

ELUCIDATING THE EVOLUTIONARY ORIGINS OF COLLECTIVE ANIMAL
BEHAVIOR

By

Randal S. Olson

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Computer Science - Doctor of Philosophy

2015

ABSTRACT

ELUCIDATING THE EVOLUTIONARY ORIGINS OF COLLECTIVE ANIMAL BEHAVIOR

By

Randal S. Olson

Despite over a century of research, the evolutionary origins of collective animal behavior remain unclear. Dozens of hypotheses explaining the evolution of collective behavior have risen and fallen in the past century, but until recently it has been difficult to perform controlled behavioral evolution experiments to isolate these various hypotheses and test their individual effects. In this dissertation, I outline a relatively new method using digital models of evolution to perform controlled behavioral evolution experiments. In particular, I use these models to directly explore the evolutionary consequence of the selfish herd, predator confusion, and the many eyes hypotheses, and demonstrate how the models can lend key insights useful to behavioral biologists, computer scientists, and robotics researchers. This dissertation lays the groundwork for the experimental study of the hypotheses surrounding the evolution of collective animal behavior, and establishes a path for future experiments to explore and disentangle how the various hypothesized benefits of collective behavior interact over evolutionary time.

TABLE OF CONTENTS

LIST OF TABLES	v
LIST OF FIGURES	vi
Chapter 1 Introduction	1
Chapter 2 Background and Related Work	4
2.1 Hypotheses explaining the evolution of collective animal behavior	4
2.1.1 Selfish herd	5
2.1.2 Predator confusion	7
2.1.3 Many eyes	8
2.2 Digital models of evolution for collective animal behavior	9
2.3 Markov Networks	10
2.3.1 How Markov Networks Function	11
2.3.2 Genetic Encoding of Markov Networks	14
2.3.3 Visualization of Markov Networks	16
2.4 Particle swarm optimization and swarm robotics	17
Chapter 3 Selfish Herd Hypothesis	19
3.1 Model of predator-prey interactions	19
3.2 Artificial Predation	23
3.2.1 Random Attacks	25
3.2.2 Random Walk Attacks	26
3.2.3 Outside Attacks	28
3.2.4 Density-Dependent Predation	30
3.2.5 High-Density Area Attacks	31
3.3 Predator-Prey Coevolution	34
3.4 Evolved Prey Markov Network Analysis	37
3.5 Discussion	40
Chapter 4 Predator Confusion Hypothesis	42
4.1 Model of predator-prey interactions	42
4.1.1 Predator and prey agents	44
4.1.2 Simulation environment	48
4.2 Effects of predator confusion	49
4.3 Evolved predator and prey behavior	54
4.4 Eco-evolutionary dynamics	55
4.5 Coevolution between predator visual systems and prey behavior	57
4.6 Discussion	62

Chapter 5	Many Eyes Hypothesis	66
5.1	Model of predator-prey interactions	67
5.1.1	Simulation of predators and prey	68
5.1.2	Evolutionary process	69
5.1.3	Group size	70
5.1.4	Genetic relatedness	70
5.1.5	Reproductive strategy	72
5.1.6	Explicit cost of grouping	73
5.2	Forced grouping	73
5.3	Optional grouping	77
5.4	Tragedy of the commons in heterogeneous groups	78
5.5	Explicit cost of grouping	80
5.6	Discussion	81
Chapter 6	Conclusion	85
References		89

LIST OF TABLES

Table 2.1	An example MG that could be used to control a prey agent which avoids nearby predator agents. “PL” and “PR” correspond to the predator sensors just to the left and right of the agent’s heading, respectively, as shown in Figure 3.2. The columns labeled $P(X)$ indicate the probability of the MG deciding on action X given the corresponding input pair. MF = Move Forward; TR = Turn Right; TL = Turn Left; SS = Stay Still.	12
Table 2.2	Typical mutation rates for experiments evolving Markov Networks. .	15
Table 3.1	Possible actions encoded by the agent’s output. Each output pair encodes a discrete action taken by the agent. The agent’s MN changes the values stored in output states L and R to indicate the action it has decided to take in the next simulation time step.	22
Table 3.2	Genetic algorithm and experiment settings.	23
Table 3.3	High-density area attack (HDAA) experiment treatments. The values listed for each treatment are the handling times for the corresponding predator attack mode.	31

LIST OF FIGURES

Figure 2.1	Example “domains of danger” (DODs) from Hamilton’s selfish herd hypothesis. Each triangle represents a prey in the group, and the area around each triangle is its DOD. Prey on the inside of the group have smaller DODs, which means/they are less likely to be targeted when a predator attacks. As a consequence, “selfish” prey that move inside the group to minimize their DOD will have an evolutionary advantage.	6
Figure 2.2	An example Markov Network (MN) with four input states (white circles labeled 0-3), two hidden states (light grey circles labeled 4 and 5), two output states (dark grey circles labeled 6 and 7), and two Markov Gates (MGs, white squares labeled “Gate 1” and “Gate 2”). The MN receives input into the input states at time step t , then performs a computation with its MGs upon activation. Together, these MGs use information about the environment, information from memory, and information about the MN’s previous action to decide where to move next.	13
Figure 2.3	A zoomed in view of the Markov Gate (MG) labeled “Gate 1” in Figure 2.2. Gate 1 has three binary inputs and two binary outputs, and is thus composed of a $2^3 \times 2^2$ probabilistic state transition table which encodes its logic. For example, p_{52} in the probabilistic state transition table is the probability of the input set 101 (state 0 is 1, state 2 is 0, state 4 is 1) mapping to the output set 10 (state 6 is 1, state 4 is 0). The probabilities across each row must sum to 1.0. . .	14
Figure 2.4	Example circular byte strings encoding the two Markov Gates (MGs) in Figure 2.2, denoted Gene 1 and Gene 2. The sequence (42, 213) represents the beginning of a new MG (white blocks). The next two bytes encode the number of input and output states used by the MG (light grey blocks), and the following eight bytes encode which states are used as input (medium grey blocks) and output (darker grey blocks). The remaining bytes in the string encode the probabilities of the MG’s logic table (darkest grey blocks).	15
Figure 2.5	A causal graph of the node connections for the Markov Network (MN) in Figure 2.2. The only states displayed are states that provide input to or receive output from the Markov Gates of the MN. Arrows between the nodes indicate the flow of binary information between the states.	17

Figure 3.1	A depiction of the simulation environment in which the agents interact. Black dots are prey agents, the black triangle is a predator agent, and the lines around the predator agent indicate its field of view. Agents wrap around the edges of the toroidal simulation environment.	20
Figure 3.2	An illustration of the agents in the model. Light grey triangles are prey agents and the dark grey triangles are predator agents. The agents have a 360° limited-distance visual system (200 virtual meters) to observe their surroundings and detect the presence of other agents. The current heading of the agent is indicated by a bold arrow. Each agent has its own Markov Network, which decides where to move next based off of a combination of sensory input and memory. The left and right actuators (labeled “L” and “R”) enable the agents to move forward, left, and right in discrete steps.	21
Figure 3.3	An illustration of the four artificial predator attack modes. A) Random attacks, B) Random walk attacks, C) Outside attacks, and D) High-density area attacks.	25
Figure 3.4	Mean swarm density over all replicates over evolutionary time, measured by the mean number of prey within 30 virtual meters of each other over a lifespan of 1,000 simulation time steps. Prey in groups attacked randomly (light grey triangles) took much longer to evolve cohesive swarming behavior than prey in groups attacked by a predator that follows a random walk (dark grey circles) or always from the outside of the group (black squares). When prey experience no attacks, they do not evolve swarming behavior at all (light grey stars). Error bars indicate two standard errors over 100 replicates.	27
Figure 3.5	Mean swarm density over all replicates over evolutionary time, measured by the mean number of prey within 30 virtual meters of each other over a lifespan of 1,000 simulation time steps. Even when experiencing density-dependent predation, prey in groups attacked randomly (light grey triangles) took much longer to evolve swarming behavior than prey in groups attacked by a persistent artificial predator (dark grey circles) or always from the outside of the group (black squares). Error bars indicate two standard errors over 100 replicates.	29

Figure 3.6	Mean swarm density over all replicates over evolutionary time, measured by the mean number of prey within 30 virtual meters of each other over a lifespan of 1,000 simulation time steps. Swarm density was measured from evolved populations that were not experiencing predation during measurement, eliminating any possible effects of attack modes that kill more prey faster. Prey in groups attacked only by outside attacks (light grey triangles) evolved cohesive swarming behavior. Increasing the relative frequency of high-density area attacks from infrequent (dark grey circles) to frequent (black squares) caused the prey to evolve increasingly dispersive behavior. Error bars indicate two standard errors over 100 replicates.	33
Figure 3.7	Mean swarm density over all replicates over evolutionary time, measured by the mean number of prey within 30 virtual meters of each other over a lifespan of 2,000 simulation time steps. Prey in groups experiencing density-dependent predation (black circles) evolved cohesive swarming behavior, whereas prey in groups not experiencing density-dependent predation (light grey triangles) evolved dispersive behavior. Error bars indicate two standard errors over 100 replicates.	35
Figure 3.8	Number of sensory input connections from 100 evolved prey Markov Networks mapped onto a prey agent. Only causal connections from the sensory inputs to the actuators are shown. The arrow indicates the facing of the agent. The prey Markov Networks evolved a strong preference for connecting to prey sensors in front and a slight preference for sensors behind the prey agent, but tended to not connect to the sensors on the sides.	38
Figure 4.1	An illustration of the predator and prey agents in the model. Light grey triangles are prey agents and the dark grey triangle is a predator agent. The predator and prey agents have a 180° limited-distance visual system (100 virtual meters for the prey agents; 200 virtual meters for the predator agent) to observe their surroundings and detect the presence of the predator and prey agents. Each agent has its own Markov Network, which decides where to move next based off of a combination of sensory input and memory. The left and right actuators (labeled “L” and “R”) enable the agents to move forward, left, and right in discrete steps.	45

Figure 4.2	Relation of predator attack efficiency (# successful attacks / total # attacks) to number of prey. The solid line with triangles indicates simulated predator attack efficiency as a function of the number of prey within the visual field of the predator (A_{NV}). Similarly, the dashed line with error bars shows the measured simulated predator attack efficiency given the predator attacks a group of swarming prey of a given size, using the A_{NV} curve to determine the per-attack predator attack success rate. Error bars indicate two standard errors over 100 replicate experiments.	47
Figure 4.3	Screen captures of (A) dispersed prey in a swarm hunted by a predator without predator confusion, (B) prey forming a single elongated swarm under attack by a predator with predator confusion, and (C) prey forming multiple cohesive swarms to defend themselves from a predator with predator confusion after 1,200 generations of evolution. Black dots are prey, the triangle is the predator, the lines projecting from the predator represent the predator's frontal 180° visual field, and the star denotes where a prey was just captured.	50
Figure 4.4	Mean swarm density (A), swarm dispersion (B), and survivorship (C) within the swarm over all replicates over evolutionary time. The swarm density was measured by the mean number of prey within 30 virtual meters of each other over a lifespan of 2,000 simulation time steps. Swarm dispersion was measured by the mean distance to the nearest prey for every living prey over a lifespan of 2,000 simulation time steps. Survivorship within the swarm was measured as the mean number of surviving prey (out of an initial total of 50) at the end of the simulation at a given generation. Prey hunted by a predator with predator confusion (black circles with a full line) evolved to maintain significantly higher swarm density and significantly less dispersed swarming behavior than prey in the swarms hunted by a predator without predator confusion (grey triangles with a dashed line). As a result, significantly more prey survived in the swarms hunted by a predator with predator confusion than the swarms hunted by a predator without predator confusion. Error bars indicate two standard errors across 180 replicate experiments.	53
Figure 4.5	Functional response curves of cohesive swarms hunted by a predator with predator confusion (black circles with a full line) and dispersed swarms hunted by a predator without predator confusion (grey triangles with a dashed line). The evolved, cohesive swarms hunted by a predator with predator confusion result in a Type II functional response with a lowered plateau. Error bars indicate two standard errors across 180 replicate experiments.	56

Figure 4.6	Mean swarm density at generation 1,200 as a function of predator view angle. Swarming to confuse the predator was an ineffective behavior if the predator's visual field covered only the frontal 60° or less, due to the predator's focused visual system. As the predator's visual field was incrementally increased to cover the frontal 90° and beyond, predator confusion via swarming again became an effective anti-predator behavior, as evidenced by the swarms exhibiting significantly higher swarm density at generation 1,200. Error bars indicate two standard errors across 180 replicate experiments.	58
Figure 4.7	Swarm density and predator view angle from the LOD of a single coevolution experiment. The predator and prey populations appear to continually cycle between different states of view angles and behaviors.	59
Figure 4.8	Pearson's r between swarm density and predator view angle from the LODs of 30 coevolution experiments. All coevolution experiments have a negative correlation between swarm density and predator view angle, indicating that when swarm density goes up, predator view angle goes down and vice versa. $P \leq 0.001$ for all correlations. . .	60
Figure 4.9	Number of simulation time steps that prey are present anywhere in an evolved predator's visual system depending on the predator's view angle. The predator is competed against dispersive prey. Predators with higher view angles are more likely to have prey anywhere in their visual system at a given time. $P \leq 0.001$ between all view angles, Kruskal-Wallis multiple comparison.	61
Figure 4.10	Number of simulation time steps that prey are visible in a portion of an evolved predator's visual system that it pays attention to, depending on the predator's view angle. The predator is competed against dispersive prey. Predators with higher view angles are more likely to spot prey at a given time, which increases their foraging efficiency. $P \leq 0.001$ between all view angles except 180 vs. 210, Kruskal-Wallis multiple comparison.	61
Figure 4.11	Fitness of an evolved predator when competed against dispersive prey, depending on the predator's view angle. Predators with higher view angles forage for prey more efficiently, thus capturing more prey in their lifetime and improving their fitness. $P \leq 0.001$ between all view angles except 150 vs. 180 and 180 vs. 210, Kruskal-Wallis multiple comparison.	62
Figure 4.12	Diagram depicting the observed coevolutionary cycle between the predator and prey in the presence of the predator confusion effect. .	64

Figure 5.1	Depiction of the disembodied simulation. Prey seek to forage as much as possible while avoiding being captured by the predator. If none of the prey in the group are vigilant, the target prey is captured 100% of the time.	67
Figure 5.2	Treatment comparison when prey are forced to forage in groups. Both group homogeneity and a semelparous reproductive strategy select for high levels of vigilance. However, only homogeneous groups experience an increase in fitness as group size increases. In contrast, vigilance behavior breaks down in larger, heterogeneous groups of semelparous prey. Error bars indicate bootstrapped 95% confidence intervals over 100 replicates; some error bars are too small to be visible.	74
Figure 5.3	Treatment comparison when prey can choose to forage in groups. Allowing prey to decide whether they wish to be in the group produces similar results compared to when they are forced to group. In homogeneous groups, prey choose to spend most of their time in the group. However, grouping breaks down (alongside vigilance) in heterogeneous groups of semelparous prey. This occurs despite there being no direct penalty assessed for choosing to group. Error bars indicate bootstrapped 95% confidence intervals over 100 replicates; some error bars are too small to be visible.	76
Figure 5.4	Vigilance in prey with and without the option to forage in groups. In homogeneous groups, prey with forced and optional grouping evolve similar vigilance behaviors. In contrast, individualistic (non-grouping) prey evolve vigilance behaviors that maximize individual fitness. Meanwhile, individuals in heterogeneous/semelparous populations with the option to group evolve to be less vigilant than either of the other two treatments. Error bars indicate bootstrapped 95% confidence intervals over 100 replicates; some error bars are too small to be visible.	79
Figure 5.5	Fitness for prey with and without the option to forage in groups. In heterogeneous/semelparous groups, prey with the option to group have lower fitness than prey that are forced to group. Error bars indicate bootstrapped 95% confidence intervals over 100 replicates; some error bars are too small to be visible.	79

Figure 5.6 **Grouping behaviors in prey experiencing grouping penalties.**
 Even with a small grouping penalty ($M = 1.0$), all treatments except homogeneous/semelparous no longer evolve grouping behavior. Prey in the homogeneous/semelparous treatment evolve only slightly lower levels of grouping behavior, even with extreme penalties to foraging in a group ($M = 1,000$). Error bars indicate bootstrapped 95% confidence intervals over 100 replicates; some error bars are too small to be visible. 80

Chapter 1

Introduction

Over the past century, researchers have devoted considerable effort into studying collective animal behavior due to its important implications for social intelligence, collective cognition, and potential applications in artificial intelligence, swarm robotics, and distributed systems [1]. Indeed, collective behaviors are pervasive across all forms of life. For example, European starlings (*Sturnus vulgaris*) are known to form murmurations of millions of birds while performing awe-inspiring displays of coordinated movement [2, 3]. Western honeybees (*Apis mellifera*) communicate the location of food and nest sites to other bees in their group via a complex dance language [4]. Even relatively simple bacteria exhibit grouping behavior, such as *Escherichia coli* forming biofilms, which allow their group to survive in hostile environments [5].

Despite the abundance of examples of collective behavior in nature, the *why* of collective behavior remains difficult to ascertain to this day [6]. What are the advantages of working and living together in a group? More importantly, which of these advantages are the underlying reason collective behavior evolved in the multitude of group-living species we observe today? The objective of this dissertation is to identify the individual hypotheses explaining how and why collective animal behavior has evolved, then directly explore these hypotheses in isolation to determine how, when, and *if* each of the hypothesized benefits actually selects for the evolution of collective behavior.

Unfortunately, performing long-term evolution experiments is especially challenging in

most biological systems because many species that exhibit collective behavior take months or even years to produce offspring. These long generation times make it extremely difficult to experimentally determine which of the aforementioned benefits are sufficient to select for collective behavior as an evolutionary response, let alone study the behaviors as they evolve [7, 8]. To complicate matters even further, any studies attempting to isolate individual selection pressures in a biological system are often muddled by the multitude of selection pressures present in the system, making it challenging to experimentally establish a direct cause-and-effect relationship between the selection pressure and the evolved behavior.

