



**HAL**  
open science

# Automatic Calibration of Artificial Neural Networks for Zebrafish Collective Behaviours using a Quality Diversity Algorithm

Leo Cazenille, Nicolas Bredeche, José Halloy

► **To cite this version:**

Leo Cazenille, Nicolas Bredeche, José Halloy. Automatic Calibration of Artificial Neural Networks for Zebrafish Collective Behaviours using a Quality Diversity Algorithm. Living Machines Conference, 2019, Nara, Japan. hal-03175254

**HAL Id: hal-03175254**

<https://hal.sorbonne-universite.fr/hal-03175254v1>

Submitted on 19 Mar 2021

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Automatic Calibration of Artificial Neural Networks for Zebrafish Collective Behaviours using a Quality Diversity Algorithm

Leo Cazenille<sup>1</sup>, Nicolas Bredeche<sup>2</sup>, and José Halloy<sup>3</sup>

<sup>1</sup> Ochanomizu University, Department of Information Science, Tokyo, Japan

<sup>2</sup> Sorbonne Université, CNRS, ISIR, F-75005 Paris, France

<sup>3</sup> Univ Paris Diderot, Sorbonne Paris Cité, LIED, UMR 8236, 75013, Paris, France

**Keywords:** collective behaviour, neural networks, QD-algorithms, CVT-MAP-Elites, bio-hybrid systems, biomimetic, robot, zebrafish, fish

**Abstract.** During the last two decades, various models have been proposed for fish collective motion. These models are mainly developed to decipher the biological mechanisms of social interaction between animals. They consider very simple homogeneous unbounded environments and it is not clear that they can simulate accurately the collective trajectories. Moreover when the models are more accurate, the question of their scalability to either larger groups or more elaborate environments remains open. This study deals with learning how to simulate realistic collective motion of collective of zebrafish, using real-world tracking data. The objective is to devise an agent-based model that can be implemented on an artificial robotic fish that can blend into a collective of real fish. We present a novel approach that uses Quality Diversity algorithms, a class of algorithms that emphasise exploration over pure optimisation. In particular, we use CVT-MAP-Elites [32], a variant of the state-of-the-art MAP-Elites algorithm [25] for high dimensional search space. Results show that Quality Diversity algorithms not only outperform classic evolutionary reinforcement learning methods at the macroscopic level (i.e. group behaviour), but are also able to generate more realistic biomimetic behaviours at the microscopic level (i.e. individual behaviour).

## 1 Introduction

Many models have been proposed for fish collective behaviours and motion [24, 30, 13]. At an early stage, they were developed to model realistic collective motion in computer simulation [28]. Nowadays, most of the models are developed to decipher the interaction rules of the animals and not to replicate their behaviour in autonomous agents be them robots or simulations. It is not clear that they can be used to produce a realistic description of fish collective interactions with collective trajectories [19] similar to the observations. Moreover, most of the models consider an unbounded homogeneous space that could be the case in pelagic conditions but not in bounded and in-homogeneous environments.

Only a few models consider the walls of the tanks that have a important effect on the fish [21, 9, 3]. In the robotic context, developing bio-mimetic and realistic fish behavioural models that can be implemented in robots are difficult to develop [6, 5]. These issues are related: (i) how can we develop models producing good descriptions of fish collective behaviours and (2) that, when used as controllers, allow fully autonomous agents (robots, simulations) to cope with bounded inhomogeneous environments and social interactions?

For this type of question, currently two kind of modelling methods are pursued to simply take into account the tank walls and the social context. The first one is equation-based. Equations for the motion of the individuals are developed and calibrated on experimental data [21, 3]. It has been shown that they give excellent results for groups of two fish (*Hemmigramus blerei*) in a circular bounded environment [3]. It remains to demonstrate that such method is scalable for groups made of more than two individuals and more elaborate set-ups. The second kind of modelling technique is agent based. For example, we have developed agent based models that take into account bounded in-homogeneous environment and the social context of the fish [9, 6]. However, agent based models become rapidly complicated as the number of variables and parameters increases. The scalability of this modelling technique remains also an issue.

