Collective Evolution Learning Model for Vision-Based Collective Motion with Collision Avoidance

David Krongauz[1]◗*, Teddy Lazebnik[2]◗,

1 Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel
2 Department of Cancer Biology, Cancer Institute, University College London, UK

◗These authors contributed equally to this work.
* kingkrong@gmail.com

Abstract

Collective motion takes many forms in nature; schools of fish, flocks of birds, and swarms of locusts to name a few. Commonly, during collective motion the individuals of the group avoid collisions. These collective motion and collision avoidance behaviors are based on input from the environment such as smell, air pressure, and vision, all of which are processed by the individual and defined action. In this work, a novel vision-based collective motion with collision avoidance model (i.e., VCMCA) that is simulating the collective evolution learning process is proposed. In this setting, a learning agent obtains a visual signal about its environment, and throughout trial-and-error over multiple attempts, the individual learns to perform a local collective motion with collision avoidance which emerges into a global collective motion with collision avoidance dynamics. The proposed algorithm was evaluated on the case of locusts' swarms, showing the evolution of these behaviors in a swarm from the learning process of the individual in the swarm. Thus, this work proposes a biologically-inspired learning process to obtain multi-agents multi-objective dynamics.

Introduction

Collective motion (CM) is a phenomenon occurring in homogeneous populations of interacting individuals [1, 2, 3, 4]. It is manifested in multiple forms in nature such as locusts aggregate into groups [5, 6], birds flocking in an organized manner [7, 8], and fish schools responding to predators and showing migratory movements [9, 10]. During the last few decades, researchers seek to elucidate the mechanisms allowing animals to achieve such a remarkably synchronized motion, while lacking centralized control. Empirical data from experimental observations [11, 12] is analysed with the aim to uncover individual's behaviour using mathematical and computational models [13, 14, 15].

A large body of work can be categorized as agent-based models, where the collective dynamics is explained by the action on the individual-level [16]. One of the leading approaches in agent-based simulation for CM is Self-Propelled-Particles. In this group of models [17, 18, 19, 20, 21], agents are moving with constant speed [17], aligning their velocity direction in response to inter-agent interactions. Moreover, a notable subgroup of models is the Avoidance-Attraction-Alignment group [5, 18, 22] where agents are governed by three types of superimposed forces: aligning force (For example see [17]), separating force for collision avoidance, and cohesion to the center of mass of the group.

Although these models reproduce aspects of the macro-dynamics of natural swarms [3, 23, 24], they do not address the problem of information acquisition by an individual. This is a results of the models relying on globally available information. For instance, accurate positions and velocities of agents in the environment are commonly assumed to be fully known to the focal agent [17, 18, 19]. However, in reality, individuals' perception of the environment is mediated by their sensory systems (e.g. vision, hearing, echolocation, magnetic field sensing etc.) [25, 26, 27, 28]. Specifically, objects in the environment are reflected on the retina of the focal agent's eye and later translated to the brain allowing an individual to act upon the visual stimuli it observed [29, 30]. Indeed, many animals use vision as the main sense for coordination and navigation [31, 32, 33]. Hence, to obtain more biologically accurate modeling of collective motion, agents in the swarm are assumed to estimate an object's location and velocity using vision-based sensing.

Multiple models that encompass visual projection in CM modelling were proposed [34, 35, 36, 37]. Strandburg-Peshkin et al. [36] explored visual sensor networks and their role in information transfer between schools of fish. Bastien and Romanczuk [38] used elementary occlusions of visual rays to create avoidance and self-organizations in disc-like shaped agents. The authors demonstrated several group behaviors with steering by angular visual occlusions. A follow-up work introduced an additional collision avoidance and demonstrated collective behavior in the presence of static obstacles in several settings [39]. Both models are oblivious to object recognition, thus any agent-specific features are ignored. In the present work, we propose and demonstrate a different approach inwhich the shape and size of the agents in the swarm is effective the decision-making process due to the vision-based input.

Several attempts to model CM using learning models are presented [40, 41, 42]. Durve et al. [40] developed both Vicsek CM teacher-based reinforcement learning and concurrent learning. The authors conclude that the obtained results show that the velocity alignment mechanism allows spontaneously emerge from the minimization of neighbor-loss (e.g., agents that are not in the visual radius) rate a strategy to keep group cohesion when the perception is limited to the velocity of neighboring agents (e.g., agents that can be seen from the focal agent) [40]. They assume the focal agent is accurately aware of the location and velocities of the neighboring agents. Into the bargain, Young and La [41] presented a hybrid system that achieves an efficient CM behavior for escaping attacking predators while maintaining a flocking formation using multiple reinforcement learning methods. The authors show that independent learning with and without function approximation proved to be unreliable in learning to flock together towards the same target. Similarly, the authors of [41], assume the focal agent is accurately aware of the location and velocities of the neighboring agents while solving a multi-objective task. Nonetheless, they were not able to obtain a collective behavior by learning a policy at the individual level as obtained by Durve et al. Furthermore, Lopez-Incera et al. [42] studied the collective behavior of artificial learning agents driven by reinforcement learning action-making process and with abstract sensing mechanism, that arises as they attempt to survive in foraging environments. They design multiple foraging scenarios in one-dimensional worlds in which the resources are either near or far from the region where agents are initialized and show the emergence of CM [42]. While the authors took into consideration a sensing-based approach they investigated a single objective behaviour in an extremely simplistic scenario.

The novelty of the work presented here, lies in the development of a collective evolution learning (CEL) algorithm that achieves both CM and CA using the vision-based approach, donated by VCMCA. The VCMCA model is obtained by simulating the evolution process of an animal population that learns over time using trial-and-error. We show that the proposed algorithm can learn a multi-agent

multi-objective behavior by optimizing an individual's objective to achieve a collective (swarm) objective. Furthermore, we evaluate the performance of the algorithm to obtain a VCMCA model in different settings, showing robust results.

