Modelling zebrafish collective behaviours with multilayer perceptrons optimised by evolutionary algorithms

Leo Cazenille, Nicolas Bredeche, José Halloy

Modelling zebrafish collective behaviours with multilayer perceptrons optimised by evolutionary algorithms

Leo Cazenille[1,2]**, Nicolas Bredeche**[2]**, and José Halloy**[1,*]

[1]Univ Paris Diderot, Sorbonne Paris Cité, LIED, UMR 8236, 75013, Paris, France
[2]Sorbonne Université, CNRS, ISIR, F-75005 Paris, France
[*]jose.halloy@univ-paris-diderot.fr

ABSTRACT

Collective movements are pervasive behaviours among social organisms and have led to the development of many models. However, modelling animal trajectories and social interactions in simple bounded environments remains a challenge. Moreover, advances in the understanding of the sensory-motor loop and the information processing by animals are leading to revisions of the traditional assumptions made in decision-making algorithms. In this context, we develop a methodology based on artificial neural networks (ANN) to describe the collective motion of small zebrafish groups in a bounded environment. Although ANN models are commonly used in artificial systems they are still under-explored to model animal collective behaviours. Here, we present a methodology to calibrate Multilayer Perceptrons by learning from real fish experimental data. The ANNs are trained using either supervised learning or various forms of evolutionary reinforcement learning methods (using the CMA-ES and NSGA-III algorithms). We reveal that ANN models trained using evolutionary methods are capable of generating realistic collective motions for groups of 5 zebrafish including the tank wall effects, a feature that is lacking in previous models. Finally, we also discuss the benefits of optimised ANNs as candidates for driving robotic lure with biologically realistic behaviour, a method that is becoming increasingly popular to gather data and validate assumptions on collective behaviours.

Introduction

Collective motion is one of the most ubiquitous collective behaviour displayed by fish and can lead to collective decision-making[1]. An interest to study the link between individual behaviours and collective patterns has arisen out of these numerous observations and led to the development of models simulating agents performing collective movement and behaviours. A large variety of models exist to describe collective motion in biological systems[2].

Modelling the collective motion of fish groups remains a challenging task, in particular for generating accurate collective trajectories[3]. The first models were computer simulations describing the motion of a collection of individuals taking into account a collision avoidance and speed alignment mechanisms[4]. In this kind of simulations, the time step updating can be synchronous[5,6] or asynchronous[7,8]. Recent updated version of this kind of modelling has also been used[9,10]. In physics, models of self-propelled particles (SPP) have been developed to analyse the type of transitions between different group patterns from a phase transition point of view[11–15]. Inspired by mathematical models of random walk, kinematic models modelling fish trajectories with stochastic differential equations have also being developed to describe animal collective motions[16–19]. Most of these models consider animals as simple particles submitted to some kind of "social" forces that allow the emergence of collective motion patterns.

Probabilistic models based on a mechanism that takes into account visual perception and motion direction have been proposed[20–22]. The model considers the three dimensional visual sensory system of fish that adjust their trajectory according to their perception field. It introduces a stochastic process based on a probability distribution function to move in targeted directions rather than on a summation of influential vectors as it is assumed by previous models.

If the considered behaviours are more complex than aligned collective motion in unbounded environments more elaborate models are needed. Modelling realistic bounded and structured environments requires multilevel models that take into account perception of the environment and the other individuals and that allow decision-making[23–25]. It is also an important issue to calibrate multilevel and spatially dependent social behaviour model because it involves trade-offs between social tendencies (*e.g.* aggregation, group splitting), and response to the environment (*e.g.*

wall-following, zone occupation). We have previously shown that an optimisation method based on multi-objective evolutionary computation gives good results for complex models of collective behaviour[26–28].

Objectives

Zebrafish is a classic animal model in genetics and neurosciences. It has been used extensively to study individual and collective behaviours, through experimental observations and modelling. Our goal is to automatically build models able to generate zebrafish trajectories that capture social interactions of individual fish within a small group of 5 fish placed in a bounded environment (i.e. with walls). We aim at generating models for autonomous fish agents that perform realistic collective behaviours for lasting duration of at least 30 minutes. This work participates to another general objective, namely to design robotic fish that can socially interact with live fish in realistic environments[29, 30].

Contributions

We consider groups of 5 agents including n simulated autonomous agents driven by the MLP model and $(5 - n)$ "real" agents with trajectories replayed from fish trajectories acquired from experimental observations. Varying the number n of simulated agents makes it possible to tune the level of difficulty. A value of $n = 1$ (easy setup) implies that the simulated agent should blend into a group performing realistic trajectories while a value of $n = 5$ (difficult setup) requires that realistic trajectories are endogenous to the simulated agents.

Our method is summarised in Fig. 1. We rely on a multi-agents approach where each simulated agent is controlled by a Multilayer Perceptrons (MLP)[31] using high-level sensory inputs and actuators that are similar to what is available to zebrafish (see Table 1). MLP weights are trained using either supervised learning or stochastic optimisation techniques so as to generate relevant trajectories w.r.t. the expected behaviours. To do so, we define an objective function to maximise, which quantifies simulated agent individual trajectories (e.g. linear and angular speeds), group dynamics (e.g. inter-individual distances, polarisation) and relation with the environment (e.g. distance to wall). By using experimental data as a reference, the performance of MLP-driven simulated agents can be assessed.

Our results show that it is possible to design and automatically calibrate neural network models to act as simulated agent controllers. In particular, we reveal that evolutionary optimisation methods, namely CMA-ES[32] and NSGA-III[33], provide a great improvement over more classic supervised learning methods. This is due to the ability of evolutionary computation to act as a policy-search reinforcement learning method[34], where sequential decisions are considered. This is very different than with supervised learning, where the goal is to devise effector outputs based solely on immediate sensory information without any regards from the sequence of actions (past or future).
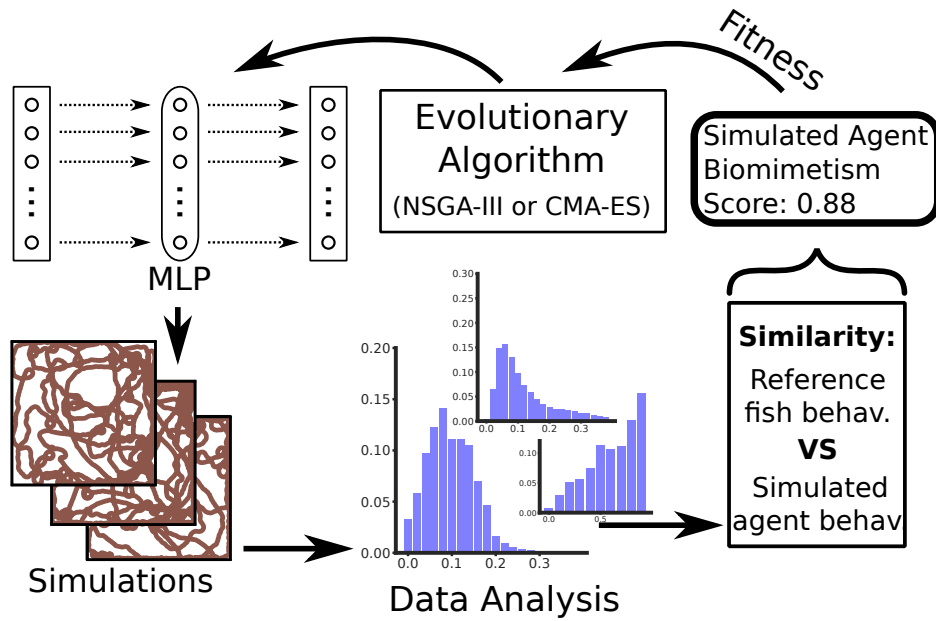
Methods

Animal handling and experimental set-up

Our experiments involved 10 groups of 5 adults (6-12 months old) wild-type AB zebrafish (*Danio rerio*) in ten 30-minutes trials as in[28, 30, 35–37]. Rearing and handling of the fish are described in the supplementary file. We use the experimental set-up that consists of a walled arena of $1 \times 1 \times 0.1$ m described in [28, 30, 35–37] and detailed in the supplementary.

Artificial neural network model

Black-box models, like artificial neural networks (ANN), can model phenomena with few *a priori* information. To our knowledge, they were only used to model fish collective behaviour in preliminary studies[36, 38]. We extend the methodology from[36] based on experimental data, and we show that it is relevant to model zebrafish individual and collective behaviours. This methodology (Fig. 1) makes use of Multilayer Perceptron (MLP)[31] artificial neural networks that are calibrated through the use of evolutionary algorithms to model the behaviour of simulated fish-like agents in a group of 5 individuals. Three cases with increasing difficulty are considered with groups composed of respectively 1,3,5 simulated agents driven by the calibrated MLP models and 4,2,0 agents following the experimental trajectories of experimental fish.

MLP are a type of feedforward artificial neural networks. They are very popular in artificial intelligence to solve a large variety of real-world problems[39]. They have the capability to universally approximate functions[40] which

Figure 1. Workflow of the methodology to optimise artificial neural networks to model fish behaviour. An evolutionary algorithm is used to evolve the weight of Multilayer Perceptron (MLP, 1 hidden layer, 100 neurons) that serve as behavioural models of fish-like agents in multi-agents simulations. The fitness function is computed through data-analysis of these simulations and represent the biomimetism metric that quantifies the realism of the simulated fish behaviour compared to the behaviour exhibited by the experimental fish. Two evolutionary algorithms are tested: CMA-ES (mono-objective) and NSGA-III (multi-objective).

makes them suitable to model control problems[39]. We consider MLP with only one hidden layer of 100 neurons (using a hyperbolic tangent function as activation function) as in[36].

For the considered focal agent, the neural network model takes the 20 parameters of Table 1 as input. This set of inputs is typically used in multi-agent modelling of animal collective behaviour[2,41]. We consider that it is sufficient to model fish behaviour with neural networks.

The neural network has 2 outputs (Table 1) corresponding to the change in linear and angular speeds to apply from the current time-step to the next time-step. Our approach models fish trajectories resulting from social interactions in a homogeneous environment bounded by walls. Very few models of fish collective behaviours take into account the presence of walls[22,23,30,36].

Data analysis

Zebrafish behaviour is multi-level (*i.e.* with different dynamics at the level of individual trajectories *vs.* at the level of group behaviour) and multi-modal (*i.e.* with several kind of exhibited behaviours)[30,35,42]. In a bounded environment, zebrafish group exhibit a trade-off between aggregation and wall-following behaviours. Groups tend to be short-lived with individuals continuously leaving and joining the groups[30].

We quantify zebrafish behaviour by analysing each trial e (experiments or simulations) and use the tracked positions of agents to compute the following behavioural metrics:

at the level of individual trajectories the distributions of *instant linear speeds* (L_e) and the distributions of *instant angular speeds* (A_e). These metrics assess the biomimetism of agent individual trajectories.

at the level of group dynamics the distribution of *inter-individual distances* between agents (D_e) and the distribution of *polarisation* of the agents in the group (P_e). The inter-individual distance evaluates the aggregative behaviour of the fish groups. The polarisation of an agent group measures how aligned the agents in a group are, and is

Inputs

Name	# of param.	Description
Linear speed	1	Instant linear speed of the focal agent at the previous time-step
Angular speed	1	Instant angular speed of the focal agent at the previous time-step
Distance towards agents	4	Linear distance from the focal agent towards each other agent
Angle towards agents	4	Angular distance from the focal agent towards each other agent
Alignment (angle)	4	Angular distance between the focal agent heading and each other agent heading
Alignment (linear speed)	4	Difference of linear speed between the focal agent and each other agent linear speed
Distance to nearest wall	1	Linear distance from the focal agent towards the nearest wall
Angle towards nearest wall	1	Angular distance from the focal agent towards the nearest wall

Outputs

Name	# of param.	Description
Delta linear speed	1	Change of instant linear speed of the focal agent from the previous time-step
Delta angular speed	1	Change of instant angular speed of the focal agent from the previous time-step

Table 1. List of the 20 parameters used in inputs and of the 2 parameters used as outputs of the neural network models of agent behaviour.

defined as the absolute value of the mean agent heading: $P = \frac{1}{N}\left|\sum_{i=1}^{N} u_i\right|$ where u_i is the unit direction of agent i and $N = 5$ is the number of agents[43].

in relation with the environment the spatial (2D) distribution of *probability (density) of presence in the arena* (E_e) and the distribution of *distances of agents to their nearest wall* (W_e). These metrics assess the spatial repartition of the fish in arena.

