Research

Authors for correspondence:
Vaios Papaspyros
e-mail: Vaios.Papaspyros@epfl.ch
Clément Sire
e-mail: Clement.Sire@univ-tlse3.fr

Predicting the long-term collective behaviour of fish pairs with deep learning

Vaios Papaspyros[1], Ramón Escobedo[3], Alexandre Alahi[2], Guy Theraulaz[3], Clément Sire[4] and Francesco Mondada[1]

[1]Mobile Robotic Systems (Mobots) group, Institute of Electrical and Micro Engineering, and [2]VITA group, Civil Engineering Institute, École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland
[3]Centre de Recherches sur la Cognition Animale, Centre de Biologie Intégrative, CNRS, and
[4]Laboratoire de Physique Théorique, CNRS, Université de Toulouse III – Paul Sabatier, 31062 Toulouse, France

(iD) VP, 0000-0001-8257-1891; RE, 0000-0002-9945-0460; AA, 0000-0002-5004-1498; GT, 0000-0002-3062-4215; CS, 0000-0003-4089-4013; FM, 0000-0001-8641-8704

Modern computing has enhanced our understanding of how social interactions shape collective behaviour in animal societies. Although analytical models dominate in studying collective behaviour, this study introduces a deep learning model to assess social interactions in the fish species *Hemigrammus rhodostomus*. We compare the results of our deep learning approach with experiments and with the results of a state-of-the-art analytical model. To that end, we propose a systematic methodology to assess the faithfulness of a collective motion model, exploiting a set of stringent individual and collective spatio-temporal observables. We demonstrate that machine learning (ML) models of social interactions can directly compete with their analytical counterparts in reproducing subtle experimental observables. Moreover, this work emphasizes the need for consistent validation across different timescales, and identifies key design aspects that enable our deep learning approach to capture both short- and long-term dynamics. We also show that our approach can be extended to larger groups without any retraining, and to other fish species, while retaining the same architecture of the deep learning network. Finally, we discuss the added value of ML in the context of the study of collective motion in animal groups and its potential as a complementary approach to analytical models.

1. Introduction

Collective behaviour in animal groups is a very active field of research, studying the fundamental mechanisms by which individuals coordinate their actions [1–3] and self-organize [4,5]. One of the most common forms of collective behaviour can be observed in schools of fish and flocks of birds that have the ability to coordinate their movements to collectively escape predator attacks or improve their foraging efficiency [6,7]. This coordination at the group level mainly results from the social interactions between individuals. Important steps to understand these collective phenomena consist in characterizing these interactions and understanding the way individuals integrate interactions with other group members [8–12].

New tracking techniques and tools for behavioural analysis have been developed that have greatly improved the quality of collective motion data [13–19]. In particular, advances in computing have allowed the development of computationally demanding data-oriented model generation techniques [12,20–24] and the subsequent simulation of biological models [25]. This has resulted in more realistic models that attempt to recover the social interactions that govern collective behaviours. Yet, the bottleneck with most of these approaches is that they rely on demanding and laborious mathematical work to obtain the interactions from experimental data.

An alternative to such analytical models is to exploit machine learning (ML) techniques and let an algorithm learn the interactions directly from data. The know-how required to use these techniques is different from the one needed to design analytical models. Nevertheless, the structure of ML algorithms, here a neural network, has an impact on the modelling performance, and requires specific expertise [26]. Once the architecture of an ML algorithm is set, ML can often process data for different species without structural adaptation, and generate new models quickly. This is very different from analytical models, where each new species requires redefining the model nearly from scratch. The downside of this flexibility is that ML models are usually less explainable (black box). Yet, recent ML algorithms provide higher-level information mappable to more tangible formats, such as force maps, which show the strength and direction of behavioural changes experienced by an individual when interacting with other individuals in a moving group [23,24]. Despite their limited explainability, ML algorithms require only a few biological assumptions. They offer an almost hypothesis-free procedure [27] that can even outperform human experts in detecting subtle patterns [28], making ML a very appealing complementary approach to analytical models.

For both analytical and ML models, several studies evaluate models over *short timescales* and through instantaneous quantities such as speed, acceleration, distance and angle to objects [22,29], or by measuring the error between predictions and ground truth [23,30,31]. Only more recently, long timescales have also been considered [21]. However, a model that performs well at short timescales compared with experiment does not necessarily perform well at long timescales. This is especially true for models that try to reproduce complex collective phenomena in living systems. To our knowledge, the predictive capacity of ML models in this context has not been evaluated over both short *and* long timescales, that is, their ability to generate synthetic data that replicate the outcomes of social interactions over both timescales.

Here, we demonstrate that ML models can generate realistic synthetic data with minimal biological assumptions, and that they allow to accelerate and generalize the process of collective behaviour modelling. More specifically, we present a social interaction model using a deep neural network that captures both the short- and long-term dynamics observed in schooling fish. We apply our approach to pairs of rummy-nose tetra (*Hemigrammus rhodostomus*) swimming in a circular tank, and show that it can also be applied to fish species with similar burst-and-coast swimming (zebrafish; *Danio rerio*). Our ML model is benchmarked against the state-of-the-art analytical model for this species [32], showing that it performs as well as the latter, even for very subtle quantities measured in the experiments. Moreover, we also introduce a systematic methodology to stringently test the results of an analytical or ML model against experiment, at different timescales, and in the context of animal collective motion.

2. Methods

2.1. Experimental data

The trajectory data used in this study were originally published in [12] for *Hemigrammus rhodostomus* swimming either alone or in pairs in a circular tank of radius 25 cm. This species is characterized by a burst-and-coast swimming mode, where the fish perform a succession of sudden and short acceleration periods (of typical duration 0.1 s), each followed by a longer gliding period almost in a straight line, resulting in a mean total duration of the kicks of 0.6 s. The instant of the kicks, when heading changes take place, are assimilated to decision instants [12].

The dataset corresponds to 15 h of video recordings at 25 Hz. Fish are tracked with idTracker [17], an image analysis software which extracts the two-dimensional trajectories of all individuals. Occasionally, the tracking algorithm is temporarily unable to report positions accurately. This can be due to small fluctuations in lighting conditions, fish standing still or moving at very low speed, fish swimming very close to the surface, to the border or to each other. These instances are corrected using several filtering processes. Since our analyses focus on social interactions, we remove the periods during which fish are inactive. Fish body length (BL) is approximately 3.5 cm, and the intervals of time during which fish velocity is less than 1 BL s^{-1} are removed. Large leaps in fish trajectories during which fish move by more than 1.5 BL $\approx$ 5.25 cm between two consecutive frames, meaning that fish move at almost 65 cm s^{-1}, are also identified and removed, as they result from tracking errors. Finally, missing points are filled by linear interpolation. The final dataset used in this work represents approximately 4 h of trajectory data for pairs of *H. rhodostomus*.

Moreover, trajectories of the original dataset have been resampled with a timestep of $\Delta t = 0.12$ s instead of the original 0.04 s provided by the camera, and data points have been converted from pixel space to a normalized $[-1, 1]$ range to facilitate the training of the networks. This subsampling helps to reduce the random noise between subsequent camera frames at the very short timescale of 0.04 s (especially for measuring fish headings and speeds), while maintaining a sufficiently small timestep to study and model the social interactions. The new timestep $\Delta t = 0.12$ s is of the same order as the sudden acceleration period of a kick and approximately one-fifth of the average total kick duration [12]. In addition to reducing the noise, the subsampling also reduces the dimension of the input vector and of the effective size of the training dataset and, as a result, of the training time for the ANN models presented in this work.

2.2. Quantification of individual and collective behaviour in pairs of fish

We use a set of observables to quantify how close the results of the models are from the measures obtained in the experiments [12,20,21]. These observables constitute a stringent benchmarking and validation when designing and testing a model. In the case of deep learning techniques, those observables also serve as means to partially explain what the algorithm has learned.

Let us first define the temporal variables characterizing the individual and collective behaviour of the fish. Figure 1a shows two fish swimming in a circular tank of radius $R = 25$ cm. The position vector of a fish i at time t is given by its Cartesian coordinates $\boldsymbol{u}_i(t) = (x^i(t), y^i(t))$ in the system of reference, centred at the centre of the tank $C(0, 0)$. The components of the velocity vector $\boldsymbol{v}_i(t) = (v_x^i(t), v_y^i(t))$ are given by $v_x^i(t) = (u_x^i(t) - u_x^i(t - \Delta t))/\Delta t$ (similar expression for v_y^i). The heading angle of the fish is assumed to indicate its direction of motion and is therefore given by the angle that the velocity vector forms with the horizontal, $\phi_i(t) = \texttt{ATAN2}(v_y^i(t), v_x^i(t))$.