This paper is organized as follows: Section , outlines the mathematical formalization for the CM with CA in a two-dimensional setting with agents that based their decisions on a vision mechanism. In Section 3, we propose a new computational model for vision-based CM with CA, called VCMCA. In Section 4, we evaluate the performance of VCMCA using several *in silico* experiments. In Section 5, we discuss the main advantages and limitations of the model and propose future work.

Vision-Based Collective Motion with Collision Avoidance

The model (M) is defined by a tuple $M := (E, P)$ where $E \subset \mathbb{R}^2$ is the environment and P is a population of agents operating in the environment. Formally, the environment, E, is a continuous space with height $H \in \mathbb{N}$ and width $W \in \mathbb{N}$ with a toroidal geometrical configuration (i.e., a two-dimensional box with periodic boundary conditions). Each agent, $p \in P$, is defined by a tuple $p := (l, v, a, h, w, r, \gamma)$ where $l = (l_x, l_y) \in E$ is the location of the agent's center of mass, $v \in \mathbb{R}^2$ is the current velocity vector of the agent, $a \in \mathbb{R}^2$ is the current acceleration vector of the agent, h is the agent's height, w is the agent's width, r is the vision-based sensing radius of the agent, and $\gamma := (\psi, \dot{\psi}, \theta, \dot{\theta})$ is the vision mechanism of the agent such that ψ is the subtended angle caused by the reflection of a neighbor agent on the visual sensor (e.g. retina, camera sensor) of the focal agent, $\dot{\psi}$ is the time derivative of ψ, θ is bearing angle of a neighbor agent and the focal agent, and $\dot{\theta}$ is the time derivative of θ. The agents have a rectangle geometrical configuration. This formalization is inspired by [43].

Based on the sensing radius, r, we define the *local environment* of an agent $p_i \in P$ to be the subset of agents P_i^* that are located in a sensing distance from the focal agent, formally defined as $\{p_j \in P_i^* : |l_j - l_i| < r_i\}$. Each agent, $p_j \in P_i^*$, can be sensed using the vision mechanism if it is not fully latent by other agents $P_s \subset P_i^*$. The vision-based properties (γ) are computed for each neighbor agent in the *local environment* of the focal agent, resulted in a vector $\nu \in \mathbb{R}^{|\gamma| \cdot |P^*|}$, such that $\nu := [\gamma_1, \ldots, \gamma_{|P^*|}]$. The parameter ψ for each neighbor agent is obtained as follows
$\psi := \max(\arccos \frac{\chi_1 \cdot \chi_2}{\|\chi_1\| \|\chi_2\|}, \ \arccos \frac{\chi_3 \cdot \chi_4}{\|\chi_3\| \|\chi_4\|})$ where $\chi_{i=0}^4 = \rho_i - l$ and
$\rho_i = l + (-1)^i \frac{v_n \cdot w}{2\|v_n\|} + (-1)^{i+1} \frac{v_n^\perp \cdot h}{2\|v_n^\perp\|}$ such that $v_n^\perp$ is the perpendicular vector to the velocity vector v of a neighbor. In addition, the parameter θ is defined by
$\theta = \arccos \frac{\mathbf{v_o} \cdot (\mathbf{v} - \mathbf{l_n})}{\|\mathbf{v_o}\| \|\mathbf{v} - \mathbf{l_n}\|}$ where l_n is the position of the nth neighbor's center of mass. The definitions of ψ and θ are formalization of the vision parameters proposed by [38, 43, 44].

An agent can be fully latent to the focal agent by other agents in its local environment. In such case, it is not taken into consideration in the vision-based computation. A schematic view of the focal's agent vision-based sensing (γ) is shown in Fig. 1a.

The model has a discrete, synchronized clock that all the agents follow. At the beginning of the simulation (t_0), the locations and velocities of the agents' population (P) are chosen at random. Then, for each clock tick (marked by t_i for the i_{th} tic) the following happens: first, each agent $p \in P$ senses its local environment and computes its action $\xi(t) \in \mathbb{R}^2$, in a random order. Second, each agent moves according to:

$$a(t_{i+1}) = \xi(t_i), \ v(t_{i+1}) = v(t_i) + \delta a(t_{i+1}), \ l(t_{i+1}) = l(t_i) + v(t_{i+1}), \qquad (1)$$

where $\delta \in \mathbb{R}^+$ is a weighted variable indicates the influence of the desired direction on the current direction such that if $l \notin E$ then $l_x^{new} = mod(l_x, W)$ and $l_y^{new} = mod(l_y, H)$.

In case two agents $p_1 \in P$, $p_2 \in P$ collide, a collision protocol is employed: First, the simulation identifies the both colliding agents to be either *active* or *passive* colliders, e.g., in a head-tail collision the head colliding agent is regarded as the active and the tail colliding as the passive (since the latter could not have actively avoided the collision). In the case of a frontal collision both parties are regarded as active colliders. The active colliding agent(s) are stopped (i.e., $v(t_{i+0.5}) = 0$) while the passive colliding agent's velocity follows purely-elastic collision:

$$v(t_{i+0.5}) = v(t_i) + v_{colliding}(t_i), \tag{2}$$

where $v_{colliding}(t_i)$ is the velocity of the colliding agent at the time of the collision (t_i). A schematic view of the perception-action making loop of a single agent is presented in Fig. 1b.

(a) Spatial view from a single agent's point of view.

(b) Computation interaction from a single agent's point of view.

