Evolving Predator-Prey Dynamics in Agent-Based Models of Collective Motion

Matthew Bowditch, Luke Murray Kearney, Sam Turley, Nathan Van Der Riet

Mathematics Institute, University of Warwick.

Abstract

We propose a bio-inspired, force-based model of agents to replicate the dynamics of a group of predators attacking a swarm of prey in a bounded space. Our model uses a first-order Euler method with a set of local metric-based interaction rules to implement a discrete time update process on a continuous space. Novel approaches to boundary conditions and predator evasion are implemented. We show that this set of simple update rules can generate complicated group chasing and evading strategies. We then implement an original evolutionary adaptation mechanism on the prey and predator behavioural parameters, to minimise or maximise the proportion of prey killed respectively. The parameter optimisation process was carried out using the BIPOP-CM-AES algorithm sequentially. Optimisation resulted in continual oscillation between distinct strategies, without the emergence of a dominant strategy for either species.

Keywords: Collective Motion, Predator-Prey, Agent-Based, Optimisation

1 Introduction

Collective motion is a broad term that encompasses many unique phenomena, where order emerges from chaos. Examples include schooling fish [1], flocking birds [2], swarming locusts [3] and even human festival goers [4]. The most famous model of active matter was proposed by Viscek et al. in 1995 [5]. In this model, each particle updates its direction at every time step, by aligning with its neighbours in a radius of interaction. The particle then moves at a constant speed in the new direction, providing an exceedingly simple heuristic to explain the emergence of order, which defined a starting point for many subsequent analyses.

Later, force-based models grew in popularity in which agents and the environment impose forces on one another much like interacting charged particles [6–9]. In particular, this allows agents to vary their speed, unlike the original model proposed by Viscek. These more complicated models were applied to a variety of situations, including predators hunting prey [6, 8, 9].

The dynamics of predators and prey play a crucial role in collective motion, especially through the methods by which prey evade their predators. Research by Chen et al. has shown that when there are multiple prey and a single predator, larger groups of prey necessitate a stronger predator in order to ensure a successful kill [9]. This may be attributed to the confusion caused by having multiple targets, making it difficult for a predator to focus on a single target. Another similar problem is modelling scenarios with multiple predators pursuing a single prey. In one such case, with the assumption that the prey was considerably faster than the predators, it was demonstrated that a simple set of forces allowed predators to

coordinate, greatly improving the success rate in capturing prey [7].

Prior to development of a model, it is important to characterise distances in which particles interact. *Metric interaction*, distance based interactions, and *topological interaction*, a nearest neighbours approach, are the most common approaches in the literature. Kumar et al., provided an analysis of both approaches, which motivated the derivation of our model [10]. The addition of a vision cone is widely used, reducing the agent's field of view to emulate a species of choice [11]. Our model is inspired by the dynamics of Lemon Sharks hunting, which have been shown to have very large vision cones, close to 360° [12], with their prey also having near 360° vision [13]. For this reason, and for simplicity we do no not include vision cones in our model.

It is believed that travelling collectively can allow for large groups to transfer information between pack members quickly. This could be information regarding a food source [14], migration [15, 16], or to escape a predator [17, 18]. A simplistic exploration of this phenomenon can be found in Couzin et al., where they show that *only a very small proportion of informed individuals are required to achieve great accuracy* in group guidance [19].

Whilst much work has been done on multiple prey, single predator models [9] and vice versa [7], research on multiple predator, many prey systems is sparse. This situation is often observed in nature, an example being the previously mentioned Lemon Shark hunting. The aim of this paper is to develop a force-based bio-inspired predator-prey model to describe these scenarios. Following this, we employ an evolutionary approach to optimise evasion/capture for prey and predators respectively.

In Section 2, we derive our model by defining forces that act on each agent and the properties of the space in which they interact. One such property of the space which is particularly interesting are the soft boundary conditions, which are a deviation from the periodic boundary conditions used by most models in the literature. Sections 3 and 4 describe the initial conditions of our model and the method used to simulate it. Section 5 provides an overview of the different behaviours of agents observed within our model, which change based on our choices of force coefficients. Various measures are also provided to help understand these behaviours, building an intuition of the dynamics within the parameter space of the model. In Section 6, we define a strategy to allow prey and predators to update their behaviour, minimising or maximising the amount of prey captured in a simulation. The results of these optimisations are discussed in Section 7, providing interesting behaviour from our simple set of update rules, such as predators lining up to shepherd prey into the boundaries to get kills. We end the paper with a conclusion of our results and discussion future work in Section 8.

2 Model

We employ a two-dimensional model and consider it comparable to a two-dimensional projection of hunting in three-dimensional shallows. The lack of depth does not allow for significant vertical chases [7]. Each agent has a position $\mathbf{x}_i(t)$ and a velocity $\mathbf{v}_i(t)$ and experiences an endogenous force $\mathbf{F}_i(t)$ due to local interactions over some radius r. Our agents have full 360° vision and are point particles. By Newton's law

$$m\frac{d\mathbf{v}_i(t)}{dt} + \mu\mathbf{v}_i(t) = \mathbf{F}_i(t), \tag{1}$$

$$\frac{d\mathbf{x}_i(t)}{dt} = \mathbf{v}_i(t) \tag{2}$$

where m is the mass of the agent and μ is the coefficient of friction within the surrounding medium. For simplicity, Chen, et al [9] and consequently Chakraborty, et al [8] use the assumption that $m \ll \mu$ and reduce the system to first order. This makes analytical results easier to derive. Instead, we opt for a model where agents have momentum and we drop the friction term, setting $\mu = 0$. In real-world systems, predators and prey are often vastly different masses. If friction were to be kept, predators and prey would likely have different mass-to-friction ratios, so picking these representative parameters to be impactful on the model whilst maintaining realism may be difficult. The aim of the project is to explore a simple abstract model for predator-prey dynamics, rather than developing a high-level model for a specific system. Since we drop the friction terms, the mass

m for each of predator and prey can be absorbed into the force coefficients.

In the previous approaches, agents only move while experiencing a force and will slow down to stationary when no force is applied; here, agents maintain their velocity in the absence of a force. We also enforce maximum velocities $\mathbf{v}_{\max}^{\text{prey}}$, $\mathbf{v}_{\max}^{\text{pred}}$ and accelerations $\mathbf{a}_{\max}^{\text{prey}}$, $\mathbf{a}_{\max}^{\text{pred}}$ a natural limitation. For any real-world system characteristic length and time scales can be chosen as desired to arrive at our non-dimensionalised model. For the purposes of this report we have chosen the characteristic length scale to be equal to the prey's vision radius and characteristic time scale to be defined by their maximum velocity, as such we always take these values to be one. Other choices are available, however, and as long as they are consistent, this model with dimensionless terms will apply.

2.1 Prey Forces

Let E and C be the sets of prey (**E**scapers) and predator (**C**haser) agents within our model respectively. For a prey agent, $i \in E$, we break down our force term into seven components,

$$\begin{aligned}\mathbf{F}_i = &\mathbf{F}_{i,E}^{\text{align}} + \mathbf{F}_{i,E}^{\text{attract}} + \mathbf{F}_{i,E}^{\text{repulse}} \\ &+ \mathbf{F}_{i,C}^{\text{disalign}} + \mathbf{F}_{i,C}^{\text{repulse}} \\ &+ \mathbf{F}_i^{\text{bound}} + \mathbf{F}_i^{\text{noise}}\end{aligned} \tag{3}$$

Many of these contributing forces are common among the literature, for example, Viscek's seminal paper on a constant speed self-propelled system [5] had a similar alignment term emulating an escapers desire to travel with their kind. The attraction, prey-prey and prey-predator repulsion were implemented by Chen et al., as a simple force-based model for predator-prey interaction [9]. Attraction and prey-prey repulsion imitate a fishes desire to be near their kind, but not too close, and prey-predator repulsion emulates a prey desire to flee from a nearby predator. The boundary, noise and disalignment forces are inspired by the boundary zig-zag and predictive forces implemented by Janosov et al., in their analysis of a multiple predator single prey system [7]. The zig-zag force is a stochastic evading behaviour, which attempts to travel at some rotation from the fleeing direction, to prevent predator tactics, similar to a noise term. The disalignment term is novel and attempts to emulate the predictive capabilities of agents, pushing the escapers to turn away from where the chasers are going to be.

Within our model, we also include a radius of interaction for prey r_E and predators r_C. Agents only experience forces from other agents within this radius. We denote the radius of interaction of an individual agent i by r_i. While some of our forces decay in magnitude with the distance of interaction eliminating the need for this radius of interaction, limiting the radius will later allow us to take advantage of grids when simulating the model, leading to much cheaper computations. For an agent i we denote by

$$\begin{aligned}N_i &= |\{j \in E \mid |\mathbf{x}_j - \mathbf{x}_i| \leq r_i\}| \\ M_i &= |\{j \in C \mid |\mathbf{x}_j - \mathbf{x}_i| \leq r_i\}|\end{aligned}$$

the number of prey and predator agents within its vision radius respectively.