Here we explore how to develop scalable effective models to generate robot controllers producing realistic collective behaviours. We do not look for understanding specific collective behaviour mechanisms. In recent works, we explored the use of artificial neural network models (multilayer perceptrons) to generate realistic collective motion and trajectories of a group of five zebrafish in a bounded environment [7, 8]. We compared supervised learning and reinforcement learning techniques to optimise the behaviour of artificial Zebrafish, so that they would match the trajectories obtained from real-world experimental data. In this setup, learning a behavioural model is challenging because of the continuous state and action spaces as well as the lack of a world model. We showed that evolutionary reinforcement learning, *i.e.* a direct policy search method [31, 34], can be used to obtain relevant fish trajectories with respect to individual and collective dynamics, and outperforms results obtained by supervised learning. We also showed that while multi-objective evolutionary optimisation using NSGA-III [35] could provide different results over single objective optimisation using CMA-ES [1], the overall quality of trajectories generated is limited by the multiple aspects of behavioural dynamics to be captured simultaneously: wall-following, aggregation, individual trajectories and group dynamics. As a result, we showed that while the *global* biomimetic score (*i.e.* the aggregation of all criteria) is improved with these methods, there is no guarantee that *all* behavioural features will be optimised. In other words, generated trajectories may display unrealistic behaviours, such as low alignment between individuals or erratic wall-following behaviours, while matching real world data in term of inter-individual distances.

In order to improve the quality of biomimetic behavioural strategies, we propose to favour exploration over pure optimisation by using Quality-Diversity

(QD) algorithms [27, 12]. These algorithms are particularly successful in evolutionary robotics problems [25, 11, 15], either by improving diversity to overcome deceptive search spaces [23], or by generating a large repertoire of solutions instead of just one single solution [25]. In the current setup (Fig. 1), we enforce diversity to guide the search by exploring trade-offs between overall quality, which results from aggregating different criteria, and unique realistic behavioural traits, which focus on specific behavioural features, in this case: (1) inter-individual distances between agents, (2) polarisation of the agents in the group, (3) distribution of agent linear speed and (4) probability of presence in the arena. We use CVT-MAP-Elites [32], a variant of the MAP-Elites algorithm [25] using centroidal Voronoi tessellations to tackle high-dimensional feature spaces. CVT-MAP-Elites makes it possible to explore a range of both diverse and high-performing solutions by partitioning the search space into geometric regions according to features predefined by the user. It is then possible to find solutions that can be very different from one another.

We show that CVT-MAP-Elites outperforms state-of-the-art evolutionary optimisation methods (CMA-ES and NSGA-III) for revealing biomimetic behavioural strategies in a fish collective. Even more interestingly, we show that trajectories generated by individuals obtained with CVT-MAP-Elites are also more realistic (when compared to actual data from the fish) at the *microscopic* scale, with realistic behaviours at the level of the individuals. Quality Diversity algorithms offer a promising alternative to classical evolutionary optimisation and reinforcement learning algorithms with respect to learning biomimetic controller for artificial fish.


## 2 Methods

### Experimental set-up

We apply the same experimental method, fish handling and set-up as in [6, 29, 7, 8]. During experiments, fish are placed in an immersed square white plexiglass arena of  $1000 \times 1000 \times 100$  mm. An overhead camera records a video of the experiment at 15 FPS with a  $500 \times 500$ px resolution. It is then analysed to track the fish positions. Experiments were carried out with 10 groups of 5 adult (6-12 months old) wild-type AB zebrafish (*Danio rerio*) in ten 30-minutes trials as in [6, 29]. Experiments conducted in this study were performed 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.