The motion of a given fish i is then described using the three following instantaneous variables: the speed, $V_i(t) = \|\boldsymbol{v}_i(t)\|$, the distance of the fish to the wall, $r_w^i(t) = R - \|\boldsymbol{u}_i(t)\|$, and the angle of incidence of the fish to the wall, $\theta_w^i(t)$, defined by the angle formed by the velocity vector and the normal to the wall: $\theta_w^i(t) = \phi_i(t) - \texttt{ATAN2}(y^i(t), x^i(t))$, see figure 1a.
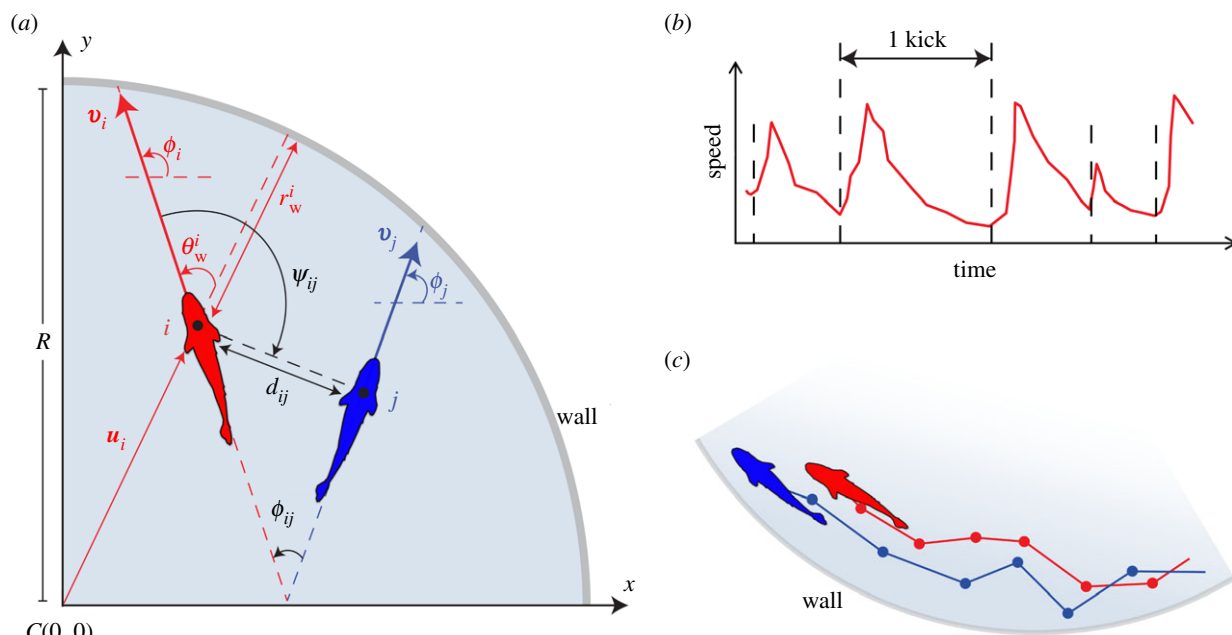
Figure 1. (a) Individual and collective variables characterizing the instantaneous state of an individual (focal fish in red) and its pairwise relation with a neighbour (blue): distance to the wall $r_w^i(t)$, angle of incidence to the wall $\theta_w^i(t)$, heading angle $\phi_i(t)$, distance between individuals $d_{ij}(t)$, difference of heading angles $\phi_{ij}(t)$ and angle of perception $\psi_{ij}(t)$. Positive angles (curved arrows) are defined in the anti-clockwise direction, starting from the positive semi-axis of abscissas. The radius of the circular set-up is $R = 25$ cm. For visualization purposes, the size of fish is not to scale with the tank. (b) Typical profile of the fish speed, $V(t)$, showing the typical sequence of kicks (abrupt accelerations followed by longer gliding phases). (c) Trajectories of two fish close to the wall due to their burst-and-coast swimming mode. The dots in the trajectories denote the instants of the kicks, where fish decision-making is assumed to take place.

When there are two fish i and j in the tank, their relative motion is characterized by means of three variables: the distance between fish, $d_{ij}(t) = \|u_i(t) - u_j(t)\|$, the difference between their heading angles, $\phi_{ij}(t) = \phi_j(t) - \phi_i(t)$, which measures the degree of alignment between both fish, and the angle of view, $\psi_{ij}(t)$, which is the angle with which fish i perceives fish j, and which is generally independent of $\psi_{ji}(t)$. See figure 1a for the graphical representation of these quantities. The angle of perception of the fish also allows us to define the notion of *geometrical leadership* for two fish: fish i is the *geometrical leader* (and therefore, j is the *geometrical follower*), if $|\psi_{ij}(t)| > |\psi_{ji}(t)|$, meaning that i has to turn by a larger angle to face j than the angle that j has to turn to face i. In practice, these definitions of the geometrical leader and follower provide a precise and intuitive characterization of a fish being ahead of the other. Note that being the leader or the follower is an instantaneous state that can change from one kick to the other.

These six quantities $V_i(t)$, $r_w^i(t)$, $\theta_w^i(t)$, $d_{ij}(t)$, $\phi_{ij}(t)$ and $\psi_{ij}(t)$ being defined, the measure of their probability distribution functions (PDF) constitutes a set of observables probing the individual and collective instantaneous fish dynamics in a fine-grained and precise manner. The PDF of $V_i(t)$, $r_w^i(t)$, $\theta_w^i(t)$ probe the behaviour of a focal fish sampled over the observed dynamics, and are hence called *instantaneous individual observables*. The PDF of $d_{ij}(t)$, $\phi_{ij}(t)$ and $\psi_{ij}(t)$ characterize the correlations between two fish *at the same time* t and are hence called *instantaneous collective observables*. These three collective observables can be easily generalized to a group of arbitrary size $N > 2$, by considering i and j as pairs of nearest neighbours, or pairs of second-nearest neighbours (or even farther neighbours), or even averaging them over all pairs in the group (then probing the size, the polarization and the anisotropy of the group). Ultimately, comparing experimental results and model predictions for these individual and collective observables constitutes a stringent test of a model.

Moreover, to characterize the *temporal correlations* arising in the dynamics, we make use of three additional observables involving quantities measured *at two different times*, for a given focal fish [21]: the mean-squared displacement $C_X(t)$, the velocity autocorrelation $C_V(t)$, and, especially challenging, the autocorrelation of the angle of incidence to the wall $C_{\theta_w}(t)$, defined, respectively, by

$$C_X(t) = \langle [u_i(t + t') - u_i(t')]^2 \rangle, \tag{2.1}$$

$$C_V(t) = \langle v_i(t + t') \cdot v_i(t') \rangle \tag{2.2}$$

and

$$C_{\theta_w}(t) = \langle \cos [\theta_w^i(t + t') - \theta_w^i(t')] \rangle, \tag{2.3}$$

where $\langle w(t) \rangle$ is the average of a variable $w(t)$ over all reference times t' (assumption of a stationary dynamics, where correlations between two times depend solely on their time separation), over all focal fish, and over all experimental runs. In principle, these correlation observables can also be generalized to probe the (collective) time correlations between the two different fish (or between nearest neighbours in a group of $N > 2$ individuals). For instance, one could consider $C_{V_{NN}}(t) = \langle v_i(t + t') \cdot v_j(t') \rangle$, where the average is now over nearest neighbour pairs. However, in the present study, we will limit ourselves to the study of the three (individual) correlation functions listed in equations (2.1)–(2.3).

2.3. Analytical and deep learning models of fish behaviour

Many species of fish like *H. rhodostomus* or *Danio rerio* move in a *burst-and-coast* manner, meaning that their swimming pattern consists of a sequence of abrupt accelerations each followed by a longer gliding period (figure 1b), during which a fish moves more or less in a straight line (figure 1c). The kicking instants observed in the curve of the speed can be interpreted as decision times when the fish potentially initiates a

change of direction. In *H. rhodostomus*, the mean time interval between kicks and the typical kick length were experimentally found to be close to 0.5 s and 7 cm, respectively [12]. When confined in circular tanks, fish tend to swim close to the curved wall because their trajectory is made of quasi straight segments with limited variance of the heading angle between kicks, hence preventing the fish from escaping from the tank wall (unless when a rare large heading angle change occurs) [12,33]. When swimming in groups, *H. rhodostomus* tend to remain close to each other, especially when the number of fish in the tank is small. In fact, the social interactions between fish reflect the combined tendency to align with and follow their neighbours while at the same time maintaining a safe distance with the wall. At a given kicking instant, only a few neighbours (one or two) have a relevant influence on the behaviour of a fish [34]. The decision-making of fish displaying a burst-and-coast swimming mode can thus be reproduced by considering only pairwise interactions. Obviously, if one only considers pairs of fish, like here, it therefore suffices to consider the relative state of the neighbouring fish (relative position and velocity) and the effect of the distance and the relative orientation to the wall [12,20].