We compute a **biomimetism score** using the similarity measure defined in[30, 36, 42] to quantify the realism of the behaviour of a group of simulated agents. It compares the behaviours of this group with the behaviours of the experimental fish averaged across all 10 experimental trials (**Control** case e_c). This score ranges from 0.0 to 1.0 and is defined as:

$$S(e, e_c) = \sqrt[5]{I(L_e, L_{e_c})I(A_e, A_{e_c})I(D_e, D_{e_c})I(P_e, P_{e_c})I(E_e, E_{e_c})} \tag{1}$$

The function $I(X, Y)$ is defined as such: $I(X, Y) = 1 - H(X, Y)$. The $H(X, Y)$ function is the Hellinger distance between two histograms[44]. It is defined as: $H(X, Y) = \frac{1}{\sqrt{2}}\sqrt{\sum_{i=1}^{d}(\sqrt{X_i} - \sqrt{Y_i})^2}$ where X_i and Y_i are the bin frequencies.

Optimisation

We calibrate the MLP models describing agent behaviour to match as close as possible the dynamics of fish in groups of 5 individuals during 30 minutes simulations (15 time-steps per seconds, *i.e.* 27000 steps per simulation) using a methodology inspired from[36, 42]. This is achieved by optimising the connection weights of the MLP using evolutionary algorithms that iteratively perform global optimisation (inspired by biological evolution) on a defined fitness function so as to find its optimal value[45, 46].

This calibration is a challenging problem because controllers have to cope with the multi-level and multi-modal nature of fish collective behaviour. It has to take into account conflicting behaviours and dynamics: dynamics at the level of the individual (individual trajectories), at the level of the group (aggregation) and in response to environmental stimuli (wall following behaviour). The optimisation process must handle an ill-defined fitness function (incorporating the biomimetic score defined earlier) that only describes partially the behaviour of experimental fish and that follow an aggregated representation (the biomimetic score aggregates several unrelated metrics).

We consider five optimisation methods (Table 2). Each case corresponds to a different way of optimising the MLP controllers and caters to different difficulties of the calibration process. We employ two global optimisers: CMA-ES[32] and NSGA-III. The CMA-ES[32] algorithm is a popular state-of-the-art global optimiser that is able to

handle problems where few assumptions need to be made on the nature of the underlying fitness function. It is able to cope with noisy, ill-defined fitness functions. It is a mono-objective optimisation algorithm (*i.e.* the fitness function is unidimensional, only one objective is optimised). The NSGA-III[33] algorithm is a popular multi-objective (*i.e.* the fitness function is multi-dimensional, several objectives are optimised) global optimiser . It is considered instead of the NSGA-II algorithm[47] employed in[42] because it is known to converge faster than NSGA-II on problems with more than two objectives[48].

In the **CMA-ES** case, we use the CMA-ES[32] algorithm with the task of maximising the biomimetism score between MLP-driven agents simulations (e_o) and experiment fish groups (e_c): S_{e_o,e_c}. This case tests the hypothesis that a mono-objective solution is sufficient to cope with the multi-modal nature of the problem.

In the **NSGA-III-SP** (Selection pressures) case, we use the NSGA-III[33] to maximise three objectives. The first objective is a performance objective corresponding to the biomimetism score function: S_{e_o,e_c}. We also consider two other objectives used to guide the evolutionary process: one that promotes genotypic diversity[49] (defined by the mean euclidean distance of the genome of an individual to the genomes of the other individuals of the current population), the other encouraging behavioural diversity (defined by the euclidean distance between the L_e, A_e, D_e, P_e and E_e scores of an individual). This case tests the hypothesis that the optimisation process needs a mechanism adapting the exploration vs exploitation tradeoff and encourage the exploration of solution with varying behaviours.

In the **NSGA-III-CM** (Combined Metrics) case, we use the NSGA-III[33] to maximise three objectives. The first objective is a performance objective corresponding to the biomimetism score function: S_{e_o,e_c}. The second and third objectives correspond to the aggregation of behavioural metrics related either to group behaviours (mean of inter-individual distance score D_e and polarisation score P_e) or related to individuals trajectories and response to environmental factors (mean of linear speeds score L_e and density of presence in the arena score E_e). This case tests the hypothesis that the optimisation process must explore solutions handling dynamics at different scales.

In the **NSGA-III-AM** (All metrics) case, we use the NSGA-III[33] to maximise five objectives: one corresponding to the biomimetim score function S_{e_o,e_c}, and the four others corresponding to the four behavioural distances metrics considered: the linear speeds score L_e, the inter-individual distance score D_e, the polarisation score P_e and the density of presence score E_e. This case tests the hypothesis that the optimisation process must explore solutions catering to each considered behavioural metrics.

In the **NSGA-III-Nov** (Novelty) case, we use the NSGA-III[33] to maximise two objectives. The first objective is a performance objective corresponding to the biomimetism score function: S_{e_o,e_c}. The second objective is a novelty measure as defined in[50,51]. It promotes the search of novel behaviours rather than only the search of efficiency (performance). This process has two benefits over optimising only performance: it can circumvent the deceptiveness present in some problems and it allows the evolutionary process to be more open-ended. Here, our problem is potentially deceptive has it involves trade-offs between different expected behaviours; their combination produces a large number of local optima in the fitness function. The novelty measure $\rho(i)$ of individual i is computed from a growing archive of previously encountered solutions (containing genomes and performances of previously evaluated individuals). It corresponds to the mean behavioural distance between i and its k nearest neighbours in the archive: $\rho(x) = \frac{1}{k} \sum_{j=0}^{k} d(x, \mu_j)$, with μ_j the j-th nearest neighbour of x, and d the euclidean distance. Our best performing results are obtained with $k = 4$.

In the **SL** (Supervised Learning) case, we employ a supervised learning approach to train MLPs to match experimental fish behaviour. The learning process is achieved by using the Adam[52] global optimiser. During training, a dropout layer (probability of dropout: 0.3) is added after the hidden layer. The training dataset corresponds to the reconstructed sets of inputs and outputs for each agent of the experimental fish trajectory at each time-step. The loss function is the mean squared error compared to the reference outputs for a given set of inputs.

In all cases, we use populations of 120 individuals and 500 generations. Each case is repeated in 10 different trials. We use CMA-ES and NSGA-III implementations based on the DEAP python library[53].

Case name	Algorithm	# of objectives	Objectives
CMA-ES	CMA-ES[32]	1	Biomimetism score
NSGA-III-SP (Selection Pressures)	NSGA-III[33]	3	Biomimetism score, genotypic diversity, behavioural diversity
NSGA-III-CM (Combined Metrics)	NSGA-III[33]	3	Biomimetism score, trajectory metrics, environment metrics
NSGA-III-AM (All Metrics)	NSGA-III[33]	5	Biomimetism score, inter-individual score, linear speeds score, polarisation score, density of presence score
NSGA-III-Nov (Novelty)	NSGA-III[33]	2	Biomimetism score, novelty
SL	Adam[52]	-	Loss function: mean squared error

Table 2. List of evolutionary algorithms methods used to optimise the MLP controllers weights. The best-performing controllers obtained through each method is compared to the experimental fish behaviours (**Control** case) of ten 30-minutes trials of five agents (simulated or experimental fish). The considered objectives are described in the Methods section. All evolutionary methods have a budget of 60 000 evaluations over 500 generations with populations of 120 individuals. Each methods is repeated through 10 different runs and fitness are computed across 10 different simulation trials per evaluation. The **CMA-ES** algorithm was used with an initial $\sigma = 0.5$. The **NSGA-III** algorithm was used with rates of crossovers of 1.0 and rates of mutations of 0.2. In the **SL** case, MLP are trained by Adam[52] over 150 epochs with a batch size of 64 and a learning rate of 0.000005.
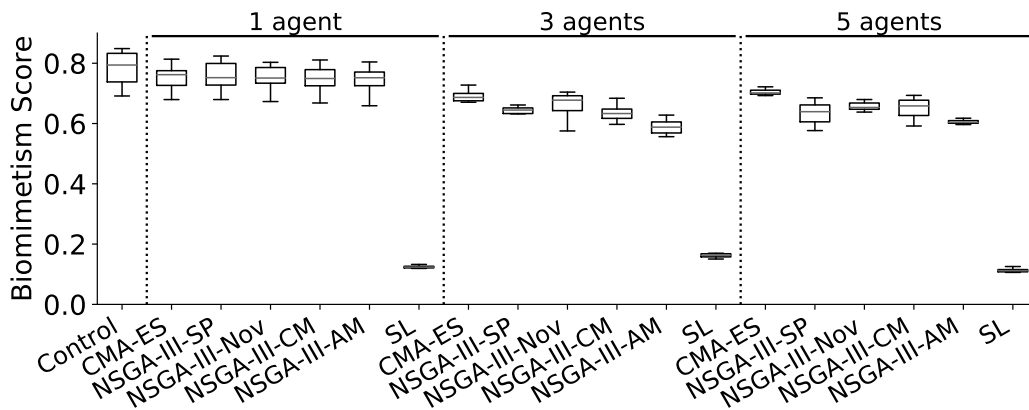
Figure 2. Quantitative analysis of the difference of behaviour between the 5-agent experimental fish groups (**Control**) and 5-agent groups formed by 1,3,5 simulated agents and respectively 4,2,0 experimental fish from the **Control** case. The simulated agents are driven by the best-performing MLP models optimised by CMA-ES or NSGA-III algorithms and according to the configurations detailed on Table 2, obtained after 10 optimisation runs. Results are obtained over 10 different trials (experiments for fish-only groups, and simulations for ANN models). We consider five behavioural features to characterise exhibited behaviours. **Inter-individual distances** corresponds to the similarity in distribution of inter-individual distances between all agents and measures the capabilities of the agents to aggregate. **Linear speed distributions** corresponds to the distributions of linear speed of the agents. **Polarisation** measures how aligned the agents are in the group. **Probability of presence** corresponds to the similarity in spatial distribution of agent density of presence in each part of the arena. It involves the comparison of 2D spatial histograms as seen in Fig. 3B. The **Biomimetic score** corresponds to the geometric mean of the other scores. Scores of simulated groups with 1 simulated agent optimised by evolutionary methods are not significantly different (Wilcoxon test, $p < 0.05$) from experimental fish groups. However, scores of simulated groups with 3 or 5 simulated agents optimised by all methods are all significantly different from experimental fish groups.
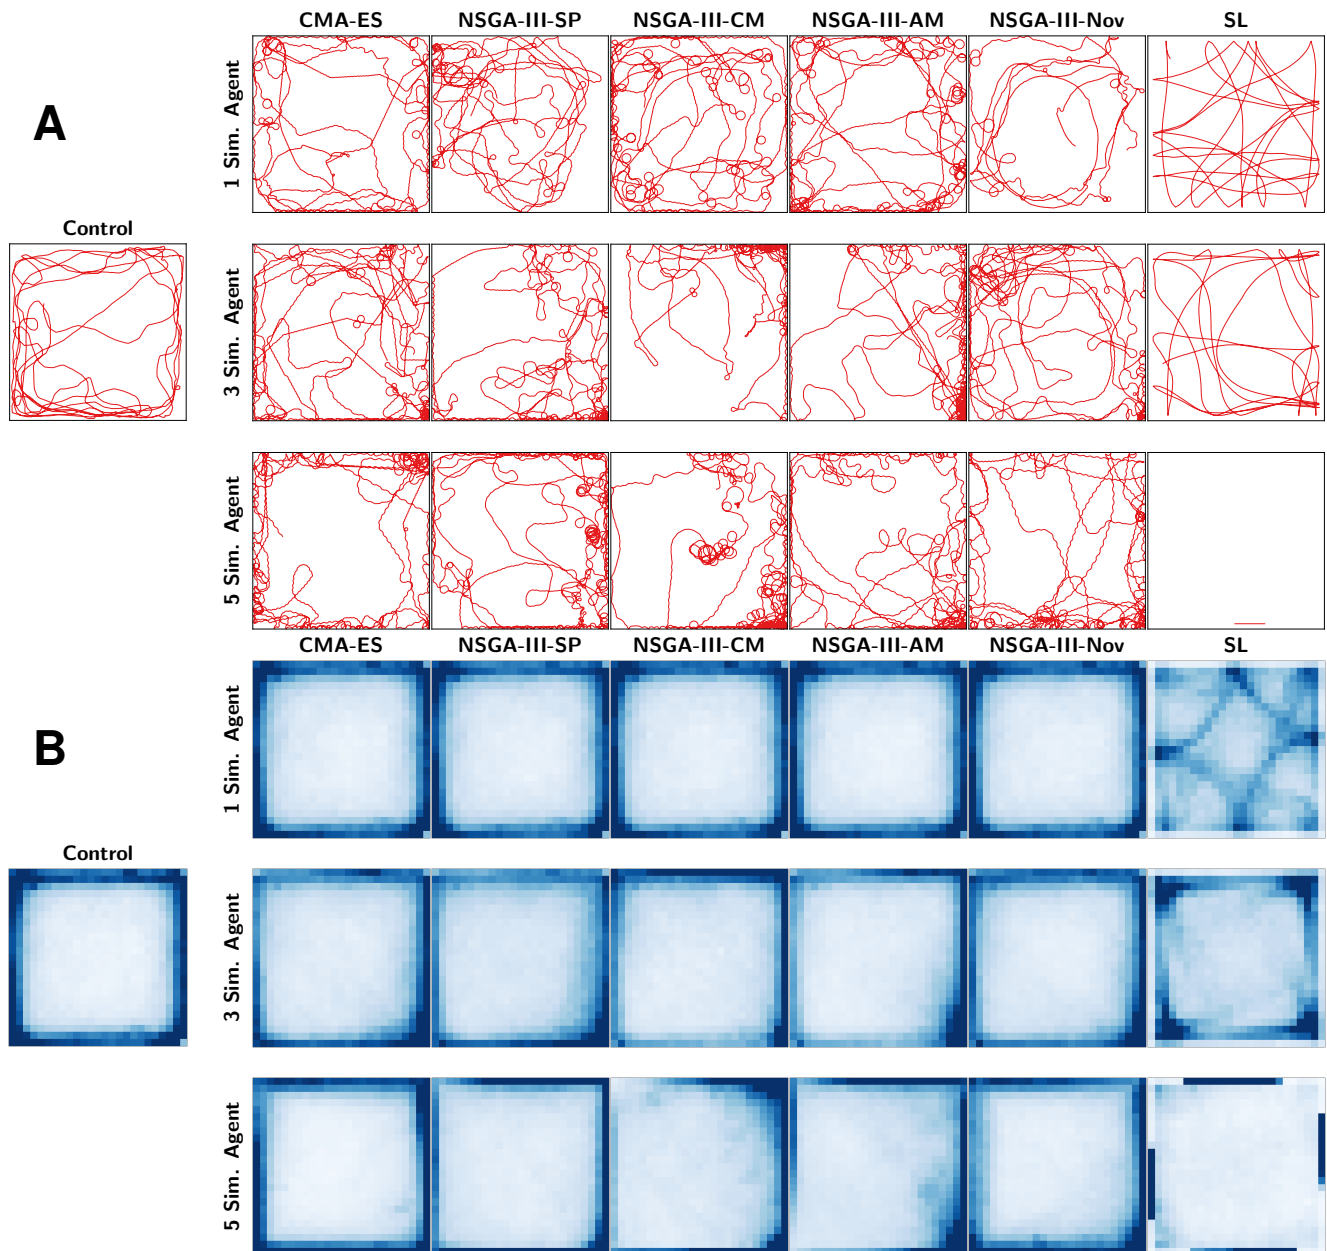
Figure 3. Kinetic behaviour of the agents (simulated or experimental fish) in the arena, using the best-performing controllers. These controllers are MLP neural networks evolved through the five considered methods. They can be compared to experimental fish (**Control** case) in 30-minutes trials in the square ($1m \times 1m$) arena, obtained as in[22, 35]. **A** Examples of individual trajectories of one simulated agent among the five making the group over one minute (900 time-steps). **B** Presence probability density of all autonomous agents in the arena represented by 2D histograms with 25×25 bins.