### Artificial neural network model

Artificial neural networks (ANN) are universal function approximators able to model phenomena with *a priori* information. They were used in previous studies [7, 8, 20] to model fish collective behaviour and generate biomimetic trajectories of fish in groups. However this problem is challenging, and it is still possible



**Fig. 1:** Description of the presented methodology to calibrate artificial neural networks to generate fish trajectories. We apply CVT-MAP-Elites [32], a quality-diversity algorithm [27], to optimise the weights of a Multilayer Perceptron (MLP, 1 hidden layer, 10 neurons) that drive 5 fish-like agents in simulations. Simulated agents trajectories are compared to experimental fish trajectories. The fitness function corresponds to the biomimetism score of simulated agent groups. CVT-MAP-Elites is compared to CMA-ES [1] as in [7, 8].

to improve upon the biomimetism of resulting trajectories. Our methodology builds on Cazenille *et al.* [8] and calibrates Multilayer Perceptron (MLP) [2] artificial neural networks to drive simulated fish-like agents in groups of 5 individuals. All simulations involve 5 simulated agents driven by the optimised MLP (see workflow on Fig 1).

MLP are a class of feedforward artificial neural networks. They can be employed in a wide variety of modelling and control tasks [26]. As in [7, 8], our approach uses MLP with one hidden layer of 10 neurons with a hyperbolic tangent activation function. We use this simple and limited ANN as a baseline for benchmarking the various optimisation algorithms.

Table 1 lists the parameters used as inputs and outputs of the MLP controllers for each simulated focal agent. The 20 inputs parameters are often used in multi-agent models of animal collective behaviour [13, 30], and can arguably be considered to be sufficient to model fish groups trajectories. As we consider fish trajectories observed in a bounded environment, we also take into account the presence of walls, which is often ignored in models of fish behaviour, and only found in a small number of recent studies [9, 3, 7, 6, 8].

## Data analysis

As in [7, 8], we analyse the tracked positions of agents in each trial  $e$  (experiments or simulations) and compute several behavioural metrics: (i) the distribution of

## Inputs

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

## Outputs

Name	#Param.	Description
Delta linear speed	1	Change of inst. linear speed of the FA from the prev. time-step
Delta angular speed	1	Change of inst. angular speed of the FA from the prev. 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. Here, FA refers to the focal agent.

*inter-individual distances* between agents ( $D_e$ ); (ii) the distributions of *instant linear speeds* ( $L_e$ ); (iii) the distribution of *polarisation* of the agents in the group ( $P_e$ ); (iv) the *probability of presence of agents in the arena* ( $E_e$ ). The polarisation of an agent group assesses the extent to which fish are aligned. It corresponds to the absolute value of the mean agent heading:  $P = \frac{1}{N} |\sum_{i=1}^N u_i|$  where  $u_i$  is the unit direction of agent  $i$  and  $N = 5$  is the number of agents [33]. Recent studies introduced more complex metrics to assess fish behaviour, like 2D features maps of neighbours compared to a focal fish used in [22, 18]. Our approach here aims to provide a simple methodological baseline, so we only take into account simple and established behavioural metrics like polarisation and inter-individual distances. While more complex metrics based on 2D features maps could describe more accurately fish collective dynamics, they may also require quantities with higher dimensionality than simple metrics, which may make their synthesis into behavioural scores more difficult.

We quantify the realism of the simulated fish groups by computing a **biomimetism score** of their behaviour, as in [6–8]. It measures the similarity between behaviours exhibited by the simulated fish group and those exhibited by 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 the geometric mean of the other behavioural scores:

$$S(e, e_c) = \sqrt[4]{I(L_e, L_{e_c})I(D_e, D_{e_c})I(P_e, P_{e_c})I(E_e, E_{e_c})} \quad (1)$$

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 [14]. 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. As opposed to [8], we do not take into account the distribution of angular speeds

in the computation of the fitness. Indeed, the distributions of angular speeds of evolved individuals was always similar to the ones from random individuals. Thus, we removed this behavioural metrics from the features taken into account to reduce the dimensionality of the feature space.

