DRAFT

Adaptive Systems Project

Flocking Under Predation

A Robust Evolutionary Strategy

Dan Sayers, EASy MSc 2008

20.04.09

Abstract

Increasing evidence points to predation as a likely source of evolutionary pressure behind flocking behaviours in animals. In simulation, flocking is shown to be a highly stable strategy for simple agents with steering behaviours, under (co-) evolutionary pressure of predation. A number of perturbations are undergone to assess the stability of the behaviour. Suggestions are made for future exploration of flocking strategies.

1 Introduction

Collective behaviours have long held our fascination - from the flocking of birds to the complexities of human social interactions. Even such simple entities as bacteria and basic physical processes have been shown to exhibit characteristics of emergent group dynamics. The relative advantages of coordinated group behaviour in flocks, schools and herds of animals have been widely debated. Do they arise as a response to predation, foraging or reproductive needs? What abilities must a species have for flocking to develop? How do purely individual needs give rise to such well-coordinated collective traits?

In recent years, technology has improved to the point where it is possible to model (in limited ways) the bodies, brains and evolutionary processes behind complex group behaviours. For the first time we are able to analyse the probable environmental requirements and exclusions in order for such traits to evolve. In the last decade, advances have been made towards exploring the simulated evolution of systems exhibiting collective locomotion. The purpose of this report is to give a brief assessment of the state of the art, as well as to present a simulation in which (despite recent studies having failed to achieve this goal) flocking is shown to be a highly robust and likely strategy - as a direct result of predation - in agents capable of perceiving each other's motion. Our contention is that it should no longer be asked under what circumstances flocking might arise in evolution - rather, under what if any circumstances flocking might *fail* to arise, when under the influence of predation.

2 Background

2.1 The Biology of Flocking

In the area of Life Sciences, much work has been done to statistically analyse flocking and its relation to predation, particularly in recent years. [1] discusses the mathematics behind increases in alertness to predation within flocks - and also raises the question of communication through sound as an important aspect of flocking and predator avoidance. Several studies support a strong link between flocking and predation in bird species - flocking has been observed to be reduced under reduced predation [2], to increase individual survival rates, and to slightly *reduce* foraging efficiency in redshanks [3], as well as akepa [4].

Other work has used recent technology to precisely measure animal positions within flocks - giving precise analyses of the dynamics within fish shoals [5], and flocks of starlings [6]. Most importantly for our current purposes, both studies confirm a strong link between biological flocking and the simulation concepts described in the next section - to be incorporated in our simulation.

2.2 Simulated Flocking

In the mid-twentieth century, pioneers of cybernetics such as Grey Walter created the first ”free goal-seeking mechanisms” [7]. Whilst robotics remained a popular dream, it was not until the 1980s that digital simulation of locomotion became a possibility. At this time, several key advances were made. Braitenberg published an influential exploration of a cybernetic thought experiment, involving simple vehicles with sensors connected directly to motors [8]. Reynolds proposed a simple system for simulating autonomous agents [9] - as well as a model for flocking based on simple locally-applicable rules [11]. A similar conception was previously described by Aoki [12]. Reynolds’ system for flocking agents, known as Boids, has proved successful - and is widely used as a model for real-world flocking in both academia (e.g. as noted above, real birds and fish have been shown to conform to the model) and entertainment (as computer graphics of flocks and swarms, for film, television and computer games). Little advance has been made on Reynolds’ three simple rules for flocking since their inception. Those rules are:

1. Collision avoidance: avoid collisions with nearby flockmates
2. Velocity matching: attempt to match velocity with nearby flockmates
3. Flock centering: attempt to stay close to nearby flockmates [11]

These are commonly referred to as the rules of separation, alignment and cohesion. In recent years, advances have been made in the mathematics [13] and algorithmic efficiency [14] of calculating boid-like flocking. Also, increasingly studies have looked at simulated flocking combined with simulated evolution, as well as simulated neurology - to explore flocking as an adaptive and cognitive phenomenon.

2.3 Evolving Intelligent Flocking

In the 1990s, the evolution of computing power brought simulation speeds to a level capable of connecting autonomous locomotion to the disciplines of artificial evolution and neural networks. A number of studies have attempted to evolve, coerce or coevolve flocking in autonomous intelligent systems - with varied results. Cliff and Miller give an early (1994) account of the challenges inherent in neurally coevolving pursuit and evasion [15], and present a successful simulation - postulating the emergence of flocking given a greater simulation scale. Reynolds presented a simulation for coevolving the game of tag with Genetic Programming [17] - again with simulation sizes too limited to explore flocking.

A number of successful studies have demonstrated the evolution and coevolution of simple group dynamics in khepera robots ([19], [20], [22]). Khepera robots are not equipped with sensors capable of determining mutual direction of movement, limiting their ability to develop flocking behaviours. In [22], khepera are given the interesting ability to communicate their position through sound, although this does not lead to flocking.

More recently, researchers have looked at evolving flocking agents with neural control systems. In [21], flocking is successfully coevolved in simulated fish with neural networks, as a response to predation. This is a useful result, although the authors do not apply any measure to the observed flocking, or assess the morphology of the fitness landscape involved. [24] describes a model in which (neural) coevolution of flocking is attempted under predation - with some success, although the authors admit the evolution required 'steering' to achieve flocking. Another study coevolves boids already able to flock, and shows that their flocking increases as a direct response to predation, rather than to foraging [25]. Very recent work has looked into complex neuro-evolution of specialisms in group behaviour [26]. Whilst there have been successes in the above approaches, we are still without a clear characterisation of the ecological niche occupied by flocking behaviours, or of the attributes necessary for agents to develop them.