We define the forces for agent-to-agent interactions used in our model as follows:

$$\begin{aligned}\mathbf{F}_{i,E}^{\text{align}} &= \frac{\alpha}{N_i}\sum_{j=1}^{N_i}\mathbf{v}_j - \mathbf{v}_i \\ \mathbf{F}_{i,E}^{\text{attract}} &= \frac{\beta}{N_i}\sum_{j=1}^{N_i}\mathbf{x}_j - \mathbf{x}_i \\ \mathbf{F}_{i,E}^{\text{repulse}} &= -\frac{\gamma}{N_i}\sum_{j=1}^{N_i}\frac{\mathbf{x}_j - \mathbf{x}_i}{|\mathbf{x}_j - \mathbf{x}_i|^2} \\ \mathbf{F}_{i,C}^{\text{repulse}} &= -\frac{\delta}{M_i}\sum_{j=1}^{M_i}\frac{\mathbf{x}_j - \mathbf{x}_i}{|\mathbf{x}_j - \mathbf{x}_i|^2} \\ \mathbf{F}_{i,C}^{\text{disalign}} &= \kappa\boldsymbol{\omega} \\ \mathbf{F}_i^{\text{noise}} &\sim \mathcal{N}(0, \sigma^2).\end{aligned} \tag{4}$$

The noise term is a draw from a normal distribution with variance σ^2. The term for disalignment with predators, $\boldsymbol{\omega}$ is a vector perpendicular to the average velocity of the nearby predators. In particular if

$$\mathbf{d} = \frac{1}{M_i}\sum_{j=1}^{M_i}\mathbf{x}_j - \mathbf{x}_i, \tag{5}$$

is the displacement vector from our prey to the centre of mass of the nearby predators and

$$\overline{\boldsymbol{v}} = \frac{1}{M_i}\sum_{j=1}^{M_i}\mathbf{v}_j. \tag{6}$$

the average velocity of the nearby predators then $\overline{\boldsymbol{v}} \perp \boldsymbol{\omega}$. Moreover, we choose the direction such that $\mathbf{d}\cdot\boldsymbol{\omega} \leq 0$, so that the prey turn in the direction away from the predators. We set the magnitude of $\boldsymbol{\omega}$ equal to that of the average predator velocity and as such average predator velocity acts as a scaling factor. We will describe the force due to boundary separately as it warrants a longer discussion. Regardless, there are still a few things to note here. For one, prey alignment would be more accurately called prey velocity matching. This is a stronger condition where agents seek to match both direction and speed [10]. The choice of prey attraction and repulsion terms are consistent with the minimal model proposed by Chen et al. [9]. Notably, when $\beta = \gamma$ the repulsion term is stronger than the attraction term within a radius of one, due to the quadratic term on the denominator, and is weaker outside of this radius. Thus, β and γ can be chosen to fix the distance where repulsion turns to attraction. The predator repulsion term is also found in the minimal model [9]. The predator disalignment term is not one discussed in the literature we have reviewed, though similar terms have been proposed, such as a zigzag term [7]. We find that the addition of this term facilitates some interesting behaviour such as better cohesion and survivability when being chased by a faster predator.

2.2 Predator Forces

For a predator, $i \in C$, we have a force term as follows

$$\begin{aligned}\mathbf{F}_i = \mathbf{F}_{i,C}^{\text{align}} + \mathbf{F}_{i,C}^{\text{attract}} + \mathbf{F}_{i,C}^{\text{repulse}} \\ + \mathbf{F}_{i,E}^{\text{attract}} + \mathbf{F}_i^{\text{bound}} + \mathbf{F}_i^{\text{noise}}\end{aligned} \tag{7}$$

Again we leave the discussion of the boundary force for later and list the forces arising due to agent interactions below.

$$\begin{aligned}\mathbf{F}_{i,C}^{\text{align}} &= \frac{a}{M_i}\sum_{j=1}^{M_i}\mathbf{v}_j - \mathbf{v}_i \\ \mathbf{F}_{i,C}^{\text{attract}} &= \frac{b}{M_i}\sum_{j=1}^{M_i}\mathbf{x}_j - \mathbf{x}_i \\ \mathbf{F}_{i,C}^{\text{repulse}} &= -\frac{c}{M_i}\sum_{j=1}^{M_i}\frac{\mathbf{x}_j - \mathbf{x}_i}{|\mathbf{x}_j - \mathbf{x}_i|^2} \\ \mathbf{F}_{i,E}^{\text{attract}} &= -\frac{d}{N_i}\sum_{j=1}^{N_i}\frac{\mathbf{x}_j - \mathbf{x}_i}{|\mathbf{x}_j - \mathbf{x}_i|^3} \\ \mathbf{F}_i^{\text{noise}} &\sim \mathcal{N}(0, \sigma^2).\end{aligned} \tag{8}$$

The predator forces are designed to mirror that of the prey, with equivalent but inverse representation for all terms apart from disalignment and prey attraction. Predators do not have a disalignment term in the interest of simplicity, although in a future addition to the model a predictive capability similar to Janosov et al., could provide an interesting extension [7]. Notice the denominator on the prey attraction is a cubic term, as found in the minimal model [9], meaning that the largest contribution to this force is from the prey closest to the predator. This is effectively making the force selective; the predator will target close lone prey over far-off packs. However, once near several prey, the direction of the closest one may change quickly over a short time frame, leaving the predator unable to focus on a single prey. This is known heuristically as a *confusion effect*. Other authors such as Mohapatra and Mahapatra handle the selection and confusion effect more explicitly [6], whereas we opt not to for simplicity. We assume that when a prey enters a given radius of a predator, k_r, they are killed and removed from the simulation. We also implement a kill cool down τ, during which a predator is forbidden from making another kill, though they still experience the same forces.

2.3 Boundary Conditions

Instead of using periodic boundary conditions typical to most flocking models, we use a bounded domain and soft boundaries in particular. This is to replicate hunting in areas close to shorelines or where indefinite escape directly away from

a predator is infeasible, but a hard wall is not realistic.

In our $L \times L$ space, we have a boundary radius r_b in which prey and predators experience a repulsive force away from the boundary. An example force, acting on the $x = 0$ boundary, is given by

$$(\mathbf{F}_i^{\text{bound}})_x = \max(a_{\max}, |\mathbf{F}_i|) \left(1 + \cos\left(\frac{\pi d}{r_b}\right)\right), \tag{9}$$

where d is the distance from the $x = 0$ boundary. The form of this is chosen such that the force becomes increasingly dominant the closer to the boundary the agent is, similar to the soft boundary defined in Janosov et al. [7]. This creates more realistic boundary scenarios, where agents start turning away from the boundary as they approach, instead of bouncing or reflecting off. To ensure our agents never collide with a boundary we can define a worst-case scenario to calculate a boundary radius sufficient to repel them fully. Consider the $x = 0$ boundary and an agent crossing $x = r_b$ experiencing $-F_i\mathbf{e}_x + 0\mathbf{e}_y$, i.e a force perpendicular to the boundary. Including boundary forces, the total force the agent experiences is

$$F_i^x = \max(a_{\max}, F_i) \left(1 + \cos\left(\frac{\pi x}{r_b}\right)\right) - F_i. \tag{10}$$

The worst-case scenario is when $|F_i| > a_{\max}$. To observe this, note that the function is continuous in F_i and consider the point $x = \frac{r_b}{2}$. Then

$$F_i^x = \max(a_{\max}, F_i) - F_i, \tag{11}$$

which is minimized when $|F_i| > a_{\max}$. In this case, our force is then given by

$$F_i^x = \cos\left(\frac{\pi x}{r_b}\right) F_i. \tag{12}$$

Until $x = \frac{r_b}{2}$ our agents can still be accelerating towards the boundary. At that point, there is no force on them. When calculating velocity, we impose the condition that $|\mathbf{F}_i| \leq a_{\max}$.

$$\frac{dv}{dt} = \min\left(\cos\left(\frac{\pi x}{r_b}\right) F_i, a_{\max}\right) \tag{13}$$

$$\frac{dx}{dt} = v \tag{14}$$

$$\implies \frac{dv}{dx} = \frac{1}{v}\min\left(\cos\left(\frac{\pi x}{r_b}\right) F_i, a_{\max}\right) \tag{15}$$

This is a separable system, which we can impose the boundary condition $v(\frac{r_b}{2}) = -v_{\max}$. The solution is given by

$$\begin{aligned} v^2 + r_b \min(\frac{F_i}{\pi}, \frac{a_{\max}}{2}) - v_{\max}^2 \\ = \min\left(\frac{r_b}{\pi}\sin\left(\frac{\pi x}{r_b}\right) F_i, a_{\max} x\right). \end{aligned}$$

The smallest boundary radius is when $v(0) = 0$, i.e. the agent just avoids touching the boundary. This yields

$$r_b = \frac{v_{\max}^2}{\min(\frac{F_i}{\pi}, \frac{a_{\max}}{2})}, \tag{16}$$

which is largest when F_i is minimized, but since we assumed $|F_i| > a_{\max}$, we get a worst-case boundary radius of

$$r_b = \frac{\pi v_{\max}^2}{a_{\max}} \tag{17}$$

3 Simulation Method

To simulate the model we use a first-order forward Euler method, for the time-stepping, subject to maximum accelerations and velocities.

$$\mathbf{v}_i(t + \Delta t) = \mathbf{v}_i(t) + \Delta t \cdot \hat{\mathbf{F}}_i(t + \Delta t) \tag{18}$$

$$\mathbf{x}_i(t + \Delta t) = \mathbf{x}_i(t) + \Delta t \cdot \hat{\mathbf{v}}_i(t + \Delta t) \tag{19}$$

where we use notation

$$\hat{x} = \frac{x}{|x|}\min\left\{|x|, x_{\max}\right\}. \tag{20}$$

Despite our model incorporating several parameters, it is important to note that they may be adjusted according to preferences and attempts could be made to match them experimentally to a real predator-prey system.

3.1 Enhancing Code Performance

At each time step, each agent must ascertain which of the other agents are within its vision area. To this end, a grid was implemented over the domain, with each grid-space having length equal

Notation	Parameter
N	Number of Prey
M	Number of Predators
N_i	Subset of N within vision radius of agent i
M_i	Subset of M within vision radius of agent i
$\alpha, \beta, \gamma, \delta, \kappa$	Constant prey force coefficients
a, b, c, d	Constant predator force coefficients
σ	Variance of noise
$\mathbf{F}_i$	Force experienced by agent i
$\mathbf{v}_i$	Agent i's velocity
$\mathbf{x}_i$	Agent i's position
Δt	Time step
L	Domain height and width
r_E	Prey vision radius
r_C	Predator vision radius
r_b	Soft boundary radius
$\mathbf{a}_{\max}^{\text{prey}}$	Max acceleration for prey
$\mathbf{a}_{\max}^{\text{pred}}$	Max acceleration for predators
$\mathbf{v}_{\max}^{\text{prey}}$	Max velocity for prey
$\mathbf{v}_{\max}^{\text{pred}}$	Max velocity for predators
τ	Kill cooldown time
r_k	Kill radius

Table 1: Nomenclature: The definitions and corresponding of all parameters in the model.

to the prey's vision radius. This vastly reduces computational complexity, as we are guaranteed to find all agents that are within our vision area just by in checking our grid-space and those neighbouring it. This strategy does add an additional step of calculating whether each agent changes it's grid-space after their position is updated.