To overcome these challenges, in this dissertation I use digital models of evolution to explore the evolutionary effect of the various hypothesized benefits of collective behavior. In these models, I simulate individual prey and the interactions between them without imposing any particular movement rules on them. Prey that survive longer and forage more efficiently produce more offspring into the next generation, allowing for the evolution of behavior in response to the environment. To observe the effect of the particular hypothesis being explored, I introduce an isolated selection pressure to the prey population (e.g., predation) and measure the change in prey behavior over evolutionary time. Digital models of evolution thus provide a decided advantage for exploring the evolution of animal behavior in response to isolated selection pressures, which I will demonstrate throughout this dissertation.

Thesis Statement *Evolutionary computation and multi-agent systems can be fruitfully combined to elucidate the evolutionary origins of collective animal behavior, which produces testable hypotheses about the evolution of collective behavior.*

The remainder of this dissertation proceeds as follows. In Section 2, I outline the several hypotheses explaining the evolution of collective animal behavior, describe the history of the digital models of evolution that I use to explore these hypotheses, and discuss related

applications of collective behavior in computer science and engineering. In Sections 3–5, I provide details on several research projects I completed during the course of my studies exploring the selfish herd, predator confusion, and many eyes hypotheses. Finally, I provide concluding remarks about using digital models of evolution to study the evolution of collective animal behavior in Section 6.

Chapter 2

Background and Related Work

In this chapter, I review the literature surrounding the evolutionary origins of collective animal behavior. I begin this chapter by outlining the various hypotheses explaining why collective animal behavior evolves. Next, I describe the history of the digital models of evolution that I use to explore the evolution of collective animal behavior. Finally, I discuss related work and potential applications of this thesis research in the fields of particle swarm optimization and swarm robotics.

2.1 Hypotheses explaining the evolution of collective animal behavior

As with many traits, collective behavior entails a variety of fitness costs, such as increased predation rates [6], requisite sharing of resources within the group [9], and heightened competition for mates [10]. With this fact in mind, significant effort has been dedicated to understanding the compensating benefits that collective behavior provides [6]. Many such benefits of collective behavior have been proposed, for example, it may improve mating success [11, 12], increase foraging efficiency [13, 14, 15, 16, 17], improve locomotion efficiency [18], or enable the group to solve problems that would be impossible to solve individually [1].

Furthermore, a major category of hypotheses propose that collective behaviors protect

group members from predators. For example, collective behavior can improve group vigilance [19, 20, 21, 22], reduce the chance of being encountered by predators [21, 23], dilute an individual’s risk of being attacked [24, 25, 26, 27], enable an active defense against predators [28], or reduce predator attack efficiency by confusing the predator [7, 29, 30].

With the multitude of potential costs and benefits of collective behavior, the need for an experimental platform to explore each hypothesis in isolation becomes abundantly clear. In the following three sections, I briefly review three of the hypothesized benefits of collective behavior that I explored in detail in this dissertation.

2.1.1 Selfish herd

Briefly, the selfish herd hypothesis states that prey in groups under attack from a predator will seek to place other prey in between themselves and the predator, thus maximizing their chance of survival. As a consequence of this selfish behavior, individuals continually move toward a central point in the group, which gives rise to the appearance of a cohesive swarm.

Hamilton’s original formulation of the selfish herd hypothesis introduced the concept of “domains of danger” (DODs, Figure 2.1), which served as a method to visualize the likelihood of a prey inside a group to be attacked by a predator [27]. Prey on the edges of the group would have larger DODs than prey on the inside of the group; thus, prey on the edges of the group would be attacked more frequently. Moreover, Hamilton proposed that prey on the edges of the group would seek to reduce their DOD by moving inside the group, thus placing other group members between themselves and the predator. Further work has expanded on this hypothesis by adding a limited predator attack range [31], investigating the effects of prey vigilance [32], considering the initial spatial positioning of prey when the group is attacked [33], exploring the role of prey body characteristics in shaping herd char-




Figure 2.1 Example “domains of danger” (DODs) from Hamilton’s selfish herd hypothesis. Each triangle represents a prey in the group, and the area around each triangle is its DOD. Prey on the inside of the group have smaller DODs, which means/they are less likely to be targeted when a predator attacks. As a consequence, “selfish” prey that move inside the group to minimize their DOD will have an evolutionary advantage.

acteristics [34, 35], and even corroborating Hamilton’s predictions in biological systems [36].

Additional studies have focused on the movement rules that prey in a selfish herd follow to minimize their DOD [37]. This line of work began by demonstrating that the simple movement rules proposed by Hamilton reduce predation risk for prey inside the group [38], then opened some parameters of the movement rules to evolution in an attempt to discover a more biologically plausible set of movement rules [39, 40]. Importantly, these studies demonstrated that it is possible for selfish herd behavior to evolve by natural selection on movement rules that rely on only local information for each agent, rather than global information about the entire group.

However, it still remains an open question of how the actions of the predator—for example, whether the predator can coevolve and adapt to the selfish herd behavior—can impact

the selfish herd hypothesis. This dissertation builds on this previous work by studying the effects of coevolving predators and predator attack mode (i.e., how predators select a prey in a group to attack) on the evolution of the selfish herd.

2.1.2 Predator confusion

In the predator confusion hypothesis, the presence of multiple individuals moving in a swarm confuses approaching predators, making it difficult for predators to successfully execute an attack [29, 7, 30, 41]. This confusion effect is hypothesized to arise from the increased cognitive processing time needed to decide on a target among multiple prey. In a recent review of predator-prey systems with swarming prey, Jeschke and Tollrian noted that predators appeared to become confused by swarming behavior in 16 of the 25 systems reviewed [7]. However, evidence that predator confusion is a seemingly widespread phenomenon still leaves open the question of how effective predator confusion could be as a selective force favoring the evolution of swarming behavior. Owing to the difficulty of disentangling the individual effects of hypotheses such as the predator confusion hypothesis in natural systems, there is a need to explore the predator confusion hypothesis in isolation in a digital model.

Predator confusion is broadly interesting for two additional reasons. First, it provides an opportunity to study how swarming behavior can in turn exert evolutionary pressures on predators, especially on the perceptual constraints that allow for predator confusion in the first place. For example, once swarming behavior evolves in prey, predator confusion may in turn provide a selective advantage for predators that are no longer confused by swarms. Second, predator confusion may influence the *functional response* describing the predator's consumption rate as prey density increases [42], as suggested in a previous study [43]. Understanding how pervasive mechanisms such as predator confusion affect functional response

relationships is critical for accurately modeling the dynamics of predator-prey interactions over ecological and evolutionary time [44].

This dissertation builds on previous work on the predator confusion hypothesis by exploring all three of the above possibilities in a digital model, where it is possible to experimentally control the effect of predator confusion.

2.1.3 Many eyes

Finally, the many eyes hypothesis concerns the trade-off between vigilance (e.g., watching for an incoming predator) and foraging for food in groups of prey. First proposed using a mathematical model [22] and explored experimentally a year later [15], the many eyes hypothesis makes two key predictions, both of which arise from the assumption that vigilance is costly because it imposes a trade-off with foraging efficiency: (a) individual prey vigilance will decline as a group size increases, and (b) because prey can more equitably divide the task of watching for predators in large groups, they will experience a fitness benefit from foraging more. Therefore, there will be a selective advantage for prey that forage in groups up to a certain group size. In the 40 years since its inception, these predictions have been examined in numerous species across hundreds of independent studies [45, 46, 47, 17, 48]. Furthermore, several game theoretical models have been applied to refine the predictions of when collective vigilance in foraging groups should evolve [49], and subsequently matched to experimental data [50].

However, it still remains an open question of whether—and under what conditions—the many eyes hypothesis provides a sufficient selective pressure to favor the evolution of collective behavior in foraging species. This dissertation builds on previous work on the many eyes hypothesis by exploring several conditions under which selection favors gregarious

foraging behavior.

2.2 Digital models of evolution for collective animal behavior

It is by no means a new idea to use digital models to study animal behavior. Digital models have previously been used to provide key insights into core evolutionary processes [51, 52], and several well-known studies have adopted digital models as a method to study collective behavior [53, 54, 55]. More recently, digital models have even been used to elucidate the emergence of prey collective behavior as a response to predation [24].

These previous studies have provided insight into the fundamental dynamics of collective behavior. However, most have not focused on isolating the evolutionary pressures that might favor the formation of groups, and few have explored the coevolution of predator and prey behavior. In fact, except for only a handful of studies, the collective behavior literature typically has not studied Darwinian evolution as a process affecting the properties of groups. It is therefore the goal of this dissertation to highlight the studies that have explored the evolution of collective animal behavior and to synthesize their work to build a solid theoretical foundation on which the evolution of collective animal behavior can be studied.

When considering the evolution of collective behavior, it is vital to take into account both the benefits *and* costs imposed by the behavior [56]. To satisfy this requirement, several researchers have recently turned to digital models of evolution to study the evolution of animal behavior [32, 13, 57]. These researchers use a digital model of evolution to evolve the behavior of a population of locally-interacting animats, enabling them to explore the evolution of behavior in complex environments that are beyond the means of mathematical

models [58, 59].

With these models, several studies have explored the evolution of selfish herd behavior in response to predation [40, 39]. Other studies have investigated the evolution of predator behavior in response to prey density [60], the evolution of prey behavior in the presence of the predator confusion effect [61, 62], the role of relative predator and prey speeds on the evolution of grouping behavior [63], and have elaborated upon the interaction between ecology and the evolution of grouping behavior [64, 65].

In this dissertation, I build upon the ideas from these studies and establish a general framework for studying the evolution of collective behavior. In each chapter of this dissertation, I construct a digital model of evolution to focus on a single hypothesized benefit of collective behavior, and control for the possible benefits introduced by other mechanisms. To perform these digital evolution experiments, I evolve prey agents with a *genetic algorithm* (GA), which is a digital model of evolution by natural selection [66]. In a GA, pools of genomes are evolved over time by evaluating the fitness of each genome at each generation and preferentially selecting those with higher fitness (e.g., from consuming more food or surviving longer) to populate the next generation. The genomes here are variable-length strings of integers that are translated into Markov Networks (MNs) during fitness evaluation. More information on MNs—including details on their genetic encoding, mutational operators, and functionality—is available in the following section.

2.3 Markov Networks

In every digital model in this dissertation, each agent is controlled by its own Markov Network (MN), which is a probabilistic controller that makes decisions about how the agent interacts

with the environment and other agents within that environment. Since a MN is responsible for the control decisions of its agent, it can be thought of as an *artificial brain* for the agent it controls. Every time step in the simulation, the MNs receive input via sensors (e.g., a visual system), perform a computation on inputs and any hidden states (i.e., memory), then place the result of the computation into hidden or output states (e.g., actuators). I note that MN states are binary and only assume a value of 0 or 1. When I evolve MNs with a GA, mutations affect (1) which states the MN pays attention to as input, (2) which states the MN outputs the result of its computation to, and (3) the internal logic that converts the input into the corresponding output.

2.3.1 How Markov Networks Function

When I embed an agent into the simulation environment, I provide sensory inputs from its visual system into its MN every simulation step (labeled “retina” and “Markov Network”, respectively). Once I provide a MN with its inputs, I activate it and allow it to store the result of the computation into its hidden and output states for the next time step. MNs are networks of Markov Gates (MGs), which perform the computation for the MN. In Figure 2.2, we see two example MGs, labeled “Gate 1” and “Gate 2.” At time t , Gate 1 receives sensory input from states 0 and 2 and retrieves state information (i.e., memory) from state 4. At time $t + 1$, Gate 1 then stores its output in hidden state 4 and output state 6. Similarly, at time t Gate 2 receives sensory input from state 2 and retrieves state information in state 6, then places its output into states 6 and 7 at time step $t + 1$. When MGs place their output into the same state, the outputs are combined into a single output using the OR logic function. Thus, the MN uses information from the environment and its memory to decide where to move in the next time step $t + 1$.

Arbitrary encodings can be used, but simpler encodings are more conducive to the evolution of effective behavior. In order for agent to be able to react to the environment, the output states must somehow meaningfully connect to the input states. Additionally, if memory about state information from the previous time step is required for more complex tasks, MNs can store state information in memory by connecting input states to hidden states, then connecting those hidden states to output states. Finally, state information can be stored in memory for longer than one time step by connecting hidden states to yet more hidden states.

In a MN, states are updated by MGs, which function similarly to digital logic gates, e.g., AND & OR. A digital logic gate, such as XOR, reads two binary states as input and outputs a single binary value according to the XOR logic. Similarly, MGs output binary values based on their input, but do so with a probabilistic logic table. Table 2.1 shows an example MG that could be used to control a prey agent that avoids nearby predator agents. For example, if a predator is to the right of the prey’s heading (i.e., $PL = 0$ and $PR = 1$, corresponding to the second row of this table), then the outputs are move forward (MF) with a 20% chance, turn right (TR) with a 5% chance, turn left (TL) with a 65% chance, and stay still (SS) with a 10% chance. Thus, due to this probabilistic input-output mapping, the agent MNs are capable of producing stochastic agent behavior.

Table 2.1 An example MG that could be used to control a prey agent which avoids nearby predator agents. “PL” and “PR” correspond to the predator sensors just to the left and right of the agent’s heading, respectively, as shown in Figure 3.2. The columns labeled $P(X)$ indicate the probability of the MG deciding on action X given the corresponding input pair. MF = Move Forward; TR = Turn Right; TL = Turn Left; SS = Stay Still.

PL	PR	P(MF)	P(TR)	P(TL)	P(SS)
0	0	0.7	0.05	0.05	0.2
0	1	0.2	0.05	0.65	0.1
1	0	0.2	0.65	0.05	0.1
1	1	0.05	0.8	0.1	0.05




Figure 2.2 An example Markov Network (MN) with four input states (white circles labeled 0-3), two hidden states (light grey circles labeled 4 and 5), two output states (dark grey circles labeled 6 and 7), and two Markov Gates (MGs, white squares labeled “Gate 1” and “Gate 2”). The MN receives input into the input states at time step t , then performs a computation with its MGs upon activation. Together, these MGs use information about the environment, information from memory, and information about the MN’s previous action to decide where to move next.

While digital logic gates are deterministic, MGs can be composed of any set of probabilities in their probability table. Therefore, while the output states still depend on the input states, they can also have a degree stochasticity to their output. Figure 2.3 illustrates an example MG with three binary inputs entering the MG: 0 and 2 coming from sensory input states, while input 4 comes from a hidden state. This example MG is composed of a $2^3 \times 2^2$ state transition table (because it has three inputs and two outputs) that encodes the logic for the MG. Once provided with inputs, the MG activates and updates output state 6 and hidden state 4. Because the MG outputs to the same hidden state that it receives input from, it is forming a *recurrent connection*, i.e., memory.

The MGs in this model can receive input from a maximum of four states, and write into a maximum of 4 states, with a minimum of one input and one output state for each MG. Any state (input, output, or hidden) in the MN can be used as an input or output for a MG. MNs can be composed of any number of MGs, and the MGs are what define the internal




Figure 2.3 A zoomed in view of the Markov Gate (MG) labeled “Gate 1” in Figure 2.2. Gate 1 has three binary inputs and two binary outputs, and is thus composed of a $2^3 \times 2^2$ probabilistic state transition table which encodes its logic. For example, p_{52} in the probabilistic state transition table is the probability of the input set 101 (state 0 is 1, state 2 is 0, state 4 is 1) mapping to the output set 10 (state 6 is 1, state 4 is 0). The probabilities across each row must sum to 1.0.

logic of the MN. Thus, to evolve a MN, mutations change the connections between states and MGs, and modify the probabilistic logic tables that describe each MG. Mutations act directly on the genetic encoding of the MN, which is described next.

2.3.2 Genetic Encoding of Markov Networks

We use a circular string of bytes as a genome, which contains all the information necessary to describe a MN. The genome is composed of *genes*, and each gene encodes a single MG. Therefore, a gene contains the information about which states the MG reads input from, which states the MG writes its output to, and the probability table defining the logic of the MG. The start of a gene is indicated by a *start codon*, which is represented by the sequence (42, 213) in the genome.

Figure 2.4 depicts an example genome. After the start codon, the next two bytes describe the number of inputs (N_{in}) and outputs (N_{out}) used in this MG, where each $N = 1 + (\text{byte} \bmod N_{max})$. Here, $N_{max} = 4$. The following N_{max} bytes specify which states the MG reads from by mapping to a state ID number with the equation: $(\text{byte} \bmod N_{states})$,




Figure 2.4 Example circular byte strings encoding the two Markov Gates (MGs) in Figure 2.2, denoted Gene 1 and Gene 2. The sequence (42, 213) represents the beginning of a new MG (white blocks). The next two bytes encode the number of input and output states used by the MG (light grey blocks), and the following eight bytes encode which states are used as input (medium grey blocks) and output (darker grey blocks). The remaining bytes in the string encode the probabilities of the MG’s logic table (darkest grey blocks).

Table 2.2 Typical mutation rates for experiments evolving Markov Networks.

Parameter	Value
Per-gene mutation rate	1%
Gene duplication rate	5%
Gene deletion rate	2%
Crossover	None

where N_{states} is the total number of input, output, and hidden states. Similarly, the next N_{max} bytes encode which states the MG writes to with the same equation as N_{in} . If too many inputs or outputs are specified, the remaining sites in that section of the gene are ignored, designated by the # signs. The remaining $2^{N_{\text{in}}+N_{\text{out}}}$ bytes of the gene define the probabilities in the logic table.