2.4 A Broader Perspective

In the continuing struggle to understand how flocking arises, the number of observations and applications of group behaviours grows. Bacteria exhibit group morphologies [27], and can coordinate aggregate behaviours through chemical communication [28]. In Statistical Physics, phase transitions have been studied in particle systems under a 'rule of alignment' - with applications to various physical and biological systems [29]. And in Information Theory and Artificial Intelligence, the concepts of Particle Swarm Optimisation and Swarm Intelligence have added new computational applications of flocking behaviours. A greater understanding of the ontology of coordinated group behaviours may have wide-reaching consequences in a broad variety of fields.

3 Method

A simulation environment is created in which two-dimensional mobile agents of two types - predators and prey - coevolve. The predators are rewarded for their hunting ability with reproductive success. The prey become increasingly fecund as their lifespan continues - before dying of (fixed-length) old age.

To model the agents and environment, the steering behaviours described by Reynolds were used [9]. The agents slow down in collisions with each other and with obstacles, and have weighted connections mapping sensors to steering. These weights constitute the agents' genomes, which evolve via a haploid process. The aim of the simulation is to ascertain whether, given the exact information needed to produce flocking, these simple agents will evolve flocking. And if so, how sensitive such behaviour will be to perturbation. Note that the sensory information required for flocking (average position and heading of nearby fellow prey) is also sufficient for other behaviours, e.g. avoidance, clustering, collision, dispersal.

To model the agents and environment, the OpenSteer engine for C++ was used. This library was originally developed by Reynolds, and has since been extended by a network of contributors. It allows for efficient development of autonomous vehicles following Reynolds' steering paradigm [31].

The code for the simulation is available online [10] as a 'plugin' for OpenSteer. It can thus be tested on any of the Platforms OpenSteer supports - currently Windows, OS X and POSIX (UNIX-like) systems. In order to increase the performance of OpenSteer for testing purposes, some changes have been made to the library files - the complete project for XCode in OS X is also available on the project web page.

3.1 Vehicles

Name	Type
Mass	Scalar
Position	Vector
Velocity	Vector
Max Force	Scalar
Max Speed	Scalar
Orientation	N Basis Vectors

Table 1: Attributes of Steering Agents

In Reynolds' formulation of steering for autonomous agents [9], each agent has the physical properties given in Table 1. A steering force is applied in each time step (constrained by Max Force, and also in our simulation, a maximum deviation angle from the vehicle's forward direction). This, divided by the mass, gives the vehicle's acceleration - which is added to the velocity, and constrained by Max Speed (as well as a maximum deviation angle), to give the resultant position of the vehicle in each time step.

3.2 Sensors

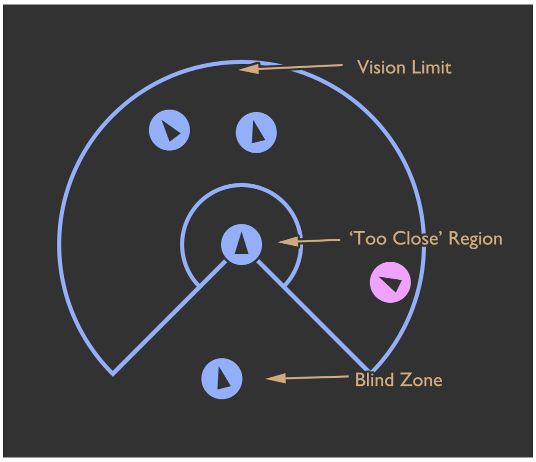

Figure 1: Prey perception areas

Type	Value
Fellow prey position	$\vec{x}$
Fellow prey velocity	$\vec{v}$
Fellow prey 'too close'	$\frac{r_{close}}{\|\vec{c}\|^2}\vec{c}$
Obstacle	$\frac{r_{max}}{\|\vec{o}\|^2}\vec{o}$
Predator	$\frac{r_{max}}{\|\vec{p}\|^2}\vec{p}$

Table 2: Prey sensors

Following the work of e.g. [12], [25], the prey vehicles have an inner and outer radius of perception - as well as a maximum viewing angle, as shown in Fig. 1. The outer radius r_{max} is the distance limit to which fellow prey, predators and obstacles can be perceived. The inner radius r_{close} activates a 'too close' sensor, allowing the vehicles to evolve the behaviour of separation. The prey sensors are as given in Table 2, where $\vec{x}$ is the average position of fellow prey between r_{close} and r_{max} (and within viewing angle), $\vec{v}$ is the average velocity of the same vehicles, $\vec{c}$ is the average position of prey in the 'too close' region, $\vec{c}$ is the average closest-points of any obstacles within vision, and $\vec{p}$ is the average position of any predators within vision. Note that the sensor vectors for 'too close' prey, obstacles and predators are inverted in length so that the vectors are greater in magnitude when the target is closer. This corresponds to the intuitive sense that these sensors are likely to be used for avoidance and hence will be more useful to the vehicles when giving greater readings for nearer targets.

