

Evolution of Swarming Behavior Is Shaped by How Predators Attack

Randal S. Olson^{*,**}

David B. Knoester[†]

Christoph Adami[‡]

Michigan State University

Abstract Animal grouping behaviors have been widely studied due to their implications for understanding social intelligence, collective cognition, and potential applications in engineering, artificial intelligence, and robotics. An important biological aspect of these studies is discerning which selection pressures favor the evolution of grouping behavior. In the past decade, researchers have begun using evolutionary computation to study the evolutionary effects of these selection pressures in predator-prey models. The selfish herd hypothesis states that concentrated groups arise because prey selfishly attempt to place their conspecifics between themselves and the predator, thus causing an endless cycle of movement toward the center of the group. Using an evolutionary model of a predator-prey system, we show that how predators attack is critical to the evolution of the selfish herd. Following this discovery, we show that density-dependent predation provides an abstraction of Hamilton's original formulation of *domains of danger*. Finally, we verify that density-dependent predation provides a sufficient selective advantage for prey to evolve the selfish herd in response to predation by coevolving predators. Thus, our work corroborates Hamilton's selfish herd hypothesis in a digital evolutionary model, refines the assumptions of the selfish herd hypothesis, and generalizes the domain of danger concept to density-dependent predation.

Keywords

Group behavior, selfish herd theory, predator attack mode, density-dependent predation, predator-prey coevolution, evolutionary algorithm, digital evolutionary model

1 Introduction

Over the past century, researchers have devoted considerable effort to studying animal grouping behavior, due to its important implications for social intelligence, collective cognition, and potential applications in engineering, artificial intelligence, and robotics [5]. Indeed, grouping behaviors are pervasive across all forms of life. For example, European starlings (*Sturnus vulgaris*) are known to form murmurations of millions of birds, which perform awe-inspiring displays of coordinated movement [12, 19]. Western honeybees (*Apis mellifera*) communicate the location of food and nest sites to other bees in their group via a complex dance language [9]. Even relatively simple bacteria exhibit

* Contact author.

** Department of Computer Science and Engineering, and BEACON Center for the Study of Evolution in Action, Michigan State University, East Lansing, MI 48824. E-mail: rso@randalolson.com

† Department of Microbiology and Molecular Genetics, and BEACON Center for the Study of Evolution in Action, Michigan State University, East Lansing, MI 48824. E-mail: dk@msu.edu

‡ Department of Microbiology and Molecular Genetics, Department of Physics and Astronomy, and BEACON Center for the Study of Evolution in Action, Michigan State University, East Lansing, MI 48824. E-mail: adami@msu.edu

grouping behavior, such as *Escherichia coli* forming biofilms that allow their group to survive in hostile environments [17].

Swarming is one example of grouping behavior, where animals coordinate their movement with conspecifics to maintain a cohesive group. Although swarmlike groups could arise by chance—as with little egrets (*Egretta garzetta*) pursuing a common resource in water pools [28]—typically swarms are maintained via behavioral mechanisms that ensure group cohesion [2]. As with many traits, swarming behavior entails a variety of fitness costs, such as increased risk of predation and the requisite sharing of resources with the group [44]. With this fact in mind, significant effort has been dedicated to understanding the compensating benefits that grouping behavior provides [30]. Many such benefits of grouping behavior have been proposed; for example, swarming may improve mating success [68, 8], increase foraging efficiency [47], or enable the group to solve problems that would be impossible to solve individually [5]. Furthermore, swarming behaviors are hypothesized to protect group members from predators in several ways. For example, swarming can improve group vigilance [57, 27, 59, 46], reduce the chance of being encountered by predators [59, 22], dilute an individual’s risk of being attacked [23, 58, 13, 18], enable an active defense against predators [30], or reduce predator attack efficiency by confusing the predator [24, 26, 29].

Unfortunately, many swarming animals 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 swarming behavior as an evolutionary response, let alone study the behaviors as they evolve [26, 3]. In this article, we use a digital model of predator-prey coevolution to explore Hamilton’s selfish herd hypothesis [18]. 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 article expands on earlier work [41] 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.

2 Related Work

Hamilton’s original formulation of the selfish herd hypothesis introduced the concept of *domains of danger* (DODs, Figure 1), which served as a method to visualize the likelihood for a prey inside a




Figure 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.

group to be attacked by a predator [18]. 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 [25], investigating the effects of prey vigilance [4], considering the initial spatial positioning of prey when the group is attacked [38], exploring the role of prey body characteristics in shaping herd characteristics [31, 20], and even confirming Hamilton's predictions in biological systems [48].