After initial testing in Python, the model was coded using the Rust programming language, which allowed for larger scale simulations, real-time graphical rendering and parallelisation for the CPU-intensive behaviour optimisation discussed in Section 6. The main performance benefit of Rust comes from compile time interpretation rather than dynamic interpretation, as used in Python. Furthermore, Rust's memory handling allows for more efficient memory usage by enforcing a system of ownership at compile time instead of relying on a garbage collector as in Python. The trade-off for this increase in performance is a language that has a steeper learning curve and longer development times, however, by making the switch from Python, we improved the performance of our code 10,000's of times.[1]

[1] A more detailed explanation of the performance benefits of Rust is unfortunately beyond the scope of this project. More information on the Rust programming language can be found in [20].

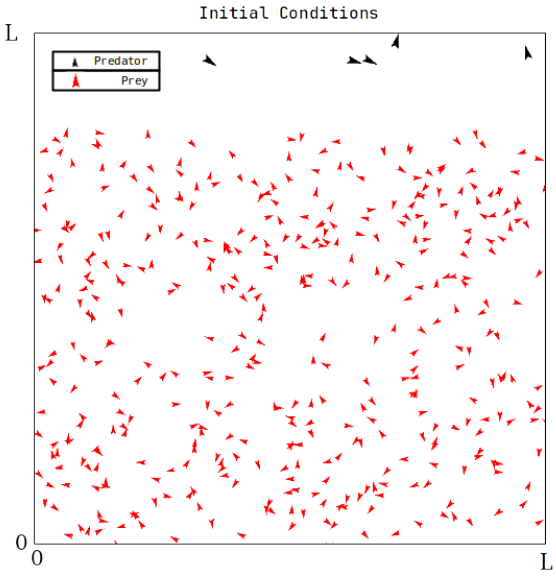

Fig. 1: One randomised set of initial conditions for a simulation of the model. Prey and predators are represented by red and black polygons respectively. The boundary length is shown along the edges.

4 Initial Conditions

In our simulations, we start with N prey is positioned randomly within $[0, L] \times [0, \frac{4L}{5}]$ and M predators positioned randomly within $[0, L] \times [\frac{9L}{10}, L]$, see Figure 1. These conditions ensure that there are very few kills within the initial time steps due to prey and predators being randomly being adjacent. It also ensures prey cannot see a predator at the start of the simulation in all settings we consider ($L \geq 10$), emulating predators finding a group of prey. Each agent is started at unit (the maximum prey) velocity, facing random directions, so as to not add any bias to our system. As a possible adaptation to the model, a smaller initial velocity may better emulate a resting shoal off fish unaware of an oncoming predator.

5 Exploring the Model

5.1 Categorization of Behaviours

In systems without predators, we can categorise different observable behaviour patterns relating to the coefficients β and γ of prey attraction and repulsion respectively. Here we assume that

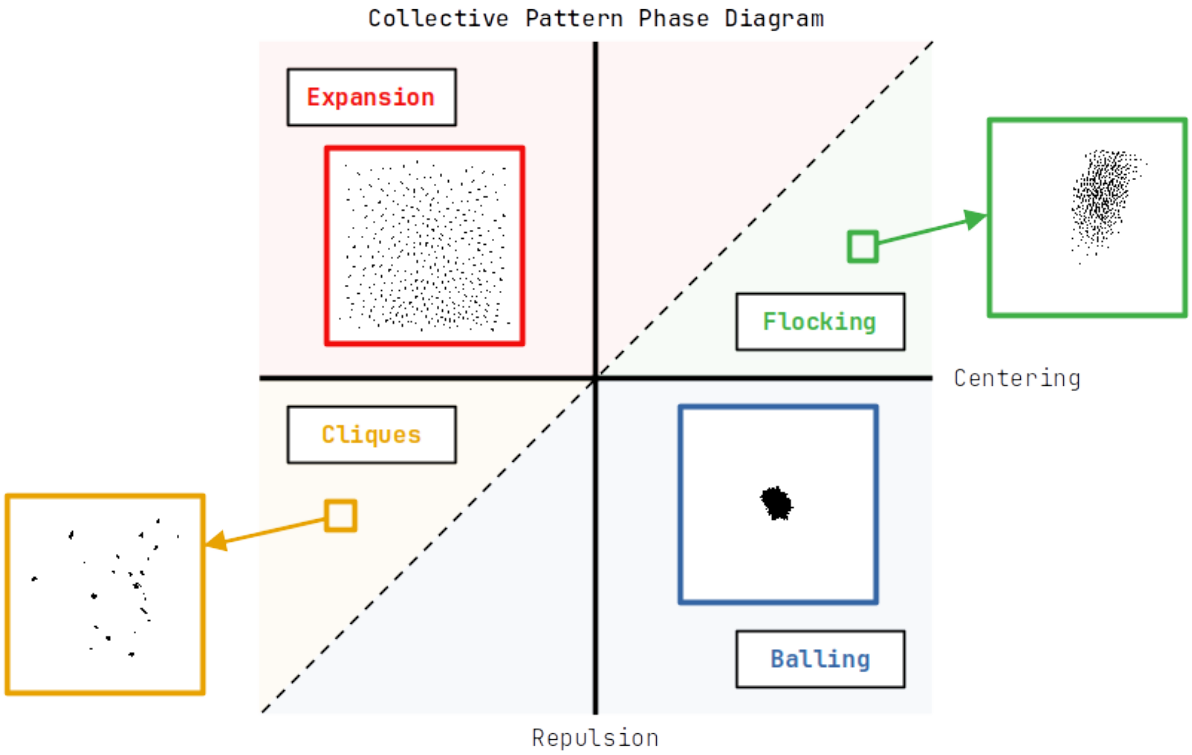

Fig. 2: Qualitative agent behaviour types observed when varying the centering and repulsion coefficients of our forces, with $r_E = 1$. Snapshots of models are displayed in each section with corresponding parameter values.

α is positive, so prey have some desire to match velocities. For $\alpha \ll 1$, we can observe *swarming*, though this is more frequently observed as an insectoid behaviour pattern [3]. We consider a periodic space and the effect of prey entering the vision radius of each other. The phase space is shown in Figure 2.

In *expansion*, agents are repelled away from each other regardless of the distance of interaction. In a bounded space, this amounts to filling the space as uniformly as possible. In *flocking*, there is a "Goldilocks" zone within the vision radii; outside the zone, prey are drawn together but inside they are repelled. Over long time periods, our flocks become circular due to the nature of the forces. In *balling*, prey wish to compact as much as possible. In reality, there should be some volume exclusion but since our prey are point particles, there is no issue with them becoming arbitrarily close together. In *cliques*, we see another "Goldilocks" zone with the opposite effect; inside the zone causes prey to be drawn in further and outside you are repelled. This leads to very tight collections of prey that act as a single cohesive unit but do not form one mass.

The dashed line in Figure 2, intercepts at [0,0] irrespective of r_{E}. The slope of this line is 1 for $r_{\mathrm{E}} = 1$, although as r_{E} varies, the slope and order of this curve changes. Our model does not produce *milling*, a common behaviour pattern observed in fish. One way this can be produced by an agent based model is to weigh visual information such that fore agents are given more precedence than aft agents, an exemplar continuous transition is given by $1 + \cos(\theta)$ [21].

5.2 Measures

To understand certain behaviours, we need metrics to discuss the model. An *order parameter* is some measure of how ordered the system is: a mapping of physical shpace to a parameter space. Usually, they are invariant under some symmetry

(reflection, rotation, translation). We have three order parameters, the average velocity, the polarization and the group number.

The average speed of the prey in the system is given by

$$V = \frac{1}{N}\sum_{i=1}^{N} |\mathbf{v}_i|, \tag{21}$$

and the polarization:

$$P = \frac{1}{N}\left|\sum_{i=1}^{N} \frac{\mathbf{v}_i}{|\mathbf{v}_i|}\right|. \tag{22}$$

Because our prey are allowed different speeds, we take the unit vector in the direction of movement. A polarization of zero means that particles are randomly oriented, whereas one means that all particles are travelling in the same direction.

The average number of groups of prey present at a given point in time is calculated using Density-Based Spatial Clustering of applications with noise (DBSCAN) [22]. DBSCAN is a widely used clustering algorithm[2], it provides an easily understood heuristic, and fast implementation in $\mathcal{O}(n \log n)$. DBSCAN uses 3 point identifiers: Core, Periphery and Noise. DBSCAN takes the inputs ϵ, the radius of a neighbourhood and n, the minimum number of agents in a group, and identifies groups using the following heuristic:

- If there are n points within a distance ϵ of a point x, x is a core point
- If a point x is within a distance ϵ of a core point, but does not fulfil the criteria of a core point it is a periphery.
- If a point x is not within a distance ϵ of any core nodes, it is referred to as a noise point.

We then use this classification to calculate the number of independent sets of cores at each time step, i.e., the number of groups.