Note that whilst we occasionally use the word 'vision' to describe these sensors, they need not be visual. We assume the vehicles are able to perceive the positions and velocities of their neighbours using some sort of detection method. We leave the specifics of the morphology and modality of these sensors as a given, in order to assess the behaviour of evolving creatures capable of perceiving such localised data.

Type	Value
Target prey position	$\vec{x}_{target}$
Target prey velocity	$\vec{v}_{target}$
Target prey pursuit	$\vec{x}_{target} + t\vec{v}_{target}$
Obstacle	$\frac{r_{max}}{\|\vec{o}\|^2}\vec{o}$
Fellow predator position	$\vec{y}$
Fellow predator velocity	$\vec{u}$
Fellow predator 'too close'	$\frac{r_{max}}{\|\vec{d}\|^2}\vec{d}$

Table 3: Predator sensors

The predators' sensors are similar to the prey's - with the addition of extra sensors for detecting prey position and heading, plus a predictive vector for pursuit given by a linear combination of prey position and velocity. The predator sensors are defined in Table 3, where t is the prediction time for the pursuit sensor and $\vec{d}$ is the average position vector of 'too close' fellow predators. Note that t is a fixed constant. The pursuit sensor was added during development as a default predatory behaviour. In fact, as we will see in section 4.2.3, it is not necessary to include this behaviour in order to stimulate the evolution of flocking.

A predator chooses a target by searching for the nearest prey. It also considers slow-moving prey as priority over faster moving nearer prey. This factor, again, was added during development in order to speed up evolution (punishing prey 'clumping'). Again it is found to be non-essential in the precipitation of group prey behaviour.

In this simulation, the predators' vision is not angle-limited as the prey's is, and can pick out prey from a distance. It is not expected that this should adversely affect the results. [25] has shown that in a simulated coevolutionary setting, predators tend to develop small vision angles while prey develop wide ones.

3.3 Interactions

As well as sensing their environment and each other, vehicles are faced with physical interactions. They can bump into each other, or into obstacles. When they collide into each other, their maximum speed is reduced to a minimum. Thereafter it increases linearly, returning to its previous value. Similarly for obstacles, a collision causes slowing. When a predator meets a prey, it 'eats' it - the prey's body is consumed (and another prey is regenerated by mutating one of it's fellows - see section

3.7). The predator slows down while ’eating’. Experiments are made modifying and removing these interactions, as detailed in section 4.2.2.

3.4 Noise

As a feature of the simulation environment, broadband noise may be added to each of the inputs and outputs of the prey control systems. This allows us to broadly model inefficiencies and imperfections of sensors or actuators (see section 4.2.5).

3.5 Obstacles

The primary obstacle is a circular perimeter wall, within which all interactions occur. Additionally, an arbitrary number of smaller circular obstacles can be placed inside, or removed from, the perimeter boundary during runtime. These obstacles, when added to the scene, appear in random non-overlapping positions and provide a method for analysing the ability of the vehicles’ behaviour to adapt to changing environmental morphology (see section 4.2.6).

3.6 Evolution

Both predators and prey have fixed population sizes. This constrains the wilder aspects of coevolution, allowing us to focus on the development of behaviours without regress of causality in cases of failure. The prey regenerate on an individual basis - whenever one dies (by being eaten, hitting into an obstacle or fellow prey too hard or from old age) a new one is created, with random position and heading within the boundary. For the connection weights of its control system, it mutates the genome of one haploid parent - chosen from the surviving prey with probability weighted to those that have lived longest.

The predators regenerate on a group basis, in situ: they have a fixed regeneration period, and each has a ’score’ of how many prey they have eaten. When the regeneration period is complete, all predator are replaced by mutations of the single highest-scoring predator from the current and previous generations. Note that this rather aggresive ”alpha male” breeding strategy for predators was found to be the most powerful for stimulating the evolution of flocking, having also tested broader-population methods of predator evolution. It does, however, create a a preponderance of drift and instability in the predator population.

3.7 Reporting

For the purpose of analysis, the simulation environment generates a data report every fixed number of timesteps. The reports include means and standard deviations of connection weights for current predators and prey, the average age of the prey, nearest neighbour distance and cosine of nearest neighbour anglular deviation. These last two give a broad measure of flocking behaviour - flocking creatures tend to keep their nearest neighbour distance to a low safe minimum (cohesion and separation) and the angular deviation of their velocities small (alignment).

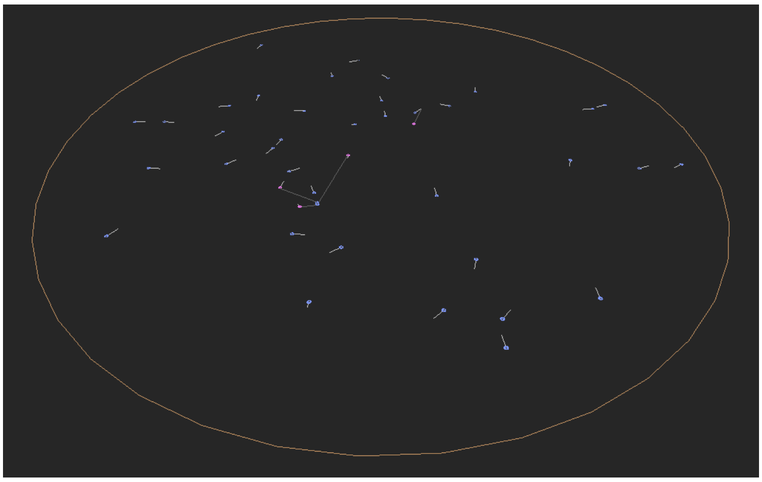

Figure 2: A typical run - start. Prey are blue, predators pink. Grey lines are velocities.