Additional studies have focused on the movement rules that prey in a selfish herd follow to minimize their DOD [62]. This line of work began by demonstrating that the simple movement rules proposed by Hamilton reduce predation risk for prey inside the group [39], then opened some parameters of the movement rules to evolution in an attempt to discover a more biologically plausible set of movement rules [49, 65]. 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. This article builds on this 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.

That said, there are many other potential causes of swarming behavior that we do not address in this study. For example, some studies have investigated the evolution of predator behavior in response to prey density [56] and the role of relative predator and prey speeds in the evolution of grouping behavior [66], elaborated upon the interaction between ecology and the evolution of grouping behavior [55, 64], and explored the role of group vigilance (i.e., the "many eyes" hypothesis) in the evolution of grouping behavior [16, 43]. Two recent studies have explored the coevolution of predator and prey behavior in the presence of the predator confusion effect [40, 32], and found that the predator confusion effect is sufficient to select for the evolution of swarming behavior in the absence of any other group benefits. Further, it has been shown that predators can adapt composite tactics to improve their efficacy against swarms of prey, which greatly reduces the defensive benefit of swarming [7]. It is therefore necessary to keep in mind that while this study investigates the evolution of swarming behavior according to Hamilton's selfish herd hypothesis, in natural populations there are often many interacting benefits and costs of swarming behavior that must be taken into account [30].

Of course, prior to this work there had been considerable research exploring the (co)evolution of animal behavior in agent-based models. Craig Reynolds' work on the evolution of prey behavior in response to simulated predation was one of the earliest articles to demonstrate that predation can directly select for aggregative behavior in prey populations [51]. Similarly, Karl Sims' work in 1994 established a new paradigm for the coevolution of agent behavior and morphology [54], which is still an object of intense study to this day. The work in this article seeks to provide a stronger biological grounding to these kinds of agent-based evolution experiments, in particular by implementing specific selection pressures on the prey that are discussed in the animal behavior literature.

More broadly, in the past decade researchers have focused on the application of locally interacting swarming agents to optimization problems, called *particle swarm optimization* (PSO) [45]. PSO applications range from feature selection for classifiers [67] to video processing [60] to open vehicle routing [34]. A related technique within PSO seeks to combine PSO with coevolving predator and prey solutions to avoid local minima [53], and has proven effective in other evolutionary computation domains as well [1]. Researchers have even sought to harness the collective problem-solving power of swarming agents to design robust autonomous robotic swarms [6]. Thus, elaborations on the foundations of animal grouping behavior have the potential to improve our ability to solve engineering problems.

3 Methods

To study the evolution of the selfish herd, we developed an agent-based model in which agents interact in a continuous, toroidal virtual environment (736×736 virtual meters), shown in Figure 2.




Figure 2. 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.

At the beginning of each simulation, we 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) [10]. We evolve the agent MNs with a genetic algorithm (GA) [11, 14] under varying selection regimes, which will be described in more detail below.

During each simulation time step, all agents read information from their sensors and take action (i.e., move) based on their effectors. In our first set of treatments, we simulate an ideal, disembodied predator by periodically removing prey agents from the environment and marking them as consumed—for example, when they are on the outermost edges of the group. Subsequent treatments introduce an embodied, coevolving predator agent that is controlled by its own MN. The data¹ and source code² from these experiments are available online for further analysis. In the remainder of this section, we describe the sensory-motor architecture of individual agents and present details related to the function and encoding of MNs.

3.1 Agent Model

Figure 3 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), pixelated retina covering its entire 360° visual field. Its retina 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 [35]. Regardless of the number of agents present in a single retina slice, the prey agent only knows whether a conspecific or a predator resides within that slice, but not how many. For example, in Figure 3, 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 1. Prey agents turn in 8° increments and move 1 virtual meter each time step.

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

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




Figure 3. An illustration of the agents in the model. Light gray triangles are prey agents, and dark gray triangles are predator agents. The agents have a 360° limited-distance retina (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 on 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.