Fig 1. A schematic view of the vision sensing mechanism of an agent. The computation is presented for a single focal agent but performed symmetrically for all the agents in the population.

We define the agents' population CM metric by the degree of alignment of the velocity headings of the simulated agents, in similar fashion to other CM models' metrics [17, 5]. Formally, the CM metric is a function $F : \mathbb{P} \to \mathbb{R}$ where $\mathbb{P}$ is the space of all possible agent populations. $F(P)$ is defined as follows:

$$F(P) := \frac{1}{2\pi} \sqrt{\frac{\Sigma_{p \in P}(v - \bar{v})^2}{|P|}}, \tag{3}$$

where $\bar{v} := \frac{\Sigma_{p \in P}(v)}{|P|}$.

In addition, we define a collision between two agents $p_1 \in P$ and $p_2 \in P$ taking place if and only if:

$$||l_1 + v_1 - l_2 - v_2|| \le max(h_1 + h_2, w_1 + w_2)/2. \tag{4}$$

Therefore, a function $C : \mathbb{P} \to \mathbb{N}$ such that $C(P)$ counts the rate of collisions between the agents of the population to the possible number of collisions in a single point of time, is used as the CA metric.

Based on these two metrics, the objective of the CM with CA takes the form:

$$\min_{\omega \in \Omega} \left(\Sigma_{i=0}^{\infty} \big(\zeta F(P) + (1 - \zeta)C(P) \big) \right), \tag{5}$$

where Ω is the space of all possible policies of the agents and $\zeta \in [0, 1]$ is the weight of the CM score compared to the CA score.

Collective Evolution Learning Algorithm

The CM behavior of locusts is believed to develop over time throughout an evolution process in order to adapt to their environment and generate the next population of locusts. During this process, multiple locusts experienced similar situations and operate slightly differently based on unique properties each locust have which yielded different outcomes. Locusts that perform better were more likely to reproduce and keep their behavior over time. As such, the wanted behavior (e.g., CM with CA) is emerged throughout the collective experience of the locust and improved over time using a natural trial-and-error approach.

Based on this motivation, we propose an algorithm (A) for vision-based collective motion with collision avoidance (VCMCA) that receives the observable parameters γ of a subset of agents $P^* \subset P/\{p\}$ in the local environment R of an agent p and returns the action ξ that best approximate the optimal action of the MC with MA objective metric (Eq. (5)) is proposed. The motivation is that the CM is a result of the attempt of each individual to optimize its local objective. A schematic view of the model training, validation, and testing phases with the interaction between them is presented in Fig. 2.

Fig 2. A schematic view of the model training, validation, and testing phases with the interaction between them.

Formally, the model is a function $A : \mathbb{R}^{4k} \to \mathbb{N}$ such that $A(\nu) \to [\xi_1, \dots, \xi_n]$ where $k \in \mathbb{N}$ is the maximum possible number of agents that an agent $p \in P$ can sense in a single point of time multiplied by the number of vision-based parameters per agent (recall, this model uses $\psi, \theta, \dot{\psi},$ and $\dot{\theta}$), $\nu \in \mathbb{R}^{4k}$ is the outcome of the vision-based sensing of an agent $p \in P$, and $\{x_i\}_{i=1}^n$ is the set of all possible actions an agent can make.

One can obtain algorithm A by simulating the CEL process of individuals in the context of their population. Intuitively, we generate at random different cases in which an agent is surrounded by other agents in its local environment. At this point, the agent is applying multiple actions and remembering the score of each one of them. Actions that outperform others derived a higher chance of the agent to pass this behavior to next generations, resulting in a collective learning process [45, 46]. Mathematically, we define a local, individual objective function, based on the collective metric Eq. (5), as follows:

$$\min_{\omega \in \Omega} \left(\Sigma_{i=t_i}^{t_{i+1}} \left(\zeta F(P^*) + (1-\zeta)C(P^*) \right) \right), \qquad (6)$$

where $P^* \subset P$ is the subset of agents in the local environment of a focal agent $p \in P$. Of note, Eq. (6) assume a semi-Markovian property [47] to the action making processes of the agents which means that the agents are aiming to perform the optimal action for the next state, taking into consideration the current and previous states.

The collective learning is performed using the Q-learning approach [48], stochastically mapping each state ν to the set of possibles actions $\{x_i\}_{i=1}^n$ by assuming a discretization of both the input and output spaces. Hence, we assume that the number of possible actions $n \in \mathbb{N}$ is large but finite and that $\nu \in \mathbb{N}^{4k}$ obtained by the discretization of each vision parameter in order to define a classification problem rather than a regression problem.

The proposed algorithm (A) is trained and evaluated in a three-phase structure, including training, evaluation, and testing phases. The training and evaluation phases are executed sequentially multiple times until a stop condition is satisfied. Afterward, the obtained VCMCA model is evaluated during the testing phase.

In the training phase, we generate at random a wide range of possible states (cases) and compute the action ξ that satisfies Eq. (6) using the genetic algorithm (GA) approach [49]. In particular, we generate a set of random action vectors as the possible population of actions (also refer to as *genes*), $\{s\}_{i=1}^x$ where $x \in \mathbb{N}$ is the size of the set of actions in the GA. We perform g iterations (generations) such that in each generation we computed a cross-over between the genes and a mutation for each gene. Afterward, we compute the fitness of each gene and using a tournament with royalty operator to obtain the next population of genes. To be exact, the cross-over operator $COO(s_1, s_2, w) \rightarrow (s_1', s_2')$ gets two actions s_1, s_2 and a weight $w \in [0, 1]$ and returns two new actions (s_1', s_2') such that $s_1' = (1 - w) \cdot s_1 + w \cdot s_2$ and $s_2' = w \cdot s_1 + (1 - w) \cdot s_2$. The mutation operator $MO(s, \mu, \phi) \rightarrow s'$ gets an action s, a mean μ and standard deviation ϕ for a normal-distributed random variable $Y \sim N(\mu, \phi)$ and returns an new action s' such obtained by the sum of s and a vector y sampled from Y. The next generation operator (NGO) leave a portion $\rho \in [0, 1]$ of the genes with the highest fitness score while the remaining genes are chosen by a distribution defined by the L_1 normalization of the fitness scores of the remaining genes. The action ξ with the smallest cosine similarity to the gene with the best fitness score after the last generation is taken be the action for the Q-learner during the training phase.