Results

We consider simulations with groups of 5 agents including 1, 3, or 5 simulated MLP-driven agents (Fig. 1) and respectively 4, 2, or 0 agents with trajectories corresponding to experimental fish data. We analyse the behaviour the MLP-driven agents and quantify how they compare to the behaviour of fish-only groups (**Control** case). The MLP controllers are calibrated using the methods described in the previous section. We only consider the best-evolved ANN controllers.

Examples of agents trajectories are found in Fig. 3A. For all simulated cases, they correspond to the trajectories of an MLP-driven agents, and can be compared to the trajectories of an actual fish in the **Control** case. We can see that the MLP-driven agents exhibit wall-following behaviour as in **Control** but their trajectories incorporate patterns not found in actual fish trajectories. Namely, small circular loops can appear in all simulated cases (especially here with **NSGA-III** cases) when simulated agents performs U-turns to catch up with the rest of the group. Previous analyses in[36] showed that this effect cannot be easily mitigated even when the fitness function takes into account agents angular speeds.

We present the presence probability density of all agents in the arena in Fig. 3B. In experiments with only one MLP-driven agent, the resulting presence probability densities of all cases match those of the **Control** case. However, experiments with 3 or 5 MLP-driven agents can exhibit excesses of density of presence in corners compared to the **Control** case, despite that presence probability density is taken into account in the fitness function computation. It is especially present in experiments with 3 MLP-driven agents, and in the **NSGA-III-SP**, **NSGA-III-CM** and **NSGA-III-AM** cases.

We compute the biomimetism scores (Fig. 2) presented in the Methods section quantifying the similarity between MLP-driven agents trajectories and **Control** case experimental fish trajectories. The similarity scores for each behavioural metric is found in Supplementary Information (Fig. 4).

The results of the **Control** case show a large variability in behaviour of the experimental fish because the 10 tested fish groups are composed of different fish (totalling 50 different fish) with different behaviour and individual preferences. This is especially the case with fish linear speeds (corroborating studies on zebrafish behaviour[35, 54]) and wall-following behaviour (measured by the probability of presence scores). This suggests that fish follow walls from a distance that vary according to group composition, with sporadic (to a varying degree) passes through the centre of the arena, possibly in small short-lived sub-groups. Conversely, the fish tend to exhibit similar degree of alignment across the 10 experimental runs and sustained aggregative tendencies (measured by inter-individual distances scores).

All results obtained using the **SL** case are sub-optimal compared to results obtained through evolutionary computation.

The similarity scores of all evolutionary cases for groups with only 1 MLP-driven agent are within the variance domain of the **Control** case, except for the polarity score: it shows that these groups exhibit relatively similar dynamics as a fish-only group (at least according to our proposed measures). However, it remains difficult for the MLP-driven agent to stay aligned with the rest of the group. This is exemplified in Fig. 3A by the tendency of the MLP-driven agents to exhibit small trajectory loops, possibly as a sustained effort to remain close to the centroid of the group. Previous results in[36] showed that it is difficult to prevent this occurrence, even when agent angular speeds distributions are taken into account in fitness computation.

However, simulations of groups with 3 or 5 MLP-driven agents correspond to far more complex problem tasks and bring about results with mostly lower mean similarity scores than the **Control** case. In both types of groups, the **CMA-ES** case displays the best-performing individuals in term of biomimetic scores, with the **NSGA-III-AM** case corresponding to individuals with the lowest biomimetic scores amongst evolutionary cases. This is also the cases for inter-individual distances scores with results of the **CMA-ES**, **NSGA-III-Nov** and **NSGA-III-CM** cases with the variance domain of the **Control** case. All evolutionary cases display lower inter-individual distances scores for groups with 3 MLP-driven agents than with groups with 5 MLP-driven agents: it suggests that it is easier to aggregate with autonomous simulated agents capable of reacting to other agents (groups of 5 MLP-driven agents) rather than aggregate with the 2 non-responsive agents present in groups of 3 MLP-driven agents.

The situation is reversed with linear speeds scores, where most cases exhibit lower similarity scores with groups

of 5 MLP-driven agents than in groups of 3 MLP-driven agents. Indeed, groups of 3 MLP-driven agents includes 2 agents with trajectories taken from experimental fish which would bias the scores to be closer to the **Control** case.

Depending on the behavioural metrics, different cases have varying degree of similarity to the **Control** case because they each have a different objective space, with several way of representing the target task. As such, some behaviours are easier to reach depending on the methodological case. With groups of 5 MLP-driven agents, the **CMA-ES** case has the highest scores of inter-individual distances, linear speeds and probability of presence. The **NSGA-III-CM** case has the highest polarity scores.

Discussion

Developing artificial neural network models for fish collective behaviours is a challenging issue because fish behaviours are multi-level (tail-beats as motor response vs individual trajectories vs collective dynamics), multi-modal (variety of behavioural patterns, input/output sources), context-dependent (behaviours depending on the spatial position and neighbours or group structure) and intrinsically stochastic (leading to individual and collectives choices and action selection)[22, 55]. Moreover, fish dynamics involve trade-offs between social tendencies (aggregation, group formation), and response to the environment (wall-following, zone occupation). They also present distinct movement patterns that allow them to move in polarised groups and react collectively to environmental and social cues.

Here we show that evolved MLP models give good results for modelling zebrafish collective behaviours taking into account the boundaries of their environment. We quantify the realism of these MLP-driven autonomous agents compared to experimental fish by computing biomimetism scores.

Our methodology involves the calibration of these MLP models by using evolutionary computation to match experimental fish behaviour, using a biomimetism score as a performance objective. We employ and test both a mono-objective (CMA-ES[32]) and a multi-objective (NSGA-III[33]) evolutionary algorithms to calibrate the weights of the MLP models. Here, we consider 1 to 5 autonomous MLP-driven agents. This allows to test models that do not need any fish present and still exhibit similar dynamics as fish groups. When considering more than one autonomous agent the task of optimising MLP models become far more difficult but it is more realistic as all modelled agents are capable of reacting to one another. We compare the results obtained through evolutionary computation to a supervised learning methodology where we train the MLP models on the reconstructed dataset of inputs and outputs for each time-step of the experimental data. This supervised learning method only provides sub-optimal results (similar to the results from[38]), and are largely dominated by results obtained using evolutionary computation. The latter method differs mostly by the ability to pursue global optimisation, rather than local optimisation (supervised learning with gradient descent) which is prone to premature convergence.

Most of the models developed in behavioural biology describe only a small and specific part of the behavioural repertoire. Often, they do not completely describe animals as autonomous agents in their environment even in laboratory conditions. They are designed to answer specific biological questions. Moreover, the larger the behavioural repertoire, the larger the agent models, such that it becomes difficult to calibrate them on experimental data as the number of parameters explodes. The same issue arises to describe them by kinetic equations for the same reasons. In stark contrast, evolved ANN models can take into account behavioural details that are usually difficult to take into account in current behavioural models.

ANN models make it possible to ask relevant biological questions. Firstly, they can be used to find in the experimental data the necessary information needed to reproduce animal behaviours. They also open the door at using ANN to model the functional neural working of the information processing by the animal during the sensory-motor loop. They do not need to be biomimetic at the lowest level *i.e.* the individual neural physiology, but can model how information is processed. Finally, they open new possibilities to design models that can be transferred to robotics systems. Such system-based models allow to describe a fully autonomous agent in a realistic environment contrary to most models that describe only a sub-part of the agent behaviours[56].

Indeed, research on animal and robot interactions also need bio-mimetic formal models as robot behavioural controllers for robots to behave like animals[30, 42, 57–59]. Robots controllers have to deal with a whole range of behaviours to allow them to take into account not only other individuals but also environment particularities, such

as walls[27, 30]. Robot behavioural models must be multilevel as they have to cope with low level hardware events and high level social behaviours. In this study, we showed that ANN are good candidates to model individual and collective fish behaviours, which is particularly relevant in the context of social bio-hybrid systems composed of animals and robots.

References

1. Lopez, U., Gautrais, J., Couzin, I. D. & Theraulaz, G. From behavioural analyses to models of collective motion in fish schools. *Interface focus* **2**, 693–707 (2012).

2. Deutsch, A., Theraulaz, G. & Vicsek, T. Collective motion in biological systems. *Interface Focus.* **2**, 689 (2012).

3. Herbert-Read, J. E., Romenskyy, M. & Sumpter, D. J. A turing test for collective motion. *Biol. letters* **11**, 20150674 (2015).

4. Reynolds, C. Flocks, herds, and schools: a distributed behavioral model. *Comput. Graph.* **21**, 25–34 (1987).

5. Aoki, I. A simulation study on the schooling mechanism in fish. *Bull. Jpn. Soc. for Sci. Fish* **48**, 1081–1088 (1982).

6. Aoki, I. Internal dynamics of fish schools in relation to inter-fish distance. *Bull. Jpn. Soc. for Sci. Fish* **50**, 751–758 (1984).

7. Bode, N., Faria, J., Franks, D., Krause, J. & Wood, A. How perceived threat increases synchronization in collectively moving animal groups. *Proc. Royal Soc. Lond. B* **277**, 3065–3070 (2010).

8. Bode, N., Franks, D. & Wood, A. Limited interactions in flocks: relating model simulations to empirical data. *J. Royal Soc. Interface* **8**, 301–304 (2011).

