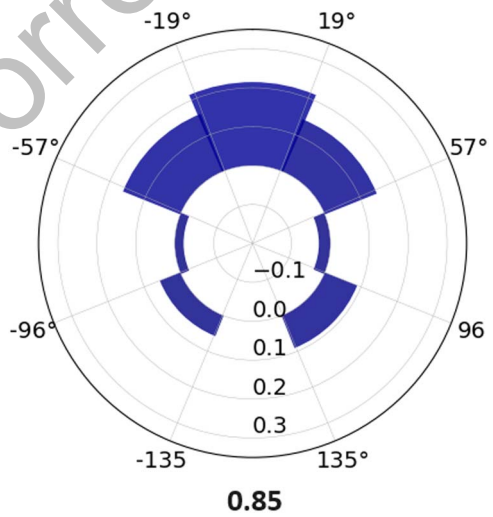

Figure 8. Plot of the parameters that determine the contribution of each of the seven topological* neighbors located in seven corresponding visual sectors to the determination of the steering action. The data are averaged over 10 solutions obtained through 10 corresponding replications of the experiment. The number in bold at the bottom indicates the contribution of the predator.

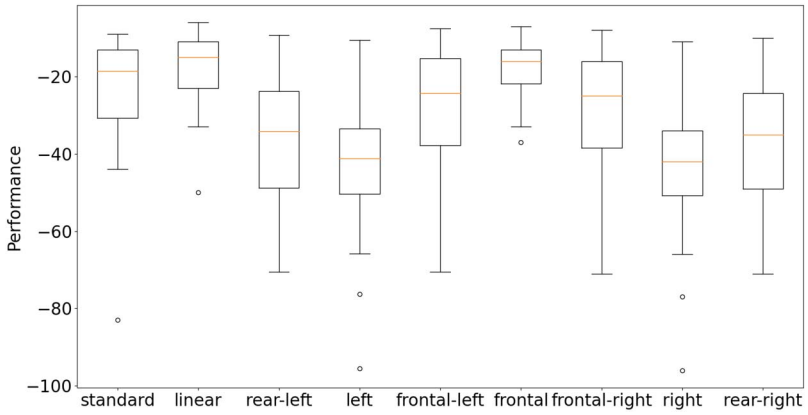

Figure 9. Performance of the standard model, linear adapted models, and seven variations of the linear adapted model in which individuals react to a single topological* neighbor located in the rear-left, left, frontal-left, frontal, frontal-right, right, or rear-right sector of the visual field. Boxes represent interquartile ranges of the data, and horizontal lines inside boxes mark median values. Whiskers extend to the most extreme data points within 1.5 times interquartile range from the box. Circles indicate outliers.

parameters determining the reaction to each topological* neighbor were set to 0.0 and the remaining parameter was set to the sum of the original values of the seven parameters. As shown in Figure 9, swarms relying on the frontal neighbor outperformed swarms relying on frontal-left or frontal-right neighbors only (Wilcoxon nonparametric test p-value < 0.001), which outperformed swarms relying on rear-left or rear-right neighbors only (Wilcoxon nonparametric test p-value < 0.001), which outperformed swarms relying on left or right neighbors only (Wilcoxon nonparametric test p-value < 0.001). The performance of swarms that rely on the frontal neighbor is equally good as the performance of solutions evolved with the linear model (Wilcoxon nonparametric test p-value > 0.05) and significantly better than the performance of the standard model (Wilcoxon nonparametric test p-value < 0.001).

These results demonstrate that, for agents that have a limited view range, the utility of information provided by topological* neighbors varies. It is maximum in the case of the frontal neighbor and decreases for neighbors located toward the lateral side of the visual field. Moreover, these results show that the strategy of the standard model is suboptimal and is outperformed by a simpler strategy that relies on the frontal neighbor only.

In a previous work, Gauci et al. (2014) demonstrated that perceiving only the frontal neighbor is sufficient to produce aggregation behavior. Here we demonstrated that the information provided by a frontal neighbor alone is sufficient to achieve optimal performance, that is, that considering also the information provided by the other neighbors does not yield better performance.

Figure 10 shows the topological and metrical aggregation exhibited by swarms. The former is calculated by counting the number of sectors including at least one conspecific located within the perceptual range, averaged over individuals and evaluation steps. The latter is calculated by counting the number of conspecifics located within the perceptual range. As expected, the condition leading to higher performance also leads to higher aggregation levels with respect to both measures.