The generation of states is repeated $Z_{train} \in \mathbb{N}$ times. Afterward, during the evaluation phase, $Z_{validation} \in \mathbb{N}$ states are generated and evaluated using the model trained up to this point. Based on the evaluation scores of the states, we compute a distribution over the state space such that the sampling probability of a state S is defined to be a L_1 norm distance, weighted average, of the three closest states sampled during the evaluation phase. The training and evaluation phases are repeated until the best ($max(e)$) and worse scores ($min(e)$) in the evaluation processes satisfies $max(e) - min(e) < \tau$ where $\tau \in \mathbb{R}$ is a pre-defined threshold value.

Once the VCMCA model is obtained, as the testing phase, $Z_{testing} \in \mathbb{N}$ simulations with different initial conditions are computed and used to evaluate the model's performance using Eq. (5).

Results

To evaluate the ability of the CEL algorithm to obtain a VCMCA model and explore the influence of several properties of the agents on the model's performance, we conduct three experiments. First, measure the performance of the proposed algorithm compared to two other approaches - vision-based CM without CA and random walk. Second, we analyze the influence of population density, agent's height to width ratio, and sensing radius on the VCMCA model's performance. Finally, we approximate the VCMCA

model obtained by the CEL algorithm using an explainable Random Forest [50] model
to extract rule-base structure for the agent's action-making process, allowing the user to
interpret the connection between the visual stimuli and the reaction of the agent.

Setup

We implemented the VCMCA model and CEL algorithm as a simulator for the case of
locust. We chose locust since at the beginning of the flocking, they move on the ground,
which is approximated by a 2D plan. However, the model can be adapted to other
species as well. The simulation is divided into two phases: 1) agent's individual
VCMCA model training using the CEL algorithm; and 2) usage of the obtained
VCMCA model in the locust swarm simulation. The default hyperparameters' values
used in the simulation are picked to represent a population of locusts that becoming
locusts. The parameters with their descriptions and values are summarized in Table 1.
If not stated otherwise, these hyperparameters' values are used across all experiments.
In particular, the agent's height to width is set to be 7/1 by calculating the average
ratio of 50 random, upper-view locust images that resulted in 6.78 ± 1.48. The agents'
dimensions were chosen based on the *Schistocerca cancellata* and *Schistocerca gregaria*
locusts' species that are known to exhibit swarming behaviours and density-dependent
phase transition [51, 52]. In addition, group population sizes were set with accordance
to available biological records of locust swarms [53, 52]. Though the exact density
values are one order of magnitude larger than the simulations, we concentrate on the
emerging phases of the flock when the density is yet to reach its maximal numbers. The
model uses abstract discrete time steps. To calibrate the simulation's abstract time step
into a real one, we define each time step to be the average duration taking the average
locust (agent) to make a decision and move [54].

ECL algorithm training

We train the locusts' VCMCA model using the ECL algorithm for $n = 30$ times in order
to evaluate the average relationship between the number of training samples and the
obtained VCMCA model's performance due to the stochastic nature of the algorithm.
In Fig. 3 the blue line indicates the training fitness' function value (i.e., Eq. (6)), and
the CMCA score (i.e., Eq. (5)). The results are shown as mean $\pm$ standard deviation.
The vertical dashed-yellow lines indicate the validation phase during the training
process as described in Fig. 2.

We fitted the average training fitness and CMCA score using the least mean square
method, obtaining $F_1 : fitness = 0.189ln(samples) + 0.053$ and
$F_2 : cmca = 0.196ln(samples) - 0.268$ with coefficient of determination $R^2 = 0.989$ and
$R^2 = 0.972$, respectively. Hence, the functions F_1 and F_2 are well-fitting the dynamics.
Since both functions are monotonically increasing, it shows that there is a connection
between the CMCA score of an individual and the population's CMCA score. In a
similar manner, the Pearson correlation between the training fitness and CMCA scores
in the validation stages is $R^2 = 0.971$ with a p-value of 0.0057. Thus, the scores are
strongly linearly correlated as well.

Screenshots of the simulator at the beginning and end of the simulation of an
average performance VCMCA model after 50 million training samples are shown in Figs.
4a and 4b, respectively. Each agent is represented by a black rectangle with a
surrounding gray circle that represents the sensing radius of the agent and their index
number. The agents marked in dashed blue indicate that they have collided. One can
see that in Fig. 4a that the agents are heading to random directions while in Fig. 4b
the agents heading the same direction (top-right corner), obtaining CM.