We sequentially fill the logic table row-by-row with bytes from the genome. Once the logic table is filled, I convert the bytes into the corresponding probabilities (p_{ij}) with the following equation:

$$p_{ij} = \frac{1 + \text{byte}_{ij}}{\sum_{j=1}^{2^{N_{\text{out}}}} (1 + \text{byte}_{ij})} \quad (2.1)$$

where byte_{ij} is the byte in the genome corresponding to the probability in the table at row i and column j , and N_{out} is the number of outputs used by the MG. Because I use bytes to

specify the values in the table, I normalize the values for each row in the probability table so the sum of each row is 1.0. I apply the modulo operator on the number of inputs, the number of outputs, and the IDs of the states used as inputs and outputs in order to keep them within the allowed ranges.

The maximum number of states allowed and which states are used as inputs and outputs are specified as constants by the user. Combined with these constants, the genome described above unambiguously defines a MN. All evolutionary changes such as point mutations, duplications, deletions, or crossover are performed on the byte string genome, with probabilities as shown in Table 2.2. During a point mutation, a random byte in the genome is replaced with a new byte drawn from a uniform random distribution. If a duplication event occurs, two random positions are chosen in the genome and all bytes between those points are duplicated into another part of the genome. Similarly, when a deletion event occurs, two random positions are chosen in the genome and all bytes between those points are deleted. Crossover for MNs is not implemented in this experiment.

2.3.3 Visualization of Markov Networks

MNs can be visualized in several ways. Because the visualization of the MN in Figure 2.2 shows many states that are not even used, and the MGs are less important than how states causally depend on each other, I usually only display a graph similar to Figure 2.5 showing the causal relations between the states.




Figure 2.5 A causal graph of the node connections for the Markov Network (MN) in Figure 2.2. The only states displayed are states that provide input to or receive output from the Markov Gates of the MN. Arrows between the nodes indicate the flow of binary information between the states.

2.4 Particle swarm optimization and swarm robotics

In the past decade, researchers have focused on the application of locally-interacting swarming agents to optimization problems, called Particle Swarm Optimization (PSO) [67]. PSO applications range from feature selection for classifiers [68], to video processing [69], to open vehicle routing [70]. A related technique within PSO seeks to combine PSO with coevolving “predator” and “prey” solutions to avoid local minima [71]. Thus, elaborations on the foundations of collective animal behavior has the potential to improve our ability to solve engineering problems.

Further, researchers have sought to harness the collective problem solving power of swarming agents to design robust autonomous robotic swarms [72]. Given that most swarm control algorithms—such as the popular Boids algorithm [73]—require significant computing power and global information about the swarm that is typically unavailable in the real world, it is vital to develop swarm control algorithms that can produce reliable swarming behavior with *individual-based* control mechanisms that require only localized information.

Chapter 3.4, in particular, will focus on how such individual-based control algorithms can be discovered using digital models of evolution.

Chapter 3

Selfish Herd Hypothesis

In this chapter, I use a digital model of predator-prey coevolution to explore Hamilton's selfish herd hypothesis [27]. Briefly, the selfish herd hypothesis states that prey in groups under attack from a predator will seek to place other prey in between themselves and the predator, thus maximizing their chance of survival. As a consequence of this selfish behavior, individuals continually move toward a central point in the group, which gives rise to the appearance of a cohesive swarm. This chapter expands on my earlier work [74] by studying the long-term evolutionary effects of differing attack modes, exploring a new attack mode that directly selects against swarming behavior, and providing an analysis of the control algorithms that evolved in the swarming prey.

This chapter begins with the details of the digital model that I used to explore the evolution of selfish herd behavior in groups of prey. Next, I describe the results from the model and how predator attack mode affects the evolution of selfish herd behavior. Finally, I conclude the chapter by discussing some of the broader implications of the findings in this chapter.

3.1 Model of predator-prey interactions

To study the evolution of the selfish herd, I developed an agent-based model in which agents interact in a continuous, toroidal virtual environment (736×736 virtual meters), shown in




Figure 3.1 A depiction of the simulation environment in which the agents interact. Black dots are prey agents, the black triangle is a predator agent, and the lines around the predator agent indicate its field of view. Agents wrap around the edges of the toroidal simulation environment.

Figure 3.1. At the beginning of each simulation, I place 250 agents in the environment at uniformly random locations. These agents are treated as “virtual prey.” Each agent is controlled by a *Markov Network* (MN), which is a probabilistic controller that makes movement decisions based on a combination of sensory input (i.e., vision) and internal states (i.e., memory). I evolve the agent MNs with a genetic algorithm (GA) [75, 66] under varying selection regimes, which will be described in more detail below. More information on MNs—including details on their genetic encoding, mutational operators, and functionality—is available in Chapter 2.3.

During each simulation time step, all agents read information from their sensors and take action (i.e., move) based on their effectors. In my first set of treatments, I simulate




Figure 3.2 An illustration of the agents in the model. Light grey triangles are prey agents and the dark grey triangles are predator agents. The agents have a 360° limited-distance visual system (200 virtual meters) to observe their surroundings and detect the presence of other agents. The current heading of the agent is indicated by a bold arrow. Each agent has its own Markov Network, which decides where to move next based off of a combination of sensory input and memory. The left and right actuators (labeled “L” and “R”) enable the agents to move forward, left, and right in discrete steps.

an ideal, disembodied predator by periodically removing prey agents from the environment and marking them as consumed, e.g., when they are on the outermost edges of the group. Subsequent treatments introduce an embodied, coevolving predator agent which is controlled by its own MN. The data¹ and source code² from these experiments are available online for further analysis.

Figure 3.2 depicts the sensory-motor architecture of the agents used for this study. A prey agent can sense predators and conspecifics with a limited-distance (200 virtual meters),

¹Data: <http://dx.doi.org/10.6084/m9.figshare.663680>

²Code: <https://github.com/adamilab/eos-selfish-herd>

Table 3.1 Possible actions encoded by the agent’s output. Each output pair encodes a discrete action taken by the agent. The agent’s MN changes the values stored in output states L and R to indicate the action it has decided to take in the next simulation time step.

Output L	Output R	Encoded Action
0	0	Move forward
0	1	Turn right
1	0	Turn left
1	1	Stay still

pixelated visual system covering its entire 360° visual field. Its visual system is split into 24 even slices, each covering an arc of 15° , which is an abstraction of the broad, coarse visual systems often observed in grouping prey [76]. Regardless of the number of agents present in a single retina slice, the prey agent only knows whether a conspecific or predator resides within that slice, but not how many. For example, in Figure 3.2, the fourth retina slice to the right of the agent’s heading (labeled “A”) has both the predator and prey sensors activated because there are two predator agents and a prey agent inside that slice. Once provided with its sensory information, the prey agent chooses one of four discrete actions, as shown in Table 3.1. Prey agents turn in 8° increments and move 1 virtual meter each time step.

In my coevolution experiments, the predator agents can detect only nearby prey agents using a limited-distance (200 virtual meters), pixelated visual system covering its frontal 180° that works just like the prey agent’s visual system (Figure 3.2). Similar to the prey agents, predators make decisions about how to move next using their MN, as shown in Table 3.1, but move $3\times$ faster than the prey agents and turn correspondingly slower (6° per simulation time step) due to their higher speed. Finally, if a predator agent moves within 5 virtual meters of a prey agent that is anywhere within its visual system, the predator agent makes an attack attempt on the prey agent. If the attack attempt is successful, I remove the prey agent from the simulation and mark it as consumed.

Table 3.2 Genetic algorithm and experiment settings.

GA Parameter	Value
Selection	Fitness proportionate
Population size	250
Per-gene mutation rate	1%
Gene duplication rate	5%
Gene deletion rate	2%
Crossover	None
Generations	40,000
Replicates	100

3.2 Artificial Predation

In my first set of experiments, I observe the evolution of prey behavior in response to various forms of artificial predation. This experimental setup enables me to control the specific modes of predation and observe their effect on the evolution of the selfish herd. I evolve the prey genomes with a GA with the settings described in Table 3.2. I begin the evolutionary process by seeding the prey genome pool with a set of randomly-generated ancestor genomes of length 5,000. Following this, I evaluate the relative fitness of each prey genome by translating the genome into its corresponding MN, embodying each MN in a prey agent, and competing the prey agents in a simulation environment for 1,000 simulation time steps. This evaluation period is akin to the agents’ lifespan, hence each agent has a potential lifespan of 1,000 time steps. I assign each prey genome an individual fitness according to how long its corresponding prey agent survived, following the equation:

$$W_{\text{prey}} = T \tag{3.1}$$

where T is the number of time steps the prey agent survived in the simulation environment. Thus, individual prey genomes are rewarded for their agent surviving longer than other agents

in the group. Once all of the prey genomes are assigned fitness values, I perform fitness-proportionate selection on the population of genomes via a Moran process [77], increment the generation counter, and repeat the evaluation process on the new population of genomes until the final generation (40,000) is reached.

In all cases, I give the prey an initial 250 simulation time steps without predation to move around, so that prey starting on the outside of the group have the chance to move toward the center of the group if they wish to. Once the initial 250 simulation time steps elapse, I apply artificial predation every 4 simulation time steps by simulating an ideal predator that attacks the group according to a specific attack type. Artificial predators succeed with their attacks every time. I limit the artificial predator attack rate to one attack attempt every 4 simulation time steps, which is called the *handling time*. The handling time represents the time it takes the simulated predator to consume and digest a prey after successful prey capture, or the time it takes to refocus on another prey in the case of an unsuccessful attack attempt. I selected a handling time of 4 because it reduces the herd of prey down to 25% of its original size by the end of the simulation, therefore applying strong selection pressure for survivorship in the herd.

For each experiment, I characterize the grouping behavior by measuring the *swarm density* of the entire prey population every generation [78]. I measure the swarm density as the mean number of prey within 30 virtual meters of each other over a lifespan of 1,000 simulation time steps, which I have experimentally shown to differentiate between swarming and non-swarming behavior in previous published experiments [79]. Qualitatively, a swarm density of ≥ 15 indicates cohesive swarming behavior, between 15 and 5 loosely grouping behavior, and ≤ 5 random, non-grouping behavior. Thus, swarm density captures how cohesively the prey are swarming, or if the prey are even grouping at all.




Figure 3.3 An illustration of the four artificial predator attack modes. A) Random attacks, B) Random walk attacks, C) Outside attacks, and D) High-density area attacks.

In the following sections, I study the effect of four different attack modes on the evolution of swarming behavior: uncorrelated random attacks (Figure 3.3A), correlated random attacks (random walk attacks, Figure 3.3B), peripheral attacks (Figure 3.3C), and attacks that target the most dense area of the swarm (Figure 3.3D).

3.2.1 Random Attacks

My initial study sought to verify Hamilton’s selfish herd hypothesis by modeling evolving prey under attack by predators that ambush prey from a random location in the simulation environment. If the selfish herd hypothesis holds, I expect prey to minimize their “domain of danger” to the predators by placing as many conspecifics as possible around them [27]. Similar to previous models studying the selfish herd [40], a random attack proceeds by selecting a uniformly random location inside the simulation space, then attacking the prey closest to that location, as shown in Figure 3.3A.

As seen in Figure 3.4, swarming behavior is weakly selected for when the predators make uniformly random attacks on the prey³ (light grey triangles). Particularly, I found that prey took upwards of 5,000 generations to evolve cohesive swarming behavior when experiencing

³Evolution of prey behavior under Random Attack treatment: <http://dx.doi.org/10.6084/m9.figshare.658857>

random attacks, compared to fewer than 1,000 generations with the other attack modes. However, even random attacks selected for more cohesive swarming behavior than no attacks at all, which resulted in completely dispersive behavior (Figure 3.4, light grey stars).

This finding has important implications, namely that one of the original assumptions of the selfish herd hypothesis—that the predator attack mode has no important impact on the evolution of swarming behavior—is not corroborated by this model. Following this discovery, I hypothesized that the *directionality* of the predators’ attacks play a critical role in the evolution of the selfish herd. To test this hypothesis, I next explore two different predator attack modes, each with their own distinct directionality of predation.

3.2.2 Random Walk Attacks

My next experiment alters the mode of predation from a predator that attacks randomly selected locations to a predator that follows a random walk within the simulation environment. Shown in Figure 3.3B, after each attack made by this predator, it is then moved to a random location within 50 virtual meters of its previous location. This models a predator that persistently feeds on a group of prey, rather than ambushing.

Figure 3.4 shows that swarming evolved quickly when the prey were attacked by a predator following a random walk⁴ (dark grey circles). Notably, even by generation 40,000, prey experiencing random walk attacks formed significantly more cohesive swarms than prey experiencing random attacks. Thus, the random walk predator attack mode appears to capture an important aspect of predation that selects for swarming behavior.

⁴Evolution of prey behavior under Random Walk treatment: <http://dx.doi.org/10.6084/m9.figshare.658856>




Figure 3.4 Mean swarm density over all replicates over evolutionary time, measured by the mean number of prey within 30 virtual meters of each other over a lifespan of 1,000 simulation time steps. Prey in groups attacked randomly (light grey triangles) took much longer to evolve cohesive swarming behavior than prey in groups attacked by a predator that follows a random walk (dark grey circles) or always from the outside of the group (black squares). When prey experience no attacks, they do not evolve swarming behavior at all (light grey stars). Error bars indicate two standard errors over 100 replicates.

3.2.3 Outside Attacks

In the last of my initial artificial predation experiments, I simulate a predator that always approaches from the outside of the group and attacks the prey nearest to it, as in [80]. This predator attack mode effectively has the predators consistently attacking prey on the outer edges of the group. As shown in shown in Figure 3.3C, I simulate this predator attack mode by first choosing a random angle outside of the group for the predator to approach from. Once an angle is chosen, I convert the angle into a location on the edge of the visible simulation space and attack the prey nearest to that location.

As shown in Figure 3.4, this form of predation has the most significant impact on the evolution of the selfish herd so far. When attacked by predators that consistently target prey on the edges of the group, prey quickly evolve cohesive swarming behavior⁵ (black squares). Taken together, the results of these artificial predation experiments demonstrate another discovery of this work: The more predators attack prey on the outside of the group, the faster the selfish herd will evolve.

One translation of this finding is that in order for the selfish herd to evolve, prey must experience a higher predation rate on the outside of the group than in the middle of the group. While this phenomenon can be explained by each prey having a “domain of danger” (DOD) influenced by its relative position in the group [27, 31, 38], an alternative hypothesis is that of density-dependent predation.




Figure 3.5 Mean swarm density over all replicates over evolutionary time, measured by the mean number of prey within 30 virtual meters of each other over a lifespan of 1,000 simulation time steps. Even when experiencing density-dependent predation, prey in groups attacked randomly (light grey triangles) took much longer to evolve swarming behavior than prey in groups attacked by a persistent artificial predator (dark grey circles) or always from the outside of the group (black squares). Error bars indicate two standard errors over 100 replicates.

3.2.4 Density-Dependent Predation

To study the impact of density-dependent predation on the evolution of the selfish herd, I impose a constraint on the predator which reduces its attack efficiency when it attacks areas of the group with high prey density. This reduced attack efficiency is meant to represent the increased predation rate that prey on edges of the group are expected to endure [27, 31, 38], and such density-dependence can also be thought of as a proxy for group defense. I compute the predator’s probability of capturing a prey during a given attack (P_{capture}) with the following equation:

$$P_{\text{capture}} = \frac{1}{A_{\text{density}}} \quad (3.2)$$

where A_{density} is the number of prey within 30 virtual meters of the target prey, including the target prey itself. For example, if the predator attacks a prey with 4 other prey nearby ($A_{\text{density}} = 5$), it has a 20% chance of successfully capturing the prey. A biological analogue of this mechanism would be, for example, lions having more success capturing buffalo on the edge rather than in the middle of the herd. As a consequence of this mechanism, the prey experience density-dependent predation.

Figure 3.5 demonstrates the effect of density-dependent predation on the previous artificial predation experiments. Just as before, when predators did not preferentially attack prey on the outside of the group, as in the random attack experiment (light grey triangles), swarming behavior took much longer to evolve. In contrast, when the predators followed a random walk (dark grey circles) or always attacked from the outside of the group (black squares), the prey experiencing density-dependent predation again quickly evolved swarm-

⁵Evolution of prey behavior under Outside Attack treatment: <http://dx.doi.org/10.6084/m9.figshare.658854>

Table 3.3 High-density area attack (HDAA) experiment treatments. The values listed for each treatment are the handling times for the corresponding predator attack mode.

HDAA?	Outside Attack Frequency	HDAA Frequency
No	10	N/A
Infrequent	10	250
Frequent	10	25

ing behavior. The most noticeable effect of density-dependent predation is on the random attack treatment, where the swarm density measurement at generation 5,000 increased from 11.19 ± 2.58 (mean $\pm$ two standard errors) to 17.61 ± 2.72 , indicating significantly stronger selection for swarming.

3.2.5 High-Density Area Attacks

Thus far, I have explored attack modes that select for the evolution of swarming behavior. It is not surprising that there are also attack modes exhibited by natural predators that must select against swarming behavior in their prey. For example, blue whales (*Balaenoptera musculus*) are known to dive into the densest areas in swarms of krill, consuming hundreds of thousands of krill in the middle of the swarm in a single attack [81]. I call this kind of attack mode a *high-density area attack*. Such an attack clearly selects against swarming behavior because it targets the prey that swarm the most. If krill swarms consistently experience these high-density area attacks, then why do they still evolve swarming behavior?

It is important to note that krill swarms are also fed on by smaller species, such as crabeater seals (*Lobodon carcinophagus*), that consistently attack the krill on the outside of the swarm [82]. Thus, krill swarms are experiencing two forms of attack modes simultaneously: High-density area attacks from whales and outside attacks from crabeater seals. Thus, it is possible that the selection pressure to swarm from outside attacks (Figure 3.4)

could outweigh the selection pressure to disperse from high-density area attacks.