To further check the utility of distance information in experiments performed with nonlinear distance adapted models, we postevaluated the evolved swarm in a control condition in which the state of distance sensors is always set to 0.0. The fact that performance does not statistically differ in the two conditions (Wilcoxon nonparametric test p-value > 0.05; Figure 11) demonstrates that evolved solutions neglect distance information, that is, that information provided by distance sensors is unnecessary.

Finally, we investigated the role of swarm size in the case of the standard model and linear adapted model, that is, simpler models that achieved better performance. More specifically, we

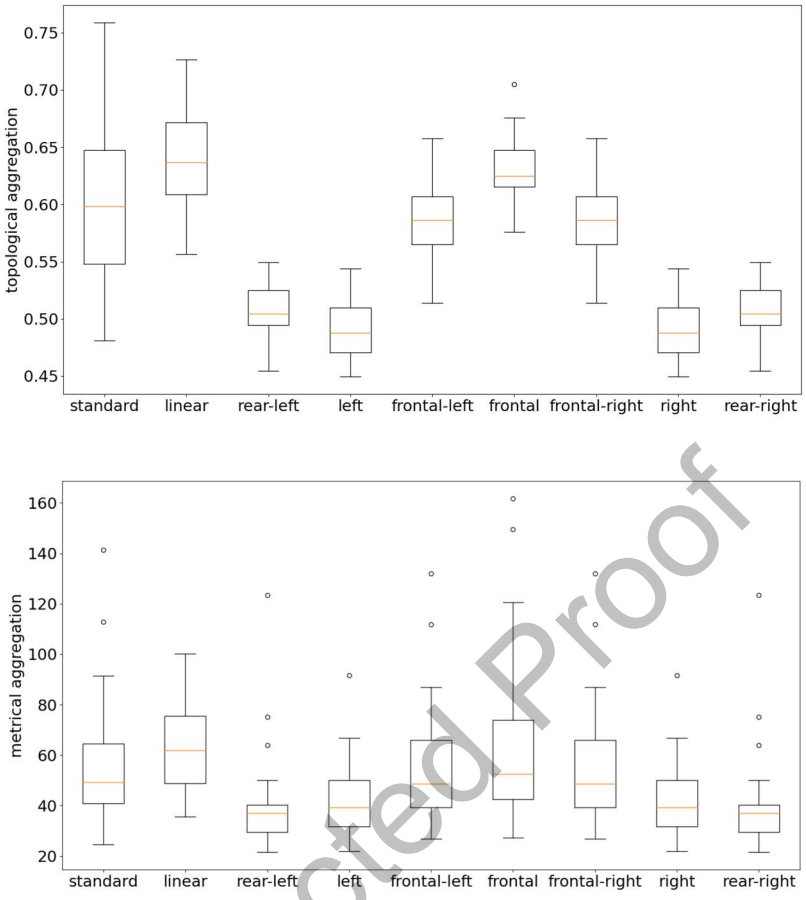

Figure 10. Average aggregation displayed by swarms obtained with the standard model, with the linear adapted models, and with seven variations of the linear adapted model in which individuals react to a single topological* neighbor: results obtained by using (top) topological and (bottom) metrical measures of aggregation.

replicated experiments in four conditions in which swarm size was set to 64, 128, 256, and 512 individuals (Figure 12). The size of the swarm in the experiments reported earlier is 512. Flocking is observed in all cases (see videos described in the appendix). As expected, the ability to escape the predator increases in smaller swarms. In addition, the advantage of the linear adapted model with respect to the standard model decreases in smaller swarms. The latter effect is explained by the fact that the adaptive advantage of improved aggregation ability achieved by the linear adaptive model decreases when the number of the agents diminishes.

5 Discussion

Complex collective behavior can originate from the interaction of multiple individuals exhibiting simple behavior. This is possible thanks to the indirect relationship between the rules regulating the interactions of individuals with their physical and social environments and the behavior of the group. Such an indirect relationship implies that anticipating how the characteristics of the behavior exhibited by single individuals affect the resulting collective behavior is far from trivial. In this article, we explored such a relationship in the case of flocking behaviors. More specifically, we

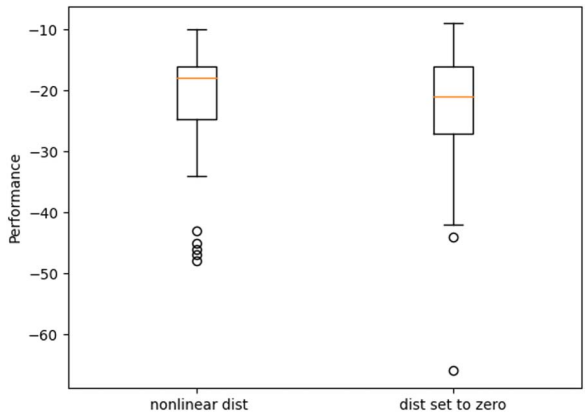

Figure 11. Performance of swarms obtained with the nonlinear distance adapted model in a standard condition and in a control condition in which the state of distance sensors is always set to 0.0. The data were obtained by postevaluating the 10 evolved swarms obtained in the 10 corresponding replications for 5 episodes for a total of 50 evaluation episodes. Boxes represent the interquartile ranges of data, and horizontal lines inside boxes mark the median values. Whiskers extend to the most extreme data points within 1.5 times the interquartile range from the box.