5.3 Without Predators

We can now use these measures to help us understand the different behaviour categories shown in Figure 2, for both soft and periodic boundary conditions. Figure 3 shows each measure over time for each behaviour category and boundary condition. From this, we can observe the following differences and similarities between periodic and soft boundary conditions:

1. In general, measures on simulations with periodic boundary conditions have much smoother lines. This implies that periodic boundary conditions result in much less stochastic simulations compared to soft boundary conditions.
2. Polarization increases at a slower rate for soft boundary conditions. This is due to the boundaries forcing the agents to turn around, decreasing the polarization. We also see a periodicity to the dips as large groups all hit the boundary at similar times.
3. The average speed increases over time, at a similar rate, for both periodic and soft boundary conditions.
4. The group number for each behaviour is consistent between periodic and soft boundary conditions. This implies that neither boundary condition provides a significant advantage to help agents form groups.

We can also use Figure 3 to observe the following about the different behaviour categories:

1. The group number for Expansion converges to 0 in both cases. This is due to the prey expanding to fill the space and not being considered a group as there is too much space between them.
2. Expansion is the only behaviour to have an average speed that is decreasing over time. This is due to the prey uniformly filling the space and then being held in place by repulsive forces from all of the prey surrounding them.
3. Balling is the behaviour that sees the steepest increase in average speed. Balling has a negative repulsion and a positive centering as seen in Figure 2, meaning these two forces are pointing in the same direction. This cohesion may be part of the reason we see the prey increasing their average speed at a faster rate, as there is less cancellation between the force components.

5.4 Difference between Disalignment and Repulsion

In Figure 4, we observe a single predator attacking a flock. In the different simulations, the prey only have one of the escape forces active and we can

[2]While the efficiency and efficacy have been questioned in high-dimensional complex spaces recently [23] this is not an issue in our case of a simple 2D space.

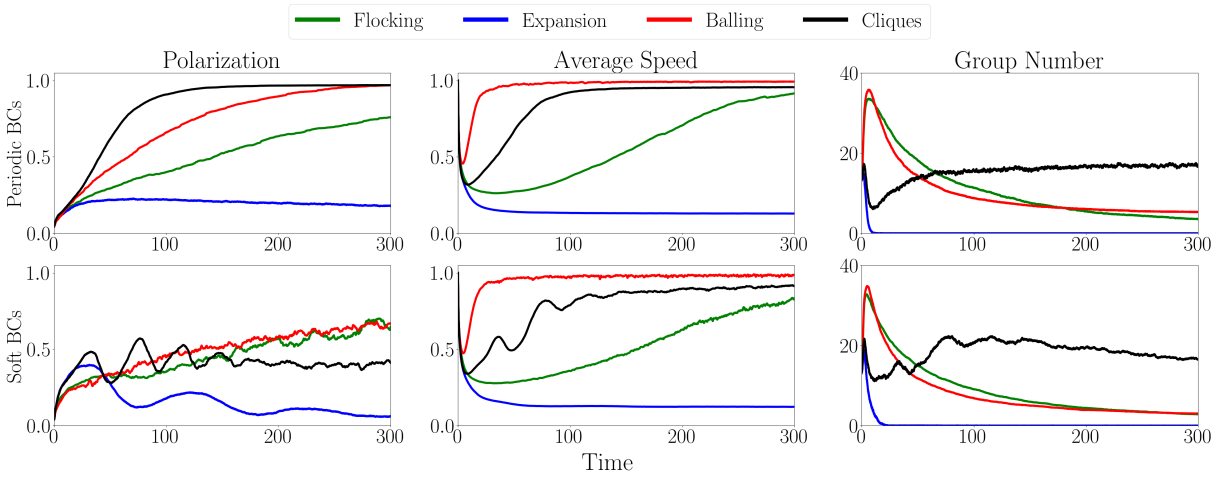

Fig. 3: Examples of our metrics, simulated for different behaviours outlined in the collective patterns phase diagram 2, with periodic and soft boundary conditions. The metrics included averaged polarization, speed and group number plots from 100 simulations in periodic boundary conditions (top), and soft boundary conditions (bottom) with $N = 400, M = 0, L = 30, r_b = 2, t_{\max} = 300, \Delta t = \frac{1}{20}$. The parameter sets for the behaviours was $\alpha = 1, \beta = \pm 0.5, \gamma = \pm 0.05$.

observe different behaviour. With only repulsion, the escaping prey form a circular fan around the predator. The prey directly in front of the predator will continue escaping parallel to the predator and so would likely be caught by a faster predator. The wake produced does not widen rapidly, and for a larger r, may allow the flock to reform. In contrast, the disalignment term splits a flock into two; no prey run directly away from the predator. This has a stronger impact on the adjacent prey and produces a more lasting effect as the wave of high-speed prey involves different agents over time: the number of prey involved is different between the strategies. Against a predator with a large turning circle, this could be a highly efficient strategy. Conversely, wise predators may co-operate to encircle a group of prey akin to a sheepdog, where the nature of travelling perpendicular may allow prey to be compressed and herded.

We can introduce a measure of success for predators, the proportion of prey killed within a certain time frame. This a better metric than the time taken to kill a fixed proportion of prey as that may never occur.

In Figure 5(a), we first notice that when a predator-evasive strategy is implemented, the prey are harvested at a slower rate than when no strategy is used. This is a useful model validation. Next, we see that the rate of predation over time is mostly linear until the number of prey alive is considerably smaller. This is likely due to the remaining prey being hard to locate: when there is more prey targets are abundant. Finally, we can note the strength of each strategy, repulsion is better than disalignment, but even better is the combination, they are complementary strategies.

When these terms are combined, situations can occur like Figure 5(b). There are no prey in the immediate vicinity of the predator and a small wake behind that is getting refilled by prey. One observed behaviour is *fountaining*, where prey peel off and away from predators, and re-fill the space behind.

5.5 Predator Flocking

The only behaviour turned on for the predator was attraction to prey. This can lead to situations in which predators chase the same prey, reducing their efficacy. Instead, we can introduce the alignment, attraction and repulsion between predators, i.e, allow them to flock. This is the main novelty of our group research project: flocking predators may be able to herd prey in a confined space.

In Figure 5(d) and 5(c), we see shepherding is an observable behaviour in our system. Notice the different separation distance between predators affects the dynamics of the prey's escape

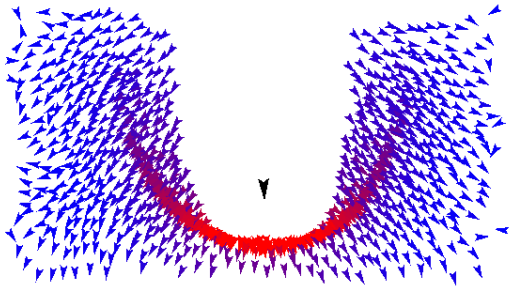

(a) Repulsion

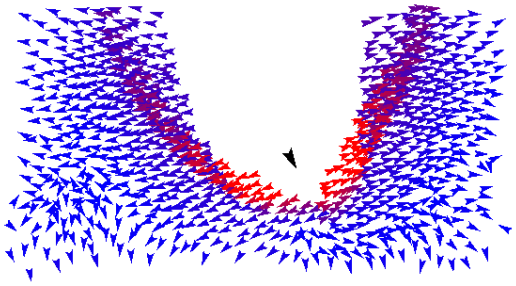

(b) Disalignment

Fig. 4: Repulsion causes the prey to fan out and disalignment causes group splitting in front of the predator. Predators are larger and depicted in black, prey follow a speed-based colour mapping where pure blue means they are stationary and pure red represents travelling at maximum speed. This colour scheme is retained throughout the report. In (a) $\alpha = 1,\ \delta = 1$ and all other prey coefficients equal 0. In (b) $\alpha = 1,\ \kappa = 1$ and all other prey coefficients equal 0.

trajectories. There are situations in which the prey becomes trapped between two predators and there is confusion between prey: cooperation is difficult as there are multiple competing messages. An interesting question is what proportion of prey should be allowed to "filter" through the gaps of the predators. If they are too close, then prey may fountain and peel off towards the edge, splitting the large flock into smaller ones, leaving the predators to choose which to follow, or split and be less effective. Excessively wide gaps mean there is no benefit of shepherding: prey can escape through the gaps leaving predators flocking towards a pitiful meal.

We can glimpse the situation at the boundary in Figure 5(e). The prey wish to avoid beaching themselves on the edge of the area whilst also avoiding being caught and so can slip into the soft boundary. However, now there are competing signals about which direction to travel and information propagation becomes hard, it is every prey for themselves. This situation is where we expect more deaths to occur when prey and predators flock.

6 Learning Behaviour

We are now prepared to ask a natural question: what parameters minimize or maximize predation? In this section, we attempt to find a set of behaviour parameters from which both predators and prey would not choose to deviate given the other species' behaviour. We fix the spatial and physical parameters of our system, and alternatively allow predators and prey to 'learn' an optimal behavioural parameter set to improve the predation in their favour. This repeated process allows us to understand the transient behaviour of our system, the trajectory through which our agents evolve.

Our attempt to mimic natural evolutionary dynamics uses the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [24], a stochastic, derivative-free method. Our lack of prior knowledge on the shape of our solution space makes CMA-ES's flexibility to non-linear and non-convex problems very useful. This algorithm belongs to a family known as evolutionary strategies, which are inspired by natural selection. Parameter sets (individuals) are randomly defined, simulated and are allowed to reproduce based on their relative fitness, i.e., their objective function evaluation. This process allows evolutionary strategies to optimise for the given system, by employing mutation and recombination operators of the given parameter set. A random sampling approach is implemented on the state space of our parameters in CMA-ES and a maximum likelihood approach is employed on these samples to create a multi-dimensional Gaussian distribution on our parameter set. The resultant distribution can then improve sampling efficiency.