Shown in Figure 3.3D, I model high-density area attacks as an artificial attack that always targets the prey at the most dense area of the swarm (i.e., highest A_{density}). I note that this attack mode is the opposite of the density-dependent mechanism explored in the previous section, which favors predators that target prey in the *least* dense area of the swarm. Once the target is selected, I execute the attack by removing the target prey and all other prey within 30 virtual meters of the target prey. Outside attacks are modeled as described above. To study the effect of high-density area attacks on the evolution of swarming behavior, I allow the prey to evolve while experiencing both attack modes simultaneously. I vary the relative handling times of both attacks (Table 3.3) to explore whether relative attack frequency could explain why some swarming animals evolved swarming behavior despite the fact that they experience high-density area attacks.

As shown in Figure 3.6, prey experiencing only outside attacks quickly evolve cohesive swarming behavior (light grey triangles). However, when I introduce infrequent high-density area attacks (dark grey circles), the selection pressure for prey to swarm is reduced. Finally, when I introduce frequent high-density area attacks (black squares), the prey do not evolve swarming behavior at all. Thus, one possible explanation for animals evolving swarming behavior despite experiencing high-density area attacks is that the high-density area attacks are too infrequent relative to other attack types to exert a strong enough selection pressure for prey to disperse.

In summary, the artificial predation experiments provided me with two important findings regarding the evolution of the selfish herd: (1) attacks on prey on the periphery of the herd exert a strong selection pressure for prey to swarm and (2) prey in less dense areas, such as those on the outside of the herd, must experience a higher predation rate than in areas of




Figure 3.6 Mean swarm density over all replicates over evolutionary time, measured by the mean number of prey within 30 virtual meters of each other over a lifespan of 1,000 simulation time steps. Swarm density was measured from evolved populations that were not experiencing predation during measurement, eliminating any possible effects of attack modes that kill more prey faster. Prey in groups attacked only by outside attacks (light grey triangles) evolved cohesive swarming behavior. Increasing the relative frequency of high-density area attacks from infrequent (dark grey circles) to frequent (black squares) caused the prey to evolve increasingly dispersive behavior. Error bars indicate two standard errors over 100 replicates.

dense prey, such as those in the center of the herd.

3.3 Predator-Prey Coevolution

Building upon the artificial predation experiments, I implemented density-dependent predation in a predator-prey coevolution experiment. Adding predators into the simulation environment enables me to observe how embodied coevolving predators affect the evolution of the selfish herd.

For this experiment, I coevolve a population of 100 predator genomes with a population of 100 prey genomes using a GA with settings described in Table 3.2. Specifically, I evaluate each predator genome against the entire prey genome population for 2,000 simulation time steps each generation. During evaluation, I place 4 clonal predator agents inside a 512×512 virtual meters simulation environment with all 100 prey agents and allow the predator agents to make attack attempts on the prey agents. The prey genome population size, simulation environment area, and total number of GA generations were decreased in this experiment due to computational limitations imposed by predator-prey coevolution. I assigned the prey individual fitness values as in the previous experiments, and evaluated predator fitness according to the following equation:

$$W_{\text{predator}} = \sum_{t=1}^{t_{\text{max}}} (S_0 - A_t) \quad (3.3)$$

where t is the current simulation time step, t_{max} is the total number of simulation time steps (here, $t_{\text{max}} = 2,000$), S_0 is the starting group size (here, $S_0 = 100$), and A_t is the number of prey alive at update t . Thus, predators are selected to consume more prey faster, and prey are selected to survive longer than other prey in the group. Once all of the predator




Figure 3.7 Mean swarm density over all replicates over evolutionary time, measured by the mean number of prey within 30 virtual meters of each other over a lifespan of 2,000 simulation time steps. Prey in groups experiencing density-dependent predation (black circles) evolved cohesive swarming behavior, whereas prey in groups not experiencing density-dependent predation (light grey triangles) evolved dispersive behavior. Error bars indicate two standard errors over 100 replicates.

and prey genomes are assigned fitness values, I perform fitness proportionate selection on the populations via a Moran process [77], increment the generation counter, and repeat the evaluation process on the new populations until the final generation (1,200) is reached.

To evaluate the coevolved predators and prey quantitatively, I obtained the line of descent (LOD) for every replicate by tracing the ancestors of the most-fit prey MN in the final population until I reached the randomly-generated ancestral MN with which the starting population was seeded (see [51] for an introduction to the concept of a LOD in the context of digital evolution). I again characterized the prey grouping behavior by measuring the swarm density of the entire prey population every generation.

Figure 3.7 depicts the prey behavior measurements for the coevolution experiments with

density-dependent predation⁶ (black circles; mean swarm density at generation 1,200 $\pm$ two standard errors: 26.2 $\pm$ 2.3) and without density-dependent predation (light grey triangles; 3.9 $\pm$ 0.8). Without density-dependent predation, the prey evolved purely dispersive behavior as a mechanism to escape the predators, even after 10,000 generations of evolution ([83], Figure S1). In contrast, with density-dependent predation, the prey quickly evolved cohesive swarming behavior in response to attacks from the predators within 400 generations. As a caveat, density-dependent predation only selects for cohesive swarming behavior when the predators are faster than the prey ([83], Figure S2), which corroborates earlier findings exploring the role of relative predator-prey speeds in the evolution of swarming behavior [63].

Here I see that density-dependent predation provides a sufficient selective advantage for prey to evolve the selfish herd in response to predation by coevolving predators, despite the fact that swarming prey experience an increased attack rate from the predators due to this behavior ([79], Figures S3 & S4). Accordingly, these results uphold Hamilton's hypothesis that grouping behavior could evolve in animals purely due to selfish reasons, without the need for an explanation that involves the benefits to the whole group [27]. Moreover, the discoveries in this work refine the selfish herd hypothesis by clarifying that the predator's attack mode has a significant effect on the evolution of selfish herd behavior: in particular, that selfish herd behavior is much more strongly selected for when the predator consistently attacks the edges of the group rather than attacking randomly.

⁶Prey behavior from predator-prey coevolution treatment: <http://dx.doi.org/10.6084/m9.figshare.658855>

3.4 Evolved Prey Markov Network Analysis

Now that I have evolved emergent swarming behavior in an agent-based model under several different treatments, I can analyze the resulting Markov Networks (MNs) to gain a deeper understanding of the individual-based mechanisms underlying swarming behavior. For this analysis, I chose the most-abundant prey MN from each of the Outside Attack artificial predation experiment replicates, resulting in 100 MNs that exhibit swarming behavior.

First, I analyze the structure of the 100 MNs by looking at the specific retina sensors that the MNs evolved to connect to. Shown in Figure 3.8, the prey MNs show a strong bias for connecting to the prey-specific retina sensors in front of the prey, but not to the sides. Additionally, some of the prey MNs show a preference for connecting to the prey-specific retina sensors behind the prey. From this analysis alone, I can deduce that the retina sensors that are most conducive to swarming behavior are in front of the prey agent.

To understand how prey make movement decisions based on their sensory input, I map every possible input combination in the prey’s visual system to the corresponding movement decision that the prey made. Due to the stochastic nature of Markov Networks, the prey agents do not always make the same movement decision when given the same input. Thus, I take the most-likely output out of 1,000 trials as the representative decision for a given sensory input combination. Effectively, this process produces a truth table that maps every possible sensory input to its corresponding movement decision. I input this truth table into the logic minimization software *espresso* [84], which outputs the minimal representative logic of the truth table. This process results in a truth table that is reduced enough to make the evolved prey behavior comprehensible by humans.

Surprisingly, the individual-based mechanisms underlying the emergent swarming behav-




Figure 3.8 Number of sensory input connections from 100 evolved prey Markov Networks mapped onto a prey agent. Only causal connections from the sensory inputs to the actuators are shown. The arrow indicates the facing of the agent. The prey Markov Networks evolved a strong preference for connecting to prey sensors in front and a slight preference for sensors behind the prey agent, but tended to not connect to the sensors on the sides.

ior are remarkably simple. Most of the prey MNs evolved to make their movement decisions based off of only one prey sensor in front of the prey agent. If the prey sensor does not detect another prey agent, the agent repeatedly turns in one direction until it detects another prey agent in that sensor. Once the agent detects another prey agent in the sensor, it moves forward until the agent is no longer visible. This mechanism alone proved sufficient to produce cohesive swarming behavior in the majority of my experiments. Interestingly, this discovery corroborates the findings in earlier studies suggesting that complex swarming behavior can emerge from simple movement rules when applied over a population of locally-interacting agents [79, 85, 73].

In a small subset of the evolved prey MNs, I observe MNs that occasionally connect to one of the prey sensors behind them. These MNs watch for a prey agent to appear in a single prey sensor behind the agent and turn repeatedly in one direction until a prey agent is no longer visible in that sensor. Once a prey agent is no longer visible in the back sensor, the MN moves forward or turns depending on the state of the frontal sensor. I note that this mechanism *only* evolved in prey MNs that already exhibited swarming behavior using one of the frontal sensors, which suggests that this mechanism does not play a major role in swarming behavior. Instead, this mechanism seems to cause the prey agent to turn toward the center of the swarm instead of swarming in a circle with the rest of the prey agents. This mechanism can be thought of as a “selfish herd” mechanism that attempts to selfishly move the agent toward the center of the swarm to avoid predation.

3.5 Discussion

In this chapter, I demonstrated Hamilton’s selfish herd hypothesis in a digital model of evolution and highlighted that it is the attack mode of the predator that critically determines the evolvability of swarming behavior. Further, we showed that density-dependent predation is sufficient for the selfish herd to evolve as long as the predators cannot consistently attack prey in the center of the group. Finally, I showed that density-dependent predation is sufficient to evolve grouping behavior in prey as a response to predation by coevolving predators. Consequently, future work exploring the evolution of the selfish herd in animals should not only consider the behavior of the prey in the group, but the attack mode of the predators as well. Following these experiments, I analyzed the evolved control algorithms of the swarming prey and identified simple, biologically-plausible agent-based algorithms that produce emergent swarming behavior, including a mechanism that produces “selfish” behavior that drives the prey toward the center of the swarm.

The findings in this chapter point to two general conclusions regarding the evolution of collective animal behavior. First and foremost, the fact that seemingly cooperative behavior evolved for purely selfish reasons (i.e., to survive longer than other conspecifics in the group) suggests that the advantages offered by grouping do not necessarily need to benefit the entire group: Indeed, swarming behavior evolved in these simulations simply because the prey found swarming to be a viable tactic to survive longer than other prey in the group, even when swarming didn’t provide any inherent advantage to the entire group. Second, the relatively simple control algorithms that evolved in the swarming prey suggests that the control algorithms used in top-down swarm modeling approaches—such as the popular Boids model [73]—are likely overly complicated and require sensory information that is not

realistically available to real swarming animals. Consequently, further analysis of the evolved control algorithms from future agent-based swarming experiments should prove enlightening for uncovering the true mechanisms that real animals use to swarm.

Chapter 4

Predator Confusion Hypothesis

In the predator confusion hypothesis, the presence of multiple individuals moving in a swarm confuses approaching predators, making it difficult for predators to successfully execute an attack [29, 7, 30, 41]. Despite the inherent advantage that predator confusion offers swarming groups of prey, it remains unclear whether predator confusion can provide a sufficient selective advantage for swarming behavior to evolve in the first place.

In this chapter, I use digital models of evolution to explore the evolutionary consequence of the predator confusion hypothesis on populations of coevolving predators and prey. First, I describe the details of the digital model that I used in this project. Next, I explain the results from the model and how predator confusion affects the evolution of prey behavior. Following that, I briefly describe an additional experiment I performed to explore how predator confusion can in turn affect the evolution of predator visual systems. Finally, I conclude the chapter by discussing some of the broader implications of the findings in this chapter.

4.1 Model of predator-prey interactions

To study the effects of predator confusion on the evolution of swarming, I created an agent-based simulation in which predator and prey agents interact in a continuous two-dimensional virtual environment. Each agent is controlled by a *Markov Network* (MN), which is a stochastic state machine that makes control decisions based on a combination of sensory input (i.e.,

vision) and internal states (i.e., memory) [86]. I coevolve the predator and prey with a *genetic algorithm* (GA), which is a digital model of evolution by natural selection [66]. In a GA, pools of genomes are evolved over time by evaluating the fitness of each genome at each generation and preferentially selecting those with higher fitness to populate the next generation. The genomes here are variable-length strings of integers that are translated into MNs during fitness evaluation. More information on MNs—including details on their genetic encoding, mutational operators, and functionality—is available in Chapter 2.3.

To perform this coevolution, I create separate genome pools for the predator and prey genomes. Next, I evaluate the genomes’ fitness by selecting pairs of predator and prey genomes at random without replacement, then place each pair into a simulation environment and evaluate them for 2,000 simulation time steps. Within this simulation environment, I generate 50 identical prey agents from the single prey genome and compete them with the single predator agent to obtain their respective fitness. This evaluation period is akin to the agents’ lifespan, hence each agent has a potential lifespan of 2,000 time steps (enough time for the prey to travel approximately 400 body lengths). The fitness values, calculated using the fitness function described below, are used to determine the next generation of the respective genome pools. Parameters describing the operation of this GA are summarized in Table S1. At the end of the lifetime simulation, I assign the predator and prey genomes separate fitness values according to the fitness functions:

$$W_{\text{predator}} = \sum_{t=1}^{2,000} S - A_t \quad (4.1)$$

$$W_{\text{prey}} = \sum_{t=1}^{2,000} A_t \quad (4.2)$$

where t is the current simulation time step, S is the starting swarm size (here, $S = 50$), and A_t is the number of prey agents alive at simulation time step t . It can be shown that the predator fitness (Eq. 4.1) is proportional to the mean kill rate k (mean number of prey consumed per time step), while the prey fitness (Eq. 4.2) is proportional to $(1 - k)$. Thus, predators are awarded higher fitness for capturing more prey faster, and prey are rewarded for surviving longer. I simulate only a portion of the prey’s lifespan where they are under predation because I am investigating swarming as a response to predation, rather than a feeding or mating behavior.

Once I evaluate all of the predator-prey genome pairs in a generation, I perform fitness-proportionate selection on the populations via a Moran process, allow the selected genomes to asexually reproduce into the next generation’s populations, increment the generation counter, and repeat the evaluation process on the new populations until the final generation (1,200) is reached.

I perform 180 replicates of each experiment, where for each replicate I seed the prey population with a set of randomly-generated MNs and the predator population with a pre-evolved predator MN that exhibits rudimentary prey-tracking behavior. Seeding the predator population in this manner only serves to speed up the coevolutionary process, and has negligible effects on the outcome of the experiment ([79], Figure S1).

4.1.1 Predator and prey agents

Figure 4.1 depicts the sensory-motor architecture of predator and prey agents in this system. The visual system sensors of both predator and prey agents are logically organized into “layers,” where a layer includes 12 sensors, with each sensor having a field of view of 15° and a range of 100 virtual meters (200 virtual meters for predators). Moreover, each layer is




Figure 4.1 An illustration of the predator and prey agents in the model. Light grey triangles are prey agents and the dark grey triangle is a predator agent. The predator and prey agents have a 180° limited-distance visual system (100 virtual meters for the prey agents; 200 virtual meters for the predator agent) to observe their surroundings and detect the presence of the predator and prey agents. Each agent has its own Markov Network, which decides where to move next based off of a combination of sensory input and memory. The left and right actuators (labeled “L” and “R”) enable the agents to move forward, left, and right in discrete steps.

attuned to sensing a specific type of agent. Specifically, the predator agents have a single-layer visual system that is only capable of sensing prey. In contrast, the prey agents have a dual-layer visual system, where one layer is able to sense conspecifics, and the other senses the predator. (I note that there is only a single predator active during each simulation, hence the lack of a predator-sensing retinal layer for the predator agent.)

Regardless of the number of agents present in a single retina slice, the agents only know the agent type(s) that reside within that slice, but not how many, representing the wide, relatively coarse-grain visual systems typical in swarming birds such as Starlings [76]. For example in Figure 4.1, the furthest-right retina slice has two prey in it (light grey triangles), so the prey sensor for that slice activates. Similarly, the sixth retina slice from the left has both a predator (dark grey triangle) and a prey (light grey triangle) agent in it, so both the predator and prey sensors activate and inform the MN that one or more predators *and* one or more prey are currently in that slice. Lastly, since the prey near the 4th retina slice from the left is just outside the range of the retina slice, the prey sensor for that slice does not activate. I note that although the agent’s sensors do not report the number of agents present in a single retina slice, this constraint does not preclude the agent’s MN from evolving and making use of a counting mechanism which reports the number of agents present in a set of retina slices. Once provided with its sensory information, the prey agent chooses one of four discrete actions: (1) stay still; (2) move forward 1 unit; (3) turn left 8° while moving forward 1 unit; or (4) turn right 8° while moving forward 1 unit.

Likewise, the predator agent detects nearby prey agents using a limited-distance (200 virtual meters), pixelated visual system covering its frontal 180° that functions just like the prey agent’s visual system. Similar to the prey agents, predator agents make decisions about where to move next, but the predator agents move 3x faster than the prey agents and turn




Figure 4.2 Relation of predator attack efficiency ($\#$ successful attacks / total $\#$ attacks) to number of prey. The solid line with triangles indicates simulated predator attack efficiency as a function of the number of prey within the visual field of the predator (A_{NV}). Similarly, the dashed line with error bars shows the measured simulated predator attack efficiency given the predator attacks a group of swarming prey of a given size, using the A_{NV} curve to determine the per-attack predator attack success rate. Error bars indicate two standard errors over 100 replicate experiments.

correspondingly slower (6° per simulation time step) due to their higher speed.

4.1.2 Simulation environment

I use a simulation environment to evaluate the relative performance of the predator and prey agents. At the beginning of every simulation, I place a single predator agent and 50 prey agents at random locations inside a closed 512×512 unit two-dimensional simulation environment. Each of the 50 prey agents are controlled by clonal MNs of the particular prey MN being evaluated. I evaluate the swarm with clonal MNs to eliminate any possible effects of selection on the individual level, e.g., the “selfish herd” effect [40].