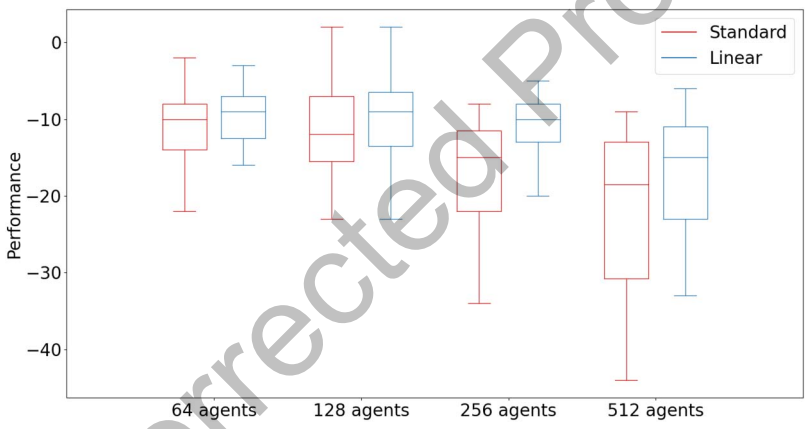

Figure 12. Performance of swarms obtained with standard and linear adaptive models in the case of swarms composed of 512, 256, 128, or 64 individuals. Performance corresponds to fitness obtained in 50 evaluation episodes. In the case of adapted models, 50 episodes are realized by performing 5 episodes with each of the 10 solutions obtained by evolving 10 replications of the experiment. Boxes represent the interquartile ranges of data, and horizontal lines inside boxes mark the median values. Whiskers extend to the most extreme data points within 1.5 times interquartile range from the box.

analyzed through simulation studies and through artificial evolutionary experiments whether the control rules used in standard models are optimal.

In standard models (Ballerini et al., 2008; Camperi et al., 2012; Olson et al., 2016; Vicsek et al., 1995), interactions of individuals with their environment is regulated by uniform, linear, and independent control rules that make individuals steer toward the relative orientation of topological* neighbors distributed over a view field of ∼270°. The term *uniform* refers to the fact that the weight that each topological* neighbor has on the steering of the agent is the same. The term *independent* refers to the fact that the offset of each neighbor contributes to steering action independently from the directional offset of other neighbors. Moreover, only observational information that affects steering is directionally offset. This implies, for example, that the relative distance between individuals is neglected.

To verify how the characteristics of the interaction rules impact the resulting collective behavior, we carried out a series of experiments in which populations of agents evolved for the ability to avoid predation are allowed to use more elaborate interaction rules in which the limits described previously are released. The obtained results indicate that using linear and independent rules and neglecting distance information does not limit the efficacy of flocking behavior, measured as the ability to aggregate and to reduce the chances of being predated. Instead, using uniform control rules limits swarms' efficacy.

The analysis of solutions evolved by enabling the use of nonuniform rules surprisingly revealed that this is due to the fact that the utility of information provided by the relative orientation of neighbors is maximum for individuals located in front and smaller for individuals located on the sides. Consequently, strategies that consider frontal neighbors only, and eventually the single neighbor located in the frontal sector only, outperform the standard model. In other words, swarms formed by individuals who use less information and operate based on simpler control rules outperform swarms formed by individuals who use more information and operate based on more complex rules.

This counterintuitive result can be explained by considering that swarms composed of many tightly interacting individuals can operate based on information generated by combining the observations and the effects of the actions of multiple individuals. In the context of the flocking behavior considered, this implies that individuals can manage to orient themselves toward the direction of their topological* neighbors even if they detect a single neighbor. This is thanks to the fact that their behavior is influenced indirectly by the behavior of other individuals detecting other neighbors. The possibility to steer toward a subset of neighbors or toward the frontal neighbor only at the level of individuals permits them to focus on the most useful information.