2.3.1. Analytical burst-and-coast model

The analytical burst-and-coast model (hereafter called ABC model) quantitatively reproduces the dynamics of *H. rhodostomus* swimming alone or in pairs under the hypothesis that fish decision-making times correspond exactly to their kicking times, that is, the new direction of movement, the duration and the length of the kick are decided precisely at the end of the previous kick [12].

Given a pair of agents i and j at a respective state $(\boldsymbol{u}_j^n, \phi_j^n)$ and $(\boldsymbol{u}_i^n, \phi_i^n)$ at time t^n, the state of agent i at the next instant of time t_i^{n+1} is given by

$$t_i^{n+1} = t_i^n + \tau_i^n, \tag{2.4}$$

$$\phi_i^{n+1} = \phi_i^n + \delta\phi_i^n \tag{2.5}$$

and

$$\boldsymbol{u}_i^{n+1} = \boldsymbol{u}_i^n + l_i^n\, \boldsymbol{e}\,(\phi_i^{n+1}), \tag{2.6}$$

where $\boldsymbol{e}\,(\phi_i^{n+1})$ is the unitary vector pointing in the heading direction ϕ_i^{n+1}, τ_i^n and l_i^n are the duration and length of the nth kick of agent i, and $\delta\phi_i^n$ is the heading change of agent i. The heading angle change $\delta\phi_i^n$ is the result of three effects: the interaction with the wall, the social interactions with the other fish (repulsion/attraction and alignment), and the natural spontaneous fluctuations of fish headings (cognitive noise) [12]. The term 'cognitive noise' encapsulates the fact that fish (or humans) would not generally replicate the exact same motion when placed under identical initial conditions, namely starting at the same positions and with the same initial velocities. Hence, a behavioural model must not only describe the social interactions between individuals, but also the properties of their spontaneous fluctuations. The social interactions depend only on the relative state of both agents, determined by the triplet $(d_{ij}, \psi_{ij}, \phi_{ij})$. The derivation of the shape and intensity of the functions involved in $\delta\phi_i^n$ is based on physical principles of symmetry of angular functions and a data-driven reconstruction procedure detailed in [12] for the case of *H. rhodostomus* and in [20] for the general case of animal groups.

Starting from the initial condition $(\boldsymbol{u}_i^0, \phi_i^0)$ of fish i, the length and the duration of its next kick, l_i^0 and τ_i^0, are sampled from the experimental distributions obtained in [12]. Then, the timeline t_i^1 of fish i is updated with equation (2.4), the heading angle of the next kick ϕ_i^1 is calculated with equation (2.5), and the position of the fish at the end of the kick $\boldsymbol{u}_i^1$ is obtained with equation (2.6). As kicks of different fish are asynchronous, the next kick can be performed by either of the two fish. Each fish has thus it own timeline, but is subject, at each of its kicks, to the evolution of the other fish along its own kicks.

The ABC model is therefore a discrete model that generates kick events instead of continuous time positions. To directly compare with the deep learning interaction (DLI) model presented in the next section, which is a continuous time model, we resampled the trajectories made of kick events produced by the ABC model and build continuous time trajectories with a timestep of size $\Delta t = 0.12$ s. We produced trajectories that add up to a total of 500 000 timesteps, corresponding to approximately 16.7 h.

2.3.2. Deep learning interaction model

The DLI model consists of an artificial neural network (ANN) which is fed with a set of variables characterizing the motion of *H. rhodostomus* and which provides the necessary information to reproduce the social interactions of these fish by estimating their motion along timestep of length $\Delta t = 0.12$ s. At time t, the DLI model is designed to take sequences of states as input to capture the short- and long-term dynamics. Then, it generates predictions for the acceleration components of the fish at the following timestep $t + \Delta t$.

For the DLI model, the state of an agent i at time t is defined by

$$\mathbf{s}_i(t) = \big(\boldsymbol{u}_i(t), \boldsymbol{v}_i(t), r_{\mathrm{w}}^i(t)\big) \in \mathbb{R}^5. \tag{2.7}$$

The state of an agent includes redundant information: in a fixed geometry, r_{w} can be deduced from $\boldsymbol{u}_i$, and $\boldsymbol{v}_i^n$ from the input sequence $\boldsymbol{u}_i^{n-4}, \ldots, \boldsymbol{u}_i^n$. This redundancy is intended to facilitate the training process of the neural network. Furthermore, these redundancies are shown to significantly boost the performance of the network compared with similar ANN structures (see electronic supplementary material).

The system's state $\mathbf{S}(t)$ is then defined as the combination of both agent states, in addition to their inter-individual distance $d_{ij}(t)$ (also a redundant variable),

$$\mathbf{S}(t) = \big(\mathbf{s}_i(t), \mathbf{s}_j(t), d_{ij}(t)\big) \in \mathbb{R}^{11}. \tag{2.8}$$

Figure 2 shows the structure of the ANN, consisting of seven layers: two long-short term memory (LSTM) layers [35], and five fully connected (dense) layers.

The first LSTM layer consists of 256 neurons and is located at the input of the ANN, where it receives the sequence of the five last states of the system, i.e. a matrix of dimension 5×11: $(\mathbf{S}(t-4), \ldots, \mathbf{S}(t))$. This history length of four timesteps (0.48 s) is borrowed from the biology of the fish: as already mentioned, the time it takes for a fish to display its characteristic behaviour, a kick, is 0.5 s [12], therefore, we input the current state plus the states that correspond to the average duration of a kick. The output of the first LSTM is then gradually reduced in dimension by two successive dense layers, and then scaled up again with a second LSTM, whose configuration is also based on a history
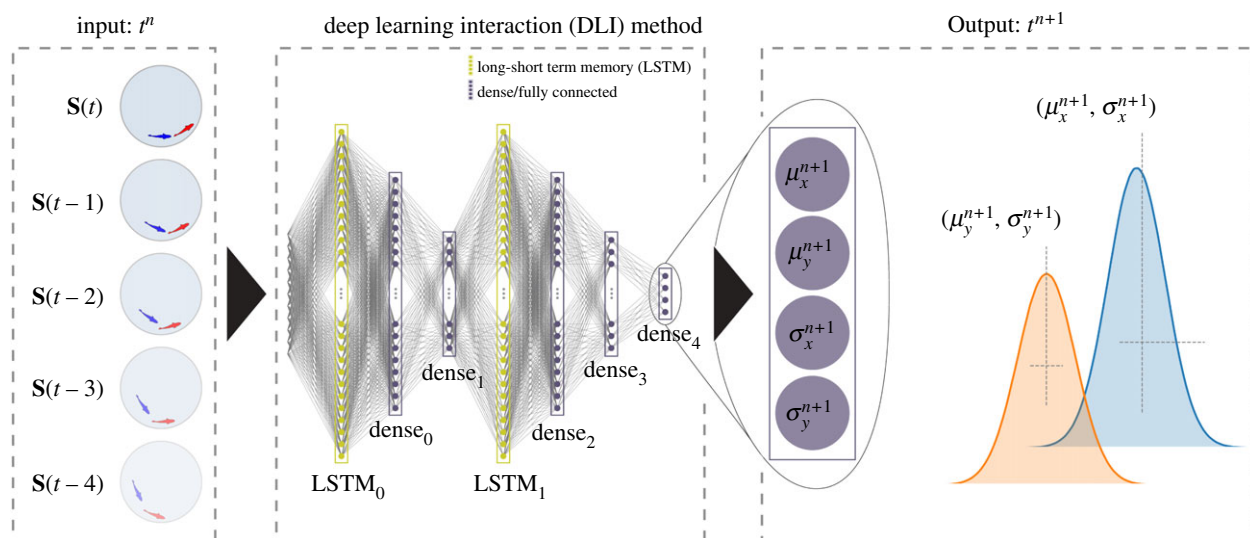
Figure 2. Structure of the artificial neural network (ANN) used in the deep learning interaction (DLI) model. From left to right: *Input* of the ANN: the five last states, $(\mathbf{S}(t-4), \ldots, \mathbf{S}(t))$ at time t. Where $\mathbf{S}(t) = (\mathbf{s}_i(t), \mathbf{s}_j(t), d_{ij}(t)) \in \mathbb{R}^{11}$ and each state is parametrized as $\mathbf{s}_i(t) = (\mathbf{u}_i(t), \mathbf{v}_i(t), r_w^j(t)) \in \mathbb{R}^5$; the seven layers (two long-short term memory, also known as LSTM, layers and five dense layers) capturing the social dynamics; *Output*: the two pairs of values (μ_x, σ_x) and (μ_y, σ_y) corresponding, respectively, to the mean and standard deviation of the probability distribution function (assumed to be Gaussian) of each component a_x and a_y of the instantaneous acceleration vector $\mathbf{a}$ at time $t+1$, constituting the prediction of the DLI model.

of five states. Then, two other dense layers are used to reduce the dimension of the output of the second LSTM, and a last dense layer is applied to provide the final output of the ANN. More details about the configuration of the ANN are given in electronic supplementary material, table S7.