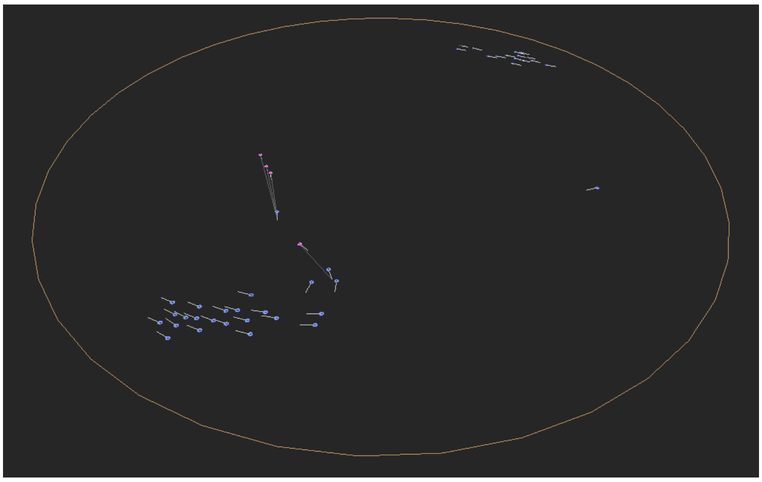

Figure 3: A typical run - later. Flocking has developed.

4 Results

The simulations was successful in evolving flocking in a majority of cases. Firstly, we discuss the results for a typical run, in which 40 prey and 4 predators coevolve to stable strategies involving flocking. We then consider the result of various perturbations of this system in order to assess the limits of this stability.

4.1 A Typical Run

A view of the simulation environment can be seen in Fig. 2. This image is taken at the start of the run, when the prey have weights set as gaussian random variables centred around zero. Over the course of approximately 1 hour (measured in a time frame making the motion appear realistic) the predators and prey have coevolved to give a scene as in Fig. 3 - the prey have settled to a state of classic flocking behaviour.

In Fig. 4 we see the development of the inter-prey distance and angle - clearly showing a tendency towards a small fixed distance (cohesion and separation) and near-unity cosine (alignment). The

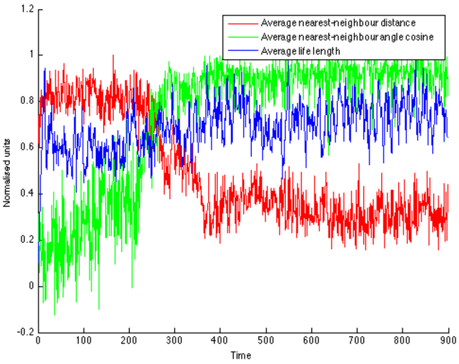

Figure 4: Evolution of flocking measures over time

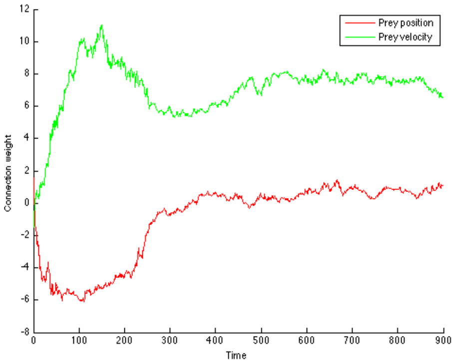

Figure 5: Evolution of weights for alignment, cohesion

averaged data is sampled once every 500 frames, and as a result the plots are quite noisy - however the trends are clear. Fig. 5 shows the connection weights relating to cohesion (position) and alignment (velocity). Again, these clearly evolve to a stable state with both weights positive. The alignment weight is of greater magnitude, with the cohesion weight being near-zero. There appears to be a phase transition around sample 250 in which the prey change from avoidance-alignment to cohesion-alignment. Previous to this transition these two connection weights appear to have a (modified) inverse relation, as can be seen in the first 250 samples of Fig. 5. This is borne out by repeated experiments.

Fig.4 shows that the average lifetime of the prey gradually increases and levels out over time. This in itself does not assure us of any superiority in fitness of the flocking behaviour over other possible behaviours in general, however - as this is a coevolutionary environment, and the predators' behaviours are changing too.