Parameter definition	Symbol	Value
Environment's (E) width [m]	W	5
Environment's (E) height [m]	H	5
Population size [1]	$|P|$	30
Agent's initial velocity [m/s]	v	$(U[0,15], U[0,15])$
Agent's initial acceleration [m/s^2]	a	$(0,0)$
Agent's height [m]	h	0.07
Agent's width [m]	w	0.01
Agent's action influence [1]	δ	0.05
Agent's sensing radius [m]	r	0.21
CM to CA weight [1]	ζ	0.5
The number of steps in time for a single simulation [1]	T	500
CEL algorithm's maximum number of agents in the local environment [1]	-	8
CEL algorithm's state (input) space discretization [1]	-	0.01
CEL algorithm's action (output) space discretization [1]	-	15 degrees
CEL algorithm's training samples [1]	Z_{train}	10,000,000
CEL algorithm's validation samples [1]	$Z_{validate}$	100,000
CEL algorithm's testing samples [1]	$Z_{testing}$	450,000
CEL algorithm's validation stop condition parameter [1]	τ	0.05
CEL algorithm's genetic algorithm mutation mean and standard deviation [1]	μ, ϕ	0.1, 0.05
CEL algorithm's genetic algorithm number of generations [1]	g	100
CEL algorithm's genetic algorithm royalty rate [1]	-	0.05

Table 1. The description of the model's parameters and their default values.

Baseline

We used the obtained VCMCA model to evaluate the CM (Eq. (3)), CA (Eq. (4)), and CMCA (Eq. (5)) scores over time. We compared these results to a model obtained in the same why but with $\zeta = 1$ rather than $\zeta = 0.5$. Namely, the CEL algorithm aims to optimize the CM while ignoring CA. In addition, we compare the proposed model with random walk. The random walk model assumes a uniform distribution probability for each action and pick one each time in random. The results of the CM, CA, and CMCA scores are shown as mean $\pm$ standard deviation of $n = 30$ simulations in Figs. 5a, 5b, and 5c, respectively.

As expected, both the CM and VCMCA models obtain flocking with a CM score of 0.24 and 0.27 while the random walk converge to 0.61, as presented in Fig. 5a. On the other hand, while the VCMCA model converge to 0.18 collision rate (e.g., 0.82 CA) the CM and random walk models obtained 0.53 and 0.62, as shown in Fig. 5b. Of note, the CM model obtain slightly worse ($\sim 17\%$) CA score due to the flocking which makes the agents closer as shown by [40].

In addition, the mean results of each model over time are shown in Fig. 6 such that the results are the mean $\pm$ standard deviation computed on the values presented in Figs. 5a, 5b, and 5c, respectively. As seen in Fig. 5a, the agents' convergence to cohesive movement, using the CM and CMCA models, emerges gradually with the collective motion index decreasing with slowing rates over time. This process is similar to displays

Fig 3. Performance of the learner as training progresses. The training fitness is the agent's score in the testing phase. The CMCA score is computed by using the trained agent in the simulation. The results are shown as mean $\pm$ standard deviation of $n = 30$ repetition. The vertical dashed-yellow lines indicate the validation phase during the training process. The training is converged after five validations for $\tau = 0.05, Z_{training} = 10000000, Z_{validate} = 100000$.

(a) At the beginning of the simulation ($t = 0$). **(b)** At the end of the simulation ($t = 500$).
Fig 4. A screenshot of the simulation. Blue-dashed agents are collided in the current simulation step.

of hopper-bands formations, in which desert locust nymphs gather together into flocks of thousands of individuals. Due to the initially disordered state, it takes time for the visual information to be mediated between the swarm-members and them changing and adapting their heading on an individual level, which later results in the swarm reaching coordinated heading direction on a group level. Another notable aspect is the collision rate (Fig. 5b) for CM and Random models reaching its final value in a relatively short period of time, in contrast to CMCA model that increases moderately over the time of the simulation, reaching a saturation value three times smaller than CM only. The gradual rise in collision rate in CMCA could be attributed to the CM component causing the agents to align and follow each other, in turn bringing them closer together.

(a) Collective motion (CM).

(b) Collision avoidance (CA).

(c) Collective motion with collision avoidance (CMCA) score.

Fig 5. The collective motion, collision avoidance, and CMCA score of the VCMCA algorithm, VCMCA algorithm without CA, and random walk over the simulation steps for $n = 30$ repetitions.

These results are aligning with the biological inspiration since flock-members facilitate coherent movement partly for their collision-avoidance abilities [55, 56], preventing them from constantly colliding into each other, which could potentially cause high mortality in flocking animals.

Parameter sensitively analysis

To estimate the sensitivity of the population density, agent's height to width ratio, and sensing radius on the VCMCA model's performance. Given the trained VCMCA model from the ECL algorithm, we computed the CMCA score over $n = 30$ simulations with different population sizes, as presented in Fig. 7. The results are shown as mean ± standard deviation.

The results clearly indicate an optimal flock size ($P = 40$), suggesting a strong link between flock density and flocking emergence. For small population sizes, the density is insufficient for phase transformation to occur since interactions between the moving agents are scarce and happen across large time intervals, thus preventing the 'directional cluster' aggregation. In other words, agents' behaviour is similar to random-walk since local interactions do not translate to the group level. An increase in population size allows for the accumulation of aligning inter-agent interactions, in which a small cluster of agents' incoherent movement gradually aligns the rest of the swarm

Fig 6. A mean ± standard deviation of the CMCA, CM, and CA scores for CMCA model, CM model, and random walk for $n = 30$ repetitions.

resulting in a large-scale coordinated motion reminiscent of natural flocks. Further increasing the population size leads to an inherent raise of the collision rate or CA maneuvers, where agents are steering away from an aligned direction in order to avoid an imminent collision with an approaching neighbor. Since our model strives to minimize the collision rate, relatively large densities hinder the aligning mechanism since CA mechanism dominates over CM.

Fig 7. A mean ± standard deviation of the CMCA score (Eq. (5)) as a function of the population size ($|P|$) for $n = 30$ repetitions.

In a similar manner, we computed the influence of the agents' height to width ratio (e.g. h/s) on the CMCA score over $n = 30$ simulations, as shown in Fig. 8.