The output of the ANN consists of two pairs of values, (μ_x, σ_x) and (μ_y, σ_y), corresponding to the expected value and standard deviation of the x and y components of the predicted acceleration, which are assumed to be Gaussian distributed [36], as actually found for *H. rhodostomus* [12]. Hence, the predicted acceleration of the agent, $\mathbf{a} = (a_x, a_y)$, can be written

$$a_x = \mu_x + \sigma_x g_x, \quad a_y = \mu_y + \sigma_x g_y, \tag{2.9}$$

where g_x and g_y are independent standard Gaussian random variables drawn from $\mathcal{N}(0, 1)$. Then, the velocity vector of the agent i at the time t^{n+1} is given by

$$\mathbf{v}_i^{n+1} = \mathbf{v}_i^n + \Delta t \, \mathbf{a}_i^n, \tag{2.10}$$

and the position of the agent is updated according to

$$\mathbf{u}_i^{n+1} = \mathbf{u}_i^n + \Delta t \, \mathbf{v}_i^{n+1}. \tag{2.11}$$

Note that in the DLI model, the predicted variance of the acceleration accounts for the fish intrinsic spontaneous behaviour exhibited during their decision process (cognitive noise), and hence translates the fact that two real (or modelled) fish will not act the same if put twice in the same given state characterized by equation (2.8).

In some rare instances, the prediction of the DLI model would move one or both fish outside the limits of the tank. To account for that, we introduce a rejection procedure: the invalid move is rejected, and we resample the Gaussian random variables drawn in equation (2.9) until a valid move is produced. Note that a similar rejection procedure is also implemented in the ABC model of [12], to strictly enforce the presence of the wall. Indeed, in the ABC model, the ABC agents would systematically escape the tank after a few seconds or very few minutes without this rejection procedure. In §3.4 and electronic supplementary material, figures S1 and S2, we show that the DLI model has, in fact, implicitly learned the presence of the wall, and that DLI agents can remain within or in the close vicinity of the tank for several dozen minutes without implementing this rejection procedure (60% chance not to escape the tank during 100 min of simulation).

The *prediction* of the ANN at time t^{n+1} is thus a vector of dimension 1×4 that can be written as $(\boldsymbol{\mu}_{\text{pred}}^{n+1}, \boldsymbol{\sigma}_{\text{pred}}^{n+1})$, where

$$\boldsymbol{\mu}_{\text{pred}}^{n+1} = \left(\mu_x^{n+1}, \mu_y^{n+1} \right) \quad \text{and} \quad \boldsymbol{\sigma}_{\text{pred}}^{n+1} = \left(\sigma_x^{n+1}, \sigma_y^{n+1} \right). \tag{2.12}$$

The ANN is then trained to approach the *real/observed* values $\boldsymbol{\mu}_{\text{real}}^{n+1}$ by means of the adaptive moment estimation optimizer (Adam) with a time-decaying learning rate $\lambda = 10^{-4}$ and a negative log-likelihood loss function ℓ defined in terms of the prediction error $\boldsymbol{\epsilon}_{n+1} = \boldsymbol{\mu}_{\text{pred}}^{n+1} - \boldsymbol{\mu}_{\text{real}}^{n+1}$ and the standard deviations as follows [37]:

$$\ell(\boldsymbol{\epsilon}_{n+1}, \boldsymbol{\sigma}^{n+1}) = \frac{1}{2} \sum_{n=1}^{N_h} [\boldsymbol{\epsilon}_{n+1}]^T \mathcal{C}^{-1}(\boldsymbol{\sigma}^{n+1}) \, \boldsymbol{\epsilon}_{n+1} + \frac{N_h}{2} \log \|\mathcal{C}(\boldsymbol{\sigma}^{n+1})\|, \tag{2.13}$$

where N_h is the number of timesteps in the history of the input of the ANN (here $N_h = 5$) and $\mathcal{C}$ is a diagonal covariance matrix with the values of $\boldsymbol{\sigma}_{\text{pred}}^{n+1}$ in the diagonal and zeroes elsewhere.

The training of the ANN is carried out with a subset of the experimental dataset. More specifically, the training process is given a budget of 45 epochs with a batch size of 512 samples on a dataset that was split 80%, 15% and 5% for training, validation and test,

respectively. Then, the DLI model is used to produce trajectories of 500 000 timesteps of size $\Delta t = 0.12$ s, as done with the ABC model. At the beginning of the simulation, each agent is given a copy of the DLI model and both agents are initialized with a random 5-timestep-long trajectory sampled from the fish dataset. At each timestep t^n, the state vector $\mathbf{S}(t^n)$ is built and introduced in the network, which provides the estimated instantaneous acceleration distributions at time t^{n+1}. Then, the acceleration is evaluated according to equation (2.9), and the next positions and velocities of the agents are obtained from the equations of motion, equations (2.10) and (2.11).

Designing the DLI model. Designing and selecting an appropriate ANN structure to model a system is for the most part non-trivial and requires either an extensive search through automatic methods (e.g. neuro-evolution [38–40]) or an exhaustive number of empirical attempts for very specific applications [22–24]. Here, we followed a hybrid approach consisting of empirically designing an ANN based on biological insight and automatically searching for its optimal structure by bootstrapping the search. Once we established this initial model, we performed an automated search for similar neural networks using the same input and output for different combinations of (i) the number of layers, (ii) the size of the layers, and (iii) the activation functions (i.e. transfer functions tasked with mapping the inputs of a neuron to a single weighted output value passed to the next layer). The search included a total of 82 neural network structures, trained with the same budget of iterations and stopping criteria, and out of which the ANN shown above is the best performing. The best performing ANN is selected according to the metrics presented in the following section.

Three notable categories of networks were considered: (i) non-probabilistic networks that only generate (μ_x^{n+1}, μ_y^{n+1}) (and hence, not explicitly including the cognitive noise), (ii) probabilistic networks that do not have memory cells (hence, missing the fact that fish are gliding passively on a timescale of order 0.5 s), and (iii) probabilistic networks that implement memory thanks to LSTM layers. Non-probabilistic networks (i) provide the mean value of the components of the acceleration for the next timestep with high accuracy, but miss the essential variability that is intrinsic to the spontaneous behaviour of fish and which allows for the emergence of social interactions. Probabilistic networks without memory (ii) are able to partly capture this intrinsic variability, but do not fully capture the nonlinear nature of the problem (see electronic supplementary material, figure S6 and video S4). Finally, probabilistic networks with memory (iii) performed generally well, and we found that the structure used in the DLI model consistently provides the best results for the number of epochs set for training and for the ANNs considered by the automatic search.

Our search approach revealed the existence of two crucial ingredients that must be considered in the model, both accounting for biological characteristics of fish behaviour observed experimentally. First, the neural network must be fed with information covering the typical timescale along which relevant changes take place in the behaviour of the fish. Since real fish kicks last 0.5–0.6 s on average, the NN needs information about the fish behaviour over time intervals of at least this duration (i.e. four to five timesteps of 0.12 s). However, we found that using longer vector lengths (up to 10 timesteps) for the case of *H. rhodostomus* does not lead to any significant improvement in the results, while considerably increasing the training time. Second, the output of the network must contain a sufficiently wide diversity of predictions so that the agents reproduce the high variability of responses that fish display when behaving spontaneously and reacting to external stimuli.

ANNs without memory tend to make too similar predictions, and agents do not initiate the typical direction changes that are observed in the experiments. A possible solution could be to add some phenomenological noise to the predictions of the NN. However, this would result in an unrealistic behaviour, albeit an improvement over not adding noise at all. For example, when a fish swims close to the wall, it does not have the same liberty to turn toward or away from the wall, which would not be captured by a too crude implementation of the fish cognitive noise. Our approach accounts for this behavioural uncertainty for each state (position, velocity, distance to the neighbour and to the wall) and for both degrees of freedom during the training phase of the ANN, being therefore able to capture these complex behavioural patterns. The performance of the two variants is depicted in electronic supplementary material, figure S4.