9. Couzin, I., Krause, J., James, R., Ruxton, G. & Franks, N. Collective memory and spatial sorting in animal groups. *J. theoretical biology* **218**, 1–11 (2002).

10. Lopez, U., Gautrais, J., Couzin, I. & Theraulaz, G. From behavioural analyses to models of collective motion in fish schools. *Interface Focus.* **2**, 693–707 (2012).

11. Vicsek, T., Czirok, A., Ben-Jacob, E., Cohen, I. & Shochet, O. Novel type of phase transition in a system of self-driven particles. *Phys. Rev. Lett.* **75**, 1226–1229 (1995).

12. Bertin, E., Droz, M. & Grégoire, G. Boltzmann and hydrodynamic description for self-propelled particles. *Phys. Rev. E* **74**, 022101 (2006).

13. Chaté, H., Ginelli, F., Grégoire, G. & Raynaud, F. Collective motion of self-propelled particles interacting without cohesion. *Physicel Rev. E* **77**, 046113 (2008).

14. Chaté, H., Ginelli, F., Grégoire, G., Peruani, F. & Raynaud, F. Modeling collective motion: variations on the vicsek model. *The Eur. Phys. J. B* **64**, 451–456 (2008).

15. Nagai, K., Sumino, Y., Montagne, R., Aranson, I. & Chaté, H. Collective motion of self-propelled particles with memory. *Phys. Rev. Lett.* **114**, 168001 (2015).

16. Gautrais, J. *et al.* Analysing fish movement as a persistent turning walker. *J. Math. Biol.* **58**, 429–445 (2009).

17. Gautrais, J. *et al.* Deciphering interactions in moving animal groups. *PLoS Comput. Biol.* e1002678 (2012).

18. Zienkiewicz, A., Barton, D., Porfiri, M. & di Bernardo, M. Data-driven stochastic modelling of zebrafish locomotion. *J. Math. Biol.* (2014).

19. Mwaffo, V., Anderson, R., Butail, S. & Porfiri, M. A jump persistent turning walker to model zebrafish locomotion. *J. Royal Soc. Interface* **12**, 20140884 (2014).

20. Lemasson, B., Anderson, J. & Goodwin, R. Collective motion in animal groups from a neurobiological perspective: the adapive benefits of dynamic sensory loads and selection attention. *J. Theor. Biol.* **261**, 501–510 (2009).

21. Lemasson, B., Anderson, J. & Goodwin, R. Motion-guided attention promotes adaptive communications during social navigation. *Proc. Royal Soc. Lond. B* **280**, 20122003 (2013).

22. Collignon, B., Séguret, A. & Halloy, J. A stochastic vision-based model inspired by zebrafish collective behaviour in heterogeneous environments. *Royal Soc. Open Sci.* **3**, 150473 (2016).

23. Calovi, D. S. *et al.* Disentangling and modeling interactions in fish with burst-and-coast swimming reveal distinct alignment and attraction behaviors. *PLoS computational biology* **14**, e1005933 (2018).

24. Jiang, L. *et al.* Identifying influential neighbors in animal flocking. *PLoS computational biology* **13**, e1005822 (2017).

25. Lecheval, V. *et al.* Domino-like propagation of collective u-turns in fish schools. *bioRxiv* 138628 (2017).

26. Cazenille, L., Bredeche, N. & Halloy, J. Multi-objective optimization of multi-level models for controlling animal collective behavior with robots. In *Biomimetic and Biohybrid Systems, Living Machines 2015*, vol. 9222 of *Lecture Notes in Computer Science*, 379–390 (Springer, 2015).

27. Cazenille, L., Bredeche, N. & Halloy, J. Automated calibration of a biomimetic space-dependent model for zebrafish and robot collective behaviour in a structured environment. *arXiv preprint* (2017).

28. Cazenille, L., Bredeche, N. & Halloy, J. Automated optimisation of multi-level models of collective behaviour in a mixed society of animals and robots. *arXiv preprint arXiv:1602.05830* (2016).

29. Romano, D., Donati, E., Benelli, G. & Stefanini, C. A review on animal–robot interaction: from bio-hybrid organisms to mixed societies. *Biol. Cybern.* 1–25 (2018).

30. Cazenille, L. *et al.* How mimetic should a robotic fish be to socially integrate into zebrafish groups ? *Bioinspiration & biomimetics* (2017).

31. Bishop, C. *Pattern Recognition and Machine Learning (Information Science and Statistics)* (Springer-Verlag, Berlin, Heidelberg, 2006).

32. Auger, A. & Hansen, N. A restart CMA evolution strategy with increasing population size. In *Evolutionary Computation, 2005. The 2005 IEEE Congress on*, vol. 2, 1769–1776 (IEEE, 2005).

33. Yuan, Y., Xu, H. & Wang, B. An improved nsga-iii procedure for evolutionary many-objective optimization. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, 661–668 (ACM, 2014).

34. Sutton, R. & Barto, A. *Reinforcement Learning, Second Edition* (MIT Press, Cambridge, MA, USA, 2018), 2nd edn.

35. Séguret, A., Collignon, B., Cazenille, L., Chemtob, Y. & Halloy, J. Loose social organisation of ab strain zebrafish groups in a two-patch environment. *arXiv preprint arXiv:1701.02572* (2017).

36. Cazenille, L., Bredeche, N. & Halloy, J. Evolutionary optimisation of neural network models for fish collective behaviours in mixed groups of robots and zebrafish. In *Conference on Biomimetic and Biohybrid Systems*, 85–96 (Springer, 2018).

37. Cazenille, L. *et al.* How to blend a robot within a group of zebrafish: Achieving social acceptance through real-time calibration of a multi-level behavioural model. *arXiv preprint arXiv:1805.11371* (2018).

38. Iizuka, H., Nakamoto, Y. & Yamamoto, M. Learning of individual sensorimotor mapping to form swarm behavior from real fish data. In *Artificial Life Conference Proceedings*, 179–185 (MIT Press, 2018).

39. Norgaard, M., Ravn, O., Poulsen, N. & Hansen, L. *Neural networks for modelling and control of dynamic systems: a practitioner's handbook. Advanced textbooks in control and signal processing* (Springer, Berlin, 2000).

40. Cybenko, G. Approximation by superpositions of a sigmoidal function. *Math. control, signals systems* **2**, 303–314 (1989).

41. Sumpter, D. J., Mann, R. P. & Perna, A. The modelling cycle for collective animal behaviour. *Interface focus* **2**, 764–773 (2012).

42. Cazenille, L. *et al.* Automated calibration of a biomimetic space-dependent model for zebrafish and robot collective behaviour in a structured environment. In *Conference on biomimetic and biohybrid systems*, 107–118 (Springer, 2017).

43. Tunstrøm, K. *et al.* Collective states, multistability and transitional behavior in schooling fish. *PLoS computational biology* **9**, e1002915 (2013).

44. Deza, M. & Deza, E. *Dictionary of distances* (Elsevier, 2006).

45. Salimans, T., Ho, J., Chen, X. & Sutskever, I. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864* (2017).

46. Jiang, F., Berry, H. & Schoenauer, M. Supervised and evolutionary learning of echo state networks. In *International Conference on Parallel Problem Solving From Nature*, 215–224 (Springer, 2008).

47. Deb, K., Pratap, A., Agarwal, S. & Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *Evol. Comput. IEEE Transactions on* **6**, 182–197 (2002).

48. Ishibuchi, H., Imada, R., Setoguchi, Y. & Nojima, Y. Performance comparison of nsga-ii and nsga-iii on various many-objective test problems. In *Evolutionary Computation (CEC), 2016 IEEE Congress on*, 3045–3052 (IEEE, 2016).

49. Mouret, J.-B. & Doncieux, S. Encouraging behavioral diversity in evolutionary robotics: An empirical study. *Evol. computation* **20**, 91–133 (2012).

50. Lehman, J. & Stanley, K. O. Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*, 329–336 (2008).

51. Mouret, J. Novelty-based multiobjectivization. In *New horizons in evolutionary robotics*, 139–154 (Springer, 2011).

52. Kingma, D. & Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

53. Fortin, F.-A., De Rainville, F.-M., Gardner, M.-A., Parizeau, M. & Gagné, C. Deap: Evolutionary algorithms made easy. *J. Mach. Learn. Res.* **13**, 2171–2175 (2012).

54. Collignon, B., Séguret, A., Chemtob, Y., Cazenille, L. & Halloy, J. Collective departures in zebrafish: profiling the initiators. *arXiv preprint arXiv:1701.03611* (2017).

55. Sumpter, D. J. T., Szorkovszky, A., Kotrschal, A., Kolm, N. & Herbert-Read, J. E. Using activity and sociability to characterize collective motion. *Phil. Trans. R. Soc. B* **373**, 20170015 (2018).

56. Webb, B. Animals versus animats: Or why not model the real iguana? *Adapt. Behav.* **17**, 269–286 (2009).

57. Bonnet, F., Binder, S., de Oliveria, M., Halloy, J. & Mondada, F. A miniature mobile robot developed to be socially integrated with species of small fish. In *Robotics and Biomimetics (ROBIO), 2014 IEEE International Conference on*, 747–752 (IEEE, 2014).

58. Bonnet, F. *et al.* Design of a modular robotic system that mimics small fish locomotion and body movements for ethological studies. *Int. J. Adv. Robotic Syst.* **14**, 1729881417706628 (2017).

59. Bonnet, F., Cazenille, L., Gribovskiy, A., Halloy, J. & Mondada, F. Multi-robots control and tracking framework for bio-hybrid systems with closed-loop interaction. In *to be published in Robotics and Automation (ICRA), 2017 IEEE International Conference on* (IEEE, 2017).

60. Pérez-Escudero, A., Vicente-Page, J., Hinz, R. C., Arganda, S. & de Polavieja, G. G. idtracker: tracking individuals in a group by automatic identification of unmarked animals. *Nat. methods* **11**, 743–748 (2014).

Acknowledgements

Author contributions statement

L.C. designed the study and performed the data analysis, ran the simulations and wrote the manuscript. N.B. contributed with artificial neural networks and evolutionary computation expertise, wrote and revised the manuscript. J.H. designed and coordinated the study, wrote and revised the manuscript. All authors gave final approval for publication.

Additional information

Accession codes (where applicable)

Competing financial interests: The authors declare no competing financial interests.

Supplementary Information

Ethical statement

The experiments performed in this study were conducted under the authorisation of the Buffon Ethical Committee (registered to the French National Ethical Committee for Animal Experiments #40) after submission to the French state ethical board for animal experiments.

Fish rearing and handling

The fish were 6-12 months old at the time of the experiments. We kept the fish under laboratory conditions, $27°C$, $500\mu S$ salinity with a 10:14 day:night cycle. The fish were reared in housing facilities ZebTEC and fed two times a day (Special Diets Services SDS-400 Scientific Fish Food). The water pH level was maintained at 7, and Nitrites (NO^{-2}) were below 0.3 mg/l.

Experimental set-up

We use the experimental set-up described in [28,30,35–37]: it consists of an arena of $1000 \times 1000 \times 100$ mm made of white plexiglass and placed in a $1200 \times 1200 \times 300$ mm experimental tank. The tank is filled with water up to a level of 60 mm. The whole set-up is exposed to diffused light and confined behind white sheets to isolate experiments and homogenise luminosity. We use an overhead acA2040-25gm monochrome GigE CCD camera (Basler AG, Germany) to grab video frames of the experiments (with a resolution of 2040×2040 pixels, 15 frames per second). This camera is equipped with low distortion lenses CF12.5HA-1 (Fujinon, Tokyo, Japan). The experimental videos are then analysed and tracked off-line to retrieve the position of the fish[60].

Supplementary results

Figure 4 shows the behavioural scores of the best-performing controllers obtained through the studied optimisation methods.

Figures 5,6,7,8,9,10 provide analyses of simulated fish groups behaviour obtained through different optimisation methods (Table 2): respectively **CMA-ES**, **NSGA-III-SP**, **NSGA-III-CM**, **NSGA-III-AM**, **NSGA-III-Nov** and **SL**. These figures take into account the following metrics: **Inter-individual distances** corresponds to the similarity in distribution of inter-individual distances between all agents and measures the capabilities of the agents to aggregate. **Linear speed distributions** corresponds to the distributions of linear speed of the agents. **Polarisation** measures how aligned the agents are in the group. **Probability of presence** corresponds to the similarity in spatial distribution of agent density of presence in each part of the arena. It involves the comparison of 2D spatial histograms as seen in Fig. 3B. The **Biomimetic score** corresponds to the geometric mean of the other scores.