## Optimisation and illumination

We calibrate the weights of the MLP models driving agent behaviour to approximate as close as possible the trajectories and behaviours of groups of 5 fish-like agents, as in [7, 8, 5]. Simulations have a duration of 30 minutes (15 time-steps per seconds, *i.e.* 27000 steps per simulation).

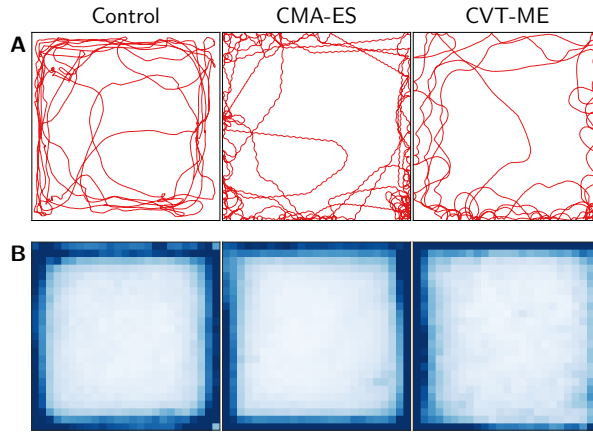
In previous studies [7, 8], we optimised these MLP controllers using evolutionary algorithms: CMA-ES [1] and NSGA-III [35].

Here, we use the CVT-MAP-Elites [32] QD algorithm, a variant of the popular MAP-Elites [25] algorithm, to search for interesting MLP controllers matching experimental fish trajectories across a user-provided space of features. The family of Map-Elites algorithms is based on the idea of exploring a clustered search space, retaining the best candidate solutions for each cluster. Clusters correspond to specific range of values for pre-defined features and each candidate solution is stored in a cell of a so-called map, which corresponds to its cluster. The seminal MAP-Elites algorithm uses a pre-defined clustering of the feature space, with the number of clusters (or "bins") quickly exploding as the number of feature dimensions considered grows. In order to tackle high-dimensional feature space, the CVT-MAP-Elite algorithm defines clusters as centroids of Voronoi tessellation, where centroids can be automatically positioned during exploration.


In our case, these features correspond to the four behavioural metrics  $L_e$ ,  $D_e$ ,  $P_e$ ,  $E_e$  presented earlier. CVT-MAP-Elites is capable of handling high dimensional feature spaces (like our case) by using centroidal Voronoi tessellations to reduce the dimensionality of the feature space. Here, the CVT-MAP-Elites case only consider 32 bins of elites, which is far lower as what would be used with MAP-Elites in a reasonable configuration (*e.g.* with 32 bins per features, it would correspond to a grid with  $32 \times 32 \times 32 \times 32 = 33554432$  bins of elites). We selected empirically 32 bins of elites in the CVT-MAP-Elites methods because it produced the best-performing results among tested numbers of bins.

We compare the generated trajectories using CVT-MAP-Elites with previous results from [7, 8] where MLP controllers were optimised by the CMA-ES [1]. CMA-ES is a popular mono-objective global optimiser capable of handling problems with noisy, ill-defined fitness function.

In all cases, the algorithms aim to maximise the biomimetism score ( $S_{e_o, e_c}$ ) of MLP-driven agents in simulations ( $e_o$ ) compared to experimental fish groups ( $e_c$ ). Both cases are tested in 10 different trials with the same budget of objective function evaluation (one simulation corresponds to one function evaluation): 60000 evaluations. The CVT-MAP-Elites case involves 6000 evaluations in the initial batch, and 450 batches of 120 individuals. The CMA-ES case involves 500 generations of 120 individuals.