3. Results

When fish swim in a circular tank (here, of radius $R = 25$ cm), they interact with each other and with the tank wall. The resulting collective dynamics can be finely characterized by exploiting the nine observables introduced and described in the Methods section. As explained there, these observables probe (i) the instantaneous individual behaviour, (ii) the instantaneous collective behaviour, and (iii) the temporal correlations of the dynamics.

Hereafter, we analyse three trajectory datasets: the first one corresponds to pairs of *H. rhodostomus* in our experiment (4 h of data), the second one to the ABC model (16.7 h) and the third one to the DLI model (16.7 h). Electronic supplementary material, video S1 shows typical trajectories for these three conditions. The aim of this section is to quantitatively validate the qualitative agreement observed in this video.

3.1. Quantification of the instantaneous individual behaviour

The individual fish behaviour is characterized by three observables: the PDF of the speed V, of the distance to the wall r_w and of the angle of incidence to the wall θ_w. When swimming in pairs, fish tend to adopt a typical speed of approximately 7 cm s^{-1} (see the peak of the PDF in figure 3a), but can also produce high speeds up to $25-30$ m s^{-1}. In fact, we observe that both the leader and follower fish produce very similar speed profiles (thus omitted in figure 3a). Both fish remain close to the wall of the tank (a consequence of the fish burst-and-coast swimming mode [12]), the leader being closer to the wall (typically, at approx. 0.5 BL) than the follower (at approx. 1.2 BL; see figure 3b). This feature is due to the follower fish trying to catch up with the leader fish by taking a shortcut while taking the turn. Moreover, fish spend most of the time almost parallel to the wall: see the peaks of both PDFs at $\theta_w \approx \pm 90°$ in figure 3c. A slight asymmetry is observed in the PDF of θ_w, showing that, in the experiments, fish have turned more frequently in the counterclockwise direction. Values of the mean and the standard deviation of the PDFs presented in this section are given in electronic supplementary material, tables S1, S2 and S3.

Both ABC and DLI models produce agents that move at the same mean speed as fish in the experiments, and figure 3a shows that the speed PDF for both models are in excellent agreement with the one observed in real fish. Moreover, the agents of the ABC model are as close to the wall and as parallel to it as fish are. The PDF of the ABC leader is in good agreement with that of the fish leader (figure 3b). However, the PDF for the ABC follower has a peak at approximately the same distance to the wall as that of the

Figure 3. Probability density functions (PDF) of observables characterizing individual behaviour: (*a*) Speed *V*, (*b*) distance to the wall r_w and (*c*) angle of incidence to the wall θ_w. Black lines: experimental fish data. Blue lines: agents of the analytical burst-and-coast model (ABC). Red lines: agents of the deep learning interaction model (DLI). In (*b*) and (*c*), dashed lines: geometrical leader; dotted lines: geometrical follower.


Figure 4. Probability density functions (PDF) of observables characterizing collective behaviour: (*a*) distance between individuals d_{ij}, (*b*) difference in heading angles ϕ_{ij} and (*c*) angle of perception of the geometrical leader and follower ψ_{ij}. Black lines: experimental fish data. Blue lines: agents of the analytical burst-and-coast model (ABC). Red lines: agents of the deep learning interaction model (DLI). Only in (*c*), dashed lines: geometrical leader; dotted lines: geometrical follower.

leader, while the corresponding peaks are more separated for real fish. Yet, the PDF for the ABC follower is broader than for the leader, showing that the ABC follower tends to be farther from the wall than the leader, as observed for real fish. For the DLI model, the peaks of both leader and follower PDFs are at approximately the same position as for real fish, although their height is smaller than for fish, meaning that DLI-agents tend to explore more frequently the interior of the tank (observe the thicker tails of the PDF of r_w for the DLI model in figure 3*b*). Alignment with the wall is also well reproduced by both models (figure 3*c*), including the asymmetry in the direction of rotation around the tank: their peak at $\theta_w > 0$ is higher than the one at $\theta_w < 0$. As already seen in the PDF of r_w, DLI-agents visit more often the interior of the tank, and are hence less aligned with the wall than the real fish and ABC agents. Note that the tendency of DLI-agents to rotate more frequently in the counterclockwise direction is learned from the training set, while this asymmetry has to be explicitly implemented in the ABC model, by introducing an asymmetric term in the analytical expression of the wall repulsion function. A closer look at figure 3*c* shows that fish actually follow the wall with a most likely angle of incidence $|\theta_w|$ that is slightly smaller than 90°, a feature resulting from the burst-and-coast swimming mode inside a tank with positive curvature: fish are found more often going toward the wall than escaping it.

We have also computed the Hellinger distance (HD) between the experimental PDF probing the individual behaviour and the corresponding PDF produced by the DLI and ABC models. The HD (see the caption of electronic supplementary material, tables S10 and S11, for more details) quantifies the (dis)agreement between two PDF for the same variable. The results of electronic supplementary material, tables S10 and S11 for both models confirm their good performance: the DLI model HD is slightly better than that of the ABC model for the speed PDF, as good for the PDF of r_w and not quite as good for the PDF of θ_w.

3.2. Quantification of the instantaneous collective behaviour

Hemigrammus rhodostomus is a social species, and figure 4*a* shows that the two fish remain most of the time close to each other, with the PDF of their distance d_{ij} presenting a peak around $d_{ij} \approx 7$ cm ≈ 2 BL (mean and standard deviation of the PDFs presented in this section are given in electronic supplementary material, tables S1, S2 and S3. The PDF of d_{ij} produced by the DLI model is slightly wider than for the experiment and the ABC model, and in particular, presents too much weight at small distances.

The fish have a strong tendency to align with each other, as shown in figure 4*b*, with the PDF of their relative heading ϕ_{ij} being sharply peaked at 0°. In addition, the PDF of the viewing angle ψ_{ij} reveals that the fish are swimming one behind the other rather than side-by-side. This is illustrated in figure 4*c* by the sharp difference in the PDF of the viewing angle for the leader and the

Figure 5. Observables quantifying temporal correlations in the system. (*a*) Mean squared displacement $C_X(t)$, (*b*) velocity temporal autocorrelation $C_V(t)$, (*c*) temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Black lines: experimental fish data. Blue lines: agents of the analytical burst-and-coast model (ABC). Red lines: agents of the deep learning interaction model (DLI). Dashed lines: geometrical leader; dotted lines: geometrical follower; full lines: average over the two fish or agents.

follower. The PDF of ψ_{leader} is peaked around $\pm 160°$, meaning that the follower fish is almost right behind the leader fish, but slightly shifted to the right or left. A slight left–right asymmetry in the PDF of the viewing angles is also visible, the follower being more frequently on the left side of the leader, a consequence of the fact that the fish in the experiment follow the wall by turning more often counterclockwise (figure 3*c*), with the follower swimming farther from the wall than the leader (figure 3*b*).

All these features are well reproduced by both models, with only some small quantitative deviations. The ABC model reproduces almost perfectly the experimental PDF of the distance between fish, whereas the PDF for the DLI model is only slightly wider and presents slightly more weight at very small distance than found for real fish or in the ABC model (figure 4*a*). The DLI model is in turn better than the ABC model at reproducing the PDF quantifying the alignment of the fish, the latter producing more weight near $0°$ than for real fish (figure 4*b*). Both models fail at reproducing the small weight in the PDF at $\phi_{ij} \approx \pm 180°$, which corresponds to sudden U-turns that real fish sometime perform. The PDF of the viewing angles for the leader and the follower (figure 4*c*) are also fairly reproduced by both models, including the slight left–right asymmetry observed in real fish, although the peak in the PDF at $\psi_{\text{follower}} = 0°$ (and to a lesser extent at $\psi_{\text{leader}} \approx -160°$) is not quite as sharp as in the experiment.

Again, we have computed the HD between the experimental PDF probing the collective behaviour and the corresponding PDF produced by the DLI and ABC models. The results of electronic supplementary material, tables S10 and S11 for both models confirm their good performance: as anticipated above, the DLI model HD for the PDF of the distance between agents is higher than for the ABC model (and is the highest found for all six PDF presented here, with $\text{HD}_{d_{ij}} = 0.13$). However, electronic supplementary material, tables S10 and S11 also confirm that the DLI model reproduces quantitatively the PDF of ϕ_{ij} and ψ_{ij}.

3.3. Quantification of temporal correlations

Figure 5 shows the three observables defined in equations (2.1)–(2.3) and probing the emerging temporal correlations in the system: the mean-squared displacement $C_X(t)$, the velocity autocorrelation $C_V(t)$ and the autocorrelation of the angle of incidence to the wall $C_{\theta_w}(t)$, as function of the time difference t between observations. The figure reveals that both models fail to fully reproduce quantitatively these very non-trivial observables, which indeed constitute the most challenging benchmark characterizing the correlations emerging from the fish behaviour.