Figure 11 shows the evolution of biomimetic score in fitness computation per evaluation.
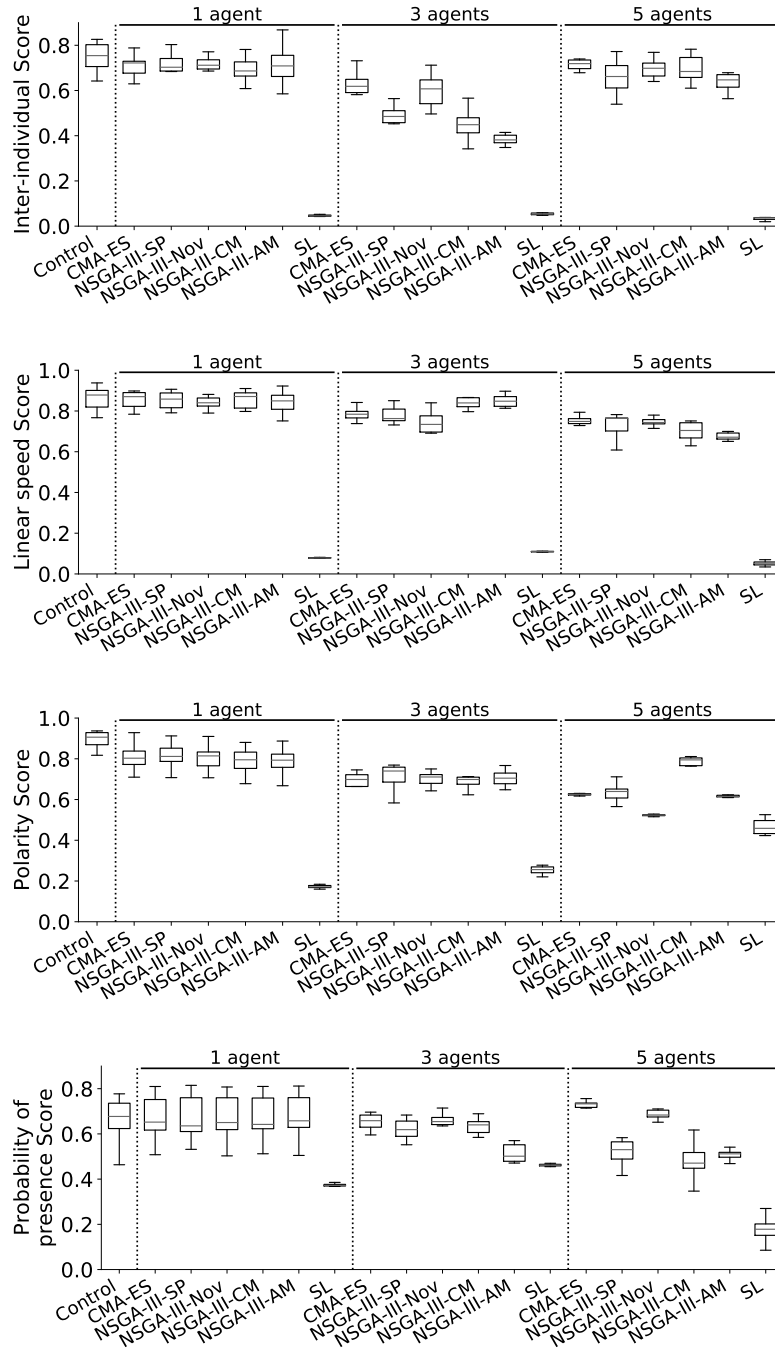
Figure 4. Quantitative analysis of the difference of behaviour between the 5-agent experimental fish groups (**Control**) and 5-agent groups formed by 1,3,5 simulated agents and respectively 4,2,0 experimental fish from the **Control** case. The simulated agents are driven by the best-performing MLP models optimised by CMA-ES or NSGA-III algorithms and according to the configurations detailed on Table 2, obtained after 10 optimisation runs. Results are obtained over 10 different trials (experiments for fish-only groups, and simulations for ANN models). We consider five behavioural features to characterise exhibited behaviours. **Inter-individual distances** corresponds to the similarity in distribution of inter-individual distances between all agents and measures the capabilities of the agents to aggregate. **Linear speed distributions** corresponds to the distributions of linear speed of the agents. **Polarisation** measures how aligned the agents are in the group. **Probability of presence** corresponds to the similarity in spatial distribution of agent density of presence in