Our result is also in line with studies that investigated the impact of communication range in swarms of cooperating individuals. Indeed, the results of these studies indicate that using short communication ranges achieves better performance than using longer communication ranges (Pirrone et al., 2022; Talamali et al., 2021).

Funding Information

This research was supported by the PNRR MUR project (grant PE0000013-FAIR).

References

Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, E., Giardina, I., Lecomte, V., Orlandi, A., Parisi, G., Procaccini, A., Viale, M., & 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 of the United States of America*, *105*(4), 1232–1237. https://doi.org/10.1073/pnas.0711437105, PubMed: 18227508

Biswas, R., Ofria, C., Bryson, D., & Wagner, A. (2014). Causes vs benefits in the evolution of prey grouping. In H. Sayama, J. Rieffel, S. Risi, R. Doursat, & H. Lipson (Eds.), *Artificial Life 14: Proceedings of the fourteenth international conference on the synthesis and simulation of living systems* (Vol. 14, pp. 641–648). MIT Press. https://doi.org/10.1162/978-0-262-32621-6-ch103

Camperi, M., Cavagna, A., Giardina, I., Parisi, G., & Silvestri, E. (2012). Spatially balanced topological interaction grants optimal cohesion in flocking models. *Interface Focus*, *2*(6), 715–725. https://doi.org/10.1098/rsfs.2012.0026, PubMed: 24312725

Couzin, I. D., Krause, J., James, R., Ruxton, G. D., & Franks, N. R. (2002). Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, *218*(1), 1–11. https://doi.org/10.1006/jtbi.2002.3065, PubMed: 12297066

Demšar, J., Hemelrijk, C. K., Hildenbrandt, H., & Bajec, I. L. (2015). Simulating predator attacks on schools: Evolving composite tactics. *Ecological Modelling*, *304*, 22–33. https://doi.org/10.1016/j.ecolmodel.2015.02.018

Demšar, J., Štrumbelj, E., & Lebar Bajec, I. (2016). A balanced mixture of antagonistic pressures promotes the evolution of parallel movement. *Scientific Reports*, *6*(1), 1–12. https://doi.org/10.1038/srep39428, PubMed: 27995967

Gauci, M., Chen, J., Li, W., Dodd, T. J., & Groß, R. (2014). Self-organized aggregation without computation. *International Journal of Robotics Research*, *33*(8), 1145–1161. https://doi.org/10.1177/0278364914525244

Ginelli, F., & Chaté, H. (2010). Relevance of metric-free interactions in flocking phenomena. *Physical Review Letters*, *105*, 168103. https://doi.org/10.1103/PhysRevLett.105.168103, PubMed: 21231019

Gregoire, G.,Chaté, H., & Tu, Y. (2003). Moving and staying together without a leader. *Physica D*, *181*(3–4), 157–170. https://doi.org/10.1016/S0167-2789(03)00102-7

Inada, Y., & Kawachi, K. (2002). Order and flexibility in the motion of fish schools. *Journal of Theoretical Biology*, *214*(3), 371–387. https://doi.org/10.1006/jtbi.2001.2449, PubMed: 11846596

Kunz, H., Züblin, T., & Hemelrijk, C. K. (2006). On prey grouping and predator confusion in artificial fish schools. In L. M. Rocha (Ed.), *Artificial Life X: Proceedings of the tenth international conference on the simulation and synthesis of living systems*. MIT Press.

Milano, N., Carvalho, J. T., & Nolfi, S. (2019). Moderate environmental variation across generations promotes the evolution of robust solutions. *Artificial Life*, *24*(4), 277–295. https://doi.org/10.1162/artl_a_00274, PubMed: 30681913

Milano, N., & Nolfi, S. (2021). Automated curriculum learning for embodied agents: A neuroevolutionary approach. *Scientific Reports*, *11*(1), 8985. https://doi.org/10.1038/s41598-021-88464-5, PubMed: 33903698

Morrell, L. J., Greenwood, L., & Ruxton, G. D. (2015). Consequences of variation in predator attack for the evolution of the selfish herd. *Evolutionary Ecology*, *29*, 107–121. https://doi.org/10.1007/s10682-014-9743-6

Olson, R. S., Haley, P. B., Dyer, F. C., & Adami, C. (2015). Exploring the evolution of a trade-off between vigilance and foraging in group-living organisms. *Royal Society Open Science*, *2*(9), 150135. https://doi.org/10.1098/rsos.150135, PubMed: 26473039

Olson, R. S., Hintze, A., Dyer, F. C., Knoester, D. B., & Adami, C. (2013). Predator confusion is sufficient to evolve swarming behaviour. *Journal of the Royal Society Interface*, *10*(85), 20130305. https://doi.org/10.1098/rsif.2013.0305, PubMed: 23740485