Fish data present three distinct regimes: a quasi-ballistic regime at short timescale ($t \lesssim 1.5\,\text{s}$) where $C_X(t) \approx \langle v^2 \rangle t^2$, followed by a second short diffusive regime ($1.5\,\text{s} \lesssim t \lesssim 5\,\text{s}$) where $C_X(t) \approx Dt$, which is limited by the finite size of the tank, ultimately leading to a third regime of saturation ($t > 5\,\text{s}$) characterized by slowly damped oscillations since fish are guided by the wall (figure 5*a*). Accordingly, the velocity correlation function starts from $C_V(t=0) = \langle v^2 \rangle$ at short time and also presents damped oscillations (figure 5*b*). The negative minima of the oscillations in $C_V(t)$ correspond to times when the focal fish is essentially at a position diametrically opposite to its position at the reference time $t = 0$, its velocity then being almost opposite to that at $t = 0$. Similarly, positive maxima correspond to times when the fish returns to almost the same position it had at $t = 0$, with a similar velocity, guided by the tank wall. Of course, these oscillations are damped as correlations are progressively lost, and the velocity correlation function $C_V(t)$ ultimately vanishes at large time $t \gg 20\,\text{s}$, due to the actual stochastic nature of the trajectories at this timescale (possible U-turns, or the fish randomly crossing the tank). Note that $C_X(t)$ is markedly different for the leader and follower fish, with a higher saturation value for the leader, which swims closer to the wall, as mentioned above.

The ABC model is able to fairly reproduce the short and intermediate regimes for $C_X(t)$ (figure 5*a*), as well as the position of its first peak, reached only 1 s later than for fish. The ABC model also reproduces the experimental saturation value of $C_X(t)$ averaged over the two fish. As for the DLI model, its predictions are only slightly worse than that of the ABC model, since the DLI agents are moving a bit farther to the wall compared with ABC agents and real fish. Yet, both models equally fail at producing more than one oscillation, and the correlations are damped faster compared with the experiment.

As for the velocity autocorrelation $C_V(t)$ (figure 5*b*), the ABC model reproduces almost perfectly the short and intermediate regimes and the position of the first negative minimum (hence, up to $t = 6\,\text{s}$), while the DLI model underestimates the depth of this first minimum. But again, both models fail at reproducing the persistence of the correlations, producing a too fast damping of the oscillations (an effect slightly stronger in the DLI model).

Both models struggle at reproducing the correlation function $C_{\theta_w}(t)$ of the angle of incidence to the wall (figure 5c), where the fish curve first sharply decreases up to $t = 6$ s and then remains close to $C_{\theta_w} \approx 0.2$. The ABC model is clearly unable to reproduce both the decreasing range (clearly diverging before $t = 2$ s) and the correct saturation value (never falling below $C_{\theta_w} \approx 0.6$). As for the DLI model, it produces a slightly sharper decay of $C_{\theta_w}(t)$ than for real fish, up to $t \approx 6$ s, but fails to reproduce the non-negligible remaining persistence of the correlation observed in fish for $t > 7$ s, with $C_{\theta_w}(t)$ in the DLI model decaying rapidly to zero. In fact, both models fail to reproduce the experimental $C_{\theta_w}(t)$ for opposite reasons. The ABC model exhibits a too high persistence of the correlations of θ_w compared with real fish, presumably because real fish indeed often follow the wall but can also produce sharp U-turns, as observed in figure 3c. On the other hand, the failure of the DLI model in reproducing $C_{\theta_w}(t)$ stems from the fact that DLI agents move farther from the wall and cross through the tank more often than real fish and ABC agents (see the discussion of figure 3b above), hence leading to a too fast, and ultimately total, loss of correlation for θ_w.

3.4. Complementary analyses

In order to test whether the DLI model has correctly learned the presence of the wall, we have run 30 simulations of duration 6000 s to check whether the DLI agents would stay within the area of the tank, even without enforcing its presence by the rejection procedure mentioned in the second paragraph below equation (2.11). We found that the DLI agents indeed remain in or very near the tank during the entire time of the simulation in 60% of runs. In the other 40% of runs, the DLI agents would ultimately escape the tank after a mean time of order 3000 s. These results are summarized in electronic supplementary material, figure S1, where we present the time series of the distance to the wall $r_w(t)$ for the 10 first runs, and in electronic supplementary material, figure S2, where we report the survival probability (i.e. the probability that the DLI agents remain within the tank up to a given time). These results indicate that the DLI model has convincingly learned the presence of the wall, and is able to maintain the agents within the wall for several dozen minutes without the need of an explicit rejection procedure.

We have also conducted several other complementary tests of our approach. First, the DLI model yields better results in generating social interactions than a similarly purposed ANN for human trajectory forecasting [30,31] (D-LSTM model; see electronic supplementary material, figures S3 and S4, tables S4, S5, S6 and S12, and video S2). In particular, the results for the HD (HD$_{r_w}$ = 0.30 and HD$_{\theta_w}$ = 0.40) show that this D-LSTM model completely fails at capturing the interaction of the fish with the tank wall. While this is expected due to the missing inputs (compared with the DLI; see electronic supplementary material, S1), these results confirm that there exist models that do indeed capture the short-term dynamics without being able to reproduce the long-term dynamics, presumably due to non-Markovian effects. In addition, we also trained a multi-layered perceptron interaction (MLI) model *without any memory cells*, and found that it fails to reproduce all six PDF (see electronic supplementary material, S1, figure S6), resulting in high values of the corresponding HD (see electronic supplementary material, S1, table S13).

Moreover, we have analysed the performance of the DLI model when varying the fraction of the dataset used in its training. The performance is quantified by using the HD between the experimental PDF and that produced by the DLI model, and electronic supplementary material S1, table S15 reports the resulting HD values. When only using 75%, 50%, or even 37.5% of the dataset, the DLI model has a similar performance as when trained with the full dataset (4 h of pair trajectories). However, the performance sharply drops when only using 25%, 12.5% and 5% of the dataset. In fact, using 25% or less of the dataset, we also found that the performance significantly depends on the training sample (we ran four training sessions in each case). Finally, we also found that without enforcing the presence of the wall with our rejection procedure, the median escape time of the fish computed over 30 runs of 6000 s when using 25%, 12.5% and 5% of the dataset are of order 500 s, 75 s, 50 s, compared with 3000 s when using 100% or even 50% of the dataset. These results show that our DLI network (and its size) is coherent with the size of the training dataset, and that its predictions remain robust when restricting the data at least down to half of the original dataset.

Finally, we have trained the DLI model with data for pairs of zebrafish (*D. rerio*), and found that it yields fair results for this species too, without any structural modification in its architecture (see electronic supplementary material, figure S5 and tables S8, S9 and S14). While acquiring a functional model of a new species' interactions proved straightforward with the DLI, the same would not be generally true for analytical models.

Following the completion of the present work, we have exploited the DLI model to study groups with more than two fish, *without any retraining*. Indeed, *H. rhodostomus* [34], like many other group-living species [7], effectively only interact with a few influential neighbours, at a given time. Thus, for a given agent in a group of $N > 2$ agents, the DLI for *H. rhodostomus* should only retain the influence of typically the two agents leading to the highest acceleration [34,41], as predicted by the DLI model. Electronic supplementary material, video S3 illustrates this procedure for $N = 5$ agents, resulting in a cohesive and aligned group, in qualitative agreement with experimental observation [34]. In addition, the present DLI model has also been recently exploited in [42] to command a robot fish initially introduced in [43] (where it was commanded by the ABC model), and moving alone in the tank, or reacting in a closed-loop to one or four real fish.

4. Conclusion and discussion

Studying social interactions in animal groups is crucial to understand how complex collective behaviours emerge from individuals' decision-making processes. Very recently, such interactions have been extensively investigated in the context of collective motion by exploiting classical computational modelling [12,20,21] and automated ML-based methods [23,24]. Although ML algorithms have been shown to provide insight into the interactions of hundreds of individuals at short timescales [23,24], their ability to reproduce the complex dynamics in animal groups at long timescales has not yet been assessed.

Here, we have presented a DLI model which reproduces the behaviour of fish swimming in pairs. The DLI model's good performance can be primarily ascribed to its memory related to a biologically relevant timescale (fish kicks of typical duration 0.5–0.6 s), and to a carefully crafted input/feature vector. Indeed, the MLI model without memory cells performs very poorly, while the D-LSTM model, characterized by a different input/feature vector, demonstrates markedly lower performance than the DLI model.