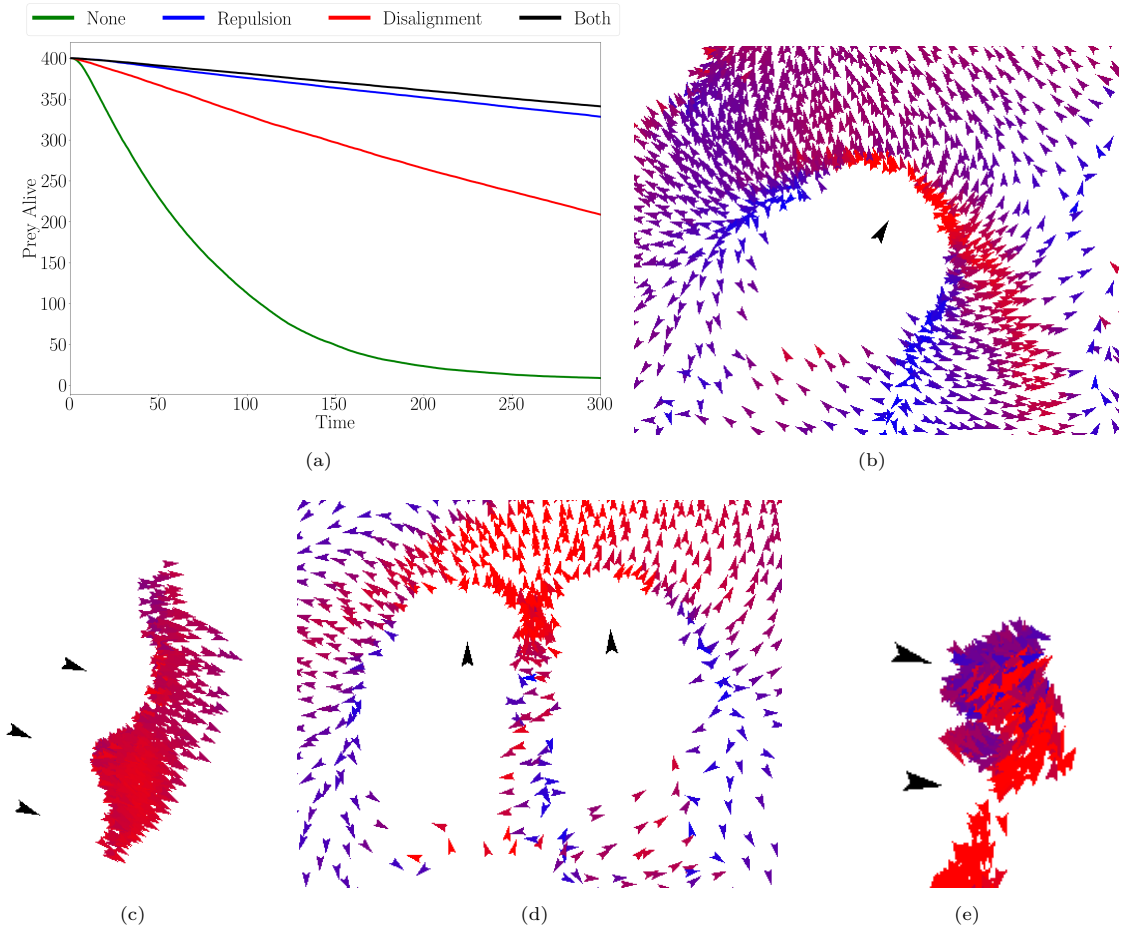

Fig. 5: (a) Number of prey left alive over time, showing the value of the prey-predator repulsion and disalignment coefficients, with each coefficient providing a positive effect in evasion. For $N = 400$, $M = 27$, $L = 30$, $r_b = 2$, $r_k = 0.05$, $t_k = 0.5$, with prey parameters $\alpha = 1$, $\beta = 0.5$, $\gamma = 0.05$, $\delta \in \{0, 1\}$, $\kappa \in \{0, 1\}$ and predator parameters $a = 1$, $b = 0$, $c = 0$. Fountaining is a behaviour seen in shoals of fish reacting to a predator, which our model displays in Figure (b). Predators Shepherding the prey by moving in parallel through the shoal can be seen in Figure (c) and (d). The simulation in (c) was run with a lower value of γ than (d), decreasing the distance between aligning predators. When prey who are being shepherded hit the boundary, they spread due to opposing boundary and predator repulsion forces. Multiple groups spreading synchronously on the boundary causes confusion from opposing forces, allowing predators to attack, as in Figure (e).

Black-box optimisers such as CMA-ES may converge to local optima in the presence of multi-modality or high levels of uncertainty [25]. To mitigate this, we employ the bi-population (BIPOP) restart strategy [26], which uses two interlaced restarting regimes with varying sample sizes. This alternative strategy has been shown to produce between linear and quadratic convergence rates, depending on modality [26], with a strong emphasis on generating a global approximation to our solution.

Testing Regimes and Scenarios

In preliminary testing, we found that prey would sometimes optimise to repel from any agent inside its vision radius (expansion behaviour as in Figure 2). In a high-density setting, despite this strong

Regime	Parameters	
Slow Predators	$\mathbf{v}_{\max}^{\text{pred}}$	0.75
	$\mathbf{a}_{\max}^{\text{pred}}$	1
Equal Predators	$\mathbf{v}_{\max}^{\text{pred}}$	1
	$\mathbf{a}_{\max}^{\text{pred}}$	2
Fast Predators	$\mathbf{v}_{\max}^{\text{pred}}$	1.25
	$\mathbf{a}_{\max}^{\text{pred}}$	3

Scenario	Parameters	
High prey density	L	10
	Predators	3
Med prey density	L	20
	Predators	12
Low prey density	L	30
	Predators	27

Table 2: The regimes and scenarios we employ, where we consider every scenario for each regime.

Parameter	Value
N	400
r_b	2
$\mathbf{v}_{\max}^{\text{prey}}$	1
$\mathbf{a}_{\max}^{\text{prey}}$	2
r_E	1
r_C	2
τ	0.5
r_k	0.05
Δt	0.05
σ	0.05

Table 3: Fixed parameter values chosen for strategy evolution, please refer to Table 1 for an explanation of notation.

repulsion, prey will fill the space and still interact with one another as there are too many of them for the areas of interaction to be mutually exclusive. This means that they can still propagate information about the predator throughout the space. We hypothesised that this strategy may falter in a low-density setting, in which prey spread out in such a way that they no longer interact, preventing information transfer. As such, we define three different *scenarios* in which to train our agents, high, medium and low-density scenarios, (4, 1, $\frac{4}{9}$ prey per unit area). This is achieved by fixing the $N = 400$ and varying the domain length, $L = 10, 20$, and 30. Predator density was kept constant, 3 predators per 100 units of area. When exploring the model, we found that predators could catch prey despite having a much lower maximum speed $\mathbf{v}_{\max,\text{ pred}}$. As such, we define three *regimes*, in which predators' maximum velocity and acceleration are changed relative to the prey, which remains fixed. This leads to a total of nine different settings for optimisation. See Table 2 for a summary of the changing variables and Table 3 for a summary of what we fixed.

In each setting prey and predators have their parameters uniformly set to 0. Simulations are ran for $T = 300, \Delta t = \frac{1}{20}$ and the average predation over 30 simulations is recorded. This is considered to be one objective function evaluation. We bound our parameter space for each parameter to $[-1, 1]$, which nullifies an agent's ability to unboundedly increase a parameter to artificially maintain an acceleration of $a_{\max}$. The BIPOP-CMA-ES algorithm performs constrained optimisation alternatively between maximizing predation for predator parameters and minimizing predation for prey parameters. We allow our optimiser 1000 total function evaluations per optimisation step of a parameter step. We call this parameter set update an *epoch*, and iterate each regime and scenario for 4 epochs. The aim is to provide an insight into the effect space and predator physical constraints have on chasing and escaping strategies. This method does not quite mirror reality: predator-prey interactions usually evolve simultaneously, yet we still hope our alternative update approach can provide some interesting qualitative behaviour.

7 Results

7.1 Sequential Optimisation of Predator-Prey Strategies

Interestingly, no dominating strategy is observed across any regimes or scenarios. For only the regime with faster predators and lower prey density does the level of predation appear to converge to a steady state, Figure 6. All other scenarios display sporadic behaviour with high variance, as in the first graph of Figure 6, when $L = 10$, or seemingly periodic behaviour, as in the second graph in Figure 6, when $L = 20$. It is possible that a larger

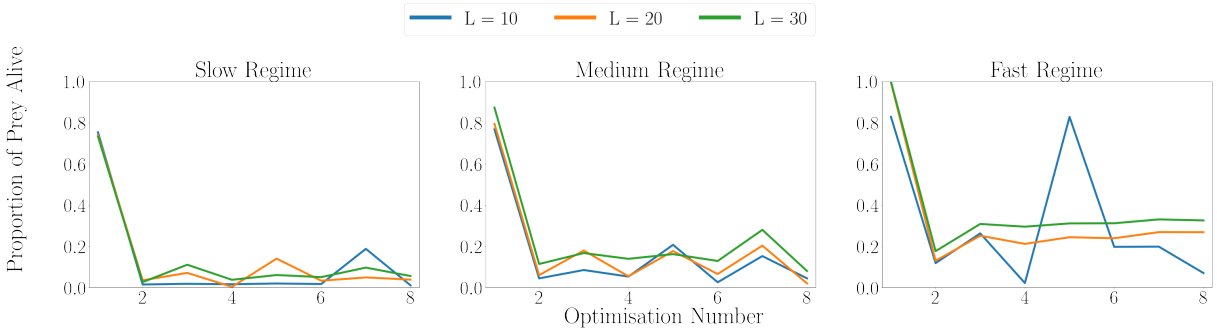

Fig. 6: Final predation levels after each optimisation step, showing transient behaviour. Predation refers to the proportion of prey remaining after 300 units of time of simulation, with the optimised parameter sets, averaged over 30 simulations. After each optimisation step this predation is shown, for the regime of slow predators, medium-speed predators, and fast predators.

number of epochs will ensure convergence. Nevertheless, the optimisations' transient behaviour shows interesting dynamics. A prime example is the slow predator regime being optimised in the large state space ($L = 30$):

- **Predator Optimise 1:** The predators favour chasing the prey as the prey have no force to evade[3].
- **Prey Optimise 1:** The prey then increase repulsion to fill the space and improve the transfer of information through the flock (expansion as in Figure 2).
- **Predator Optimise 2:** To counteract this, the predators align and reduce attraction to the prey. This sweeping and shepherding behaviour as in Figure 5(d) then allows groups of predators to catch prey at the boundaries.
- **Prey Optimise 2:** The prey increase attraction to each other, creating tight packs of prey. These packs can slip between sweeping predators, unlike when prey had a higher repulsion term.
- **Predator Optimise 3:** The predators revert back to chasing prey to counteract the balling strategy and the cycle repeats.