Olson, R. S., Knoester, D. B., & Adami, C. (2016). Evolution of swarming behavior is shaped by how predators attack. *Artificial Life*, *22*(3), 299–318. https://doi.org/10.1162/ARTL_a_00206, PubMed: 27139941

Pirrone, A., Reina, A., Stafford, T., Marshall, J. A., & Gobet, F. (2022). Magnitude-sensitivity: Rethinking decision-making. *Trends in Cognitive Sciences*, *26*(1), 66–80. https://doi.org/10.1016/j.tics.2021.10.006, PubMed: 34750080

Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. In M. C. Stone (Ed.), *Proceedings of the 14th annual conference on Computer graphics and interactive techniques* (pp. 25–34). ACM. https://doi.org/10.1145/37401.37406, PubMed: 10284632

Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., & Chen, X. (2016). Improved techniques for training GANs. In D. D. Lee, U. von Luxburg, R. Garnett, M. Sugiyama, & I. Guyon (Eds.), *Advances in neural information processing systems* (Vol. 29, pp. 2226–2234). Curran Associates.

Salimans, T., Ho, J., Chen, X., Sidor, S., & Sutskever, I., (2017). *Evolution strategies as a scalable alternative to reinforcement learning*. ArXiv. https://doi.org/10.48550/arXiv.1703.03864

Talamali, M. S., Saha, A., Marshall, J. A., & Reina, A. (2021). When less is more: Robot swarms adapt better to changes with constrained communication. *Science Robotics*, *6*(56), eabf1416. https://doi.org/10.1126/scirobotics.abf1416, PubMed: 34321345

Tang, Y., Tian, Y., & Ha, D. (2022). EvoJAX: Hardware-accelerated neuroevolution. In *GECCO'22 companion: Proceedings of the 2022 Genetic and evolutionary computation conference* (pp. 308–311). ACM. https://doi.org/10.1145/3520304.3528770

Tegeder, R. W., & Krause, J. (1995). Density dependence and numerosity in fright stimulated aggregation behaviour of shoaling fish. *Philosophical Transactions of the Royal Society of London, Series B*, *350*(1334), 381–390. https://doi.org/10.1098/rstb.1995.0172

Tosh, C. R. (2011). Which conditions promote negative density dependent selection on prey aggregations? *Journal of Theoretical Biology*, *281*(1), 24–30. https://doi.org/10.1016/j.jtbi.2011.04.014, PubMed: 21540037

Van Diggelen, F., Luo, J., Karagüzel, T. A., Cambier, N., Ferrante, E., & Eiben, A. E. (2022). Environment induced emergence of collective behavior in evolving swarms with limited sensing. In *Proceedings of the 2022 Genetic and evolutionary computation conference* (pp. 31–39). ACM. https://doi.org/10.1145/3512290.3528735

Vicsek, T., Czirok, A., Ben-Jacob, E., Cohen, I., & Shochet, O. (1995). Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, *75*, 1226–1229. https://doi.org/10.1103/PhysRevLett.75.1226, PubMed: 10060237

Wood, A. J., & Ackland, G. J. (2007). Evolving the selfish herd: Emergence of distinct aggregating strategies in an individual-based model. *Proceedings of the Royal Society B*, *274*(1618), 1637–1642. https://doi.org/10.1098/rspb.2007.0306, PubMed: 17472913

Appendix

The hyperparameters of the evolutionary algorithm were set as follows: The distribution of the perturbations of parameters was set to 0.02, the step size of the Adam optimizer was set to 0.01, the size of the population was set to 40, the observation vectors were normalized by using virtual batch normalization (Milano et al., 2019; Milano & Nolfi, 2021; Salimans et al., 2016, 2017), and the connection weights were normalized by using weight decay.

The swarms evolved were homogeneous; that is, the parameters encoded the connection weights of a single neural controller that was then duplicated 512 times to create the 512 corresponding neural network controllers that were embedded in the individuals forming the swarm.

The evolving swarms were evaluated for 16 episodes. Each episode lasted 500 steps. The evolutionary process was continued until a total number of steps was performed. The total number of evaluation steps was set to 50, 100, and 100 million in the case of the linear, nonlinear, and nonlinear distance adaptive model, respectively.

The code to replicate the experiments can be found at https://github.com/milnico/InteractionRules. Videos of sample behavior of swarms evolved under different experimental conditions can be found at https://www.youtube.com/playlist?list=PLcNe-IgRAL8X-fTH4c4NcPlev4GRf1iXm.