Expectedly, larger height-to-width ratios yield lower CMCA score. As mentioned in Section , elongated agent morphologies, opposed to circular, complicate the processing of visual stimuli. Neighbor's orientations greatly affect the subtended angle ψ, especially in longer agents, since the same agent with the same distance to its center could have highly differing values of ψ. For example, observing an agent's head or bottom vs. observing its side will greatly change the input information and introduce significant noise into the state domain. Moreover, Fig. 8 shows that square-shaped agents got the

Fig 8. A mean $\pm$ standard deviation of the CMCA score (Eq. (5)) as a function of the agent's ratio between the height and the width ($\frac{h}{w}$) for $n = 30$ repetitions.

Fig 9. A mean $\pm$ standard deviation of the CMCA score (Eq. (5)) as a function of the agent's sensing radius (r) for $n = 30$ repetitions.

highest CMCA score. One can conclude that square, with its similarity to a disc, can artificially improve the convergence of a vision-based model [38, 39], since it eliminates the problem of orientational noise with its one-to-one function of distance to the angular area. However, such agent morphologies are biologically inaccurate, especially when addressing flocking animals as locusts, starlings, fish, etc [57, 58, 59].

In a similar manner, we computed the influence of the agents' sensing radius (r) on the CMCA score over $n = 30$ simulations, as shown in Fig. 9, such that the x-axis indicates both the sensing radius in absolute size and as the multiplier of the agent's height.

Explainablility of the model

The proposed model is the outcome of multiple stochastic processes, converging to collective and intelligent behavior. In this section, we aimed to measure the ability of a simple model such as Decision Tree (DT) or even its more complex but still explainable extension, the Random Forest (RF) model. First, we generate 1000000 cases at random, in the same way we have done for the testing phase of the CEL algorithm (see Section) and query the model, resulting in a corresponding list of actions for each state. Based on this data, we train RF models with 1000 trees in order to allow a large number of parameters, allowing to fit complex dynamics [50]. The obtained RF models' accuracy relative to the VCMCA model with a restriction of the number of leaf nodes and tree's maximum depth are shown in Figs. 10b and 10a, respectively.

(a) Decision tree's depth.

(b) Decision tree's leaf nodes count.

Fig 10. Random Forest model with 1000 trees approximated the proposed algorithm on the same set of samples. The percent of accuracy to reproduce the proposed model on the default configuration as a function of tree's depth and number of leaf nodes as mean ± deviation of each decision tree.

Unsurprisingly, as more leaves are available as well as more depth to the DTs, the better the RF model approximates the VCMCA model. However, after 320 leaf nodes (or max-depth of 8) there is no significant improvement in the model's performance over 93.37%, indicating that 6.63% of the cases are anomalies. Moreover, there is a sharp improvement between 70 leaf nodes and 80. This sharp improvement is associated with the RF model's to differ between keeping the current vector of movement or rotating 15 degrees (e.g., the smallest turn possible in the default configuration of the model) to either the right or the left. These three actions seem to be very common and the correct separation between them causes a large improvement in performance.

Discussion

In this work, we proposed a vision-based collective motion with collision avoidance (VCMCA) model in a two-dimensional space. Furthermore, we proposed a collective evolution learning (CEL) algorithm that simulates the evolution process of an individual in the context of the group (population) aiming to optimize a local VCMCA task. Individuals better perform in the task at each generation contributing to the collective knowledge.

We used the VCMCA model obtained from the CEL algorithm in the case of locust, showing that multi-objective swarm behavior, such as CM with CA, emerges from the optimization of the individual-level objective as shown in Figs. 3 and 6. Moreover, we show that the results are robust over the population's density, agents' shape, and sensing radius as shown in Figs. 7, 8, and 9, respectively. In particular, we show that disc-like agents with vision-based input are obtaining better results than agents with more complex topology, reproducing [38]'s results.

In addition, by approximating the obtained VCMCA model using explainable machine-learning-based models we show that small rule-based models are not able to capture the complexity of the VCMCA model and it is requiring several hundreds of rule combinations to capture only 93.3% of the VCMCA model's ability while considered out of the scope of the average person or expert [60]. While the proposed model is not explainable on the single decision level, the overall behavior well agrees with known behaviors such as the ones described by [42].

These results agree with previous outcomes proposed by [40] which shows that

agents who learned secondary goals such as keeping contact with neighbors were able to keep CM. These results raise the question that if CM or CA are the primary objective of the individual, as assumed in this work, or other biological objectives such as reducing their risk of being predated or increasing their probability of finding resources are the source of these behaviors.

Moreover, in this work, we tackle a swarm task that assumes a group of identical agents. However, in nature, agents in the population differ by their size, capability, and even action-making process. This diversity may result in more behaviors between the agents such as picking a leader and introducing new and complex challenges for the model.

The proposed question and many others can be addressed by introducing modifications to the proposed VCMCA model and ECL algorithm proposed in this work. As such, we believe that the proposed approach can be used as a computational background to investigate a more biologically accurate representation of swarm dynamics.

Acknowledgments

The authors wish to thank Olga Chervova and Vitaly Voloshin for the thought provoking discussions.

References

[1] H. Chate, F. Ginelli, F. Peruani, and F. Raynaud. "Modeling collective motion: variations on the Vicsek model". In: *The European Physical Journal B* 64 (2008), pp. 451–456.

[2] G. Baglietto and E. V. Albano. "Nature of the order-disorder transition in the Vicsek model for the collective motion of self-propelled particles". In: *Physical Review E* 80 (2009), p. 050103.

[3] A. Deutsch, G. Theraulaz, and T. Vicsek. "Collective motion in biological systems". In: *Interface Focus* 2.6 (2012), p. 050103.