This cycle between distinct and contrasting behaviours is prevalent throughout our regimes, with many 'jumps' in total predation coming when an easily refutable strategy arises. These refutable strategies seem to be highly unstable and can be seen to oscillate for the above case. In section 7.3, we will optimise this scenario further to attempt to understand if the oscillatory behaviour persists.

The medium-speed predators in $L = 30$ seemed to converge to a steady predation value for 2 epochs, before significantly changing, shown in the second image of Figure 6. This is likely due to a relatively flat optimisation space, where many parameter sets have a very similar predation level. Selecting a seemingly strong parameter set can allow the other species to spike predation and break out of pseudo-steady-state. This is because the optimisation is a "best" response to a given parameter set, as opposed to a best response to all parameter sets. This dynamic adds more doubt on the truth of our steady states in the fast regime, requiring more optimisations to resolve.

7.2 Observed Strategies and their Counters

The emerging cyclic pattern in predation between hyper optimised strategies is often found to be: prey expansion, predators sweeping, prey balling, predators chasing, etc. The results of this qualitative cycle can be easily observed in simulation and became viable through our discrete optimisation process, with each response being highly tailored to the current environment.

Prey Expansion

In many cases where predators weigh chasing the prey highly, a large repulsion from neighbouring

[3]Due to a lack of friction and force coefficients our prey by in large maintain their velocities and the predators get the majority of their kills in the half of the domain where they started, see figure A1 in the appendix.

prey can allow spread out flocks to pass information quickly on the whereabouts of predators.

Predator Sweeping

When prey spread to fill the space in many cases the best recourse is for predators to reduce their coefficient of prey attraction, and instead favour alignment. This causes predators to travel in a line, and shepherd prey to the boundary. At the boundary the prey now experiences repulsive forces from the boundary, predator and prey on all sides preventing them from performing collective escape. From Figure 7, we can see predators catch most of their prey at the boundaries when implementing this strategy. In this case, predators also favour repulsion with each other, ensuring they could shepherd a large quantity of fish together. Notice there are "hotter" and "colder" areas on the top and bottom boundary. This is due to the spacing of the predators.

Prey Balling

Predator sweeping took advantage of the prey's desire to repel each other. To optimise against this, the prey favour attraction with their neighbours, creating dense balls of prey who all align. Balling allows prey to slip between the sweeping prey without getting caught at the boundaries. The heatmap of death positions in Figure 8, has many more prey being caught in the middle of the

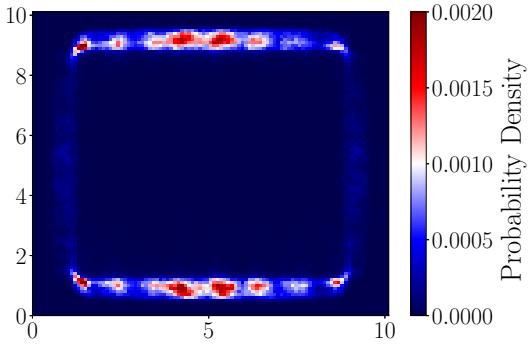

Fig. 7: Most deaths are located at the top and bottom of our boundary, while prey implement a balling strategy. The probability density of death occurring at a given location in our simulation state space, averaged from 5,000 simulations.

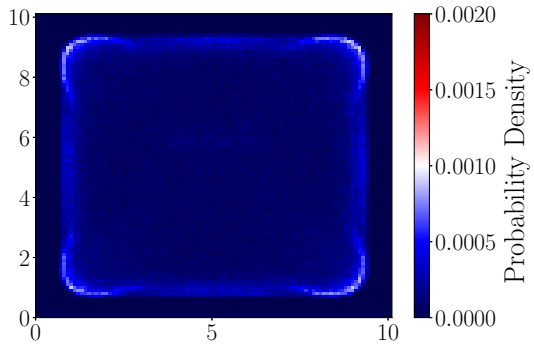

Fig. 8: Deaths more uniformly spread than sweeping, with a high proportion occurring at the corners. The probability of death occurring at each position, in a 100×100 grid of our simulation state space. This was averaged from 5,000 simulations.

space than during the sweep. Although in most cases, the flocks must get stuck in a corner to be caught by the sweeping predators.

Predator Chasing

In many cases the optimal reaction to prey balling was for prey to return to chasing and aligning to catch the dense flocks of prey. In the case, where predators are faster they can remain inside the flock and continuously catch prey. Whereas, a slower predator must use team work with aligning and repulsion terms to catch the flock. In many cases this repeats the cycle.

7.3 Continued Strategy Evolution for Interesting Scenarios

As previously mentioned we see a qualitative dichotomy in optimisation results between our fast and slow predators in $L = 30$ space; the predation oscillating in the regime with slower predators, and the predation seemingly converging in simulations with fast predators. This is something we want to explore further. Our sample size of 4 epochs is too small to understand the limiting behaviour of these systems and so we continued these interesting cases for 20 epochs. We increased the number of total objective function evaluations allowed to the optimiser to 2000 per step, and increased the simulations per function evaluation

to 80, to better ensure the optimality of the output parameters. In Figure 9 we create a box plot for the successive changes to predation relative to the mean predation from optimisation to optimisation and we do this for both the slow and fast regimes. As we can see, predation varies greatly for slow predators and much less for fast predators (and even the absolute change is greater for slow predators, see Figure A3 in the appendix). It is also the case that the predation oscillates around a mean value in both cases instead of drifting over time, see Figure A2 in the appendix. Interestingly, however, the Euclidean difference between consecutive parameter sets does not converge in either regime as shown in Figure 10. Even though predation levels show a stark contrast between convergence and oscillation, neither regime finds a dominant strategy and so continually changes its parameter set. We can attempt to illustrate this dichotomy in behaviours using several metrics, such as the previously defined average number of groups present or the average speed of prey. We choose average speed as opposed to velocity as the scenarios of interest contain 27 predators chasing 400 prey, causing a lot of disorder in all cases. The results can be found in Figure 11. We see

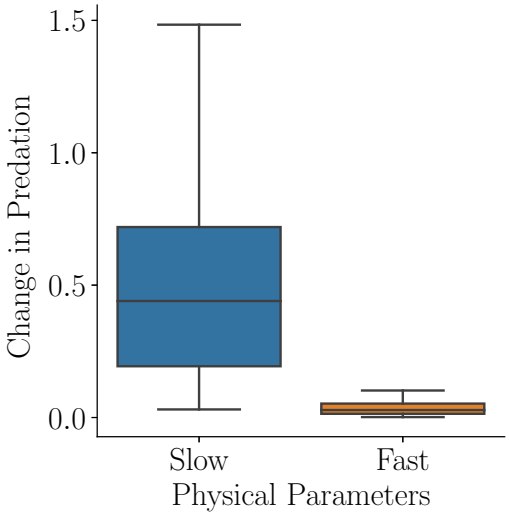

Fig. 9: A box plot of the change in predation value relative to the mean predation value for successive optimisations for both the fast and slow predator regimes. The slow regime provides a much higher variance in the final predation.

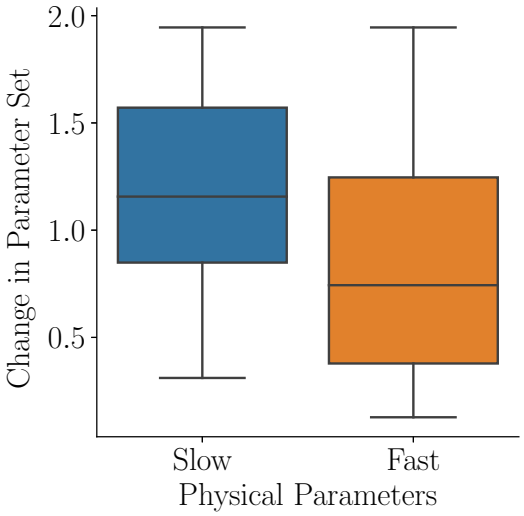

Fig. 10: A box plot of the Euclidean distance between parameter sets in successive epochs for both the slow and fast regimes. Note this is epoch to epoch for both predator and prey parameters and is *not* comparing successive prey parameter sets to predator ones. The rate of change of subsequent parameter sets is similar among both regimes.

a consistent average speed of prey across optimisations in the fast regime in contrast to the oscillatory results in the slow regime. Further, there are strong oscillations in the average group number for the slow regime and much softer variations for the fast regime. It is worth pointing out that in the slow regimes, the oscillations move in sync for both metrics, and the period of oscillation is fairly consistent, generally jumping with each prey optimisation. If we plot the prey-to-prey attraction and repulsion terms of our prey for each epoch on our phase portrait in Figure 2 we end up with Figure 12. Here we can see that in the fast regime prey always prefer to ball or flock, however in the slow regime we get two distinct groups of behaviours where about half the time the prey are balling and the other half the prey are expanding. This gives further credence to our observation that there seem to be oscillations in behaviour in the slow regimes and less so in the fast regimes. Notice that the centering when expanding is very minimal. We conjecture that increasing the centering further (to expand more strongly) would