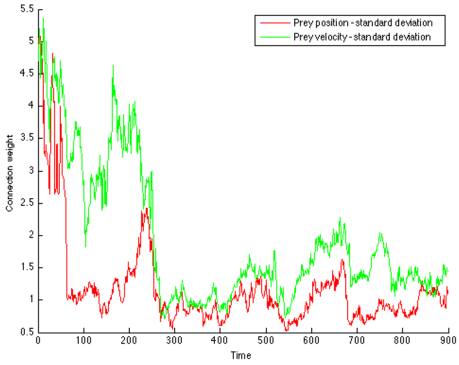

Figure 6: Weights standard deviation over time

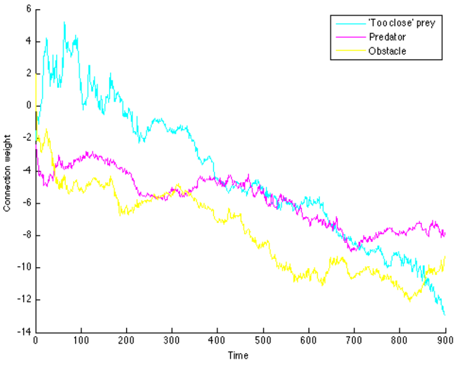

Figure 7: Evolution of weights for avoidance

In Fig. 6, we can see that the standard deviation of the prey's weights for alignment and cohesion reduces over time - with a substantial reduction at around the same time that the flocking behaviour appears. Once flocking is in place, the standard deviation for each weight stays relatively low, with small fluctuations.

Fig. 7 shows the remaining prey weights (averaged) over time. These weights give the avoidance vectors for fellow prey, obstacles (perimeter) and predators. They are all increasingly negative over the given timescale. Over longer timescales these reach more-or-less stable values (Fig. 8).

Figs 9 and 10 show the evolution of the average predator connection weights over this longer timescale. Fig. 9 shows the weights for pursuit and obstacle avoidance. Interestingly, there is a late

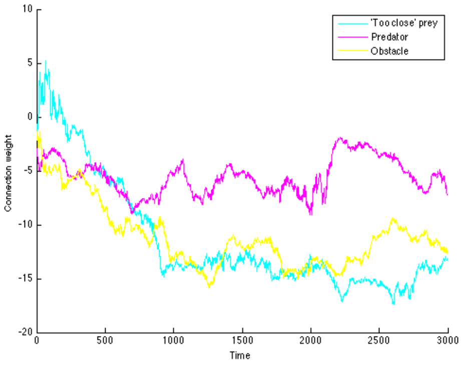

Figure 8: Avoidance weights - longer run

surge for the weight connecting prey velocity vector to predator steering. This may be a reaction to the increased alignment of the prey flock - as heading in the same direction as the prey will work well as a strategy for tracking a herd or flock. Fig. 10 shows the inter-predator behaviour weights. In this case the predators initially develop something apparently similar to flocking themselves - they have positive alignment and cohesion weights. However given the small population size (4) and since they spend more of their time chasing prey than interacting with each other, this is not obvious from their behaviour. This apparently stable state soon changes and the predators explore other behaviours. This high level of drift is clearly influenced by the aggressive winner-takes-all breeding strategy, and the small population size. However it is also an indicator that as clear a group behaviour strategy might not be available for predator as for prey. We will return to this question in section 4.2.4.

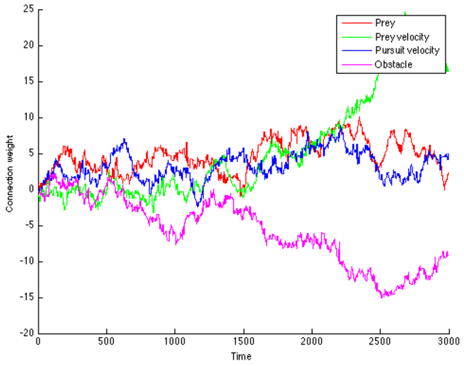

Figure 9: Predator weights for pursuit, avoidance

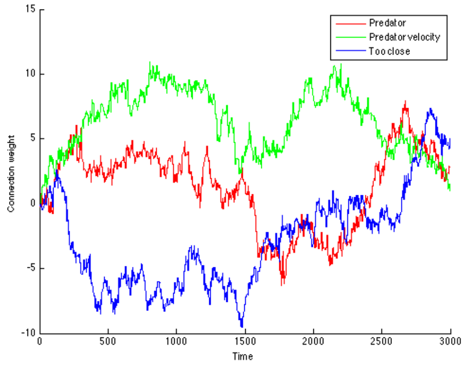

Figure 10: Predator weights for group behaviour

4.2 Perturbations

Given the successful coevolution of flocking with predation described above, we perturb the system to assess the stability of this arrangement.

4.2.1 Population sizes

Experimenting with the ratio of predators to prey, whilst keeping the combined population roughly equal, reveals that flocking is a remarkably robust strategy in this regard: the prey continue to evolve flocking until they are outnumbered roughly 2:1. At this point the realism of the simulation must be questioned, since most prey hardly live long enough to find each other, let alone to flock. Altering the ratio in the opposite direction, flocking evolves even with many prey and just one predator - see Fig. 11.