In our coevolution experiments, the predator agents can detect only nearby prey agents, using a limited-distance (200 virtual meters), pixelated retina covering its frontal 180° that works just like the prey agent’s retina (Figure 3). Similarly to the prey agents, predators make decisions about how to move next using their MN, as shown in Table 1, but move 3 times faster than the prey agents and turn correspondingly slower (6° per simulation time step) due to their higher speed. This dramatically faster predator movement speed is meant to represent predators that perform rapid attacks on groups of prey, such as a peregrine falcon dive-bombing a swarm of starlings. Finally, if a predator agent moves within 5 virtual meters of a prey agent that is anywhere within its retina, the predator agent attempts an attack on the prey agent. If the attempt is successful, we remove the prey agent from the simulation and mark it as consumed.

3.2 Markov Networks

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

Table 1. Possible actions encoded by the agent’s output.

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

Notes. 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.




Figure 4. An example Markov network (MN) with four input states (white circles labeled 0–3), two hidden states (light gray circles labeled 4 and 5), two output states (dark gray 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.

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. Although we specifically use MNs as the artificial brains in these experiments, other artificial brains such as artificial neural networks, genetic programming, or many other evolvable substrates that can produce agent-based behavior based on sensory inputs could also be used in these experiments.

Every simulation time step, the MNs receive input via sensors (e.g., visual retina), 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). We note that MN states are binary and only assume a value of 0 or 1. When we 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.

3.2.1 How Markov Networks Function

When we embed an agent into the simulation environment, we provide sensory inputs from its retina into its MN every simulation step (Figure 3, labeled “retina” and “Markov network,” respectively). Once we provide a MN with its inputs, we 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 4, 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 $t + 1$.

In a MN, states are updated by MGs, which function similarly to digital logic gates, such as AND 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 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. An example MG that could be used to control a prey agent that avoids nearby predator agents.

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

Notes. 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. 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.

The MGs in this model can receive input from a maximum of four states, and write into a maximum of four 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 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.

3.2.2 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 5 depicts an example genome. After the start codon, the next two bytes describe the numbers of inputs (N_{in}) and outputs (N_{out}) used in this MG, where each $N = 1 + (\text{byte mod } N_{max})$. 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 function $(\text{byte mod } N_{states})$, 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 $2^{N_{in}+N_{out}}$ bytes of the gene determine the probabilities in the logic table.

The maximum number of states allowed and which states are used as inputs and outputs are specified as constants by the user. In these experiments, we provided 64 states for the MNs to work with: 24 sensory inputs, 2 outputs for the actuators, and 38 hidden states for optional internal




Figure 5. Example circular byte strings encoding the two Markov gates (MGs) in Figure 4, 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 numbers of input and output states used by the MG (light gray blocks), and the following eight bytes encode which states are used as input (medium gray blocks) and output (darker gray blocks). The remaining bytes in the string encode the probabilities of the MG’s logic table (darkest gray blocks).

Table 3. Genetic algorithm and experiment settings.

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

computations. 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 3. 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, to allow for a succinct reconstruction of the line of descent of the population [33] (described further below), which is a useful tool in evolutionary studies that we harness in this article.

4 Artificial Predation

In our first set of experiments, we observe the evolution of prey behavior in response to various forms of artificial predation. This enables us to experimentally control the specific modes of predation and observe their effect on the evolution of the selfish herd. We evolve the prey genomes with a GA with the settings described in Table 3. We begin the evolutionary process by seeding the prey genome pool with a set of randomly generated ancestor genomes of length 5,000 with four random MGs. Following this, we 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' life span; hence each agent has a potential life span of 1,000 time steps. We assign each prey genome an individual fitness according to how long its corresponding prey agent survived, following the equation

$$W_{\text{prey}} = T \quad (1)$$

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, we perform fitness-proportional selection on the population of genomes via a Moran process [36], increment the generation counter, and

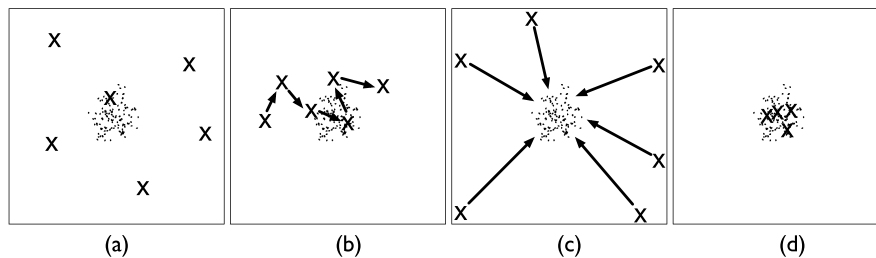


Figure 6. An illustration of the four artificial predator attack modes: (a) random attacks, (b) random walk attacks, (c) outside attacks, and (d) high-density area attacks.

repeat the evaluation process on the new population of genomes until the final generation (40,000) is reached.