During each simulation time step, I provide all agents their sensory input, update their MN, then allow the MN to make a decision about where to move next. When the predator agent moves within 5 virtual meters of a prey agent it can see, it automatically makes an attack attempt on that prey agent. If the attack attempt is successful, the target prey agent is removed from the simulation and marked as consumed. Predator agents are limited to one attack attempt every 10 simulation time steps, which is called the *handling time*. The handling time represents the time it takes to consume and digest a prey after successful prey capture, or the time it takes to refocus on another prey in the case of an unsuccessful attack attempt. Shorter handling times have negligible effects on the outcome of the experiment, except for when there is no handling time at all ([79], Figure S2).

To investigate predator confusion as an indirect selection pressure driving the evolution of swarming, I implement a perceptual constraint on the predator agent. When the predator confusion mechanism is active, the predator agent’s chance of successfully capturing its target prey agent (P_{capture}) is diminished when any prey agents near the target prey agent are visible anywhere in the predator’s visual field. This perceptual constraint is similar to

previous models of predator confusion based on observations from natural predator-prey systems [7, 29, 43], where the predator’s *attack efficiency* (# successful attacks / total # attacks) is reduced when attacking swarms of higher density. P_{capture} is determined by the equation:

$$P_{\text{capture}} = \frac{1}{A_{\text{NV}}} \quad (4.3)$$

where A_{NV} is the number of prey agents that are visible to the predator, i.e., anywhere in the predator agent’s visual field, *and* within 30 virtual meters of the target prey. By only counting prey near the target prey, this mechanism localizes the predator confusion effect to the predator’s visual system, and enables me to experimentally control the strength of the predator confusion effect. Although my predator confusion model is based on the predator’s visual system, it is qualitatively similar to previous models that are based on the total swarm size—e.g., models of predator confusion presented in [7, 29, 43, 87]—in that there is a gradual (rather than immediate) decline in predator attack efficiency as the prey group size increases (Figure 4.2, dashed line). As shown in Figure 4.2 (solid line with triangles), the predator has a 50% chance of capturing a prey with one visible prey near the target prey ($A_{\text{NV}} = 2$), a 33% chance of capturing a prey with two visible prey near the target prey ($A_{\text{NV}} = 3$), etc. As a consequence, prey are in principle able to exploit the combined effects of predator confusion and handling time by swarming.

4.2 Effects of predator confusion

Qualitatively, I observed significant differences in prey behavior over the course of evolution between swarms experiencing predators with and without predator confusion. Figure 4.3A il-




Figure 4.3 Screen captures of (A) dispersed prey in a swarm hunted by a predator without predator confusion, (B) prey forming a single elongated swarm under attack by a predator with predator confusion, and (C) prey forming multiple cohesive swarms to defend themselves from a predator with predator confusion after 1,200 generations of evolution. Black dots are prey, the triangle is the predator, the lines projecting from the predator represent the predator's frontal 180° visual field, and the star denotes where a prey was just captured.

illustrates that prey hunted by a predator without the predator confusion mechanism dispersed as much as possible to escape the predator. No replicates containing a predator without predator confusion resulted in prey behavior that resembled a cohesive swarm. Conversely, when evolution occurred with predator confusion, prey exhibited cohesive swarm behavior in the majority of the replicates (70% of my replicates). Figure 4.3B depicts one such swarm in which prey follow the conspecific directly in front of them, resulting in an elongated swarm. Similarly, Figure 4.3C shows another swarm where the prey circle around their nearest conspecific, resulting in multiple small, cohesive swarms with the prey constantly trying to circle around each other. Both of these swarms evolved as defensive behaviors to exploit the predator confusion effect.

Furthermore, predators exhibited divergent hunting behaviors when hunting prey with and without predator confusion. As seen in Figure 4.3A, predators that evolved in the absence of predator confusion, and hence had to contend with dispersed prey, simply tracked the nearest visible prey until it was captured, then immediately pursued the next nearest visible prey. On the other hand, predators that evolved in the presence of predator confusion, and hence were challenged with cohesive swarms, used a mechanism that causes them to attack prey on the outer edges of the swarm. This strategy is similar to a predatory behavior observed in many natural systems [88, 89], and effectively minimized the number of prey in the predator’s visual system and maximized its chance of capturing prey. Figure 4.3B demonstrates this behavior, where the predator just captured a prey on the top-right edge of the swarm (prey capture location denoted by a black star). Videos of the evolved swarms under predation are available in the supplementary information ([79], SI videos 1-5).

To evaluate the evolved swarms quantitatively, I obtained the line of descent (LOD) for every replicate by tracing the ancestors of the most-fit prey MN in the final population

until I reached the randomly-generated ancestral MN with which the starting population was seeded (see [51] for an introduction to the concept of a LOD in the context of digital evolution). For each ancestor in the LOD, I characterized the swarm behavior with two common behavior measurements: *swarm density* and *swarm dispersion* [78]. I measured the swarm density as the mean number of prey within 30 virtual meters of each other over a lifespan of 2,000 simulation time steps. The swarm’s dispersion was computed by averaging the distance to the nearest prey for every living prey over a lifespan of 2,000 simulation time steps. Together, these metrics captured whether or not the prey were cohesively swarming.

Figure 4.4A demonstrates that the prey hunted by a predator with only handling time (i.e., without predator confusion) moved close to each other by chance but never coordinated their movement at any point in their evolutionary history (mean swarm density ± 1 standard error across 180 replicates: 0.69 ± 0.02). In contrast, when hunted by a predator with predator confusion, the prey coordinated their movement to remain close to each other and form a swarm (mean swarm density 12.48 ± 0.8 at generation 1,200). Likewise, Figure 4.4B shows that in the absence of predator confusion, prey evolved to maximize their dispersion (mean shortest distance 46.69 ± 0.44 at generation 1,200), whereas with predator confusion, prey evolved increasingly cohesive swarm behavior (mean shortest distance 22.54 ± 1.32 at generation 1,200). Taken together, these results confirm that predator confusion provided a sufficient selection pressure to evolve cohesive swarming behavior in this model, even though the swarming prey actually experience an increased attack rate from the predator due to this behavior (see [79], Figures S3 & S4).

Figure 4.4C shows that as a result of these evolutionary trends, the cohesive swarms that evolved under predator confusion experienced significantly higher survivorship than swarms that evolved without predator confusion (34.7 ± 0.6 and 25.54 ± 0.49 prey surviving




Figure 4.4 Mean swarm density (A), swarm dispersion (B), and survivorship (C) within the swarm over all replicates over evolutionary time. The swarm density was measured by the mean number of prey within 30 virtual meters of each other over a lifespan of 2,000 simulation time steps. Swarm dispersion was measured by the mean distance to the nearest prey for every living prey over a lifespan of 2,000 simulation time steps. Survivorship within the swarm was measured as the mean number of surviving prey (out of an initial total of 50) at the end of the simulation at a given generation. Prey hunted by a predator with predator confusion (black circles with a full line) evolved to maintain significantly higher swarm density and significantly less dispersed swarming behavior than prey in the swarms hunted by a predator without predator confusion (grey triangles with a dashed line). As a result, significantly more prey survived in the swarms hunted by a predator with predator confusion than the swarms hunted by a predator without predator confusion. Error bars indicate two standard errors across 180 replicate experiments.

the simulations, respectively). This increased survivorship confirms that swarming behavior confused the predator, leading to fewer successful prey captures. I found these results robust to a variety of experimental parameters, including weaker predator confusion effects ([79], Figure S5 & S6) and applying a minimum threshold to predator attack efficiency ([79], Figure S7).

4.3 Evolved predator and prey behavior

To deduce how swarms emerge in my model from individual-level behaviors, I next determined the functionality of the evolved predator and prey MNs. I accomplished this by first visualizing the MN connectivity to discern which slices of the visual system and memory nodes of the MN were causally connected, then created a truth table from the MN mapping every possible input combination with its corresponding most-likely output from the MN. With this input-output mapping, I computed the minimal descriptive logic of the MN with Logic Friday, a hardware logic minimization program. I used the most-likely output for every input combination due to the stochastic nature of MNs, therefore the functionality I determined was the *most-likely* behavior of the predator or prey.

In all of my experiments, the prey at generation 1,200 ignored the presence of predators and instead only reacted to the presence of conspecifics in their visual system in order to follow the other prey in the swarm. This result was particularly striking because it suggested that prey can evolve swarming behavior in response to predation without the ability to sense the predators hunting them, which was suggested in a previous study [24]. I observed that the prey evolved a wide variety of simple algorithms that exhibited a diversity of emergent swarming behaviors, ranging from moderately dispersed, elongated swarms similar to Star-

ling murmurations (Figure 4.3B) to tightly-packed cohesive swarms reminiscent of fish bait balls (Figure 4.3C).

As for the predators, the evolved behavior I observed at generation 1,200 with predator confusion appeared to be rather complex: The predators avoided dense swarms and hunted prey outside, or on the edge, of the swarm. However, the algorithm underlying this behavior was relatively simple, which allowed for the predators to evolve this hunting behavior fairly early in the simulations. The predators watched only the two center retina slices and constantly turned in one direction until a prey entered one of those slices. Once a prey became visible in one of the center retina slices, the predator moved forward and pursued that prey until it made a capture attempt. This process was repeated regardless of whether the predator successfully captured the prey. The simplicity of the predator algorithm and relative simplicity of the prey algorithms supports the findings of earlier digital swarm studies that complex swarm behaviors can be described by simple rules applied over a group of locally-interacting agents [64, 73].

4.4 Eco-evolutionary dynamics

Predator confusion has been hypothesized to be not only a selective pressure favoring swarming, but also as a determinant of the *functional response* [43], i.e., the number of prey consumed by the predator as a function of prey density [90]. Figure 4.5 supports a key prediction of functional response theory: Both with and without predator confusion, the system displayed a Type II functional response (a saturating effect of prey density), but when predator confusion was present the functional response showed a lower plateau (24.01 ± 0.49 prey consumed without predator confusion; 15.18 ± 0.57 with predator confusion). The fact that




Figure 4.5 Functional response curves of cohesive swarms hunted by a predator with predator confusion (black circles with a full line) and dispersed swarms hunted by a predator without predator confusion (grey triangles with a dashed line). The evolved, cohesive swarms hunted by a predator with predator confusion result in a Type II functional response with a lowered plateau. Error bars indicate two standard errors across 180 replicate experiments.

there was a Type II functional response even in the condition without predator confusion was the result of an additional constraint present in both conditions: The handling time that was imposed on the predator after prey capture before it can attack again. Additionally, when I varied the handling time in my experiments, I found that increasing the handling time also lowers the plateau of the Type II functional response ([79], Figure S9).

Modeling functional response has been an important problem in ecology [91], and is critical for constructing accurate models that capture the dynamics of predator-prey interactions over ecological and evolutionary time [92]. I provided evidence here that predator confusion has significant effects on functional response that are not captured in traditional models [43]. Most of these traditional models, including the original formulation of Holling [42], capture the ecological interaction between predator and prey. Evolution is assumed to shape the be-

havioral strategies and constraints that influence predator-prey dynamics, but only recently have biologists begun to explicitly study the dynamics of predator-prey interactions over both ecological and evolutionary time [44]. I have shown that a Type II functional response evolves even when it is not directly selected for, and the shape of the functional response can be attributed to specific constraints such as handling time and predator confusion.

4.5 Coevolution between predator visual systems and prey behavior

In the previous sections, I implemented predator confusion by imposing a perceptual constraint that reduces the probability of successfully capturing prey if one or more prey near the target prey are visible to the predator. This is meant to simulate the difficulty, arising from attentional or cognitive limitations, that a biological predator might have in choosing among multiple available prey at the moment of attack. To examine the effect of relaxing this constraint, I coevolved the predator and prey again and experimentally reduced the size of the predator's field of view. This procedure reduces the possibility that multiple prey can be detected at the moment of attack, thereby reducing the probability of confusion. For example, experimentally decreasing the predator's field of view from 180° to 60° decreases by two-thirds the area within which the presence of multiple prey can confuse the predator.

Figure 4.6 demonstrates that when the predator's visual system only covered the frontal 60° or less, swarming to confuse the predator was no longer a viable adaptation (as indicated by a mean swarm density of 0.68 ± 0.02 at generation 1,200). In this case, the predator had such a narrow view angle that few swarming prey were visible during an attack, which minimizes the confusion effect and correspondingly increases its capture rate ([79], Figure




Figure 4.6 Mean swarm density at generation 1,200 as a function of predator view angle. Swarming to confuse the predator was an ineffective behavior if the predator's visual field covered only the frontal 60° or less, due to the predator's focused visual system. As the predator's visual field was incrementally increased to cover the frontal 90° and beyond, predator confusion via swarming again became an effective anti-predator behavior, as evidenced by the swarms exhibiting significantly higher swarm density at generation 1,200. Error bars indicate two standard errors across 180 replicate experiments.




Figure 4.7 Swarm density and predator view angle from the LOD of a single coevolution experiment. The predator and prey populations appear to continually cycle between different states of view angles and behaviors.

S8). As the predator’s visual system was incrementally modified to cover the frontal 120° and beyond, swarming again became an effective adaptation against the predator due to the confusion effect (indicated by a mean swarm density of 6.13 ± 0.76 at generation 1,200). This suggests that the predator confusion mechanism may not only provide a selective pressure for the prey to swarm, but it could also provide a selective pressure for the predator to narrow its view angle to become less easily confused.

This finding opens the possibility for coevolution between the predator’s visual system and the prey swarming behavior. To explore this possibility further, I ran the predator-prey coevolution experiments again, but this time allowing mutations (with 5% probability) to affect the view angle of the predator’s offspring. When a mutation occurs to the predator offspring’s view angle, a random number between $[-50.0, 50.0]$ is added to the offspring’s view angle. Thus, mutations can widen or shrink the predator’s view angle in large or small steps.

A typical coevolutionary experiment is depicted in Figure 4.7. Surprisingly, the predator and prey populations appear to continually cycle between different states of view angles and behaviors, respectively, such that there is a significant negative correlation between the




Figure 4.8 Pearson's r between swarm density and predator view angle from the LODs of 30 coevolution experiments. All coevolution experiments have a negative correlation between swarm density and predator view angle, indicating that when swarm density goes up, predator view angle goes down and vice versa. $P \leq 0.001$ for all correlations.

predator view angle and swarm density across all 30 coevolution experiments (Figure 4.8). This finding is surprising because the predator population can effectively “defeat” the swarming prey population by shrinking their visual system to the point that the prey will no longer evolve to swarm. Why then would the predator population evolve to widen their visual system again, allowing the prey population to again evolve swarming behavior to reduce the predators’ attack efficiency?

The answer to this question lies in the predators’ attack efficiency when the prey are no longer swarming. Shown in Figure 4.9, when predators with varying fixed view angles are competed against dispersive prey, predators with wider visual systems are more likely to find a prey anywhere in their visual system at any time. Further, Figure 4.10 demonstrates that predators with wider visual systems are also more likely to find dispersive prey in a




Figure 4.9 Number of simulation time steps that prey are present anywhere in an evolved predator's visual system depending on the predator's view angle. The predator is competed against dispersive prey. Predators with higher view angles are more likely to have prey anywhere in their visual system at a given time. $P \leq 0.001$ between all view angles, Kruskal-Wallis multiple comparison.




Figure 4.10 Number of simulation time steps that prey are visible in a portion of an evolved predator's visual system that it pays attention to, depending on the predator's view angle. The predator is competed against dispersive prey. Predators with higher view angles are more likely to spot prey at a given time, which increases their foraging efficiency. $P \leq 0.001$ between all view angles except 180 vs. 210, Kruskal-Wallis multiple comparison.




Figure 4.11 Fitness of an evolved predator when competed against dispersive prey, depending on the predator’s view angle. Predators with higher view angles forage for prey more efficiently, thus capturing more prey in their lifetime and improving their fitness. $P \leq 0.001$ between all view angles except 150 vs. 180 and 180 vs. 210, Kruskal-Wallis multiple comparison.

portion of their visual system that they pay attention to, which means they spend less time searching for prey. Thus, the increased foraging efficiency that wider visual systems offer predators against dispersive prey results in higher predator fitness (Figure 4.11). These findings point to a trade-off that natural predators likely experience when hunting prey: Wider, less-focused visual systems are more useful for initially spotting prey, but focused visual systems are better adapted for tracking an individual prey down and avoiding the effects of predator confusion when hunting prey in groups.

4.6 Discussion

I demonstrated that swarming evolves as an emergent behavior in prey when a simple perceptual constraint—predator confusion—is imposed on the predator. Further, I found that

measuring swarm density and swarm dispersion, proposed in [78], serves as an effective substitute for qualitatively assessing every swarm to determine if cohesive swarming behavior is present. A diverse collection of prey swarming behaviors evolved in my model, suggesting that predator confusion could allow for a wide range of swarming behaviors to evolve. Strikingly, most evolved prey strategies used algorithms that responded to other prey, but not to the attacking predators. This raises the interesting question of what selection pressures would favor the evolution of prey that detect and respond to the predators themselves.

In contrast to the diversity of evolutionary outcomes for prey, a common behavioral strategy emerged among the predators when evolved in the confusion condition. Namely, the evolved predators focused on attacking prey on the vulnerable edges of the swarms, which is a phenomenon commonly observed in nature [88, 89].

Modeling functional response has been an important problem in ecology [91], and is critical for constructing accurate models that capture the dynamics of predator-prey interactions over ecological and evolutionary time [92]. I provided evidence that predator confusion has significant effects on functional response that are not captured in traditional models [43]. Most of these traditional models, including the original formulation of Holling [42], capture the ecological interaction between predator and prey. Evolution is assumed to shape the behavioral strategies and constraints that influence predator-prey dynamics, but only recently have biologists begun to explicitly study the dynamics of predator-prey interactions over both ecological and evolutionary time [44]. I have shown that a Type II functional response evolves even when it is not directly selected for, and the shape of the functional response can be attributed to specific constraints such as handling time and predator confusion.