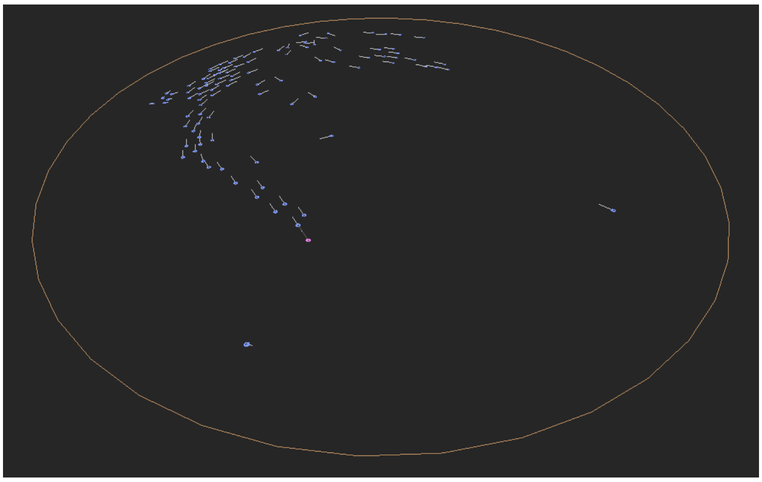

Figure 11: 100 prey, 1 predator.

In highly constrained runs, where the perimeter radius is small compared to the number of prey, flocking does not evolve. Rather a misalignment-cohesion behaviour develops where the prey hit into each other repeatedly in tight clusters.

Flocking does not occur below a certain threshold number of prey - around twelve for the perimeter size given in the typical run described in section 4.1. Given that the prey regenerate to a random position in the environment, low numbers do not allow for easy flocking. Again the system breaks down, but outside of realistic conditions.

Overall the flocking behaviour seems remarkably resilient to variations in population density and ratio.

4.2.2 Physical constraints

In the runs above, the major external physical constraints on the predators and prey occur in the actions of bumping and eating. Both slow the agents down - their speed gradually increasing afterward. If the predators are not slowed down by 'eating' prey, flocking does *not* evolve - rather, a misaligning-avoiding behaviour emerges. This suggests that the emergence of flocking is dependent on the flock being able to get away in the time taken for the predator to deal with its catch. Experimenting with gradually reintroducing the kill recovery period reveals a very sharp phase transition, in which flocking suddenly emerges at a certain threshold predator recovery time.

Agents are also internally constrained by their maximum speed and force parameters. In the simulations described above, the predators have greater values for these parameters. In fact, flocking will still evolve (albeit at a greatly reduced rate) even when the predators' maximum force and speed are set to slightly lower values than the prey's. Thus, as in nature, the predators are able to capitalize on mistakes and weaknesses (collisions) and survive on flocks of creatures normally able to outrun them. Below a (gradual) threshold, the predators are unable to kill enough prey for evolution to occur.

4.2.3 Eliminating sensors

The predators and prey are equipped with limited sensory input. We can reduce their abilities further, and still achieve flocking. Of the inter-prey sensors, removal of the cohesion sensor (nearby prey average position) results in very similar flocking patterns to when it is present. It might be argued that cohesion in this case is enforced by the presence of a circular boundary. It seems clear in any case that cohesion is a longer-range effect than alignment or separation. If an agent is near or in a flock, its effect is negligible.

On removing the alignment sensor (nearby average prey velocity) the prey evolve mostly mutually avoidant behaviour with high drift. On removing the ’too close’ sensor, the prey evolve mutual weak avoidance and strong alignment. This behaviour appears much like flocking, but the prey tend to disperse over broad areas. If only the alignment sensor is present, misalignment ensues.

If the prey’s obstacle detection sensors are removed, they evolve clumping - strong misalignment and cohesion. If the prey’s predator detection sensors are removed, remarkably, (after a lengthy delay in which they adopt a dispersion strategy) the prey still develop flocking. The minimal sensors needed for realistic flocking to emerge may thus be reduced from five to three - those giving average neighbour velocity, average neighbor position, and average obstacle position.

The predators will stimulate flocking in the prey given only one sensor - either the sensor for target prey position or the pursuit sensor is sufficient.

4.2.4 Evolution strategies

As discussed in section 4.1, the predator and prey populations evolve using different methods. The prey use probabilities based on lifetimes to breed the next generation; the predators pick a single best candidate to breed all successors. This latter method was implemented partly due to predator populations being small or individual in most runs. Notably, in no runs did stable flocking evolve in predators.

Changing the mutation variance has little effect on the evolution. This can be explained by the fact that the maximum force and maximum speed constrictions tend to make the ratio of the weights more important than their values. Increasing the mutation range merely increases the scale of all the vectors, which are then summed and clipped in simulation.

The above runs were completed with starting weights randomised on a gaussian distribution around zero - at a higher variance than attained in the ensuing runs. Increasing this starting variance to much higher values caused flocking to fail to evolve. The prey in this case are taken over by strongly-avoiding, misaligning behaviours.

4.2.5 Noise

As noted in section 3.5, the simulation allows for noise to be added to the inputs and outputs of the prey control systems. Whilst this does not necessarily capture the expected behaviour of inaccurate or low-grade sensors (which would not produce such consistent output), this does give a perspective on the stability of the observed flocking encountered above when faced with less-than-ideal physiological abilities.