In all cases, we give the prey an initial 250 simulation time steps without predation to move around, so that prey starting on the exterior 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, we apply artificial predation every 4 simulation time steps to simulate an ideal predator attacking the group. Artificial predators succeed in their attacks every time. We 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. We 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, we characterize the grouping behavior by measuring the *swarm density* of the entire prey population every generation [21]. We measure the swarm density as the mean number of prey within 30 virtual meters of each other over a life span of 1,000 simulation time steps, which we have experimentally shown to differentiate between swarming and non-swarming behavior in previous published experiments [40]. Qualitatively, a swarm density ≥ 15 indicates cohesive swarming behavior, between 15 and 5 loosely grouping behavior, and ≤ 5 random, non-grouping behavior. Thus, the swarm density captures how cohesively the prey are swarming, or whether the prey are even grouping at all.

In the following sections, we study the effect of four different attack modes on the evolution of swarming behavior: uncorrelated random attacks (Figure 6a), correlated random attacks (random walk attacks, Figure 6b), peripheral attacks (Figure 6c), and attacks that target the densest area of the swarm (Figure 6d).

4.1 Random Attacks

Our 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, we expect prey to minimize their domain of danger by placing as many conspecifics as possible around the predators [18]. Similarly to previous models studying the selfish herd [65], 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 6a.

As seen in Figure 7, swarming behavior is weakly selected for when the predators make uniformly random attacks on the prey³ (light gray triangles). In particular, we found that prey took upwards of 5,000 generations to evolve cohesive swarming behavior when experiencing random attacks, compared to less 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 7, light gray stars).

³ Video: evolution of prey under random attack treatment: <http://dx.doi.org/10.6084/m9.figshare.658857>




Figure 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 life span of 1,000 simulation time steps. Prey in groups attacked randomly (light gray triangles) took much longer to evolve cohesive swarming behavior than prey in groups attacked by a predator that followed a random walk (dark gray circles) or always from outside the group (black squares). When prey experience no attacks, they do not evolve swarming behavior at all (light gray stars). Error bars indicate two standard errors over 100 replicates.

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

4.2 Random Walk Attacks

Our 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. As shown in Figure 6b, 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 7 shows that swarming evolved quickly when prey were attacked by a predator following a random walk⁴ (dark gray 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.

4.3 Outside Attacks

In the last of our initial artificial predation experiments, we simulate a predator that always approaches from outside the group and attacks the prey nearest to it, as in [63]. This predator attack mode effectively has the predators consistently attacking prey on the outer edges of the group. As shown in Figure 6c, we simulate this predator attack mode by first choosing a random angle outside the group for the predator to approach from. Once an angle is chosen, we 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 7, this form of predation has the most significant influence on the evolution of the selfish herd so far. When attacked by predators that consistently target prey on the

4 Video: evolution of prey under random walk treatment: <http://dx.doi.org/10.6084/m9.figshare.658856>




Figure 8. 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 life span of 1,000 simulation time steps. Even when experiencing density-dependent predation, prey in groups attacked randomly (light gray triangles) took much longer to evolve swarming behavior than prey in groups attacked by a persistent artificial predator (dark gray circles) or always from outside the group (black squares). Error bars indicate two standard errors over 100 replicates.

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 in the outer parts 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 in the outer parts 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 [18, 25, 39], an alternative hypothesis is that of density-dependent predation.

4.4 Density-Dependent Predation

To study the influence of density-dependent predation on the evolution of the selfish herd, we impose a constraint on the predator that 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 [18, 25, 39], and such density dependence can also be thought of as a proxy for group defense. We compute the predator's probability of capturing a prey during a given attack (P_{capture}) with the following equation:

$$P_{\text{capture}} = \frac{1}{A_{\text{density}}} \quad (2)$$

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 four other prey nearby ($A_{\text{density}} = 5$), it has a 20% chance of successfully capturing the prey. As a consequence of this mechanism, prey experience density-dependent predation.