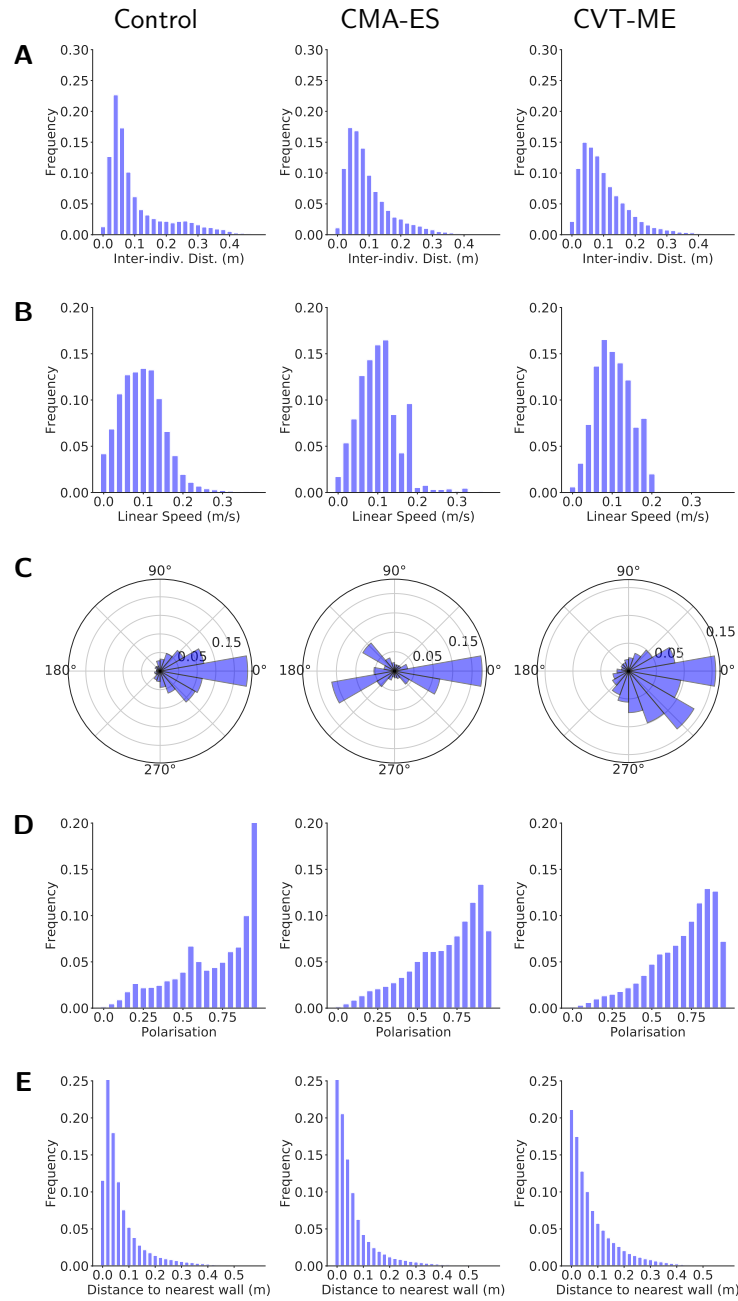


**Fig. 2:** Agent trajectories in the square (1m) experimental arena after 30-minute trials, for all considered cases: **Control** reference experimental fish data obtained as in [9, 29], **CVT-MAP-Elites** and **CMA-ES** corresponding to simulated MLP-driven agents. **A** Examples of an individual trajectory of one agent among the 5 making the group (fish or simulated agent) during 1 minute out of a 30-minute trial. **B** Presence probability density of agents in the arena.



**Fig. 3:** Similarity scores between the trajectories of the experimental fish groups (Control) and those of the best-performing simulated individuals optimised by CVT-MAP-Elites or CMA-ES. All cases are tested across 10 different trials (experiments or simulations). Four behavioural features are considered to quantify the realism of exhibited behaviours. **Inter-individual distances** measures the similarity in distribution of inter-individual distances between all agents and corresponds to the capabilities of the agents to aggregate. **Linear speed distribution** measures to the distributions of linear speeds of the agents. **Polarisation** measures how aligned the agents are in the group. **Probability of presence** corresponds to the density of agent presence in each part of the arena (cf Fig. 2B). The **Biomimetic score** is computed as the geometric mean of the other scores.