I also found that I could reduce the advantage of swarming by diminishing the predator's field of view, hence decreasing the level of confusion affecting the predator. This suggests




Figure 4.12 Diagram depicting the observed coevolutionary cycle between the predator and prey in the presence of the predator confusion effect.

that predator confusion could impose a selective pressure on the shape of the predator’s visual system: Once swarming has evolved in the prey, selection will favor predators that are no longer confused by swarms. Following the trend in Figure 4.6, I expected selection to favor predators with a narrower, more frontally focused visual system, as observed in the visual systems of many natural predators [93].

In the final section of this chapter, I directly explored the above hypothesis and found my prediction to be partly true: As demonstrated in Figure 4.7, selection does indeed favor predators with a more focused visual system once swarming has involved in prey. However, once the predators evolve a focused visual system, the prey evolve dispersive behavior in response and a coevolutionary cycle commences between the predator visual system and prey behavior in my model. Generally, researchers assume that the evolution of collective behavior is a one-way street, i.e., collective behavior is so evolutionarily advantageous that a species would only evolve increasingly collective behavior [94]. These findings demonstrate a coevolutionary cycle that could occur between natural predators and prey due to the con-

fusion effect, shown in Figure 4.12. As a part of this coevolutionary cycle, I discovered a condition under which the prey populations consistently evolved away from collective behavior. This experiment therefore highlights one possible mechanism through which collective behavior could be lost evolutionarily.

Chapter 5

Many Eyes Hypothesis

In this chapter, I focus on anti-predator vigilance (i.e., the many eyes hypothesis) as a possible selective mechanism for the evolution of gregarious foraging behavior, and control for the influence of the other benefits described in Chapter 2.1. I assume that vigilance has benefits (e.g., communicating the presence of a predator via alarm signals) but also costs (e.g., reduced foraging rates by watching for the predator). Under the many eyes hypothesis, grouping is beneficial because it reduces the cost of vigilance by sharing the cost of vigilance among the group, but it may have additional costs that must be considered, e.g., increased predation rates on larger groups [95]. Furthermore, this benefit would be diluted if some individuals can freeloader on the vigilance of others (as in heterogeneous groups), but magnified if the group members are highly related. The benefits and costs would also be affected by the life history of the prey, in particular whether their reproduction is iteroparous (i.e., repeated) or semelparous (i.e., all at once): Vigilance may be more beneficial in semelparous prey because a predation event can completely prevent them from reproducing, whereas iteroparous prey are more likely to have reproduced at least once prior to experiencing a predation event. To explore these issues, I manipulate the genetic relatedness and reproductive strategy of groups of prey that are under predation and observe the resulting behavior after thousands of generations of digital evolution have taken place. A preliminary investigation of this work was published in the ALIFE 14 conference [96], and has been significantly extended in this chapter.




Figure 5.1 **Depiction of the disembodied simulation.** Prey seek to forage as much as possible while avoiding being captured by the predator. If none of the prey in the group are vigilant, the target prey is captured 100% of the time.

This chapter proceeds as follows. First, I describe the details of the digital model that I used in this project. Next, I describe the results from the model and explain what conditions select for the evolution of gregarious foraging behavior. Finally, I conclude the chapter by discussing some of the broader implications of the findings in this chapter.

5.1 Model of predator-prey interactions

Figure 5.1 depicts my model of predator-prey interactions in a disembodied model, wherein prey must balance the trade-offs between foraging and vigilance [13]. In an embodied model, every animat is situated in the world, perceives the world via its sensors, and can act on the world via behavioral or other responses [97]. While embodied models offer more detail and can capture potentially important aspects of the real world, they are also sensitive to implementation-specific details of the sensors and actuators, which can skew the results. I therefore focus on a disembodied model ¹ for the remainder of this study, which enables me to explore several factors affecting the evolution of group vigilance in isolation.

¹Model code: <https://github.com/phaley/eos/tree/non-embodied>

In this model, prey fitness is directly related to the amount of time it spends foraging, where a single round of foraging increases prey fitness by 1.0. However, prey vigilance determines whether a predator’s attack on the prey is successful. These two options— foraging and vigilance—are assumed to be mutually exclusive. Thus, prey must evolve to maximize their food intake while remaining vigilant enough to survive the entire simulation, which is akin to the maximum possible life span of the prey.

5.1.1 Simulation of predators and prey

I designed this model to capture certain features of natural predators and to control for potentially complicating factors. First, to ensure that predator attacks are not trivially predictable I simulate predators that attack at intervals that are normally distributed around a specific attack rate. Thus, predator attacks are randomly distributed throughout the 2,000-time-step duration of the simulation. To model the observation that larger groups of prey often attract more attacks from predators—a realistic cost of group living known as the *attraction effect* [95]—I scale this attack rate with the group size, such that the group experiences 5 predator attacks for every prey initially in the group over the course of the simulation. This scaling factor also allows me to control for the *dilution effect*, which has been suggested to allow prey to survive with lower vigilance levels in larger groups only because they are less likely to be the target of a predator’s attack [17, 98, 99].

Each time a predator appears, I randomly select a target prey from the surviving prey of previous attacks. This is followed by a 10 time step delay between the appearance of the predator in the simulation and the actual attack, representing the time it takes for a predator to close the distance to the prey. It is during this time that prey vigilance becomes important. If the target prey is vigilant at any time during this interval, then it spots the predator and

the attack has only a 10% chance of success. If the target prey is not vigilant but one or more other prey in the group are vigilant, then the other prey communicate the presence of the predator via an alarm signal or other behavioral indicator and the predator will capture the target prey 30% of the time. These probabilities are chosen based on analytical models of group vigilance [13] such that group vigilance is not as effective as individual vigilance, and models the imperfect communication between members of the group [100]. Finally, if no members of the group are vigilant while the predator is closing the distance to its target, then the entire group is unaware of the predator and the attack will succeed 100% of the time. In all cases of a successful attack, the target prey is removed from the simulation and can no longer forage to increase its fitness.

Each individual prey makes the decision to forage or be vigilant every simulation time step. This decision-making process is modeled with a *Markov Network* (MN), which is an “artificial brain” that can stochastically make decisions based on sensory input, memory, and previous actions [79, 86, 101]. Every prey MN is encoded by a list of numbers known as its genotype, such that changes to the genotype can result in changes in the function of the MN. Because I do not provide any sensory input to the prey in this simulation, I am effectively modeling the probability of a prey taking an action (e.g., be vigilant or forage?) every simulation time step. More information on MNs—including details on their genetic encoding, mutational operators, and functionality—is available in Chapter 2.3.

5.1.2 Evolutionary process

At the beginning of every experiment, I create a population of 100 individuals with random Markov Networks. I repeat the evaluation procedure described above until all 100 individuals in the Genetic Algorithm (GA) population have been assigned a fitness (see, e.g., [75] for

a full description of GAs). Once all individuals have been assigned a fitness, I use fitness-proportional selection according to a Moran process [77] to produce the next generation’s population of prey. Fitness-proportional selection ensures that prey with higher fitness values generally produce more offspring. The selected prey reproduce asexually, with a small probability of mutations (0.5% per site) affecting their offspring’s genotype. I repeat this evaluation-selection-reproduction process for 2,500 generations to ensure that the GA has reached an evolutionarily stable strategy [102] and replicate the experiments 100 times for each treatment—each with a distinct random number generator seed—to verify that I am capturing evolutionary trends rather than outlier scenarios.

5.1.3 Group size

Since the many eyes hypothesis predicts an inverse relationship between individual vigilance and group size [22, 15], I study prey populations across a range of group sizes: 5, 10, 25, and 50. In my early experiments, I observe the equilibrium vigilance levels when prey are forced to group. In my later experiments, I relax this assumption and allow the prey to choose to group (or not) every time step. In the latter case, I report the group size as the maximum initial group size. To provide a baseline for the optional grouping experiment, I compare its equilibrium vigilance levels to that of experiments where prey are forced to group and experiments where prey are forced to forage individually.

5.1.4 Genetic relatedness

For all of the above experiments, I study the effect of genetic relatedness on group vigilance behavior. Given that genetically related organisms are more likely to cooperate with each

other than genetically unrelated organisms [103], I expect that genetic relatedness within the group will play a critical role in the evolution of group vigilance behavior. To explore the two extremes of genetic relatedness, I form groups in two different ways.

In **homogeneous groups**, each individual in the GA population is evaluated separately. During an individual's fitness evaluation, I fill the group in the simulation with exact copies of the individual, and the fitness for that individual is the average fitness of all of its copies at the end of the simulation. Thus, for a GA with a population size of 100 individuals, I run 100 simulations every generation to acquire the fitness for each individual.

In **heterogeneous groups**, I use a subset of the GA population (which contains many prey with different genetics) to study how the prey fare in direct competition (or cooperation) with each other. When forming a heterogeneous group, I randomly sample individuals from the GA population without replacement until I reach the desired group size for the current treatment. This group is then evaluated in the simulation, where each individual has only one copy that is assigned a fitness. Once the simulation finishes, I flag the evaluated individuals so they are not evaluated again in that generation. Since the desired group sizes (5, 10, 25, and 50) are always smaller than the GA population size (100), this procedure is repeated until all individuals have been evaluated. For example, if the desired group size is 25 and the GA population is composed of 100 individuals, then the randomly-group-and-evaluate procedure is repeated 4 times. Thus, by following this procedure, all individuals in the GA population are evaluated only once per generation in a randomly-assigned group.

Since vigilance indirectly benefits the vigilant individual in homogeneous groups by aiding its kin, I expect that group vigilance will be highly beneficial in homogeneous groups. In contrast, because the vigilance of one prey can potentially aid a rival prey in heterogeneous groups, I expect to observe lower levels of vigilance in heterogeneous groups.

5.1.5 Reproductive strategy

The benefits of making the right decision in this simulated environment are straightforward: The prey must maximize food intake by surviving the longest while minimizing the time spent being vigilant. But the cost of making the wrong decision can also depend on the life history of the prey. For example, two different reproductive strategies—semelparity and iteroparity—should incur different costs. Semelparous organisms sit on one end of the reproductive spectrum and are characterized by a single reproductive event prior to death. On the other end of the reproductive spectrum, iteroparous organisms continually reproduce throughout their lifetime. I explore these two extremes by simulating semelparous and iteroparous prey in separate treatments.

When simulating **semelparous** prey in my model, I assume that their reproductive event occurs at the end of the simulation. Therefore, if a semelparous prey is consumed by the predator before the end of the simulation, all of its gathered food counts for nothing: it will leave no offspring.

When simulating **iteroparous** prey in my model, I assume that the prey are constantly reproducing throughout their lifetime. Therefore when a predator consumes an iteroparous prey, the prey can no longer increase its fitness via foraging, but any food it gathered prior to its death counts toward its fitness for the simulation.

I note that these are highly simplified implementations of reproductive strategies and are meant to capture one key variable: the probability of reproduction occurring before a predation event. I hypothesize that the increased risk of genetic death introduced by the semelparous treatment will provide an evolutionary incentive for prey to invest in vigilance, whereas prey in the iteroparous treatment will be more likely to engage in risky,

non-cooperative behavior because their demise does not necessarily doom their genetic lineage [104].

5.1.6 Explicit cost of grouping

The model described so far includes a cost of vigilance (insofar as prey cannot forage at the same time that they are vigilant), but there is no explicit cost to choosing to group aside from the possibility of aiding a competing individual. In a final treatment, I implement such a grouping penalty in order to model the realistic constraints of limited resources and the resulting scramble competition for food [105, 50, 98, 106]. This grouping penalty is only assessed on prey who choose to forage in the group, and decreases the amount of food they receive in that simulation time step proportional to the number of prey in the group. The group foraging penalty is imposed according to the equation:

$$\text{Food} = \frac{1.0}{M * G} \quad (5.1)$$

where G is the number of prey in the group and M is the penalty multiplier that allows me to experimentally control the severity of the penalty. Given this penalty, prey foraging in larger groups receive less food every time they forage, but potentially enjoy the benefits of group vigilance.

5.2 Forced grouping

I evolved the vigilance behavior of prey by subjecting them to predation under a variety of treatments that vary reproductive strategy and group composition. Vigilance is measured as




Figure 5.2 **Treatment comparison when prey are forced to forage in groups.** Both group homogeneity and a semelparous reproductive strategy select for high levels of vigilance. However, only homogeneous groups experience an increase in fitness as group size increases. In contrast, vigilance behavior breaks down in larger, heterogeneous groups of semelparous prey. Error bars indicate bootstrapped 95% confidence intervals over 100 replicates; some error bars are too small to be visible.

the percent chance that a prey will be vigilant at a given moment in time, averaged across all of the prey in the population. These treatments are repeated across a wide range of group sizes, allowing me to study not only whether the selection for vigilance can be generalized to groups of varying sizes, but also whether I can observe the inverse relationship between group size and vigilance predicted by the many eyes hypothesis.

In my first experiment, all prey in the simulation are forced to forage in the same group, and the only trait that is evolving is the prey decision to be vigilant or not at every time step. Under these conditions, I find that prey living in homogeneous groups consistently evolve higher levels of vigilance than their counterparts living in heterogeneous groups (Figure 5.2). This finding suggests that organisms living in groups with high genetic relatedness are more likely to evolve cooperative strategies. Thus, in my model as in many natural systems, gregarious foraging is most favorable when genetic interests are aligned.

Figure 5.2 also shows that semelparous prey are more likely to evolve vigilant strategies than iteroparous prey. This is because semelparity selects more strongly than iteroparity for successful evasion of predator attacks, since prey death negates all previous foraging efforts

in semelparous prey. This effect is seen across both homogeneous and heterogeneous groups, indicating that semelparity is a strong enough selective pressure to act independently of group genetic composition. Importantly, prey vigilance does not evolve at all in the absence of predation (Figure S1), and gradually reducing the predation rate leads to a correspondingly gradual decrease in prey vigilance levels (Figure S2). Therefore, I know that the selection pressure imposed by predation is the primary driving force behind this evolved vigilance behavior.

All three treatments that evolve any level of vigilance also see the prevalence of vigilance decrease as group size increases. This pattern is important because it matches the pattern predicted by the many eyes hypothesis: As group size increases, individuals are able to rely more on collective rather than individual vigilance and can in turn devote more of their own time to foraging. Since I use a relative attack rate that scales the predator's attack frequency with group size, this phenomenon must be caused by group vigilance and not the dilution effect (i.e., fewer attacks per individual in larger groups) cited in other studies. I note that vigilance in the heterogeneous/semelparous treatment appears to evolve away almost entirely in a group size of 50. To explain why this trend might be due to something other than collective vigilance, I can instead look at trends in the fitness of the populations.

I observe several interesting trends when I look at the influence of group size on average group fitness. In both homogeneous treatments, there is a steady increase in fitness with increasing group size, suggesting that gregarious foraging behavior is under positive selection. I see no significant fitness increase with group size in the heterogeneous/iteroparous populations, where the populations do not evolve vigilance behavior (Wilcoxon rank-sum $p = 0.79$ between group size 5 and 50). Unlike the other treatments, the heterogeneous/semelparous populations actually experience a significant *decrease* in fitness with increasing group size




Figure 5.3 **Treatment comparison when prey can choose to forage in groups.** Allowing prey to decide whether they wish to be in the group produces similar results compared to when they are forced to group. In homogeneous groups, prey choose to spend most of their time in the group. However, grouping breaks down (alongside vigilance) in heterogeneous groups of semelparous prey. This occurs despite there being no direct penalty assessed for choosing to group. Error bars indicate bootstrapped 95% confidence intervals over 100 replicates; some error bars are too small to be visible.

(Wilcoxon rank-sum $p = 2.77 \times 10^{-6}$ between group size 5 and 50), which suggests that cooperative behavior is not evolutionarily stable in larger heterogeneous groups. Accordingly, these findings suggest that heterogeneous populations are much more susceptible to non-vigilant, “cheating” prey strategies that sweep the population and reduce the overall population fitness.

5.3 Optional grouping

So far I have shown that prey appear to take advantage of collective vigilance to increase their fitness when they are forced to group. We might expect from this result (and the many eyes hypothesis predicts) that grouping provides a selective advantage. To test this expectation explicitly, I relax the constraints of the previous experiment by allowing the prey to evolve whether to group or not at every simulation time step. Since there is no direct fitness trade-off for grouping in this model yet (as there was for foraging and vigilance), this allows me to study whether the evolutionary advantages of grouping are favorable enough for vigilance and grouping to co-evolve.

Figure 5.3 shows that when I allow prey to choose to group, I find nearly the same results as before. This suggests that collective vigilance provides enough of a selective advantage to favor the evolution of grouping. It is not surprising that the homogeneous treatments evolve to group nearly 100% of the time, given that the population is genetically identical and any “altruistic” action indirectly benefits the altruist as well. As in the forced grouping experiment, I observe a decline in fitness in the heterogeneous/semelparous populations as group size increases, to the point that the population is nearly driven extinct. The inability of the heterogeneous/semelparous populations to evolve consistently high levels of

vigilance further supports the hypothesis that evolution is favoring short-term competitive advantages over long-term survival. This phenomenon is commonly known as the tragedy of the commons [107, 108], where selfish actions that provide an individual short-term benefit lead to a decrease in overall group fitness.

5.4 Tragedy of the commons in heterogeneous groups

To explore this apparent tragedy of the commons scenario further, I directly compare vigilance and fitness values from the forced and optional grouping experiments alongside a third experiment where I force the population to forage and survive as individuals. Figure 5.4 shows that when given the choice to group in the homogeneous treatments, prey behavior closely mirrors the behavior observed when forced to forage in a group. This observation confirms the previous suggestion that collective vigilance in homogeneous groups provides a fitness benefit that positively selects for gregarious foraging behaviors.