[4] W. H. Warren. "Collective Motion in Human Crowds". In: *Current Directions in Psychological Science* 27.4 (2018), pp. 232–240.

[5] G. Ariel and A. Ayali. "Locust Collective Motion and Its Modeling". In: *PLOS computational Biology* 11.12 (2015), e1004522.

[6] L. Dyson, C. A. Yates, J. Buhl, and A. J. McJane. "Onset of collective motion in locusts is captured by a minimal model". In: *PLOS computational Biology* 92 (2015), p. 052708.

[7] M. Bellerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. "Empirical investigation of starling flocks: a benchmark study in collective animal behaviour". In: *Animal Behaviour* 76.1 (2008), pp. 201–215.

[8] I. L. Bajec and F. H. Heppner. "Organized flight in birds". In: *Animal Behaviour* 78.4 (2009), pp. 777–789.

[9] B. H. Seghers. "Schooling Behavior in the Guppy (Poecilia reticulata): An Evolutionary Response to Predation". In: *Evolution* 28.3 (1974), pp. 486–489.

[10] D. Weihs. "The hydrodynamics of dolphin drafting". In: *Journal of Biology* 3 (2004).

[11] A. Cavagna, A. Cimarelli, I. Giardina, G. Parisi, R. Sanagati, F. Stefanini, and R. Tavarone. "From Empirical Data to Inter-Individual Interactions: Unveiling the Rules of Collective Animal Behavior". In: *Mathematical Models and Methods in Applied Sciences* 20 (2010), pp. 1491–1510.

[12] R. Lukeman, Y-X. Li, and L. Edelstein-Keshet. "Inferring individual rules from collective behavior". In: *PNAS* 107.28 (2010), pp. 12576–12580.

[13] G. Gregorie, H. Chate, and Y. Tu. "Moving and staying together without a leader". In: *Physica D* 181 (2003).

[14] M. Bando, K. Hasebe, A. Nakayama, A. Shibata, and Y. Sugiyama. "Dynamical model of traffic congestion and numerical simulation". In: *Physical Review E* 51.2 (1995), pp. 1035–1042.

[15] W. Ebeling and L. Schimansky-Geier. "Swarm dynamics, attractors and bifurcations of active Brownian motion". In: *European Physical Journal Special Topics* 157 (2008), pp. 17–31.

[16] L. Edelstein-Keshet. "Mathematical models of swarming and social aggregation". In: *Proceedings of the 2001 International Symposium on Nonlinear Theory and Its Applications, Miyagi, Japan*. Citeseer. 2001, pp. 1–7.

[17] T. Vicsek, A. Czirok, E. Ben-Jacob, I. Cohen, and O. Shochet. "Novel type of phase transition in a system of self-driven particles". In: *Physical Review Letters* 75.6 (1995), pp. 1226–1229.

[18] I. Aoki. "A simulation study on the schooling mechanism in fish". In: *Bulletin of the Japanese Society of Fisheries* 48 (1982).

[19] A. Czirok and T. Vicsek. "Collective behavior of interacting self-propelled particles". In: *Physica A* 281 (2000), pp. 17–29.

[20] P. Romanczuk, I. D. Couzin, and L. Schimansky-Geier. "Collective motion due to individual escape and pursuit response". In: *Physical Review Letters* 102.1 (2000).

[21] A. J. Wood and G. J. Ackland. "Evolving the selfish herd: emergence of distinct aggregating strategies in an individual-based model". In: *Proceedings of Royal Society B* 274 (2007), pp. 1637–1642.

[22] C. W. Reynolds. "Flocks, herds and schools: A distributed behavioral model". In: *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*. 1987, pp. 25–34.

[23] I. Giardina. "Collective behavior in animal groups: theoretical models and empirical studies". In: *HFSP Journal* 2.4 (2008), pp. 205–219.

[24] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. "Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study". In: *PNAS* 105.4 (2008), pp. 1232–1237.

[25] D. W. E. Sankey, R. F. Storms, R. J. Musters, T. W. Russell, C. K. Hemelrijk, and S. J. Protugal. "Absence of "selfish herd" dynamics in bird flocks under threat". In: *Current Biology* 31 (2021), pp. 3192–3198.

[26] J. M. Anderson, T. M. Clegg, L. V. M. V. Q. Veras, and K. N. Holland. "Insight into shark magnetic field perception from empirical observations". In: *Scientific Reports* 7 (2017), p. 11042.

[27] G. Steiner. "Orientation and Sense of Smell in Animals". In: *Universitas* 3.1 (1959).

[28] S. Heinze and U. Homberg. "Linking the Input to the Output: New Sets of Neurons Complement the Polarization Vision Network in the Locust Central Complex". In: *Journal of Neuroscience* 29.15 (2009), pp. 4911–4921.

[29] R. H. Masland and P. R. Martin. "The unsolved mystery of vision". In: *Current Biology* 17.15 (2007), RS577–RS582.

[30] M. O'Shea and J. L. D. Williams. "The Anatomy and Output Connection of a Locust Visual lnterneurone". In: *J. comp. Physiol* 91 (1974), pp. 257–266.

[31] J. Zeli, N. Borddeker, and J. M. Hemmi. "Vision and the organization of behaviour". In: *Current Biology* 18.8 (2008), R320–R323.

[32] J. R. Serres and F. Ruffier. "Optic flow-based collision-free strategies: From insects to robots". In: *Arthropod Structure & Development* 46.5 (2017), pp. 703–717.

[33] I. G. Ros and A. A. Biewener. "Pigeons (C. livia) follow their head during turning flight: Head stabilization underlies the visual control of flight". In: *Frontiers in Neuroscience* 11 (2017), p. 655.