each part of the arena. It involves the comparison of 2D spatial histograms as seen in Fig. 3B. Similarity scores of simulated groups with 1 simulated agent optimised by all methods are not significantly different (Wilcoxon test, $p < 0.05$) from experimental fish groups (except for the polarity metric). However, Similarity scores of simulated groups with 3 or 5 simulated agents optimised by all methods are all significantly different from experimental fish groups.
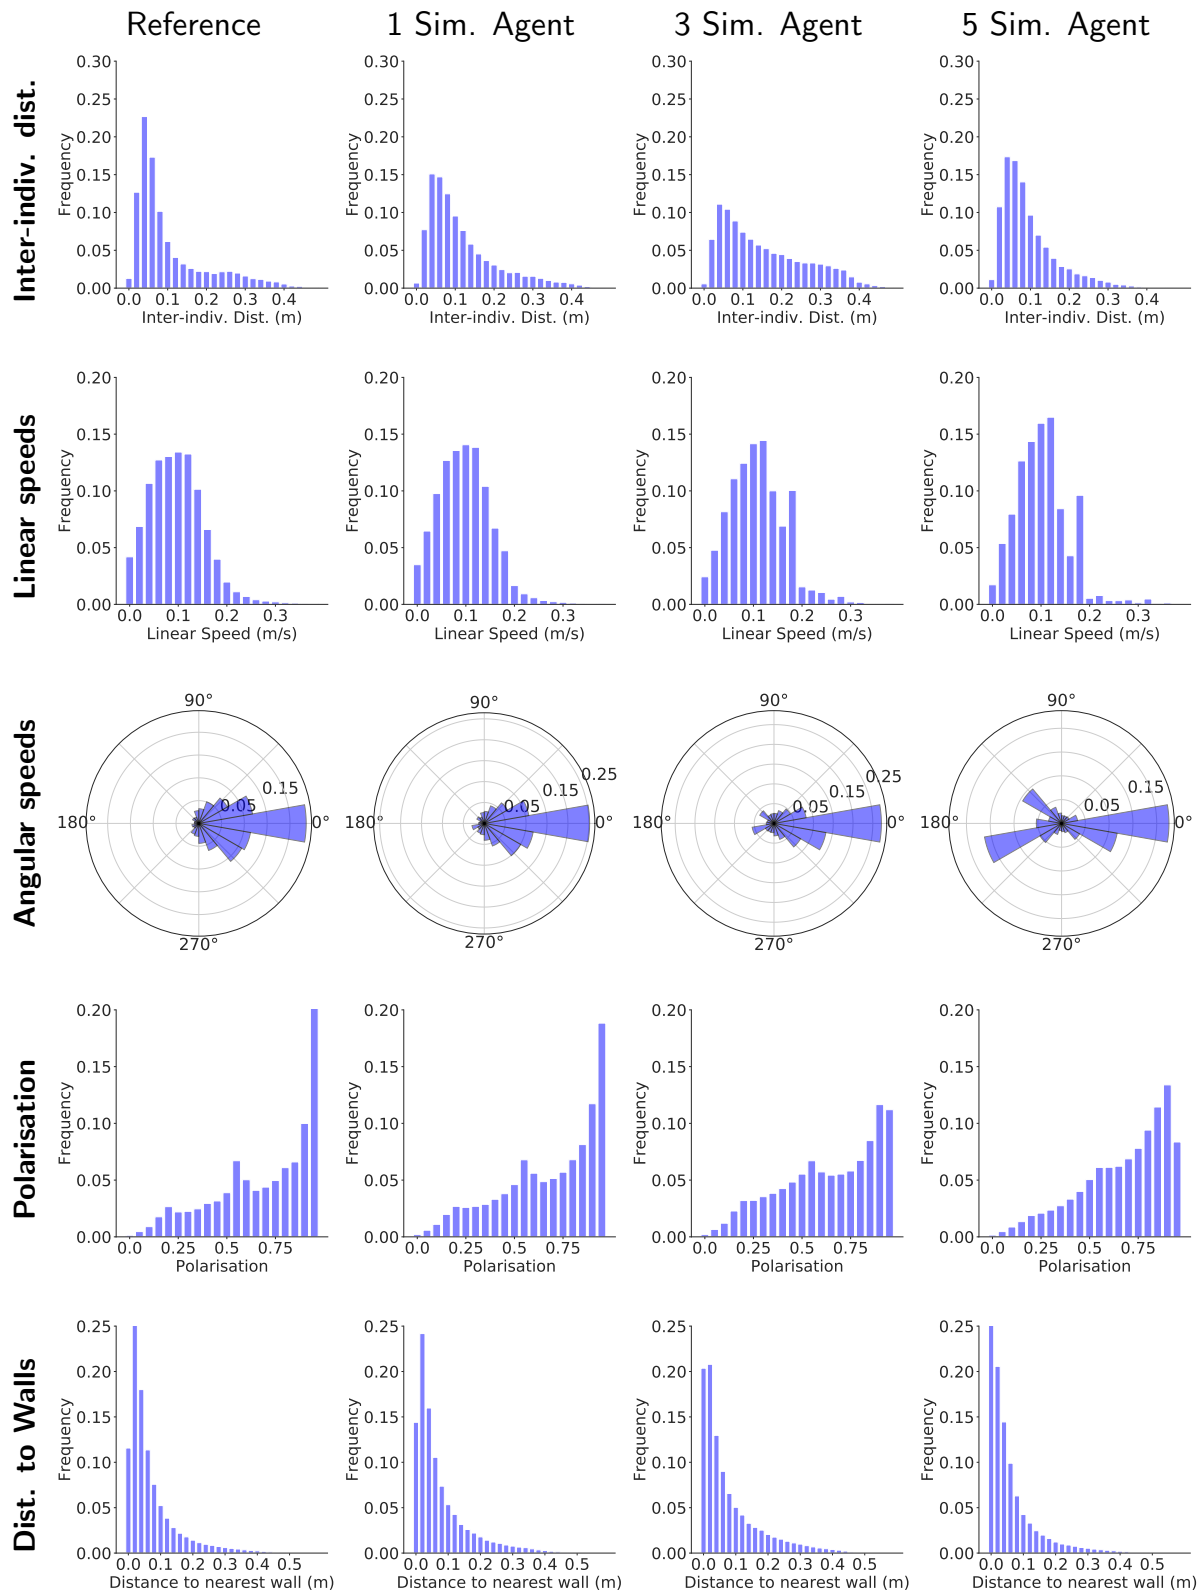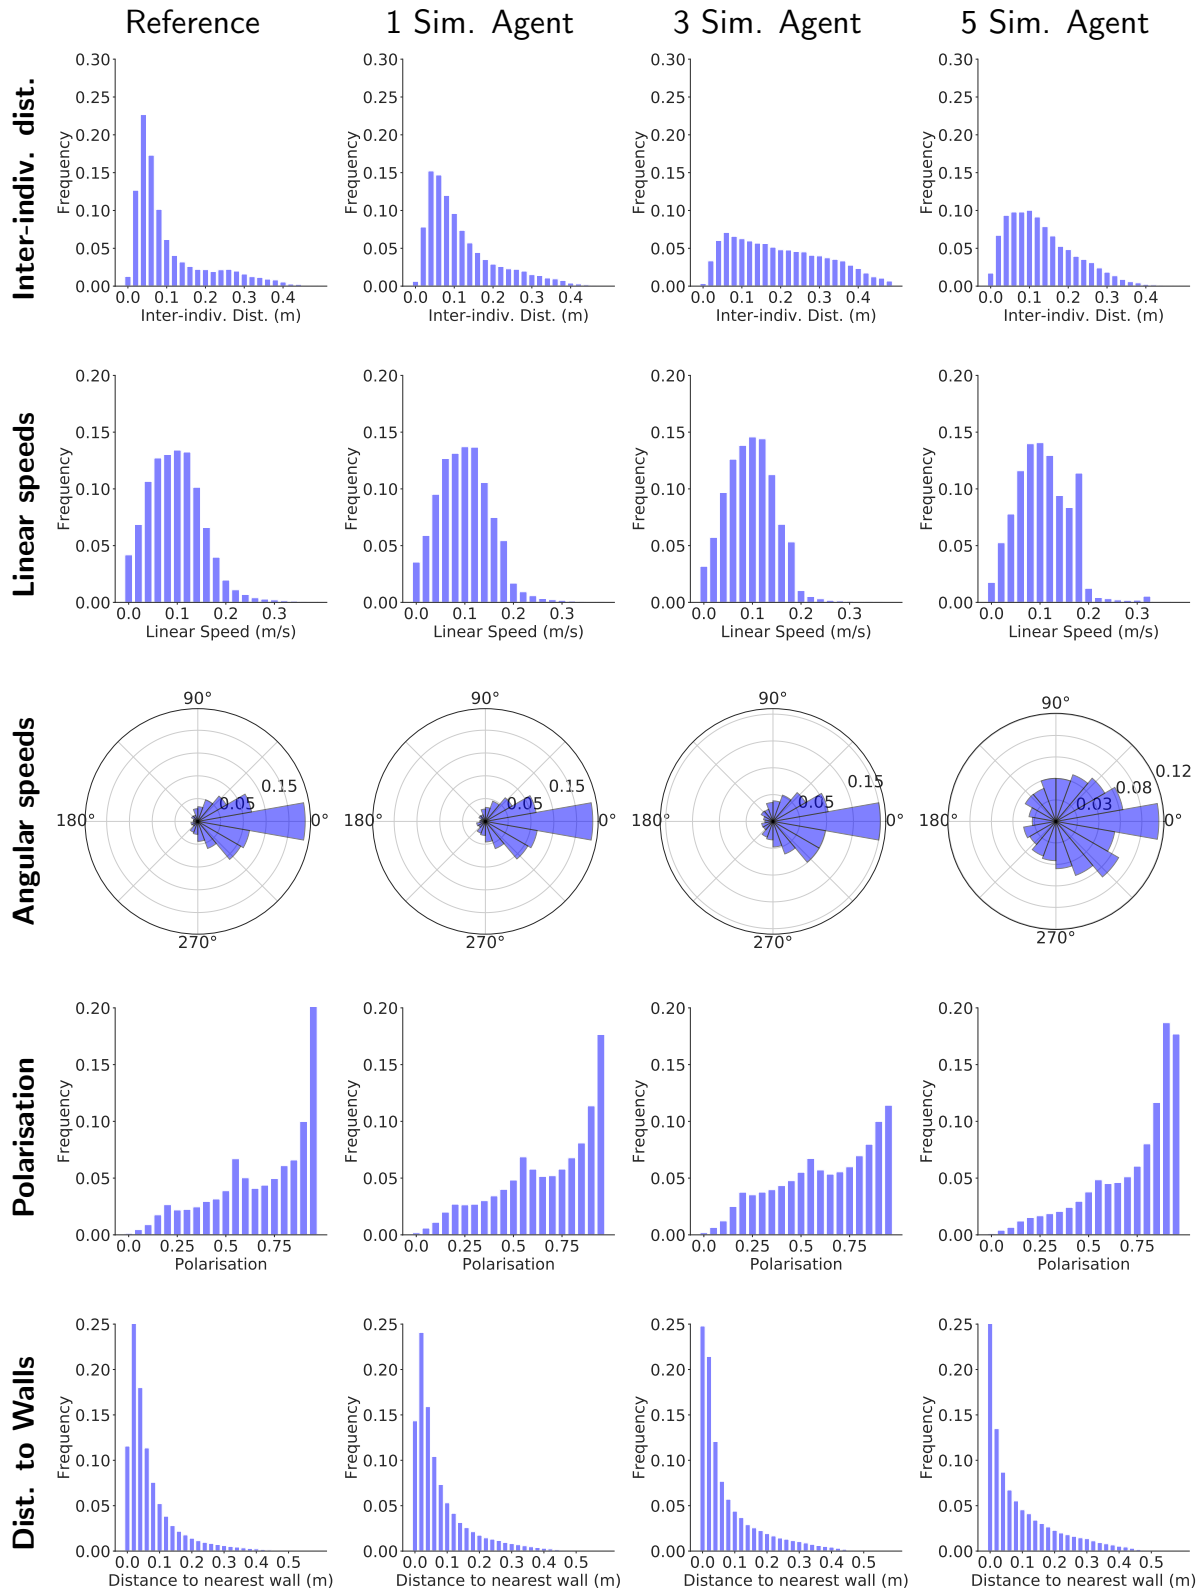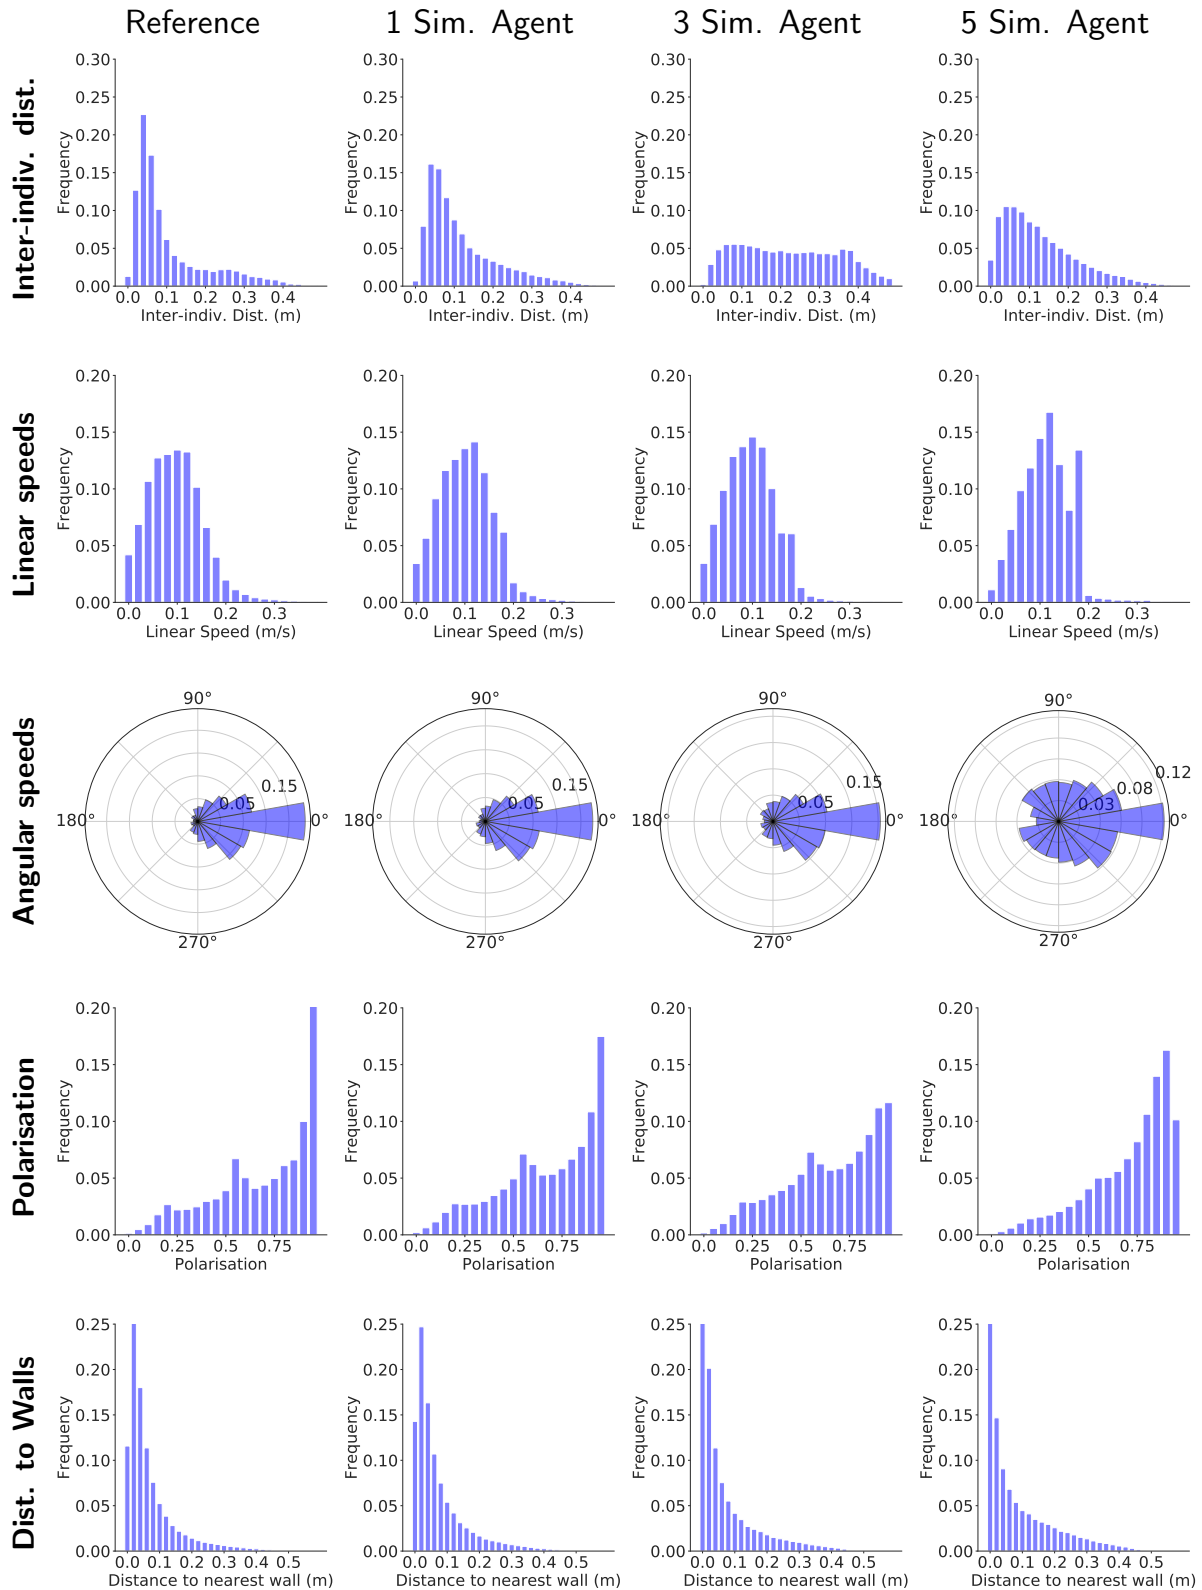
Figure 5. Comparison of agents behaviours between the **Control** case (10 trials of 30-minutes experiments involving 5 zebrafish) and the **CMA-ES** case (10 trials of 30-minutes simulations of 5 agents, with respectively 1,3,5 simulated agents driven by the best-performing MLP controller optimised through the **CMA-ES** method and 4,2,0 agents copying the trajectories of experimental fish). Agents behaviour is quantified across 5 behavioural features: inter-individual distances, linear and angular speeds distributions, polarisation, and distances to nearest wall.
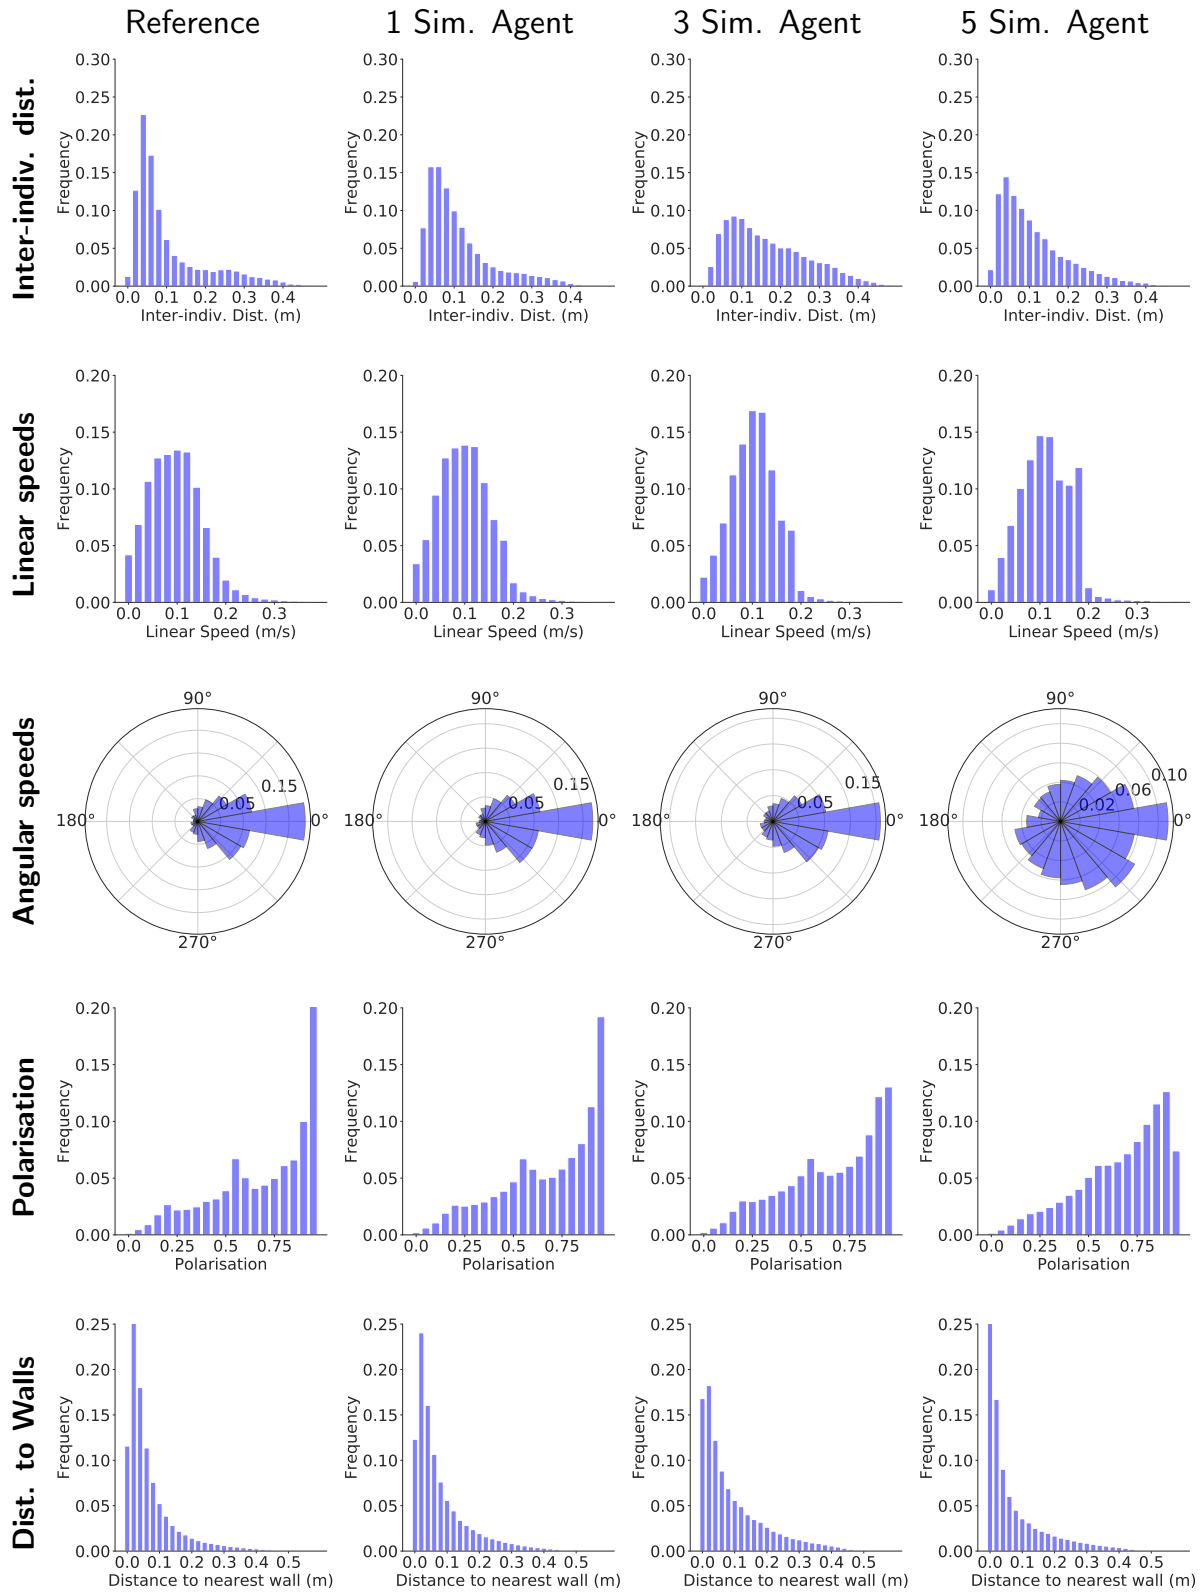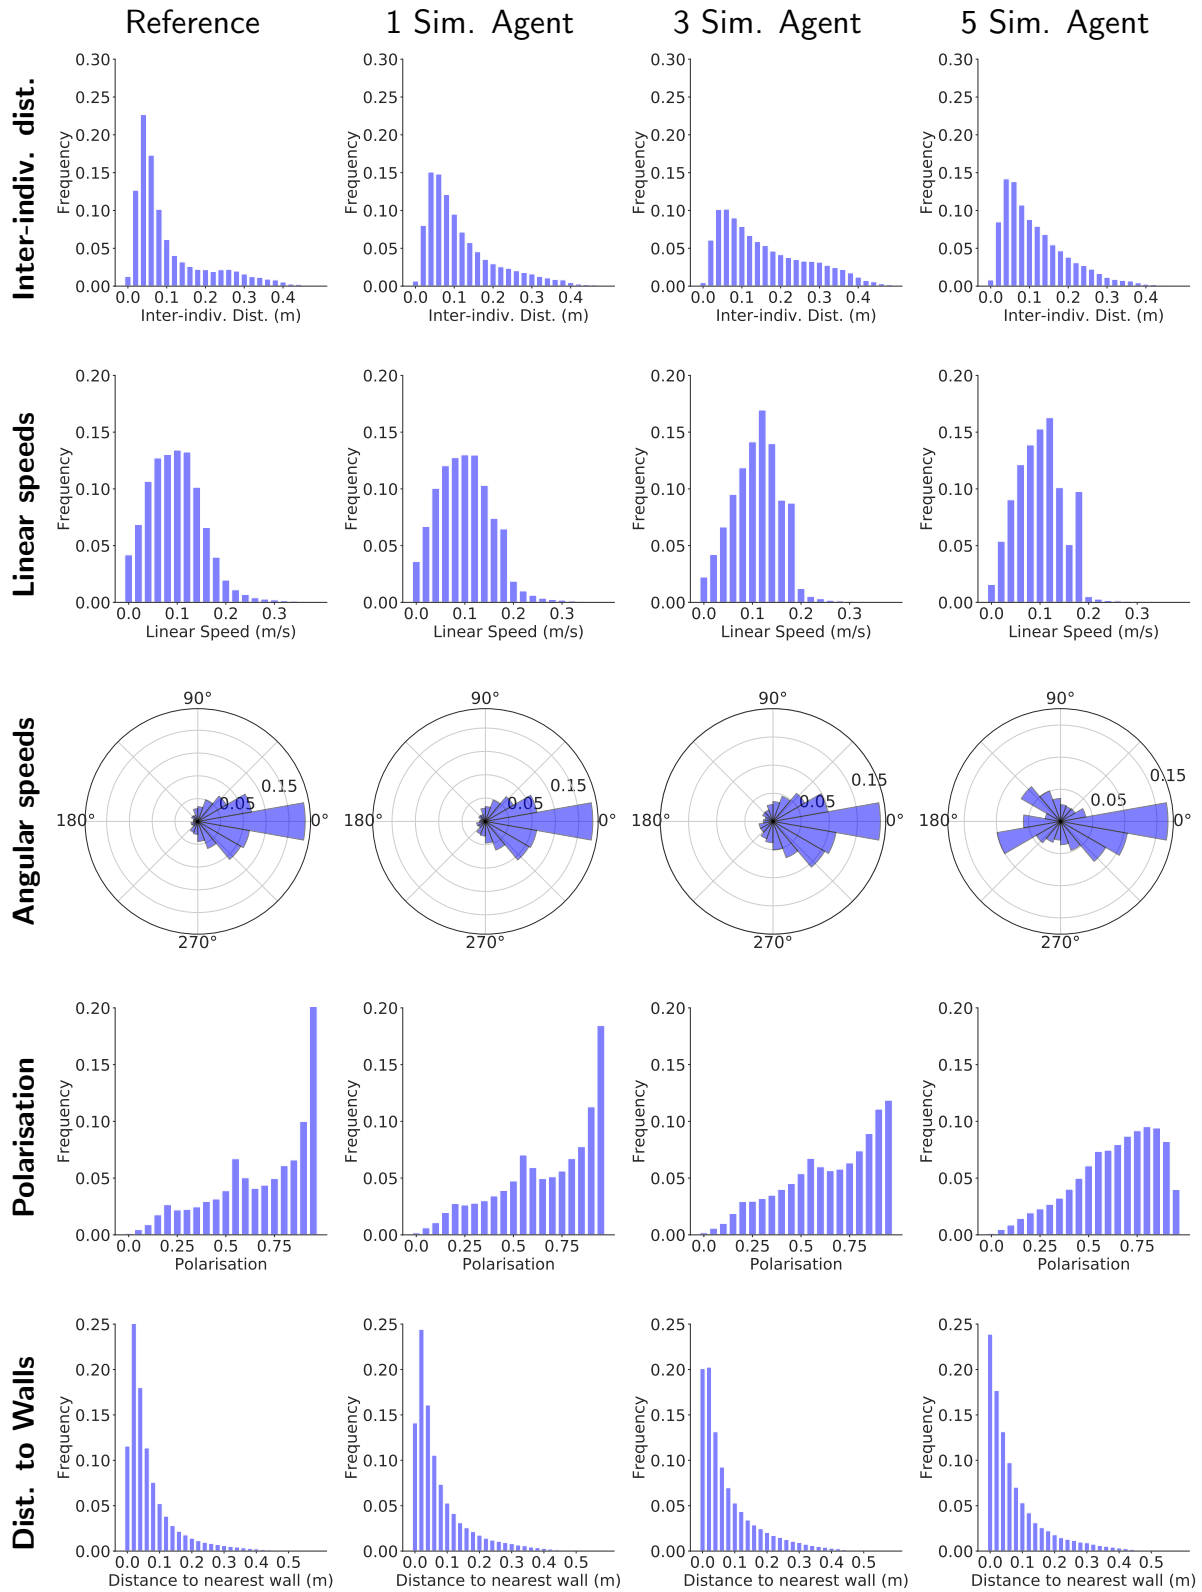