Figure 8 demonstrates the effect of density-dependent predation on the previous artificial predation experiments. Just as before, when predators did not preferentially attack prey in the outer parts of the group, as in the random attack experiment (light gray triangles), swarming behavior took much longer to evolve. In contrast, when the predators followed a random walk (dark gray circles) or always attacked from outside the group (black squares), prey experiencing density-dependent

⁵ Video: evolution of prey under outside attack treatment: <http://dx.doi.org/10.6084/m9.figshare.658854>

predation again quickly evolved swarming 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 2 standard errors) to 17.61 ± 2.72 , indicating significantly stronger selection for swarming.

4.5 High-Density Area Attacks

Thus far, we 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 [15]. We call this mode of attack 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 from outside the swarm [37]. Thus, krill swarms are experiencing two 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 7) could outweigh the selection pressure to disperse from high-density area attacks.

As shown in Figure 6d, we 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}). We 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, we 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, we allow the prey to evolve while experiencing both attack modes simultaneously. We vary the relative handling times of both attacks (Table 4) 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 9, prey experiencing only outside attacks quickly evolve cohesive swarming behavior (light gray triangles). However, when we introduce infrequent high-density area attacks (dark gray circles), the selection pressure for prey to swarm is reduced. Finally, when we introduce frequent high-density area attacks (black squares), 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 us 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 in the outer parts

Table 4. High-density area attack (HDAA) experiment treatments.

HDAA?	Outside attack frequency	HDAA frequency
No	10	N/A
Infrequent	10	250
Frequent	10	25

Notes. The values listed for each treatment are the handling times for the corresponding predator attack mode.




Figure 9. 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 life span 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 subjected only to outside attacks (light gray triangles) evolved cohesive swarming behavior. Increasing the relative frequency of high-density area attacks from infrequent (dark gray circles) to frequent (black squares) caused the prey to evolve increasingly dispersive behavior. Error bars indicate two standard errors over 100 replicates.

of the herd, must experience a higher predation rate than in areas of dense prey, such as those in the middle of the herd.

5 Predator-Prey Coevolution

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

For this experiment, we coevolve a population of 100 predator genomes with a population of 100 prey genomes, using a GA with settings described in Table 3. Specifically, we evaluate each predator genome against the entire prey genome population for 2,000 simulation time steps each generation. During evaluation, we place 4 clonal predator agents inside a 512×512 virtual meters simulation environment with all 100 prey agents and allow the predator agents to attempt attacks 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. We assigned the prey individual fitness values as in the previous experiments, and evaluated predator fitness according to the following equation:

$$W_{\text{predator}} = \sum_{t=1}^{t_{\max}} (S_0 - A_t) \quad (3)$$

where t is the current simulation time step, t_{\max} is the total number of simulation time steps (here, $t_{\max} = 2,000$), 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 and prey genomes are assigned fitness values, we perform fitness-proportional selection on the populations via a Moran process [36], 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, we obtained the line of descent (LOD) for every replicate by tracing the ancestors of the fittest prey MN in the final population until we reached the randomly generated ancestral MN with which the starting population was seeded (see [33] for an introduction to the concept of a LOD in the context of digital evolution). We again

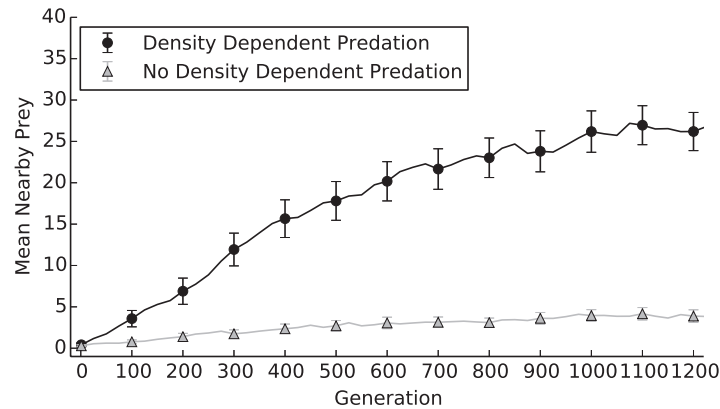


Figure 10. Mean swarm density over all replicates over evolutionary time, measured by the mean number of prey within 30 virtual meters of each other over a life span of 2,000 simulation time steps. Prey in groups experiencing density-dependent predation (black circles) evolved cohesive swarming behavior, whereas prey in groups not experiencing density-dependent predation (light gray triangles) evolved dispersive behavior. Error bars indicate two standard errors over 100 replicates.

characterized the prey grouping behavior by measuring the swarm density of the entire prey population every generation.