**Fig. 4:** Behavioural comparison between ten 30-minute trials of experimental fish in groups of 5 and the 5-sized simulated fish groups for both tested cases. The following behavioural features are examined: inter-individual distances (**A**), linear (**B**) and angular (**C**) speeds distributions, polarisation (**D**), and distances to nearest wall (**E**). Note that distributions of angular speeds and distances to nearest wall are informative and not used in the calibration process.

We use a CVT-MAP-Elites implementation from the QDpy (Quality Diversity in Python) framework [4]. The CMA-ES implementation is based on the DEAP library [16].

### 3 Results

We analyse the behaviour of the simulated agent groups for the CVT-MAP-Elites and CMA-ES cases and compare them with the behaviour of experimental fish groups (Control case). In both cases, the agents are driven by MLP controllers, calibrated either by CVT-MAP-Elites or with CMA-ES to match as close as possible the behaviour of experimental fish across the behavioural metrics presented above. Each case is repeated in 10 trials and the following statistics only consider the best-evolved MLP controllers.

Figure 2A provides examples of agents trajectories. In the control case, fish tend to follow walls but retain a capability to go to the center of arena. This is also observed in trajectories from both MLP-driven cases. However, they also incorporate patterns not found in actual fish trajectories. Small circular loops can appear in both cases. A small periodic "shaking" is present in the trajectories of the CMA-ES case. Conversely, the trajectories of the CVT-MAP-Elites appear smoother and match more closely those of the experimental fish. This suggests that CVT-MAP-Elites is more realistic at the microscopic level of agent trajectories. Figure 2B presents the mean probability of presence of all agents in the arena for all cases.

We assess the realism of the two tested cases by computing the behavioural metrics presented in Sec. 2. These metrics serve as a base to compute similarity scores between the tested cases and experimental fish behaviour (Fig. 3). Both simulated cases display lower similarity scores than the experimental fish groups. Based on a comparison of the best solutions found by both algorithms, CVT-MAP-Elites outperforms CMA-ES with statistical significance (p-value=0.0227 using the Mann-Whitney U-test). The best solution found by CVT-MAP-Elites also dominates all solutions found with CMAE-ES (best fitness: 0.724 with CVT-MAP-Elites *vs.* 0.704 with CMA-ES).

However, the controllers optimised by the two methods prioritise different features. The CVT-MAP-Elites case shows higher scores on inter-individual distances and polarisation than the CMA-ES case. In turn, CMA-ES exhibits higher probability of presence scores than the CVT-MAP-Elites case. Scores of linear speeds are roughly similar between the two cases. Overall, it means that the controllers optimised by the two methods exhibit different kind of behaviours and way of coping with the trade-offs between fish aggregative and wall-following behaviours. In term of group dynamics, the solutions of the CVT-MAP-Elites case are more cohesive than what is seen in the CMA-ES case, which evolves controllers that are more biased towards wall-following than group aggregation.

Histograms of all behavioural metrics are shown for all cases in Fig. 4, with two complementary metrics: the distribution of angular speeds (Fig. 4, related to polarisation) and distance to nearest wall (Fig. 4, related to probability of

presence). They confirm the results from Fig. 3. The distributions of angular speed (Fig. 4C) of both cases are sub-optimal in term of realism. Figure 4E displays that simulated agents of both cases tend to exhibit correctly a wall-following behaviour.

The experimental fish groups of the **Control** case display a large behavioural variability across all investigated metrics (Fig. 3 and 4). Indeed, experiments were conducted with 10 groups of 5 fish (totalling 50 different fish) displaying disparate behaviours and individual preferences. This matches results from previous zebrafish collective behaviours studies [29, 10]. Social (group composition) and environmental contexts impact fish behaviour: fish tend to aggregate in small short-lived sub-groups that follow walls from a distance that vary according to group composition. They also tend to exhibit an uniform degree of alignment within sub-groups.