[34] B. H. Lemasson, J. J. Anderson, and R. A. Goodwin. "Collective motion in animal groups from a neurobiological perspective: The adaptive benefits of dynamic sensory loads and selective attention". In: *Journal of Theoretical Biology* 261.4 (2009), pp. 501–510.

[35] B. H. Lemasson, J. J. Anderson, and R. A. Goodwin. "Motion-guided attention promotes adaptive communications during social navigation". In: *Proceedings of the Royal Society B: Biological Sciences* 280.1754 (Mar. 2013), p. 20122003.

[36] A. Strandburg-Peshkin, C. R. Twomey, N. W. F. Bode, A. B. Kao, Y. Katz, C. C. Ioannou, S. B. Rosenthal, C. J. Torney, H. S. Wu, S. A. Levin, and I. D. Couzin. "Visual sensory networks and effective information transfer in animal groups". In: *Current Biology* 23.17 (2013), R709–R711.

[37] B. Collignon, A. Séguret, and J. Halloy. "A stochastic vision-based model inspired by zebrafish collective behaviour in heterogeneous environments". In: *Royal Society open science* 3.1 (2016), p. 150473.

[38] R. Bastien and P. Romanczuk. "A model of collective behavior based purely on vision". In: *Science advances* 6.6 (2020).

[39] J. Qi, L. Bai, Y. Xiao, Y. Wei, and W. Wu. "The emergence of collective obstacle avoidance based on a visual perception mechanism". In: *Information Sciences* 582 (2022), pp. 850–864.

[40] M. Durve, F. Peruani, and A. Celani. "Learning to flock through reinforcement". In: *Physical Review E* 102 (2020), p. 012601.

[41] Z. Young and H. M. La. "Consensus, cooperative learning, and flocking for multiagent predator avoidance". In: *International Journal of Advanced Robotic Systems* (2020).

[42] A. Lopez-Incera, K. Ried, and H. J. Briegel. "Development of swarm behavior in artificial learning agents that adapt to different foraging environments". In: *PLoS ONE* 15.12 (2020), e0243628.

[43] S. B. Rosenthal, C. R. Twomey, A. T. Hartnett, H. S. Wu, and I. D. Couzin. "Revealing the hidden networks of interaction in mobile animal groups allows prediction of complex behavioral contagion". In: *Proceedings of the National Academy of Sciences* 112.15 (2015), pp. 4690–4695.

[44] J. E. Herbert-Read. "Understanding how animal groups achieve coordinated movement". In: *Journal of Experimental Biology* 219.19 (Oct. 2016), pp. 2971–2983.

[45] L. Busoniu, R. Babuska, and B. De Schutter. "Multi-agent Reinforcement Learning: An Overview". In: *Innovations in Multi-Agent Systems and Applications - 1. Studies in Computational Intelligence* 310 (2010).

[46] L. Busoniu, R. Babuska, and B. De Schutter. "A Comprehensive Survey of Multiagent Reinforcement Learning". In: *IEEE Transactions on Systems, Man, and Cybernetics* 38.2 (2008), pp. 156–172.

[47] E. Behrends. *Introduction to Markov Chains*. Vieweg and Teubner Verlag, 2000.

[48] C. J. C. H. Watkins and P. Dayan. "Q-Learning". In: *Machine Learning* 8 (1992), pp. 279–292.

[49] S. Mirjalili. "Genetic algorithm". In: *Evolutionary algorithms and neural networks*. Springer, 2019, pp. 43–55.

[50] L. Rokach. "Decision forest: twenty years of research". In: *Information Fusion* 27 (2016), pp. 111–125.

[51] M. E. Pocco, M. M. Cigliano, B. Foquet, C. E. Lange, E. L. Nieves, and H. Song. "Density-Dependent Phenotypic Plasticity in the South American Locust, Schistocerca cancellata (Orthoptera: Acrididae)". In: *Annals of the Entomological Society of America* 112.5 (2019). From appendix, pp. 458–472.

[52] A. T. Showler. "Desert Locust, Schistocerca gregaria Forskål (Orthoptera: Acrididae) Plagues". In: *Encyclopedia of Entomology*. Ed. by J. L. Capinera. Springer Netherlands, 2008, pp. 1181–1186.

[53] H. Sanchez-Arroyo. "Largest Swarm". In: *University of Florida book of insect records. University of Florida, Gainesville* 2325 (1997).

[54] L. Edelstein-Keshet, J. Watmough, and D. Grunbaum. "Do travelling band solutions describe cohesive swarms? An investigation for migratory locusts". In: *Journal of mathematical biology* 36.6 (1998), pp. 515–549.

[55] R. M. Robertson and A. G. Johnson. "Collision avoidance of flying locusts: steering torques and behaviour". In: *Journal of Experimental Biology* 183.1 (1993), pp. 35–60.

[56] F. C. Rind, R. D. Santer, and G. A. Wright. "Arousal facilitates collision avoidance mediated by a looming sensitive visual neuron in a flying locust". In: *Journal of neurophysiology* 100.2 (2008), pp. 670–680.

[57] E. G. Reis and M. G. Pawson. "Fish morphology and estimating selectivity by gillnets". In: *Fisheries research* 39.3 (1999), pp. 263–273.

[58] R. E. Snodgrass. *Principles of insect morphology*. Cornell University Press, 2018.

[59] J. P. Hume. "Systematics, morphology, and ecological history of the Mascarene starlings (Aves: Sturnidae) with the description of a new genus and species from Mauritius". In: *Zootaxa* 3849.1 (2014), pp. 1–75.

[60] A. Blanco-Justicia and J. Domingo-Ferrer. "Machine learning explainability through comprehensible decision trees". In: *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*. Springer. 2019, pp. 15–26.