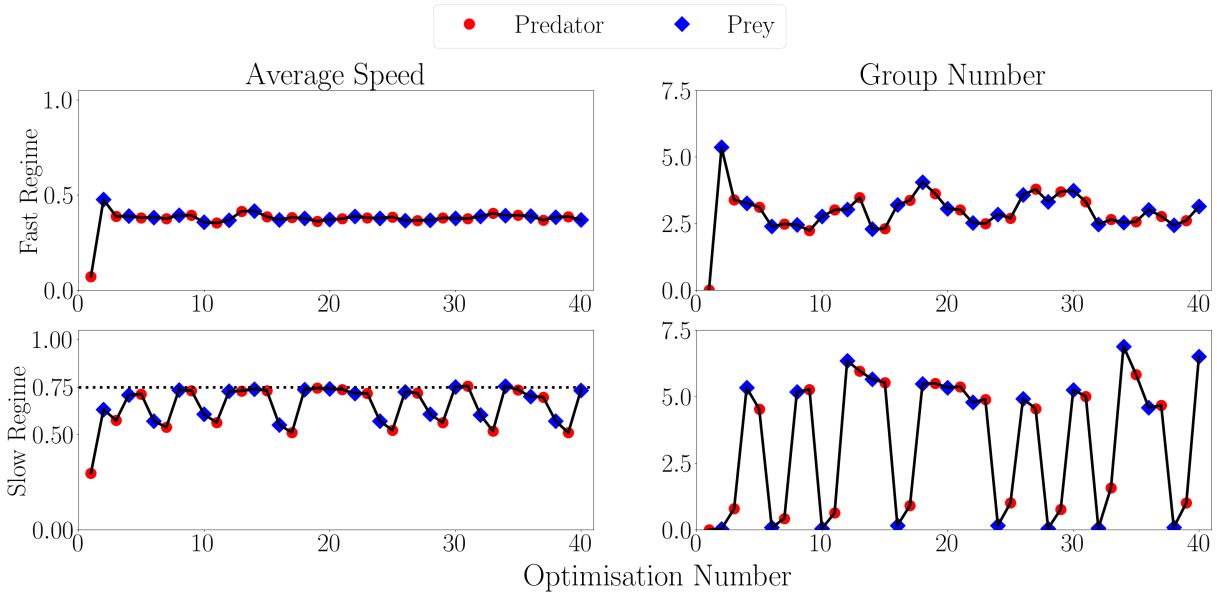

Fig. 11: We see apparent switching of dynamics between two distinct states in the slow regime and convergence with long-period, small-amplitude oscillations in the fast regime. This figure displays the average speed and group number of prey, using the parameter values after each optimisation step, from the fast and slow regimes, averaged over 100 simulations. The dashed line in the bottom left graph represents the maximum predator velocity.

be detrimental as it would overpower the alignment term used to propagate information, leading to higher predation. The distribution of balling is comparatively rather varied.

It is also surprising that in Figure 11 we can see that prey have a slower average speed in the fast regime. This is indicative of the fact that our model has momentum: it is easier to change direction (and thus react to a predator) at a lower speed. Travelling at maximum speed is not useful when a predator can outrun you. Conversely, we see that in the slow regime, the average speed of prey is at most at quick as the predators. Beyond this speed, there are no forces that will maintain acceleration for long periods of time.

7.4 Regression Analysis

In this section, we consider the two extended runs with their heterogeneous behaviour and seek to quantitatively identify consistent responses in behaviour. To this end, we wish to perform a multiple linear regression analysis for each predator coefficient, as a function of all the prey coefficients they have just optimised against and then vice versa, using the data from both extended runs. Take a for example, the coefficient of the predator attraction to prey force as in equations (8), and $\alpha, \beta, \gamma, \delta, \kappa$, the optimised prey coefficients as in equations (4). We are trying to find k_j and its corresponding significance, for $j = 0, \ldots, 5$ in the following equation;

$$a = k_0 + k_1\alpha + k_2\beta + k_3\gamma + k_4\delta + k_5\kappa \qquad (23)$$

7.4.1 Significant Results

We start by looking for correlations in how the predators optimise their parameters given the previous prey optimisation. For each of a, b, c, d, as in equations (8), we do a multiple regression on the prey parameters $\alpha, \beta, \gamma, \delta, \kappa$, as in equations (4). As predators optimise first, we cannot include their first optimisation, leaving us with a total of 38 data points to work with. The high-dimensionality of our system, with limited data, reduces the odds of finding statistically significant correlation coefficients. We deem a correlation coefficient is statistically significant if there is less than a 5% chance of observing the data.

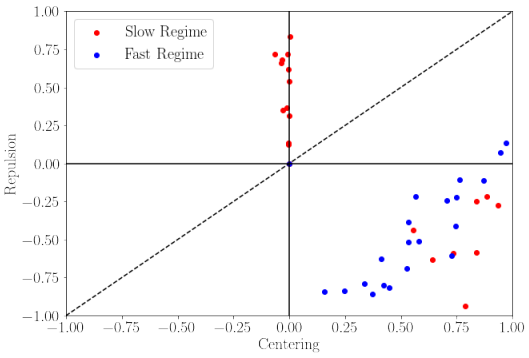

Fig. 12: Bouncing between behaviours in optimisation of the slow regime, but not in the fast. Parameter values returned from successive optimisation steps plotted on the phase plane described in Figure 2, $L = 30$.

Although, if there is no correlation between the variables in question, the 95% confidence interval for the corresponding correlation coefficient will not include zero. We include the significant correlations in Table 4. This suggests that predators are less likely to chase prey (a decrease in d), when prey repel each other more (an increase in γ). Although, predators are more likely to chase prey (an increase in d), when prey dis-align from the predators more (an increase in κ). We also see that predators are more likely to repel from one another when prey do the opposite (an increase in c, in response to increases and decreases in β and γ respectively).

Repeating this method for prey, we consider how $\alpha, \beta, \gamma, \delta, \kappa$ after a prey optimisation are correlated to a, b, c, d. Which results in the second half of the results in Table 4. Both results here can be summarised as the more predators chase, the more prey spread out. This is already enough to give a sense of back-and-forth behaviour optimisation, we see that prey centering encourages predators to chase more, which in turn encourages prey to spread out again.

7.4.2 Limitations of the Regression Analysis

Multiple linear regression assumes that input parameters are not correlated to each other and

Dependent	Independent	Coefficient	P-value
c	β	0.230	0.000
	γ	-0.257	0.000
d	γ	-0.287	0.018
	κ	0.348	0.003
β	d	-1.344	0.001
γ	d	0.965	0.000

Table 4: Force coefficients of predator and prey that are statistically significantly correlated with force coefficients of prey and predator (equations (4) and (8)) respectively following a multiple regression analysis.

therefore can be less effective at predicting the significance of the correlation between our independent and dependent variables. This issue known as co-linearity. To test for this, we can calculate the Variance Inflation Factor (VIF) of the parameters that we are fitting to [27]. The minimum possible VIF is one and a VIF of at least five is considered to be indicative of critical levels of co-linearity between the current parameter and the others. If we find values greater than five we can perform the regression without the highly co-linear variable(s). See the results in Table 5, for the prey and predator parameters. Note that two variables can be reasonably strongly correlated and still have a low VIF in regard to the entire system, this is the case for the prey-to-prey repulsion parameter and the prey-to-prey attraction parameter. A variable with a high VIF, for example, the repulsion from predator parameter is highly correlated with the other variables in the system, this could be for example, because it is desirable to maintain a certain ratio between forces. See Table A1 in the appendix, for the regression results once we restrict our independent parameters to a subset that is not highly correlated. In summary, however, the variables that were statistically significant stay so, meaning we should not feel that we need to change our previous conclusions. The fact we do not see strong trends is not surprising or a cause for concern, given both our relatively small amount of data and the fact that the parameter sets change a lot from epoch to epoch as in Figure 10. The possible disparity between both regimes may also affect this.

Coefficient	VIF		Coefficient	VIF	
α	19.9		a	11.0	
β	1.9	1.6	b	1.2	1.1
γ	4.5	4.0	c	1.8	1.8
δ	28.5		d	11.7	1.7
κ	8.4	3.9			

Table 5: VIF values for sets of parameters. The first column for all the listed parameters and the second for a restriction to some maximal subset for which the VIFs are no higher than 5 chosen to include the variables in Table 4.

8 Conclusion

During our review of collective motion, we delve into the sparsely researched dynamics of multiple predator, many prey systems. Our model is primarily based on the force-based approaches common in the literature [7, 9], but deviates by the inclusion of a soft boundary condition and predator disalignment. Following this, we provide a short characterisation of prey behaviours by the relative strength of repulsion and centering terms. We introduce various measures to analyse the differences in these behaviours and our choice of a soft boundary.

After developing an understanding of our model, we implement an evolutionary adaptation mechanism to find optimal predator-chasing and prey-survival strategies. We did this by sequentially applying the BIPOP-CMA-ES algorithm, optimising predator and prey behaviour in turn. Boad analyses varying the spatial and physical constraints on our agents provided two interesting settings, which we consider in more detail. Both of these settings are in the low-density scenario on a large domain but have different predator regimes, one using predators that were slower than the prey and the other predators that were faster. In the slow predator case, successive optimisations lead to large regular oscillations in predation, Figure 9. We observe distinct cycling in prey and predator behaviour in response to one another, made clear by tracking several order parameters, Figure 11. We also categorise behaviours on the phase diagram 12, where prey behaviour oscillates between expansion and balling. On the other hand, when predators are faster than the prey, we find prey optimisations have a negligible effect on the success of predators, Figure 9. The stability of predation values in this scenario is mirrored in the convergence of our metrics after each optimisation step, 11 and the identification of a single behaviour, balling, on our phase diagram 12. It is highly interesting that in neither scenario do we see convergence to a specific parameter set, Figure 10, although it is clear that specific behaviours are being optimised for.