We have also introduced the appropriate tools for the validation of an ANN model, when compared with experimental results and confronted with an analytical behavioural model (ABC). In fact, our study establishes a systematic methodology to assess the long-term predictive power of a model (analytical or ML), by introducing a set of fine observables probing the individual and collective behaviour of model agents, as well as the subtle correlations emerging in the system. These observables, which can be straightforwardly extended to groups of $N > 2$ agents, provide an extremely stringent test for any model aimed at producing realistic long-term trajectories mimicking that of actual animal groups. In particular, we consider that the usual validation of an ML model at a short timescale should be complemented by the type of long timescale analysis that we propose here, in order to fully assess its performance. Indeed, we have shown that a model (like the D-LSTM model) can have a good performance at very short timescales, while presenting a degraded performance at large timescales, presumably due to non-trivial non-Markovian effects.

The DLI model closely reproduces the dynamics of real fish at both the individual (speed, distance to the wall, angle of incidence to the wall) and collective (distance between individuals, relative heading angle, angle of perception) levels during long simulations corresponding to more than 16 h of fish swimming in a tank, hence successfully generating lifelike interactions between agents. When compared with experiment, the ABC model and the DLI model essentially perform equally well. Notably, the DLI model better captures the most likely distance of the leader and follower from the wall. However, the DLI model is less accurate in reproducing the temporal correlations quantified by the mean-squared displacement and the velocity autocorrelation. Yet, both ABC and DLI models fail at capturing the temporal correlations of the angle of incidence to the wall, but for very different reasons. More importantly, the DLI model convincingly infers the presence of the tank wall, and is able to keep the DLI agents within the wall boundaries for several dozen minutes, even when the rejection procedure is not enforced. In addition, we have shown that the performance of the DLI model remains robust even when only using half of the experimental training dataset, while its accuracy sharply drops when only using a quarter of the training dataset.

Our study demonstrates two advantages of ML techniques: (i) they can drastically accelerate the generation of new models (as illustrated here for zebrafish) and (ii) with minimal expertise in biology or modelling. This is especially useful in robotics, where models often act as behavioural controllers (i.e. trajectory generators) that guide the robot(s). Although there already exist many bio-hybrid experiments in the literature, most of them rely on simplified models for behavioural modulation [44–46], few of them exploit realistic models (analytical or ML) [29,47], and, to our knowledge, none of them are tested in the long term in simulations or real-life. In this context, ML has the potential to benefit multidisciplinary studies, provided such techniques are thoroughly validated in simulations.

However, accelerating the production of collective behaviour models with ML comes at a cost. Indeed, the DLI is a black-box model, and although it captures the subtle impact of social interactions between individuals, it is impossible to retrieve the interaction functions themselves. Some approaches partially address this issue by providing insight into how the network operates for specific sets of inputs [23,24]. Yet, they still do not offer explicit interaction functions. Instead, they provide insights in the form of force maps that can, to some extent, be used to interpret the underlying mechanisms of the interactions, or in the form of input/output correlation graphs, that showcase the manner in which an input state typically affects the output [48]. On the other hand, analytical models supplemented by a procedure to reconstruct social interactions [12,20] provide a concise and explicit description of the system in question. Moreover, varying the parameters of such models allows for investigating their relative impact on the dynamics, in the form of phase diagrams representing the collective observables (and the corresponding collective state of the group) as a function of the model parameters [32,41]. This is not feasible with ML models, unless they are retrained or specifically structured to allow it.

In summary, this work shows that DLI-like models may now be considered as firm candidates to shed light on groundbreaking problems such as how social interactions take place and affect collective behaviour in living groups. Yet, we have emphasized that social interaction models should be precisely tested at both short and long timescales. Future work includes the design of ANNs that provide additional information about the learned dynamics (e.g. using the framework of [48] and/or attention layers, like in [23,24]), or possibly, by exploiting symbolic regression algorithms [49,50]. We also plan to study the extension of the DLI model to larger groups, in particular, in connection with our robotic platform [42–46]. It would also be interesting to apply the DLI model in different environmental conditions, such as light intensity, as recently done for the ABC model [33]. Ultimately, a more generalized and unified version of the DLI model or similar algorithms requires extensive testing with additional social animal species (e.g. humans). We believe that these approaches could improve our understanding of the mechanisms arising in collective behaviour and allow for more precisely exploring and modulating them.

References

1. Sumpter DJ. 2010 *Collective animal behavior*. Princeton, NJ: Princeton University Press.

2. Krause J, Ruxton GD. 2002 *Living in groups*. Oxford, UK: Oxford University Press.

3. Ball P. 2011 *Flow: nature's patterns: a tapestry in three parts*, vol. 2. Oxford, UK: Oxford University Press.

4. Camazine S, Deneubourg JL, Franks NR, Sneyd J, Theraulaz G, Bonabeau E. 2020 *Self-organization in biological systems*. Princeton, NJ: Princeton University Press.

5. Couzin ID, Krause J. 2003 Self-organization and collective behavior in vertebrates. *Adv. Study Behav.* **32**, 1–75. (doi:10.1016/S0065-3454(03)01001-5)

6. Vicsek T, Zafeiris A. 2012 Collective motion. *Phys. Rep.* **517**, 71–140. (doi:10.1016/j.physrep.2012.03.004)

7. Cavagna A, Giardina I, Grigera TS. 2018 The physics of flocking: correlation as a compass from experiments to theory. *Phys. Rep.* **728**, 1–62. (doi:10.1016/j.physrep.2017.11.003)

8. Deutsch A, Friedl P, Preziosi L, Theraulaz G. 2020 Multi-scale analysis and modelling of collective migration in biological systems. *Phil. Trans. R. Soc. B* **375**, 20190377. (doi:10.1098/rstb.2019.0377)

9. Herbert-Read JE, Perna A, Mann RP, Schaerf TM, Sumpter DJ, Ward AJ. 2011 Inferring the rules of interaction of shoaling fish. *Proc. Natl Acad. Sci. USA* **108**, 18 726–18 731. (doi:10.1073/pnas.1109355108)

10. Herbert-Read JE. 2016 Understanding how animal groups achieve coordinated movement. *J. Exp. Biol.* **219**, 2971–2983. (doi:10.1242/jeb.129411)

11. Gautrais J, Ginelli F, Fournier R, Blanco S, Soria M, Chaté H, Theraulaz G. 2012 Deciphering interactions in moving animal groups. *PLoS Comput. Biol.* **8**, e1002678. (doi:10.1371/journal.pcbi.1002678)

12. Calovi DS, Litchinko A, Lecheval V, Lopez U, Escudero AP, Chaté H, Sire C, Theraulaz G. 2018 Disentangling and modeling interactions in fish with burst-and-coast swimming reveal distinct alignment and attraction behaviors. *PLoS Comput. Biol.* **14**, e1005933. (doi:10.1371/journal.pcbi.1005933)

13. Branson K, Robie AA, Bender J, Perona P, Dickinson MH. 2009 High-throughput ethomics in large groups of *Drosophila*. *Nat. Methods* **6**, 451–457. (doi:10.1038/nmeth.1328)

14. Dell AI *et al.* 2014 Automated image-based tracking and its application in ecology. *Trends Ecol. Evol.* **29**, 417–428. (doi:10.1016/j.tree.2014.05.004)

15. Gallois B, Candelier R. 2021 FastTrack: an open-source software for tracking varying numbers of deformable objects. *PLoS Comput. Biol.* **17**, e1008697. (doi:10.1371/journal.pcbi.1008697)

16. Anderson DJ, Perona P. 2014 Toward a science of computational ethology. *Neuron* **84**, 18–31. (doi:10.1016/j.neuron.2014.09.005)

17. Pérez-Escudero A, Vicente-Page J, Hinz RC, Arganda S, De Polavieja GG. 2014 idTracker: tracking individuals in a group by automatic identification of unmarked animals. *Nat. Methods* **11**, 743–748. (doi:10.1038/nmeth.2994)

18. Romero-Ferrero F, Bergomi MG, Hinz RC, Heras FJ, De Polavieja GG. 2019 Idtracker. AI: tracking all individuals in small or large collectives of unmarked animals. *Nat. Methods* **16**, 179–182. (doi:10.1038/s41592-018-0295-5)

19. Walter T, Couzin ID. 2021 TRex, a fast multi-animal tracking system with markerless identification, and 2D estimation of posture and visual fields. *Elife* **10**, e64000. (doi:10.7554/eLife.64000)

20. Escobedo R, Lecheval V, Papaspyros V, Bonnet F, Mondada F, Sire C, Theraulaz G. 2020 A data-driven method for reconstructing and modelling social interactions in moving animal groups. *Phil. Trans. R. Soc. B* **375**, 20190380. (doi:10.1098/rstb.2019.0380)

21. Jayles B, Escobedo R, Pasqua R, Zanon C, Blanchet A, Roy M, Trédan G, Theraulaz G, Sire C. 2020 Collective information processing in human phase separation. *Phil. Trans. R. Soc. B* **375**, 20190801. (doi:10.1098/rstb.2019.0801)

22. Cazenille L, Bredeche N, Halloy J. 2019 Automatic calibration of artificial neural networks for zebrafish collective behaviours using a quality diversity algorithm. In *Proc. of Biomimetic and Biohybrid Systems: 8th Int. Conf., Living Machines 2019, Nara, Japan, 9–12 July*, pp. 38–50. Springer International Publishing.