Adding broadband noise (SNR ≈ 0.5) to the prey’s input sensors results in prey’s control systems being severely disrupted. Their movement becomes palsied and agitated. They still, however, evolve over time to flock as best they can. The weight for cohesion becomes more weighted allowing the prey

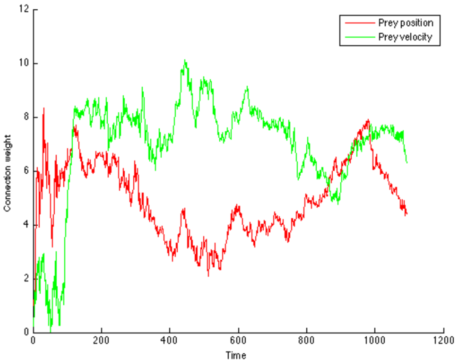

Figure 12: Prey weights - SNR $\approx$ 0.5

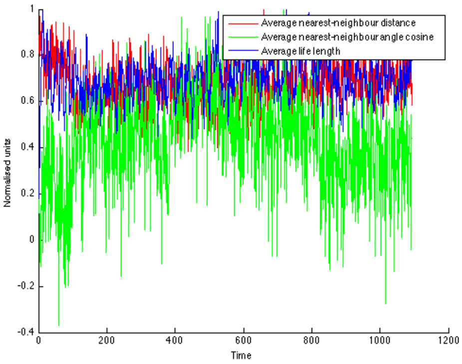

Figure 13: Flocking measures - SNR $\approx$ 0.5

to stay close together despite the fluctuations in their senses (Fig. 12). They do not achieve a very high alignment in practice (Fig. 13).

Reducing the SNR below around 0.3 results in unpredictable behaviour - some runs producing rough alignment, others producing other behaviour. Below this region, the weights continue to converge to vaiours behaviours, but the movement becomes highly erratic and uncoordinated.

Adding noise to the output produces similar behaviour to adding noise to the inputs - except that the flocking is able to persist through higher SNRs. Note the SNR values given here are very approximate, since the signal levels involved exhibit great variation.

4.2.6 Obstacles

Additional obstacles were added within the perimeter boundary after flocking had evolved (Fig. 14). As can be seen in Figs 15 & 16, after a brief period of reduced alignment, the prey develop a tendency towards increased alignment, retaining their flocking abilities in the radically altered environment. The cohesion factor remains unaffected throughout, although the prey experience a brief rise in nearest-neighbour distance after the obstacles are added.

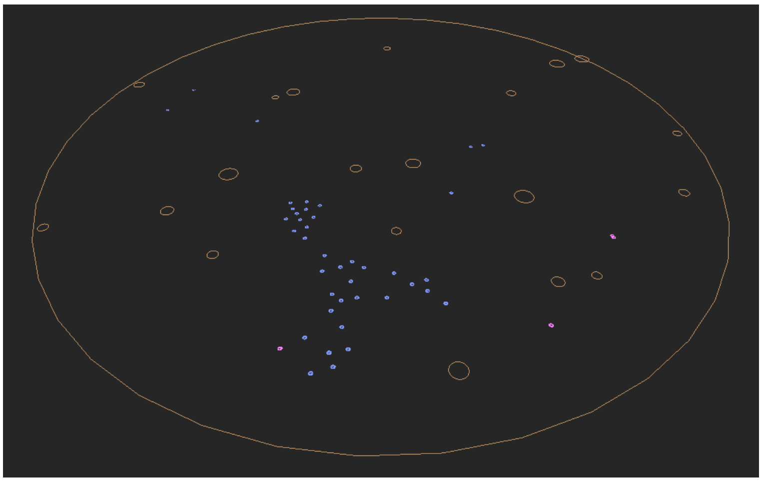

Figure 14: 20 obstacles added.

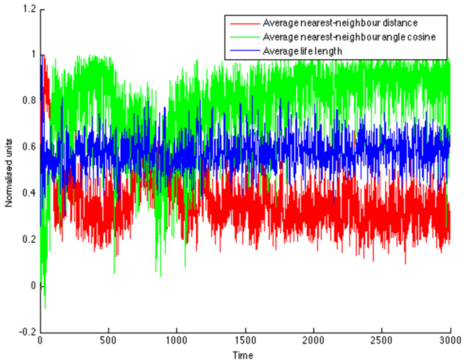

Figure 15: Weights - obstacles added at 500

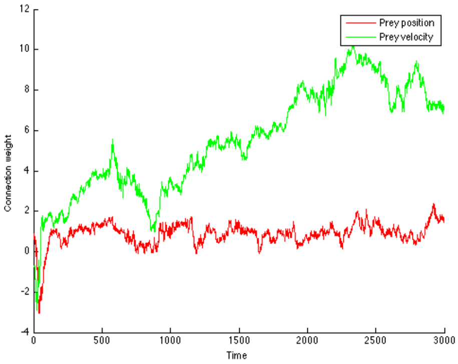

Figure 16: Flocking measures - added obstacles

5 Discussion

This work adds to the chorus of voices, and evidence, pointing to predation as a major cause of the development of flocking behaviours.

The experiments detailed above demonstrate simple coevolution of robust flocking strategies in Braitenberg-like vehicles. There has been much interest in highly complex neuro-evolutionary flocking in recent years. It is important that this exploration of complexity be balanced by a proper understanding of the mimimum requirements for such phenomena to exist. The simulation presented here shows that there is a large amount of order-for-free, to borrow Stuart Kaufmann's phrase, in an emergent form of flocking behaviour so simple it only requires minimally cognitive agents.