Three natural extensions of the model are to extend the space to thee-dimensions, reduce the vision cone to be less that 360° and to introduce friction. These would model a more realistic scenario, though the added work may not provide distinctly new results. The latter is likely the direction that will provide the greatest change. The noise term ensures the model is non-deterministic, but can arbitrarily cause agents to slowly accelerate over time. As we kept this term small, $\sigma = 0.05$, we do not believe it to have a significant effect, but one could perform a formal analysis to explore to what extent it hask. There is scope to adapt the model into a replicator system, where agents can reproduce and consequently mutate, which provides an alternative evolutionary method to CMA-ES.

Acknowledgements. We wish to extend our thanks to Gareth Alexander and Matthew Turner for nurturing our research direction and adding valuable insights. Additional thanks can be extended to Fernando Peruani and Ananyo Maitra for offering expert opinions. This research was funded by the EPSRC.

Appendix A

In Figure A1, prey have not optimised meaning their force coefficients are zero. This results in no escaping or avoiding predators meaning the majority of deaths are caused in one half of the space, corresponding to where the predators start.

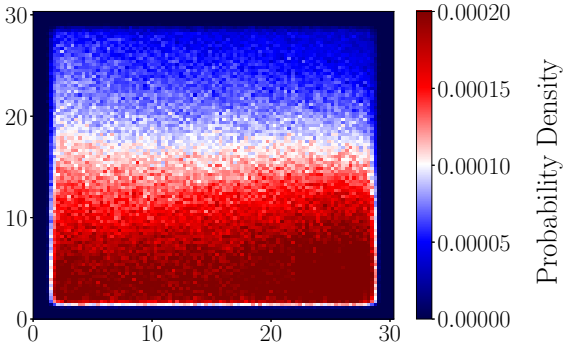

Fig. A1: Heatmap of deaths after initial predator optimisation. $L = 30$, slow regime.

In Figure A2 we can see that the relative variation of predation in the slow regime is significantly larger than that of the fast regime. Furthermore, the predation value has a very short burn-in period before reaching its transient state.

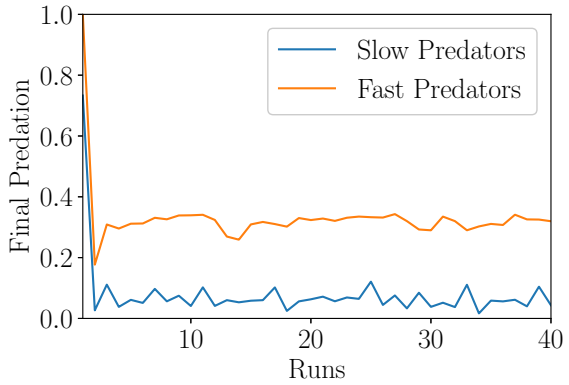

Fig. A2: Convergence of final predation in the fast and slow regimes. Slow predators are represented in blue and fast in orange.

In Figure A3, we can see that slow predators still exhibit a larger absolute change in predation between optimisations, even if not normalised by the mean, as shown in Figure 9.

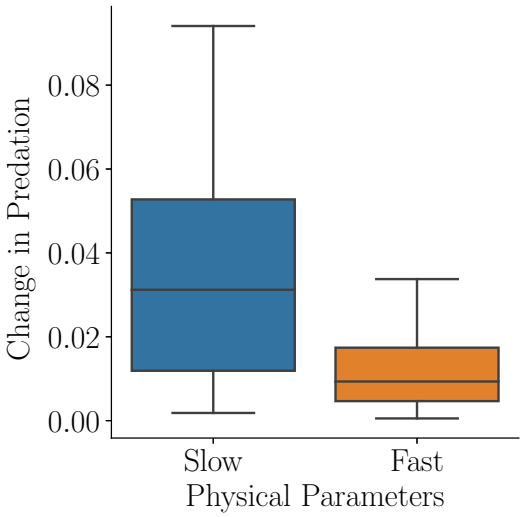

Fig. A3: A box plot of the absolute change in predation between successive optimisations. The slow and fast regimes are shown in blue and orange respectively.

Dependent	Independent	Coefficient	P-value
c	β	0.222	0.000
	γ	-0.251	0.000
d	γ	-0.296	0.010
	κ	0.344	0.002
β	d	-1.342	0.001
γ	d	0.966	0.000

Table A1: Statistically significant correlations of force coefficients (equations (4) and (8)) between predator to prey responses and vice versa. This is done using multiple linear regression analysis excluding highly colinear variables.

References

[1] Breder, C.M.: Studies on social groupings in fishes. bulletin of the amnh; v. 117, article 6 (1959)

[2] Potts, W.K.: The chorus-line hypothesis of manoeuvre coordination in avian flocks. Nature **309**(5966), 344–345 (1984)

[3] Topaz, C.M., D'Orsogna, M.R., Edelstein-Keshet, L., Bernoff, A.J.: Locust dynamics: behavioral phase change and swarming (2012)

[4] Silverberg, J.L., Bierbaum, M., Sethna, J.P., Cohen, I.: Collective motion of humans in mosh and circle pits at heavy metal concerts. Physical review letters **110**(22), 228701 (2013)

[5] Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., Shochet, O.: Novel type of phase transition in a system of self-driven particles. Phys. Rev. Lett. **75**, 1226–1229 (1995) https://doi.org/10.1103/PhysRevLett.75.1226

[6] Mohapatra, S., Mahapatra, P.S.: Confined system analysis of a predator-prey minimalistic model. Scientific Reports **9**(1), 11258 (2019)

[7] Janosov, M., Virágh, C., Vásárhelyi, G., Vicsek, T.: Group chasing tactics: how to catch a faster prey. New Journal of Physics **19**(5) (2017) https://doi.org/10.1088/1367-2630/aa69e7

[8] Chakraborty, D., Bhunia, S., De, R.: Survival chances of a prey swarm: how the cooperative interaction range affects the outcome. Scientific Reports **10** (2020) https://doi.org/10.1038/s41598-020-64084-3

[9] Chen, Y., Kolokolnikov, T.: A minimal model of predator–swarm interactions. J. R. Soc. Interface (2014) https://doi.org/10.1098/rsif.2013.1208

[10] Kumar, V., De, R.: Efficient flocking: metric versus topological interactions. Royal Society open science **8**(9), 202158 (2021)

[11] Barberis, L., Peruani, F.: Large-scale patterns in a minimal cognitive flocking model: incidental leaders, nematic patterns, and aggregates. Physical review letters **117**(24), 248001 (2016)

[12] Collin, S.P.: Scene through the eyes of an apex predator: a comparative analysis of the shark visual system. Clinical and Experimental Optometry **101**(5), 624–640 (2018)

[13] Pita, D., Moore, B.A., Tyrrell, L.P., Fernández-Juricic, E.: Vision in two cyprinid fish: implications for collective behavior. PeerJ **3**, 1113 (2015)

[14] Franks, N.R., Pratt, S.C., Mallon, E.B., Britton, N.F., Sumpter, D.J.: Information flow, opinion polling and collective intelligence in house–hunting social insects. Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences **357**(1427), 1567–1583 (2002)

[15] Lindauer, M.: Communication in swarm-bees searching for a new home. Nature **179**, 63–66 (1957)

[16] Seeley, T.D.: The Wisdom of the Hive: the Social Physiology of Honey Bee Colonies. Harvard University Press, Cambridge, Massachusetts (2009)

[17] Tambling, C.J., Druce, D.J., Hayward, M.W., Castley, J.G., Adendorff, J., Kerley, G.I.: Spatial and temporal changes in group dynamics and range use enable anti-predator responses in african buffalo. Ecology **93**(6), 1297–1304 (2012)

[18] Magurran, A.E., Higham, A.: Information transfer across fish shoals under predator threat. Ethology **78**(2), 153–158 (1988)

[19] Couzin, I.D., Krause, J., Franks, N.R., Levin, S.A.: Effective leadership and decision-making in animal groups on the move. Nature **433**(7025), 513–516 (2005)

[20] McNamara, T.F.: Rust in Action: Systems Programming Concepts and Techniques. Manning, Shelter Island (2021)

[21] Calovi, D.S., Lopez, U., Ngo, S., Sire, C., Chaté, H., Theraulaz, G.: Swarming, schooling, milling: phase diagram of a data-driven fish school model. New Journal of Physics **16**(1), 015026 (2014)

[22] Ester, M., Kriegel, H.-P., Sander, J., Xu, X., *et al.*: A density-based algorithm for discovering clusters in large spatial databases with noise. In: KDD-96 Proceedings, vol. 96, pp. 226–231 (1996)

[23] Gan, J., Tao, Y.: Dbscan revisited: Mis-claim, un-fixability, and approximation. In: Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, pp. 519–530 (2015)

[24] Hansen, N., Müller, S.D., Koumoutsakos, P.: Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es). Evolutionary computation **11**(1), 1–18 (2003)

[25] Loshchilov, I.: Cma-es with restarts for solving cec 2013 benchmark problems. In: 2013 IEEE Congress on Evolutionary Computation, pp. 369–376 (2013). Institute of Electrical and Electronics Engineers

[26] Hansen, N.: Benchmarking a bi-population cma-es on the bbob-2009 function testbed. In: Proceedings of the 11th Annual Conference Companion on Genetic and Evolutionary Computation Conference: Late Breaking Papers, pp. 2389–2396 (2009)

[27] Forthofer, R.N., Lee, E.S., Hernandez, M.: 13 - linear regression. In: Forthofer, R.N., Lee, E.S., Hernandez, M. (eds.) Biostatistics (Second Edition), Second edition edn., pp. 349–386. Academic Press, San Diego (2007). https://doi.org/10.1016/B978-0-12-369492-8.50018-2 . https://www.sciencedirect.com/science/article/pii/B9780123694928500182