Figure 6. Comparison of agents behaviours between the **Control** case (10 trials of 30-minutes experiments involving 5 zebrafish) and the **NSGA-III-SP** case (10 trials of 30-minutes simulations of 5 agents, with respectively 1,3,5 simulated agents driven by the best-performing MLP controller optimised through the **NSGA-III-SP** method and 4,2,0 agents copying the trajectories of experimental fish). Agents behaviour is quantified across 5 behavioural features: inter-individual distances, linear and angular speeds distributions, polarisation, and distances to nearest wall.
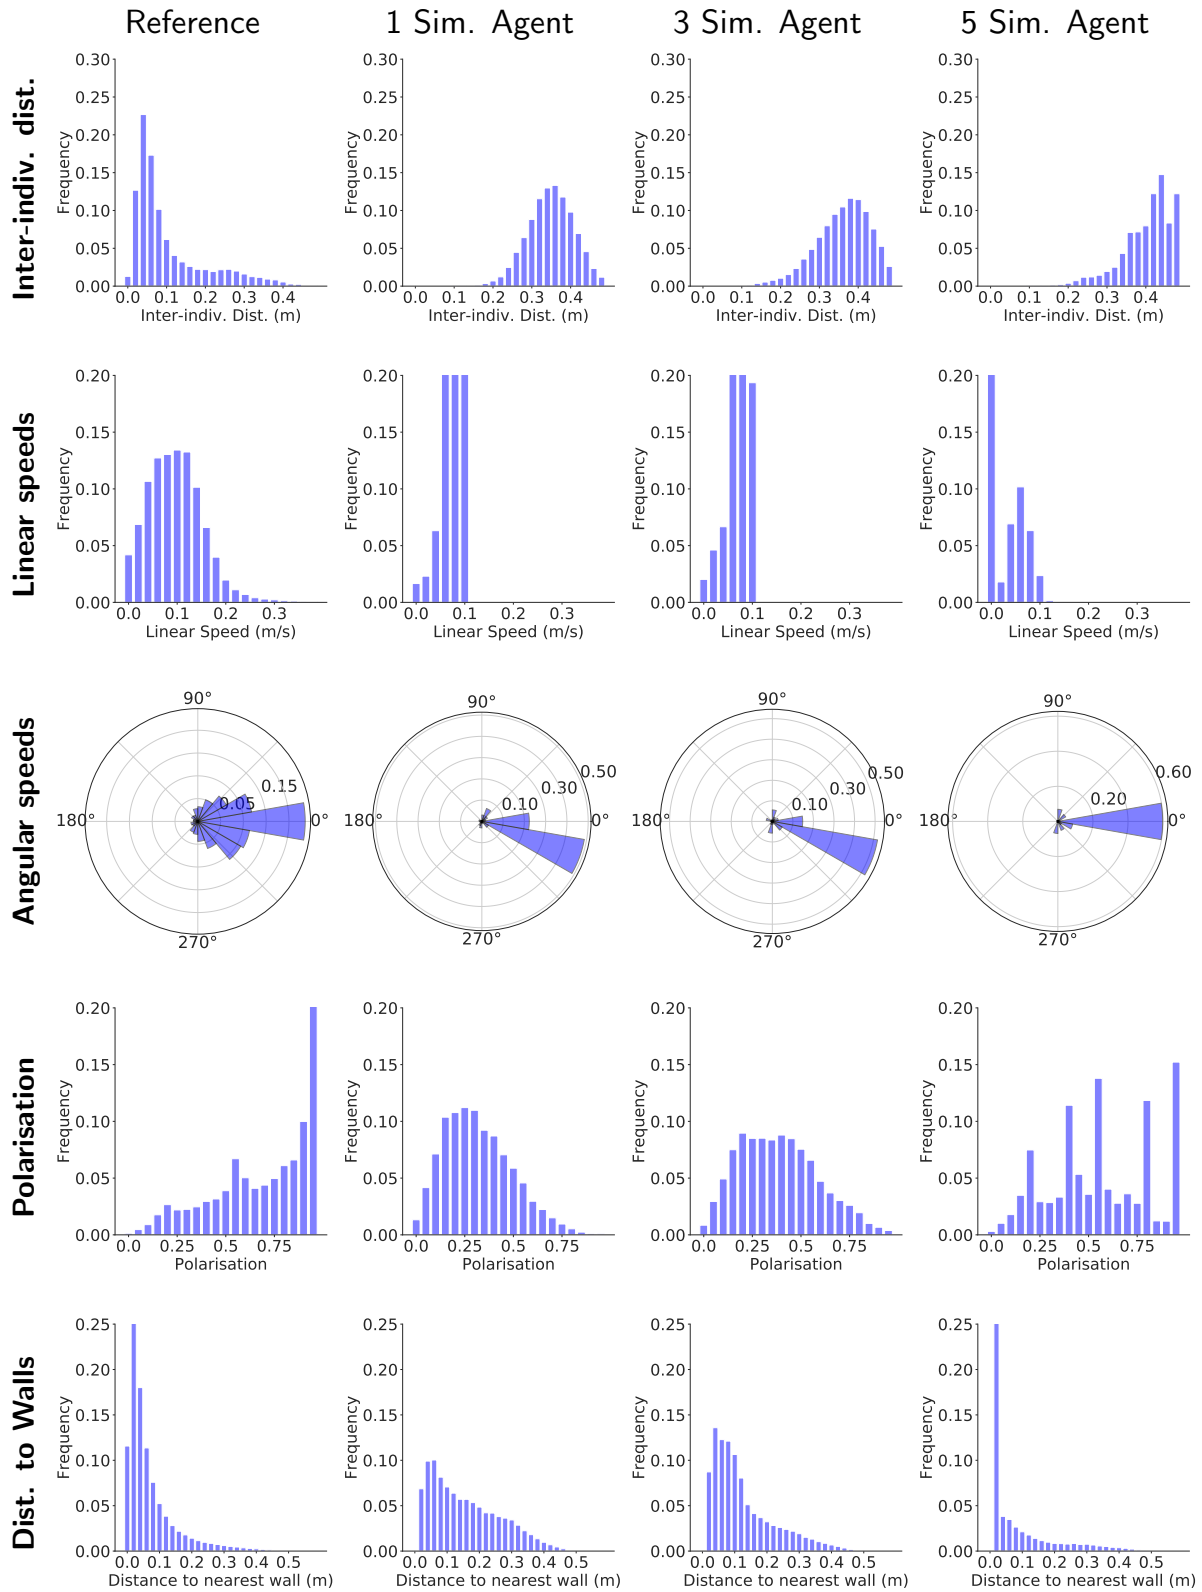
Figure 7. Comparison of agents behaviours between the **Control** case (10 trials of 30-minutes experiments involving 5 zebrafish) and the **NSGA-III-CM** case (10 trials of 30-minutes simulations of 5 agents, with respectively 1,3,5 simulated agents driven by the best-performing MLP controller optimised through the **NSGA-III-CM** method and 4,2,0 agents copying the trajectories of experimental fish). Agents behaviour is quantified across 5 behavioural features: inter-individual distances, linear and angular speeds distributions, polarisation, and distances to nearest wall.