## 4 Discussion and Conclusion

Calibrating artificial neural networks to model the collective behaviour of fish group and generate realistic fish trajectories is a challenging problem because fish behaviours involve several complementary dynamics with trade-offs between group-level dynamics (aggregative tendencies, group alignment), individual-level behaviours (agent linear speed) and response to environmental cues (wall-following behaviour, probability of presence in the arena). It is difficult to balance these conflicting behaviours during the calibration process.

Here, we show that the CVT-MAP-Elites [32], a quality diversity method that emphasises exploration over pure optimisation, calibrates controllers that are more realistic in term of agent groups polarisation and inter-individual distances when compared to previous results using stochastic optimisation methods such as the CMA-ES evolutionary method [7, 8]. Moreover, QD algorithms also have the advantage of exploring a range of diverse solutions instead of searching for a single local optimum, and could be used to decipher the interrelation between features and behavioural biomimetism in order to draw biological conclusions.

Our approach could still be improved further, either by taking into account more behavioural metrics (tangential and normal accelerations, curvature or tortuosity) or by using more complex artificial neural networks than MLP, such as recurrent neural networks or deep neural networks.

Additionally, our methodology could be adapted to make possible to derive biological conclusions from the calibrated ANN models. ANN can be used as benchmarks to find the necessary information in experimental data to replicate experimental fish behaviour. Recently, Heras *et al.* [17] hinted at the possibility of this approach to decipher the interaction mechanism in large zebrafish groups. It remains to be shown that such ANN models can also produce collective trajectories similar to those observed experimentally. If it is shown to be the case, best-performing agents optimised through such methodology could be used as controllers to drive the behaviour of robots interacting experimentally with fish to study their collective dynamics.

## Acknowledgement

This work was funded by EU-ICT project 'ASSISIBf', no 601074.

## References

1. Auger, A., Hansen, N.: A restart CMA evolution strategy with increasing population size. In: *Evolutionary Computation, 2005. The 2005 IEEE Congress on.* vol. 2, pp. 1769–1776. IEEE (2005)
2. Bishop, C.: *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg (2006)
3. Calovi, D.S., Litchinko, A., Lecheval, V., Lopez, U., Escudero, A.P., Chaté, H., Sire, C., Theraulaz, G.: Disentangling and modeling interactions in fish with burst-and-coast swimming reveal distinct alignment and attraction behaviors. *PLoS computational biology* 14(1), e1005933 (2018)
4. Cazenille, L.: Qdpy: A python framework for quality-diversity. <https://gitlab.com/leo.cazenille/qdpy> (2018)
5. Cazenille, L., Chemtob, Y., Bonnet, F., Gribovskiy, A., Mondada, F., Bredeche, N., Halloy, J.: 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*. pp. 107–118. Springer (2017)
6. Cazenille, L., Collignon, B., Bonnet, F., Gribovskiy, A., Mondada, F., Bredeche, N., Halloy, J.: How mimetic should a robotic fish be to socially integrate into zebrafish groups ? *Bioinspiration & biomimetics* (2017)
7. 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*. pp. 85–96. Springer (2018)
8. Cazenille, L., Bredeche, N., Halloy, J.: Modelling zebrafish collective behaviours with multilayer perceptrons optimised by evolutionary algorithms. *arXiv preprint arXiv:1811.11040* (2018)
9. Collignon, B., Séguret, A., Halloy, J.: A stochastic vision-based model inspired by zebrafish collective behaviour in heterogeneous environments. *Royal Society Open Science* 3(1), 150473 (2016)
10. Collignon, B., Séguret, A., Chemtob, Y., Cazenille, L., Halloy, J.: Collective departures in zebrafish: profiling the initiators. *arXiv preprint arXiv:1701.03611* (2017)
11. Cully, A., Clune, J., Tarapore, D., Mouret, J.B.: Robots that can adapt like animals. *Nature* 521(7553), 503 (2015)
12. Cully, A., Demiris, Y.: Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation* 22(2), 245–259 (2018)
13. Deutsch, A., Theraulaz, G., Vicsek, T.: Collective motion in biological systems. *Interface Focus* 2(6), 689 (2012)
14. Deza, M., Deza, E.: *Dictionary of distances*. Elsevier (2006)
15. Duarte, M., Gomes, J., Oliveira, S.M., Christensen, A.L.: Evolution of repertoire-based control for robots with complex locomotor systems. *IEEE Transactions on Evolutionary Computation* 22(2), 314–328 (2018)
16. Fortin, F.A., De Rainville, F.M., Gardner, M.A., Parizeau, M., Gagné, C.: Deap: Evolutionary algorithms made easy. *Journal of Machine Learning Research* 13(Jul), 2171–2175 (2012)