In contrast to the homogeneous populations, heterogeneous populations are much less likely to evolve gregarious foraging behaviors. Heterogeneous/iteroparous populations never evolve vigilance behavior regardless of whether the prey are forced to group or not (Figure 5.4). Similarly, heterogeneous/semelparous populations only evolve vigilance behavior in smaller groups, whereas the advantage of collective vigilance is lost in larger groups. At larger group sizes, prey with the ability to choose whether or not to forage in heterogeneous/semelparous groups instead evolve lower levels of vigilance than required to protect the group (Figure 5.4), which results in a decrease in overall group fitness relative to prey that always forage in groups (Figure 5.5).




Figure 5.4 **Vigilance in prey with and without the option to forage in groups.** In homogeneous groups, prey with forced and optional grouping evolve similar vigilance behaviors. In contrast, individualistic (non-grouping) prey evolve vigilance behaviors that maximize individual fitness. Meanwhile, individuals in heterogeneous/semelparous populations with the option to group evolve to be less vigilant than either of the other two treatments. Error bars indicate bootstrapped 95% confidence intervals over 100 replicates; some error bars are too small to be visible.




Figure 5.5 **Fitness for prey with and without the option to forage in groups.** In heterogeneous/semelparous groups, prey with the option to group have lower fitness than prey that are forced to group. Error bars indicate bootstrapped 95% confidence intervals over 100 replicates; some error bars are too small to be visible.




Figure 5.6 **Grouping behaviors in prey experiencing grouping penalties.** Even with a small grouping penalty ($M = 1.0$), all treatments except homogeneous/semelparous no longer evolve grouping behavior. Prey in the homogeneous/semelparous treatment evolve only slightly lower levels of grouping behavior, even with extreme penalties to foraging in a group ($M = 1,000$). Error bars indicate bootstrapped 95% confidence intervals over 100 replicates; some error bars are too small to be visible.

5.5 Explicit cost of grouping

In my final treatment, I investigate the impact of assessing a direct cost of foraging in a group (e.g., competition for food). Figure 5.6 shows that except in the homogeneous/semelparous treatment, an explicit grouping cost selects against gregarious foraging behavior even when the grouping penalty is small ($M = 1.0$). Conversely, prey in the homogeneous/semelparous treatment maintain some level of gregarious foraging behavior even when the penalty for foraging in groups is extreme ($M = 1,000$). Therefore, I conclude that in the presence of even a small penalty for foraging in a group and the absence of additional selection pressures that favor gregarious foraging (e.g., improved social status for sentinels), only the combination of high genetic relatedness within the group and a semelparous reproductive strategy select strongly enough for gregarious foraging behavior to evolve in my model.

5.6 Discussion

I found that gregarious foraging behavior can emerge under a variety of conditions when there is a benefit of vigilance and the spreading of information about predators. Prey that forage in homogeneous groups are more likely to evolve gregarious foraging behaviors compared to the those in heterogeneous groups. The same is true for semelparous organisms (who reproduce only once before death) compared to their iteroparous counterparts (who reproduce continually), but group homogeneity selects much more strongly for gregarious foraging behavior.

Clearly, there are numerous challenges to evolving any form of cooperative behavior in a population with unconstrained genetic relatedness. However, I have shown here that when there is strong selection for survival (as in the heterogeneous/semelparous treatment), the benefit of information sharing via being vigilant and making alarm signals is sufficient to select for cooperative behavior in heterogeneous groups. This finding demonstrates that kinship is not necessary for cooperative behavior to evolve as long as there is some benefit to information sharing within the group, e.g., reducing predator attack efficiency.

Further, my results point to a heretofore unsuspected cost of gregarious foraging that is unique to heterogeneous groups. I call this the “two-fold cost of vigilance.” In my model, vigilance behavior in heterogeneous groups is more than a trade-off with foraging on the individual level. By choosing to be vigilant, prey also risk aiding in the survival of rival prey, which then puts the vigilant prey at a fitness disadvantage because it sacrificed a round of foraging to aid the rival prey. Together, these costs could explain why prey in heterogeneous groups evolve to be less vigilant than those in homogeneous groups.

At the same time, it is also possible that there are some evolutionary advantages unique

to heterogeneous groups that I have not yet addressed. For example, my model does not currently allow for any kind of specialization in roles between individuals, which could explain the presence of multi-species groups in nature [109, 110]. If the prey could evolve to preferentially pay attention to certain “sentinel” members of the population (who, in turn, choose to be vigilant nearly always in order to receive some form of rewards, e.g., food or increased social status) then perhaps an evolutionarily stable form of gregarious foraging could be found in heterogeneous groups of all sizes. It is even possible that such a complex social structure could out-perform the relatively primitive cooperation in my homogeneous groups.

Alongside genetic relatedness, another positive selective pressure for the evolution of vigilance is a semelparous reproductive strategy. When prey must survive any and all predator attacks in order to reproduce, the impetus to be vigilant is much greater. Semelparous organisms are known to be more risk-averse than similar, iteroparous organisms [111], and the decision to forage instead of being vigilant is an example of one such risky behavior. Thus, rather than spending most of their time foraging (as iteroparous prey evolve to do in my model), semelparous prey in my model tend to devote most of their time to watching for predators. When given the opportunity to group with other prey and take advantage of collective vigilance, semelparous prey are actually able to spend less time being vigilant. Thus, when semelparous prey evolve lower levels of vigilance in larger groups, we are observing the effect of collective vigilance.

Given that many animals who rely on vigilance for survival are iteroparous, my result that vigilance is less likely to evolve in iteroparous populations may seem to be contradicted by evidence. In my experiments, I explore the two extremes of reproductive behavior: Semelparous strategies where the prey reproduce only once at the end of their lifetime,

and iteroparous strategies where the prey constantly reproduce throughout their lifetime. It is plausible that an intermediate strategy—where prey reproduce within a few breeding episodes throughout their lifetime—could select for vigilance behavior while at the same time offering the benefits of a more reliable iteroparous reproductive strategy. Even though such an intermediate strategy is not explored in this work, it would make an interesting focus for future work to explore the continuum between the two reproductive strategies.

Although my results suggest that the risk-averseness of semelparity induces semelparous prey to evolve to take advantage of collective vigilance, this selective pressure does not appear to be as strong as the pressure I observed in homogeneous groups. Proof of this observation can be found in the heterogeneous/semelparous treatment, where most group members attempt to cheat their way into collective vigilance by evolving lower levels of vigilance behavior than is observed in populations where prey are either forced to forage on their own or in the group (Figure 5.4). Ultimately, this selfish behavior results in lower fitness than the fitness of prey that are forced to forage in groups (Figure 5.5), but the constantly-present, short-term benefits of selfishness appear to be too enticing to allow a more advantageous, cooperative behavior to emerge.

The breakdown of cooperation in the heterogeneous/semelparous populations suggests that the populations are succumbing to a tragedy of the commons [107, 108]. In my experiments, all prey are competing against each other to forage as much food as possible without being captured by the predator. However, because there is an unlimited amount of food, the only depletable group resource is vigilance, which protects the entire group from the predator. As the resulting non-cooperative behavior in the heterogeneous/iteroparous populations demonstrate, absent any major selective pressures for collective vigilance, prey will evolve to selfishly forage 100% of the time. Therefore, group homogeneity and semelparity correspond

to two previously-established mechanisms for preventing a tragedy of the commons, namely kin selection and punishment for non-cooperative behaviors, respectively [107]. The relative efficacy of these mechanisms to prevent cheating merits further investigation, for example, does group homogeneity play a larger role than reproductive strategy in the evolution of collective vigilance?

In the presence of even a small penalty for foraging in groups, I observe that only prey in homogeneous groups with a semelparous reproductive strategy are capable of evolving gregarious foraging behavior (Figure 5.6). This finding suggests that, in the absence of unlimited food resources or extreme predation rates, collective vigilance (i.e., the many eyes hypothesis) is insufficient to select for gregarious foraging. However, there may be important aspects of natural systems that select for gregarious foraging that I did not model here. For example, predators have been observed to preferentially attack non-vigilant prey in groups [112], which would require prey to be vigilant even without the benefit of collective vigilance. Thus, it would be informative in future work to model such a preference for non-vigilant prey and observe the evolution of gregarious foraging under those conditions.

Chapter 6

Conclusion

In total, this dissertation thoroughly explored three of the many hypotheses explaining the evolution of collective animal behavior. So far, these projects have expanded the theory surrounding the evolutionary origins of collective behavior by producing several testable hypotheses. For example in Chapter 3.4, I discovered a simple vision-based movement algorithm that prey could use maintain cohesive swarming behavior, which contrasts with the complex algorithms that are commonly used to simulate swarming behavior in silico [85, 73]. This finding provides a much simpler movement algorithm to explain the complex swarming behavior found in nature, which can be validated in observational field studies such as [113]. Similarly in Chapter 4.5, I hypothesized that the interaction enabled by predator confusion between the predator's view angle and prey swarming behavior should select for predators with a focused visual system. If this is the case, we would expect to observe potentially unique traits and mechanisms that focus the predator's visual system in natural predators that hunt swarming prey. Finally in Chapter 5, I found that reproductive strategy plays a significant role in the evolution of cooperative group foraging, namely by showing that semelparous species will be strongly selected to forage in groups. This finding represents an altogether new discovery explaining the evolution of cooperative group foraging that could be corroborated by a meta-analysis of existing species linking the species' proclivity to cooperatively forage in groups to its life history.

During the course of this research, I came upon several surprising results that defied my

expectations. Perhaps the most notable unexpected finding was the simplicity of the evolved prey swarming mechanisms across every embodied swarming experiment that I performed. Since the majority of researchers studying swarm control mechanisms suggest that natural prey must be following some form of complex Boids rules (separation, alignment, cohesion), I expected my simulated prey to evolve some form of the Boids rules. Instead, my simulated prey evolved a simple “follow the prey in front of you” rule that results in emergent swarming behavior when applied over a group of locally-interacting prey without the need for global information about the swarm—and only minimal local information about what is in front of the prey. Although part of the simplicity of the movement algorithm may result from the fact that the prey cannot collide with one another, this unexpected finding is cause for collective behavior researchers to reevaluate and perhaps simplify the standard models that explain swarming behavior in nature.

Another surprising finding from this research was the significant impact of predator attack mode (i.e., how predators choose to attack prey) on the evolution of swarming behavior. Most research prior to the work in this dissertation had assumed that predator attack mode did not play an important role, and often simply assumed that the predator attacks prey at random. In this dissertation, I reevaluated this assumption and discovered that predators that consistently attack prey on the outside of the swarm exhibit a much stronger selection pressure to swarm than predators that attack randomly. This finding suggests that it is not safe to assume that the predator attack mode plays an unimportant role in the evolution of collective behavior, and highlights the importance of exploring the role of complex predator attack modes in models exploring the evolution of collective behavior [74, 114].

Furthermore, this work has provided insight into possible applications in Computer Science and biomimetic solutions to problems in particle swarm optimization and swarm

robotics. For example in Chapter 3.4, the same simple vision-based control algorithm can be used as a control algorithm in swarm robotics experiments. Provided that swarm robotics experiments are often limited to simple, inexpensive robots with minimal sensors due to manufacturing cost and battery life, simple control algorithms that produce emergent swarming behavior will be necessary to advance the field [72]. In Chapters 3.2.4 and 4.5, I elaborated upon the role of predator attack mode and predator confusion in the coevolutionary dynamics between predator and prey populations. Provided that a growing field of particle swarm optimization seeks to harness predator-prey coevolution during the optimization process [71], it will be critical to understand the core predator-prey coevolutionary theory underlying these optimization algorithms. Finally, it is important to note that the core of this research aims at understanding how it is possible to get a heterogeneous group of independent agents to cooperate toward a common goal, even if cooperation entails sacrificing the performance of some of the individuals in the group. There are many parallels to this problem in Computer Science, for example, it is common to see improved performance in Machine Learning classification problems by creating an ensemble of classifiers to “work together” toward better classification performance [115]. Although it is not yet common to automatically create heterogeneous ensembles of classifiers, the work presented in this dissertation will be informative for this potentially fruitful line of Machine Learning research.

Of course, this dissertation opens many new avenues of research directly following the work presented here. In Chapter 3.4, the evolved prey swarming mechanism assumed that the prey do not collide with one another. It would be instructive to follow up on this work by implementing collisions for the prey and observing the resultant behavior: Do the prey evolve a more Boids-like control algorithm, or is a simple follow-the-prey-in-front-of-you mechanism still sufficient? In Chapter 4.5, I found that the predators and prey enter a seemingly-endless

coevolutionary cycle when both prey behavior and the predator visual system are allowed to coevolve. Another fascinating venue of research would be to explore possible mechanisms for visual predators to secure their “evolutionary victory” by evolving to effectively hunt both swarming and dispersive prey—for example, by evolving a complex visual system that we often see in visual predators in nature that provides both coarse, broad vision for searching as well as narrow, focused vision for tracking prey. In Chapter 5, I discovered that the prey’s reproductive strategy plays an important role in whether it will evolve to cooperatively forage in groups. However, I only explored the two extremes of reproductive strategy—reproducing continuously and reproducing only once near the end of their lifetime—and there is an entire continuum of reproductive strategies in between left to explore.

Finally, there are many more hypothesized benefits of collective behavior that remain to be explored in future work, such as improved locomotion efficiency [18] and the feasibility of collective cognition [1]. Once all of these hypothesized benefits have been explored in isolation, it will then be possible to combine these hypothesized benefits into hybrid experiments where we can answer questions such as, “In the presence of the predator confusion effect, does collective vigilance play an important role in the evolution of collective behavior?” These hybrid experiments will bring us closer to simulating real biological systems and understanding how and why prey evolve to live in groups. By bringing us closer to understanding nature, this line of research will establish a solid basis of evolutionary theory surrounding collective behavior for researchers to draw upon, both when studying collective behavior in nature and when harnessing collective behavior in robotics and optimization problems.

REFERENCES

REFERENCES

- [1] I. Couzin, “Collective cognition in animal groups,” *Trends Cogn. Sci.*, vol. 13, pp. 36–43, 2009.
- [2] C. Feare, *The Starling*. Oxford University Press, USA, 1984.
- [3] C. Hemelrijk and H. Hildenbrandt, “Some causes of the variable shape of flocks of birds,” *PLoS ONE*, vol. 6, p. e22479, 2011.
- [4] F. Dyer, “The biology of the dance language,” *Annu. Rev. of Entomol.*, vol. 47, pp. 917–49, 2002.
- [5] L. Hall-Stoodley, J. Costerton, and P. Stoodley, “Bacterial biofilms: From the natural environment to infectious diseases,” *Nat. Rev. Micro.*, vol. 2, pp. 95–108, 2004.
- [6] J. Krause and G. Ruxton, *Living in groups*. Oxford University Press, USA, 2002.
- [7] J. Jeschke and R. Tollrian, “Prey swarming: Which predators become confused and why?,” *Anim. Behav.*, vol. 74, pp. 387–393, 2007.
- [8] G. Beauchamp, “Reduced flocking by birds on islands with relaxed predation,” *P. Roy. Soc. B*, vol. 271, pp. 1039–42, 2004.
- [9] J. Parrish, “Complexity, pattern, and evolutionary trade-offs in animal aggregation,” *Science*, vol. 284, p. 99, 1999.
- [10] M. K. McClintock, J. J. Anisko, and N. T. Adler, “Group mating among norway rats ii. the social dynamics of copulation: Competition, cooperation, and mate choice,” *Anim. Behav.*, vol. 30, pp. 410 – 425, 1982.
- [11] B. Yuval, J. Wekesa, and R. Washino, “Effect of body size on swarming behavior and mating success of male *Anopheles freeborni* (Diptera: Culicidae),” *J. Insect Behav.*, vol. 6, pp. 333–342, 1993.
- [12] A. Diabaté, A. Yaro, A. Dao, M. Diallo, D. Huestis, and T. Lehmann, “Spatial distribution and male mating success of *Anopheles gambiae* swarms,” *BMC Evol. Biol.*, vol. 11, p. 184, 2011.

- [13] G. D. Ruxton and G. Beauchamp, “The application of genetic algorithms in behavioural ecology, illustrated with a model of anti-predator vigilance,” *J. Theor. Bio.*, vol. 250, pp. 435–448, 2008.
- [14] R. Pulliam and T. Caraco, “Living in groups: Is there an optimal group size?,” in *Behavioral Ecology: An Evolutionary Approach* (J. Krebs and N. Davies, eds.), Mass.: Sinauer, 2nd ed., 1984.
- [15] G. Powell, “Experimental analysis of the social value of flocking by starlings (*sturnus vulgaris*) in relation to predation and foraging,” *Anim. Behav.*, vol. 22, pp. 501–505, 1974.
- [16] S. L. Lima, “Back to the basics of anti-predatory vigilance: the group-size effect,” *Anim. Behav.*, vol. 49, pp. 11–20, 1995.
- [17] G. Roberts, “Why individual vigilance declines as group size increases,” *Anim. Behav.*, vol. 51, pp. 1077–1086, 1996.
- [18] F. E. Fish, “Kinematics of ducklings swimming in formation: Consequences of position,” *J. Exp. Zool.*, vol. 273, pp. 1–11, 1995.
- [19] J. Treherne and W. Foster, “Group transmission of predator avoidance behaviour in a marine insect: The Trafalgar effect,” *Anim. Behav.*, vol. 29, pp. 911–917, 1981.
- [20] R. Kenward, “Hawks and doves: Factors affecting success and selection in Goshawk attacks on Woodpigeons,” *J. Anim. Ecol.*, vol. 47, pp. 449–460, 1978.
- [21] M. Treisman, “Predation and the evolution of gregariousness. I. Models for concealment and evasion,” *Anim. Behav.*, vol. 23, pp. 779–800, 1975.
- [22] H. Pulliam, “On the advantages of flocking,” *J. Theor. Bio.*, vol. 38, pp. 419–22, 1973.
- [23] A. Inman and J. Krebs, “Predation and group living,” *Trends Ecol. Evol.*, vol. 2, pp. 31–32, 1987.
- [24] C. Ioannou, V. Guttal, and I. Couzin, “Predatory fish select for coordinated collective motion in virtual prey,” *Science*, vol. 337, pp. 1212–1215, 2012.
- [25] J. Treherne and W. Foster, “Group size and anti-predator strategies in a marine insect,” *Anim. Behav.*, vol. 30, pp. 536–542, 1982.