Figure 8. Comparison of agents behaviours between the **Control** case (10 trials of 30-minutes experiments involving 5 zebrafish) and the **NSGA-III-AM** case (10 trials of 30-minutes simulations of 5 agents, with respectively 1,3,5 simulated agents driven by the best-performing MLP controller optimised through the **NSGA-III-AM** method and 4,2,0 agents copying the trajectories of experimental fish). Agents behaviour is quantified across 5 behavioural features: inter-individual distances, linear and angular speeds distributions, polarisation, and distances to nearest wall.

Figure 9. Comparison of agents behaviours between the **Control** case (10 trials of 30-minutes experiments involving 5 zebrafish) and the **NSGA-III-Nov** case (10 trials of 30-minutes simulations of 5 agents, with respectively 1,3,5 simulated agents driven by the best-performing MLP controller optimised through the **NSGA-III-Nov** method and 4,2,0 agents copying the trajectories of experimental fish). Agents behaviour is quantified across 5 behavioural features: inter-individual distances, linear and angular speeds distributions, polarisation, and distances to nearest wall.

Figure 10. Comparison of agents behaviours between the **Control** case (10 trials of 30-minutes experiments involving 5 zebrafish) and the **SL** case (10 trials of 30-minutes simulations of 5 agents, with respectively 1,3,5 simulated agents driven by the best-performing MLP controller optimised through the **SL** method and 4,2,0 agents copying the trajectories of experimental fish). Agents behaviour is quantified across 5 behavioural features: inter-individual distances, linear and angular speeds distributions, polarisation, and distances to nearest wall.
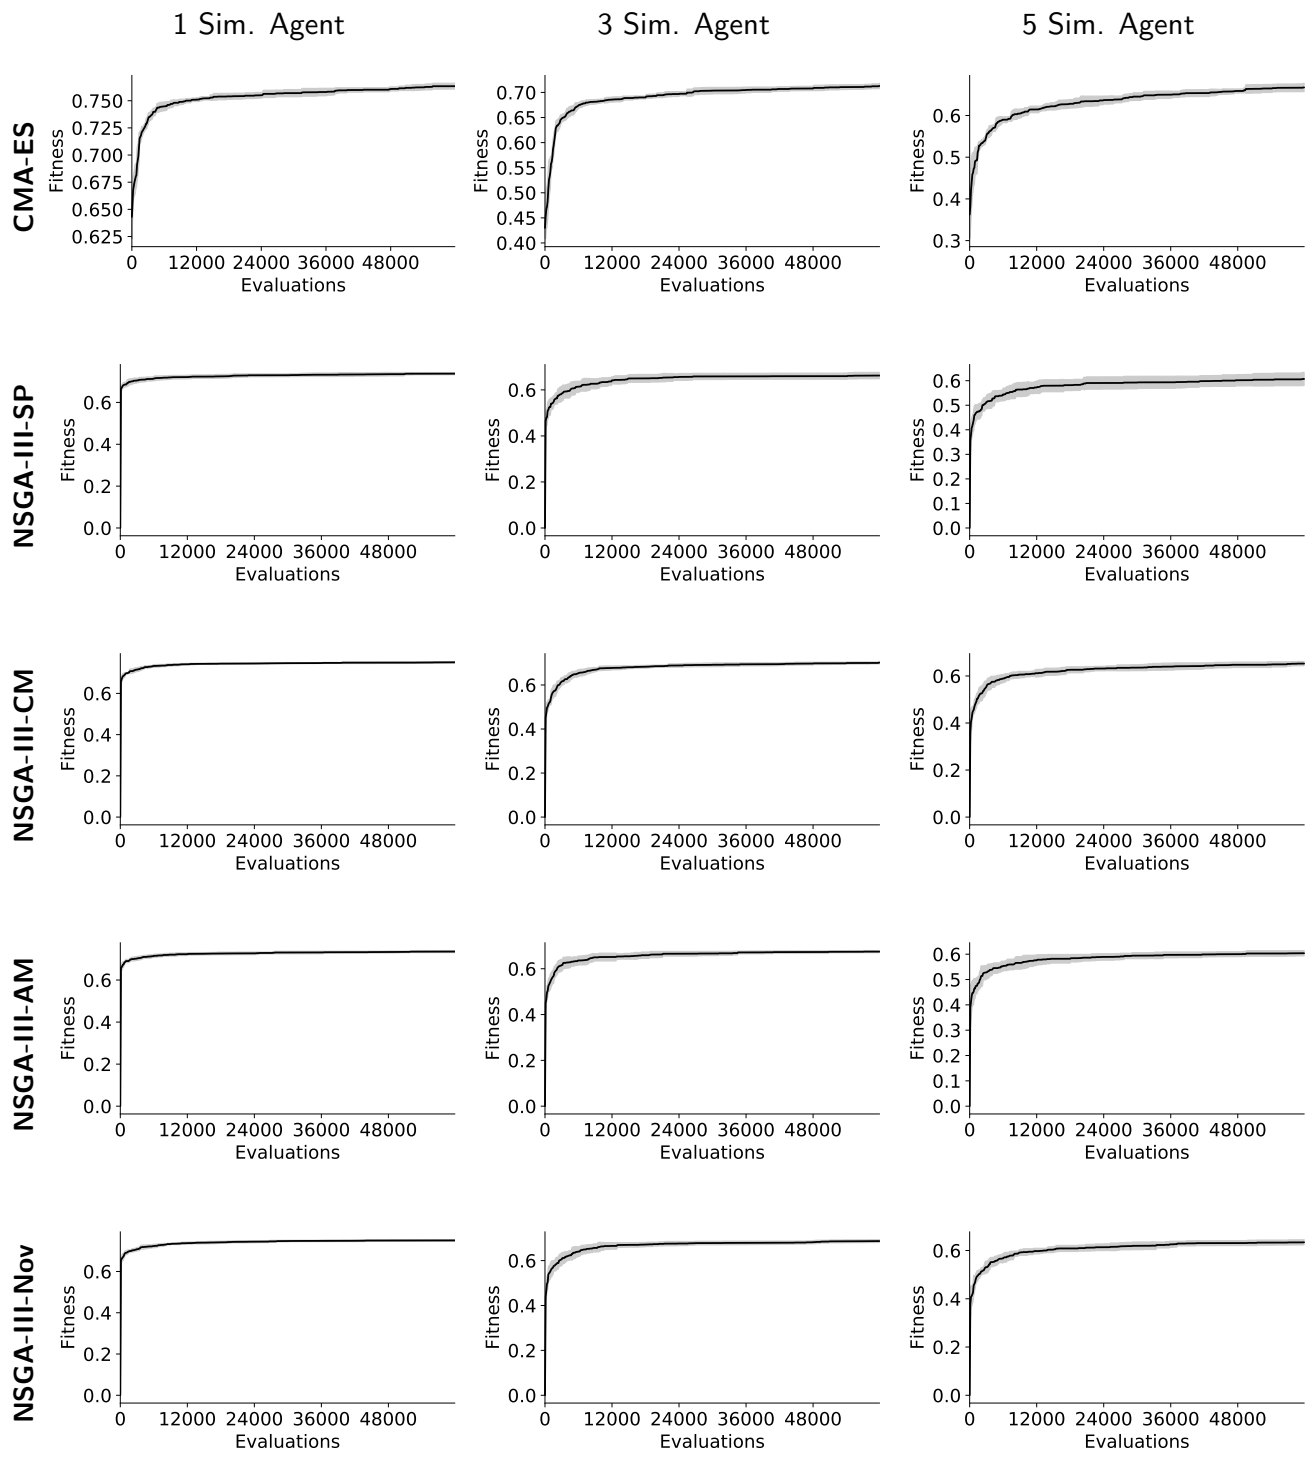
Figure 11. Evolution of the biomimetism score used in the computation of the fitness values across evaluations for all tested optimisation method. Each optimisation method is tested in 10 different runs. Each evaluation involves 10 trials of 5-minutes simulations of 5 agents, with respectively 1,3,5 simulated agents driven by the tested MLP controllers (phenotypes of the evolution process) and 4,2,0 agents copying the trajectories of experimental fish.