17. Heras, F.J., Romero-Ferrero, F., Hinz, R.C., de Polavieja, G.G.: Aggregation rule in animal collectives dynamically changes between majority and minority influence. *bioRxiv* p. 400747 (2018)
18. Heras, F.J., Romero-Ferrero, F., Hinz, R.C., de Polavieja, G.G.: Deep attention networks reveal the rules of collective motion in zebrafish. *bioRxiv* p. 400747 (2018)
19. Herbert-Read, J.E., Romenskyy, M., Sumpter, D.J.: A turing test for collective motion. *Biology letters* 11(12), 20150674 (2015)
20. Iizuka, H., Nakamoto, Y., Yamamoto, M.: Learning of individual sensorimotor mapping to form swarm behavior from real fish data. In: *Artificial Life Conference Proceedings*. pp. 179–185. MIT Press (2018)
21. Jeanson, R., Blanco, S., Fournier, R., Deneubourg, J., Fourcassié, V., Theraulaz, G.: A model of animal movements in a bounded space. *Journal of Theoretical Biology* 225(4), 443–451 (2003)
22. Jiang, L., Giuggioli, L., Perna, A., Escobedo, R., Lecheval, V., Sire, C., Han, Z., Theraulaz, G.: Identifying influential neighbors in animal flocking. *PLoS computational biology* 13(11), e1005822 (2017)
23. Lehman, J., Stanley, K.O., Miikkulainen, R.: Effective diversity maintenance in deceptive domains. In: *Proceedings of the 15th annual conference on Genetic and evolutionary computation*. pp. 215–222. ACM (2013)
24. Lopez, U., Gautrais, J., Couzin, I.D., Theraulaz, G.: From behavioural analyses to models of collective motion in fish schools. *Interface focus* 2(6), 693–707 (2012)
25. Mouret, J.B., Clune, J.: Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015)
26. 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)
27. Pugh, J.K., Soros, L.B., Stanley, K.O.: Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI* 3, 40 (2016)
28. Reynolds, C.: Flocks, herds, and schools: a distributed behavioral model. *Computer Graphics* 21, 25–34 (1987)
29. 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)
30. Sumpter, D.J., Mann, R.P., Perna, A.: The modelling cycle for collective animal behaviour. *Interface focus* 2(6), 764–773 (2012)
31. Sutton, R.S., Barto, A.G.: *Reinforcement learning: An introduction* (2018)
32. Vassiliades, V., Chatzilygeroudis, K., Mouret, J.B.: Using centroidal voronoi tessellations to scale up the multidimensional archive of phenotypic elites algorithm. *IEEE Transactions on Evolutionary Computation* 22(4), 623–630 (2018)
33. Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., Shochet, O.: Novel type of phase transition in a system of self-driven particles. *Physical review letters* 75(6), 1226 (1995)
34. Whiteson, S.: Evolutionary computation for reinforcement learning. In: *Reinforcement learning*, pp. 325–355. Springer (2012)
35. 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*. pp. 661–668. ACM (2014)