23. Heras FJH, Romero-Ferrero F, Hinz RC, de Polavieja GG. 2019 Deep attention networks reveal the rules of collective motion in zebrafish. *PLoS Comput. Biol.* **15**, 1–23. (doi:10.1371/journal.pcbi.1007354)

24. Costa T, Laan A, Heras FJ. 2020 Automated discovery of local rules for desired collective-level behavior through reinforcement learning. *Front. Phys.* **8**, 200. (doi:10.3389/fphy.2020.00200)

25. Gilpin W, Huang Y, Forger DB. 2020 Learning dynamics from large biological datasets: machine learning meets systems biology. *Curr. Opin. Syst. Biol.* **22**, 1–7. (doi:10.1016/j.coisb.2020.07.009)

26. Mammadli R, Wolf F, Jannesari A. 2019 The art of getting deep neural networks in shape. *ACM Trans. Archit. Code Optim. (TACO)* **15**, 1–21. (doi:10.1145/3291053)

27. Valletta JJ, Torney C, Kings M, Thornton A, Madden J. 2017 Applications of machine learning in animal behaviour studies. *Anim. Behav.* **124**, 203–220. (doi:10.1016/j.anbehav.2016.12.005)

28. Marques JC, Lackner S, Félix R, Orger MB. 2018 Structure of the zebrafish locomotor repertoire revealed with unsupervised behavioral clustering. *Curr. Biol.* **28**, 181–195. (doi:10.1016/j.cub.2017.12.002)

29. Cazenille L, Chemtob Y, Bonnet F, Gribovskiy A, Mondada F, Bredeche N, Halloy J. 2018 How to blend a robot within a group of zebrafish: achieving social acceptance through real-time calibration of a multi-level behavioural model. In *Proc. of Biomimetic and Biohybrid Systems: 7th Int. Conf., Living Machines 2018, Paris, France, 17–20 July*, pp. 73–84. Springer International Publishing.

30. Alahi A, Goel K, Ramanathan V, Robicquet A, Fei-Fei L, Savarese S. 2016 Social LSTM: human trajectory prediction in crowded spaces. In *2016 IEEE Conf. on Computer Vision and Pattern Recognition* (CVPR) (ed. L O'Conner), pp. 961–971.

31. Kothari P, Kreiss S, Alahi A. 2021 Human trajectory forecasting in crowds: a deep learning perspective. *IEEE Trans. Intell. Transp. Syst.* **23**, 7386-–400. (doi:10.1109/TITS.2021.3069362)

32. Calovi DS, Lopez U, Ngo S, Sire C, Chaté H, Theraulaz G. 2014 Swarming, schooling, milling: phase diagram of a data-driven fish school model. *New J. Phys.* **16**, 015026. (doi:10.1088/1367-2630/16/1/015026)

33. Xue T, Li X, Lin G, Escobedo R, Han Z, Chen X, Sire C, Theraulaz G. 2023 Tuning social interactions' strength drives collective response to light intensity in schooling fish. *bioRxiv*. (doi:10.1101/2023.04.20.537526)

34. Lei L, Escobedo R, Sire C, Theraulaz G. 2020 Computational and robotic modeling reveal parsimonious combinations of interactions between individuals in schooling fish. *PLoS Comput. Biol.* **16**, e1007194. (doi:10.1371/journal.pcbi.1007194)

35. Hochreiter S, Schmidhuber J. 1997 Long short-term memory. *Neural Comput.* **9**, 1735–1780. (doi:10.1162/neco.1997.9.8.1735)

36. Chua K, Calandra R, McAllister R, Levine S. 2018 Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in neural information processing systems* (eds S Bengio, H Wallach, H Larochelle, K Grauman, N Cesa-Bianchi, R Garnett), pp. 4754–4765. Cambridge, MA: MIT Press.

37. Kingma DP, Ba J. 2014 Adam: a method for stochastic optimization. *arXiv*. (http://arxiv.org/abs/1412.6980)

38. Martín JAH, de Lope J, Santos M. 2007 Evolution of neuro-controllers for multi-link robots. In *Innovations in hybrid intelligent systems* (eds E Corchado, JM Corchado), pp. 175–182. New York, NY: Springer.

39. Mwaura J, Keedwell E. 2015 Evolving robotic neuro-controllers using gene expression programming. In *2015 IEEE Symp. Series on Computational Intelligence* (ed. P Kellenberger), pp. 1063–1072. IEEE.

40. Sekaj I, Cíferský L, Hvozdík M. 2019 Neuro-evolution of mobile robot controller. *MENDEL* **25**, 39–42. (doi:10.13164/mendel.2019.1.039)

41. Wang W, Escobedo R, Sanchez S, Sire C, Han Z, Theraulaz G. 2022 The impact of individual perceptual and cognitive factors on collective states in a data-driven fish school model. *PLoS Comput. Biol.* **18**, e1009437. (doi:10.1371/journal.pcbi.1009437)

42. Papaspyros V, Theraulaz G, Sire C, Mondada F. 2023 Quantifying the biomimicry gap in biohybrid systems. *arXiv*. (http://arxiv.org/abs/2308.08978)

43. Papaspyros V, Burnier D, Cherfan R, Theraulaz G, Sire C, Mondada F. 2023 A biohybrid interaction framework for the integration of robots in animal societies. *IEEE Access* **11**, 67640–67659. (doi:10.1109/ACCESS.2023.3290960)

44. Papaspyros V, Bonnet F, Collignon B, Mondada F. 2019 Bidirectional interactions facilitate the integration of a robot into a shoal of zebrafish *Danio rerio*. *PLoS ONE* **14**, e0220559. (doi:10.1371/journal.pone.0220559)

45. Bonnet F, Kato Y, Halloy J, Mondada F. 2016 Infiltrating the zebrafish swarm: design, implementation and experimental tests of a miniature robotic fish lure for fish–robot interaction studies. *Artif. Life Rob.* **21**, 239–246. (doi:10.1007/s10015-016-0291-8)

46. Bonnet F, Gribovskiy A, Halloy J, Mondada F. 2018 Closed-loop interactions between a shoal of zebrafish and a group of robotic fish in a circular corridor. *Swarm Intell.* **12**, 227–244. (doi:10.1007/s11721-017-0153-6)

47. Cazenille L, Collignon B, Chemtob Y, Bonnet F, Gribovskiy A, Mondada F, Bredeche N, Halloy J. 2018 How mimetic should a robotic fish be to socially integrate into zebrafish groups? *Bioinsp. Biomimet.* **13**, 025001. (doi:10.1088/1748-3190/aa8f6a)

48. Lundberg SM, Lee SI. 2017 A unified approach to interpreting model predictions. In *Advances in neural information processing systems,* vol. 30 (eds I Guyon, U Von Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan, R Garnett). Cambridge, MA: MIT Press.

49. Quade M, Abel M, Shafi K, Niven RK, Noack BR. 2016 Prediction of dynamical systems by symbolic regression. *Phys. Rev. E* **94**, 012214. (doi:10.1103/PhysRevE.94.012214)

50. Chen Y, Angulo MT, Liu YY. 2019 Revealing complex ecological dynamics via symbolic regression. *BioEssays* **41**, 1900069. (doi:10.1002/bies.201900069)

51. Papaspyros V, Escobedo R, Alahi A, Theraulaz G, Sire C, Mondada F. 2024 Predicting the long-term collective behaviour of fish pairs with deep learning. *Zenodo*. (doi:10.5281/zenodo.7634912)

52. Papaspyros V, Escobedo R, Alahi A, Theraulaz G, Sire C, Mondada F. 2024 Predicting the long-term collective behaviour of fish pairs with deep learning. *Zenodo*. (doi:10.5281/zenodo.7634687)

53. Papaspyros V, Escobedo R, Alahi A, Theraulaz G, Sire C, Mondada F. 2024 Predicting the long-term collective behaviour of fish pairs with deep learning. Figshare. (doi:10.6084/m9.figshare.c.7090084)