- [26] W. Foster and J. Treherne, “Evidence for the dilution effect in the selfish herd from fish predation on a marine insect,” *Nature*, vol. 293, pp. 466–467, 1981.
- [27] W. Hamilton, “Geometry for the selfish herd,” *J. Theor. Bio.*, vol. 31, pp. 295–311, 1971.
- [28] B. Bertram in *Behavioral Ecology: An Evolutionary Approach* (J. R. Krebs and N. Davies, eds.), pp. 64–96, Blackwell, Oxford, 1978.
- [29] C. Ioannou, C. Tosh, L. Neville, and J. Krause, “The confusion effect—from neural networks to reduced predation risk,” *Behav. Ecol.*, vol. 19, pp. 126–130, 2008.
- [30] D. Krakauer, “Groups confuse predators by exploiting perceptual bottlenecks: A connectionist model of the confusion effect,” *Behav. Ecol. Sociobiol.*, vol. 36, pp. 421–429, 1995.
- [31] R. James, P. Bennett, and J. Krause, “Geometry for mutualistic and selfish herds: The limited domain of danger,” *J. Theoret. Biol.*, vol. 228, pp. 107–13, 2004.
- [32] G. Beauchamp, “Vigilance in a selfish herd,” *Anim. Behav.*, vol. 73, pp. 445–451, 2007.
- [33] L. Morrell, G. Ruxton, and R. James, “Spatial positioning in the selfish herd,” *Behav. Ecol.*, vol. 22, pp. 16–22, 2010.
- [34] H. Kunz and C. K. Hemelrijk, “Artificial fish schools: Collective effects of school size, body size, and body form,” *Artif. Life*, vol. 9, pp. 237–253, 2003.
- [35] C. K. Hemelrijk and H. Kunz, “Density distribution and size sorting in fish schools: an individual-based model,” *Behav. Ecol.*, vol. 16, pp. 178–187, 2005.
- [36] J. Quinn and W. Cresswell, “Testing domains of danger in the selfish herd: Sparrowhawks target widely spaced redshanks in flocks,” *P. Roy. Soc. B*, vol. 273, pp. 2521–6, 2006.
- [37] S. Viscido, M. Millery, and D. Wethey, “The dilemma of the selfish herd: The search for a realistic movement rule,” *J. Theor. Bio.*, vol. 217, pp. 183–194, 2002.
- [38] T. Morton, J. Haefner, V. Nugala, R. Decino, and L. Mendes, “The selfish herd revisited: Do simple movement rules reduce relative predation risk?,” *J. Theor. Bio.*, vol. 167, pp. 73–79, 1994.

- [39] T. Reluga and S. Viscido, “Simulated evolution of selfish herd behavior,” *J. Theor. Bio.*, vol. 234, pp. 213–25, 2005.
- [40] A. Wood and G. Ackland, “Evolving the selfish herd: Emergence of distinct aggregating strategies in an individual-based model,” *P. Roy. Soc. B*, vol. 274, pp. 1637–42, 2007.
- [41] D. Humphries and P. Driver, “Protean defence by prey animals,” *Oecologia*, vol. 5, pp. 285–302, 1970.
- [42] C. S. Holling, “Some characteristics of simple types of predation and parasitism,” *Can. Entomol.*, vol. 91, pp. 385–398, 1959.
- [43] J. M. Jeschke and R. Tollrian, “Effects of predator confusion on functional responses,” *Oikos*, vol. 111, pp. 547–555, 2005.
- [44] N. G. Hairston, S. P. Ellner, M. A. Geber, T. Yoshida, and J. A. Fox, “Rapid evolution and the convergence of ecological and evolutionary time,” *Ecol. Lett.*, vol. 8, pp. 1114–1127, 2005.
- [45] T. Caraco, “Time budgeting and group size: A test of theory,” *Ecology*, vol. 60, pp. 618–627, 1979.
- [46] A. J. W. Ward, J. E. Herbert-Read, D. J. T. Sumpter, and J. Krause, “Fast and accurate decisions through collective vigilance in fish shoals,” *Proceedings of the National Academy of Sciences*, vol. 108, pp. 2312–2315, 2011.
- [47] B. C. Bertram, “Vigilance and group size in ostriches,” *Animal Behaviour*, vol. 28, pp. 278–286, 1980.
- [48] G. Beauchamp, “What is the magnitude of the group-size effect on vigilance?,” *Behavioral Ecology*, vol. 19, pp. 1361–1368, 2008.
- [49] H. R. Pulliam, G. H. Pyke, and T. Caraco, “The scanning behavior of juncos: A game-theoretical approach,” *Journal of Theoretical Biology*, vol. 95, pp. 89–103, 1982.
- [50] J. M. McNamara and A. I. Houston, “Evolutionarily stable levels of vigilance as a function of group size,” *Animal Behaviour*, vol. 43, pp. 641–658, 1992.
- [51] R. Lenski, C. Ofria, R. Pennock, and C. Adami, “The evolutionary origin of complex features,” *Nature*, vol. 423, pp. 139–44, 2003.

- [52] C. O. Wilke, J. L. Wang, C. Ofria, L. R. E., and C. Adami, “Evolution of digital organisms at high mutation rates leads to survival of the flattest,” *Nature*, vol. 412, pp. 331–333, 2001.
- [53] R. Lukeman, Y.-X. Li, and L. Edelstein-Keshet, “Inferring individual rules from collective behavior,” *P. Nat. Acad. Sci. USA*, vol. 107, pp. 12576–12580, 2010.
- [54] I. D. Couzin, J. Krause, N. R. Franks, and S. A. Levin, “Effective leadership and decision-making in animal groups on the move,” *Nature*, vol. 433, pp. 513–516, 2005.
- [55] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks, “Collective memory and spatial sorting in animal groups,” *J. Theor. Bio.*, vol. 218, pp. 1–11, 2002.
- [56] C. W. Clark and M. Mangel, “The evolutionary advantages of group foraging,” *Theoretical Population Biology*, vol. 30, pp. 45–75, 1986.
- [57] E. Katsnelson, U. Motro, M. W. Feldman, and A. Lotem, “Evolution of learned strategy choice in a frequency-dependent game,” *Proceedings of the Royal Society B*, 2011.
- [58] C. Adami, “Boldly going beyond mathematics,” *Science*, vol. 338, pp. 1421–1422, 2012.
- [59] S. Hamblin, “On the practical usage of genetic algorithms in ecology and evolution,” *Methods in Ecology and Evolution*, vol. 4, pp. 184–194, 2013.
- [60] C. Tosh, “Which conditions promote negative density dependent selection on prey aggregations?,” *J. Theor. Bio.*, vol. 281, pp. 24–30, 2011.
- [61] R. Biswas, C. Ofria, D. Bryson, and A. Wagner, “Causes vs benefits in the evolution of prey grouping,” in *Proceedings of the International Conference on the Simulation and Synthesis of Living Systems (ALIFE)*, pp. 641–648, 2014.
- [62] H. Kunz, T. Züblin, and C. Hemelrijk, “On prey grouping and predator confusion in artificial fish schools,” in *Proceedings of the International Conference on the Simulation and Synthesis of Living Systems (ALIFE)*, pp. 365–371, 2006.
- [63] A. J. Wood, “Strategy selection under predation; evolutionary analysis of the emergence of cohesive aggregations,” *J. Theor. Bio.*, vol. 264, pp. 1102–1110, 2010.
- [64] L. Spector, J. Klein, C. Perry, and M. Feinstein, “Emergence of collective behavior in evolving populations of flying agents,” in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 61–73, 2003.

- [65] C. Ward, F. Gobet, and G. Kendall, “Evolving collective behavior in an artificial ecology,” *J. Artificial Life*, vol. 7, pp. 191–209, 2001.
- [66] D. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*. Boston, MA: Addison-Wesley, 1st ed., 1989.
- [67] R. Poli, “Analysis of the publications on the applications of particle swarm optimisation,” *J. Artif. Evol. Appl.*, vol. 2008, pp. 4:1–4:10, 2008.
- [68] B. Xue, M. Zhang, and W. Browne, “Multi-objective particle swarm optimisation (PSO) for feature selection,” in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 81–88, 2012.
- [69] E. Vellasques, R. Sabourin, and E. Granger, “Gaussian mixture modeling for dynamic particle swarm optimization of recurrent problems,” in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 73–80, 2012.
- [70] Y. Marinakis and M. Marinaki, “A honey bees mating optimization algorithm for the open vehicle routing problem,” in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 101–108, 2011.
- [71] A. Silva, A. Neves, and E. Costa, “An empirical comparison of particle swarm and predator prey optimisation,” in *Proceedings of the Irish Conference on Artificial Intelligence and Cognitive Science (AICS)*, pp. 103–110, 2002.
- [72] E. Şahin, “Swarm robotics: From sources of inspiration to domains of application,” in *Proceedings of the International Conference on Swarm Robotics (SAB)*, pp. 10–20, 2004.
- [73] C. Reynolds, “Flocks, herds and schools: A distributed behavioral model,” in *Proceedings of the Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 25–34, 1987.
- [74] R. S. Olson, D. B. Knoester, and C. Adami, “Critical interplay between density-dependent predation and evolution of the selfish herd,” in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 247–254, 2013.
- [75] A. Eiben and J. Smith, *Introduction to Evolutionary Computing*. Springer-Verlag, 2003.

- [76] G. Martin, “The eye of a passeriform bird, the European starling (*Sturnus vulgaris*): Eye movement amplitude, visual fields and schematic optics,” *J. Comp. Physiol. A*, vol. 159, pp. 545–557, 1986.
- [77] P. Moran, *The statistical processes of evolutionary theory*. Clarendon Press, 1962.
- [78] C. Huepe and M. Aldana, “New tools for characterizing swarming systems: A comparison of minimal models,” *Physica A*, vol. 387, pp. 2809–2822, 2008.
- [79] R. Olson, A. Hintze, F. Dyer, D. Knoester, and C. Adami, “Predator confusion is sufficient to evolve swarming behaviour,” *J. Roy. Soc. Interface*, vol. 10, p. 20130305, 2013.
- [80] S. V. Viscido, M. Miller, and D. S. Wethey, “The response of a selfish herd to an attack from outside the group perimeter,” *J. Theor. Biol.*, vol. 208, pp. 315–328, 2001.
- [81] J. Goldbogen, J. Calambokidis, E. Oleson, J. Potvin, N. Pyenson, G. Schorr, and R. Shadwick, “Mechanics, hydrodynamics and energetics of blue whale lunge feeding: Efficiency dependence on krill density,” *J. Exp. Biol.*, vol. 214, pp. 131–146, 2011.
- [82] M. Mori and D. Butterworth, “A first step towards modelling the krill-predator dynamics of the Antarctic ecosystem,” *CCAMLR Sci.*, vol. 13, pp. 217–277, 2006.
- [83] R. S. Olson, D. B. Knoester, and C. Adami, “Evolution of swarming behavior is shaped by how predators attack.” arXiv e-print. <http://arxiv.org/abs/1310.6012>, 2014.
- [84] R. Rudell, “Multiple-valued logic minimization for PLA synthesis,” Tech. Rep. UCB/ERL M86/65, EECS Department, University of California, Berkeley, 1986.
- [85] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet, “Novel type of phase transition in a system of self-driven particles,” *Phys. Rev. Lett.*, vol. 75, pp. 1226–1229, 1995.
- [86] J. Edlund, N. Chaumont, A. Hintze, C. Koch, G. Tononi, and C. Adami, “Integrated information increases with fitness in the evolution of animats,” *PLoS Comput. Biol.*, vol. 7, p. e1002236, 2011.
- [87] C. R. Tosh, A. L. Jackson, and G. D. Ruxton, “The confusion effect in predatory neural networks,” *Am. Nat.*, vol. 167, pp. E52–E65, 2006.

- [88] B. T. Hirsch and L. J. Morrell, “Measuring marginal predation in animal groups,” *Behav. Ecol.*, vol. 22, pp. 648–656, 2011.
- [89] W. Romey, A. Walston, and P. Watt, “Do 3-d predators attack the margins of 2-d selfish herds?,” *Behav. Ecol.*, vol. 19, pp. 74–78, 2008.
- [90] W. W. Murdoch, “The functional response of predators,” *J Appl. Ecol.*, vol. 10, pp. 335–342, 1973.
- [91] M. E. Solomon, “The natural control of animal populations,” *J Anim. Ecol.*, vol. 18, pp. 1–35, 1949.
- [92] J. M. Jeschke, M. Kopp, and R. Tollrian, “Consumer-food systems: why type i functional responses are exclusive to filter feeders,” *Biol. Rev.*, vol. 79, pp. 337–349, 2004.
- [93] V. Tucker, “The deep fovea, sideways vision and spiral flight paths in raptors,” *Journal of Experimental Biology*, vol. 203, pp. 3745–3754, 2000.
- [94] W. T. Wcislo and B. N. Danforth, “Secondarily solitary: the evolutionary loss of social behavior,” *Trends Ecol. Evol.*, vol. 12, pp. 468–474, 1997.
- [95] S. B. Ale and J. S. Brown, “The contingencies of group size and vigilance,” *Evolutionary Ecology Research*, vol. 9, p. 1263, 2007.
- [96] P. Haley, R. Olson, F. Dyer, and C. Adami, “Exploring conditions that select for the evolution of cooperative group foraging,” in *Proceedings of the International Conference on the Simulation and Synthesis of Living Systems (ALIFE)*, pp. 310–311, 2014.
- [97] R. Pfeiffer and J. C. Bongard, *How the body shapes the way we think*. MIT Press Cambridge, 2006.
- [98] G. Beauchamp, “Group-size effects on vigilance: a search for mechanisms,” *Behavioural Processes*, vol. 63, pp. 141–145, 2003.
- [99] B. Fairbanks and F. S. Dobson, “Mechanisms of the group-size effect on vigilance in columbian ground squirrels: dilution versus detection,” *Animal Behaviour*, vol. 73, pp. 115–123, 2007.
- [100] S. L. Lima, “Collective detection of predatory attack by social foragers: fraught with ambiguity?,” *Animal Behaviour*, vol. 50, pp. 1097–1108, 1995.

- [101] L. Marstaller, A. Hintze, and C. Adami, “The evolution of representation in simple cognitive networks,” *Neural Computation*, vol. 25, no. 8, pp. 2079–2107, 2013.
- [102] S. Hamblin and P. L. Hurd, “Genetic algorithms and non-ESS solutions to game theory models,” *Animal Behaviour*, vol. 74, pp. 1005–1018, 2007.
- [103] W. Hamilton, “The genetical evolution of social behaviour. I-II,” *Journal of Theoretical Biology*, vol. 7, pp. 1–52, 1964.
- [104] A. Hintze, R. S. Olson, C. Adami, and R. Hertwig, “Risk sensitivity as an evolutionary adaptation,” *Scientific Reports*, vol. 5, Feb. 2015.
- [105] M. A. Elgar, “Predator vigilance and group size in mammals and birds: a critical review of the empirical evidence,” *Biological Reviews*, vol. 64, pp. 13–33, 1989.
- [106] A. Sansom, W. Cresswell, J. Minderman, and J. Lind, “Vigilance benefits and competition costs in groups: do individual redshanks gain an overall foraging benefit?,” *Animal Behaviour*, vol. 75, pp. 1869–1875, 2008.
- [107] D. J. Rankin, K. Bargum, and H. Kokko, “The tragedy of the commons in evolutionary biology,” *Trends in Ecology & Evolution*, vol. 22, pp. 643–651, 2007.
- [108] T. Wenseleers and F. L. Ratnieks, “Tragedy of the commons in *Melipona* bees,” *Proceedings of the Royal Society B*, vol. 271, pp. S310–S312, 2004.
- [109] E. Goodale, S. W. Kotagama, and R. Holberton, “Alarm calling in Sri Lankan mixed-species bird flocks,” *The Auk*, vol. 122, pp. 108–120, 2005.
- [110] H. Sridhar, G. Beauchamp, and K. Shanker, “Why do birds participate in mixed-species foraging flocks? A large-scale synthesis,” *Animal Behaviour*, vol. 78, pp. 337–347, 2009.
- [111] P. A. Abrams, “Life history and the relationship between food availability and foraging effort,” *Ecology*, pp. 1242–1252, 1991.
- [112] J. Krause and J.-G. J. Godin, “Influence of prey foraging posture on flight behavior and predation risk: predators take advantage of unwary prey,” *Behavioral Ecology*, vol. 7, pp. 264–271, 1996.

- [113] M. Ballerini et al., “Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study,” *P. Nat. Acad. Sci. USA*, vol. 105, pp. 1232–7, 2008.
- [114] J. Demsar, C. K. Hemelrijk, H. Hildenbrandt, and I. L. Bajec, “Simulating predator attacks on schools: Evolving composite tactics,” *Ecological Modelling*, vol. 304, pp. 22–33, 2015.
- [115] R. E. Schapire, “The boosting approach to machine learning: An overview,” in *Non-linear Estimation and Classification* (D. Denison, M. Hansen, C. Holmes, B. Mallick, and B. Yu, eds.), vol. 171 of *Lecture Notes in Statistics*, pp. 149–171, Springer New York, 2003.