One major question left unanswered is the precise sensory requirements for an agent to be able to perceive the position and velocities of its neighbours. It seems likely that possession of these perceptive abilities (at least to some level of fidelity) are necessary in order for flocking to occur. In animals, flocking is achieved through the use of, e.g. vision, echolocation and lateral line senses. For artificial life purposes, we may add other modalities - as well as methods of communicating the information directly from agent to agent. The present model suggests that, if evolving agents under the influence of predation have the ability to detect this information about their neighbours, then flocking is strongly prejudiced to develop.

Such minimal models of dynamic group behaviour give us tools to understand how such behaviours might succeed or fail in other cases. If the behaviour fails to emerge in an alternative scenario, that scenario may be lacking vital aspects present in our minimal model, or conversely it may contain additions which inhibit the behaviour we seek. Given the robust development of flocking in such a simple system, we should only expect the behaviour to be enhanced by the addition of complexity in modelling biological flocking systems or synthesising artificial ones.

References

[1] Pulliam H (1973) *On the advantages of flocking* Journal of Theoretical Biology

[2] Beauchamp G (2004) *Reduced flocking by birds on islands with relaxed predation* Proceedings of the Royal Society B: Biological Sciences

[3] Cresswell W (1994) *Flocking is an effective anti-predation strategy in redshanks, Tringa totanus* Animal Behaviour

[4] Hart P, Freed L (2005) *Predator avoidance as a function of flocking in the sexually dichromatic Hawaii akepa* J. Ethol.Res

[5] Tien J, Levin S, Rubinstein D (2004) *Dynamics of fish shoals: identifying key decision rules* Evol. Ecol. Res.

[6] Ballerini M, Cabibbo N, Candelier R, Cavagna A, Cisbani E, Giardina I, Orlandi A, Parisi G, Procaccini A, Viale M, Zdravkovic V (2008) *An empirical study of large, naturally occurring starling flocks: a benchmark in collective animal behaviour* (To be published in Animal Behaviour)

[7] Holland O (1997) *Grey Walter: the pioneer of real artificial life* Artificial Life

[8] Braitenberg V (1984) *Vehicles: Experiments in synthetic psychology*

[9] Reynolds C (1999) *Steering behaviors for autonomous characters* Game Developers Conference

[10] Sayers D (2009) *Co-evolved Flocking Under Predation* http://iotic.com/critters/

[11] Reynolds C (1987) *Flocks, Herds, and Schools: A Distributed Behavioral Model* Computer Graphics

[12] Aoki I (1982) *A simulation study on the schooling mechanism in fish* Bull. Jpn Soc. Sci. Fish

[13] Toner J, Tu Y (1998) *Flocks, herds, and schools: A quantitative theory of flocking* Physical Review E

[14] Olfati-Saber R (2006) *Flocking for Multi-Agent Dynamic Systems: Algorithms and Theory* IEEE Transactions on Automatic Control

[15] Cliff D, Miller G (1994) *Co-evolution of pursuit and evasion I: Biological and game-theoretic foundations* Technical Report CSRP311, School of Cognitive and Computing Sussex, Brighton, UK

[16] Cliff D, Miller G (1996) *Co-evolution of pursuit and evasion II: Simulation methods and results* Technical Report CSRP377, School of Cognitive and Computing Sussex, Brighton, UK

[17] Reynolds C (1996) *Competition, Coevolution and the Game of Tag* Artificial Life, 1994

[18] Zaera N, Cliff D, Bruten J (1996) *(Not) evolving collective behaviours in synthetic fish*, HP Laboratories technical report HPL-96-04

[19] D Floreano, S Nol (1997) *God save the red queen! Competition in co-evolutionary robotics* Genetic Programming

[20] D Floreano, S Nolfi (1997) *Adaptive Behavior in Competing Co-Evolving Species*, Fourth European Conference on Artificial Life

[21] Ward C, Gobet F, Kendall G (2002) *Evolving Collective Behavior in an Articial Ecology* Artificial Life

[22] Baldassarre G, Nolfi S, Parisi D (2002) *Evolving Mobile Robots Able to Display Collective Behaviours* Artificial Life

[23] Spector L, Klein J, Perry C, Feinstein M (2005) *Emergence of Collective Behavior in Evolving Populations of Flying Agents* Genetic Programming and Evolvable Machines

[24] Kwasnicka H, Markowska-Kaczmar U, Mikosik M (2007) *Open-ended Evolution in Flocking Behaviour Simulation* IMCSIT

[25] Wood A, Ackland G (2007) *Evolving the selfish herd: emergence of distinct aggregating strategies in an individual-based model* Proc Biol Sci.

[26] Nitschke G, Schut M, Eiben A (2007) *Emergent specialization in biologically inspired collective behavior systems* Intelligent Complex Adaptive Systems

[27] Ben-Jacob E, Cohen I, Gutnick D (1998) *Cooperative organization of bacterial colonies: From genotype to morphotype* Annual Reviews in Microbiology

[28] Schauder S, Bassler B (2001) *The languages of bacteria* Genes & Development

[29] Vicsek T, Czirk A, Ben-Jacob E, Cohen I, Shochet O (1995) *Novel Type of Phase Transition in a System of Self-Driven Particles* Physical Review Letters

[30] Liu Y, Passino K (2000) *Swarm intelligence - Literature overview* Dept. of Electrical Engineering, The Ohio State University

[31] OpenSteer (2004) *OpenSteer - Steering Behaviors for Autonomous Characters* http://opensteer.sourceforge.net/