Figure 10 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 ± 2.3) and without density-dependent predation (light gray triangles; 3.9 ± 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 (see Figure S1 in the online supplement http://www.mitpressjournals.org/doi/suppl/10.1162/ARTL_a_00206). In contrast, with density-dependent predation, the prey quickly evolved cohesive swarming behavior in response to attacks from the predators, within 400 generations. As expected, the coevolving predators adapted to the prey swarming behavior in the density-dependent treatments by focusing on prey on the edges of the swarm, where the density of prey was lowest. As a caveat, density-dependent predation only selects for cohesive swarming behavior when the predators are faster than the prey (Figure S2 in the online supplement http://www.mitpressjournals.org/doi/suppl/10.1162/ARTL_a_00206), which corroborates earlier findings exploring the role of relative predator-prey speeds in the evolution of swarming behavior [66].

Here we 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 (see [40]; Figures S3 and S4 in the online supplement http://www.mitpressjournals.org/doi/suppl/10.1162/ARTL_a_00206). Accordingly, these results uphold Hamilton's hypothesis that grouping behavior could evolve in animals purely for selfish reasons, without the need for an explanation that involves the benefits to the whole group [18]. Moreover, the discoveries in this work refine the selfish herd hypothesis by clarifying the effect that different attack modes have on the evolution of the selfish herd.

6 Evolved Prey Markov Network Analysis

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

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




Figure 11. 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 at the front and a slight preference for sensors at the rear of the prey agent, but tended not to connect to the sensors on the sides.

First, we analyze the structure of the 100 MNs by looking at the specific retina sensors that the MNs evolved to connect to. As shown in Figure 11, 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. Strangely, there appears to be a strong preference to connect to the front right sensor but not the front left sensor, which is an artefact of the fact that the front right sensor is the only sensor that activates when other prey are directly in front of the prey. Additionally, some of the prey MNs show a preference for connecting to the prey-specific retina sensors behind the prey. From this analysis alone, we 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 inputs, we map every possible input combination in the prey's retina to the corresponding movement decision that the prey made. This mapping is accomplished by generating all 2^{24} possible input combinations (000000000000000000000000, 00000000000000000000000001, 000000000000000000000010, etc.) and passing them into the evolved MN as a simulated sensory input. Upon activating the MN, we receive an output that corresponds to the action that the prey decided to make in response to the simulated sensory input. Due to the stochastic nature of MNs, the prey agents do not always make the same movement decision when given the same input. Thus, we take the most likely output from 1,000 repeats 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. An example truth table can be seen in Table 5.

Once we have the truth table of all 2^{24} input-output mappings, we pass it to the logic minimization software *espresso* [52], which eliminates the inputs that have no effect on the outputs and provides 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. An example output minimal logic table can be seen in Table 6.

Surprisingly, the individual-based mechanisms underlying the emergent swarming behavior are remarkably simple. Most of the prey MNs evolved to make their movement decisions depend on 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

Table 5. An example truth table mapping every possible sensory input combination to the corresponding most likely movement decision from the evolved prey Markov network.

Sensory input	Corresponding output
00000000000000000000000000000000	00
00000000000000000000000000000001	10
⋮	⋮
11111111111111111111111111111111	00

Table 6. An example minimal logic table resulting from a Markov network.

Input 12	Input 19	Corresponding output
0	0	01
1	0	00
0	1	01
1	1	00

Notes. Input 12 corresponds to a frontal sensor, whereas input 19 corresponds to a back-right sensor, where 1 indicates that the sensor detects a prey agent and 0 means no prey are in that sensor. The output 01 translates into the agent turning right, and the output 00 translates into the agent moving forward. Thus, this example agent moves forward if it sees anything in its frontal sensor. Otherwise, the agent turns right if it sees another prey in its back-right sensor or if it sees nothing at all.

longer visible. This mechanism alone proved sufficient to produce cohesive swarming behavior in the majority of our 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 [40, 61, 50].

In a small subset of the evolved prey MNs, we 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 no prey agent is visible in the back sensor, the MN moves forward or turns, depending on the state of the frontal sensor. We 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.

7 Conclusions and Future Work

The contributions of this work are as follows. First, we demonstrate Hamilton’s selfish herd hypothesis in a digital evolutionary model and highlight that it is the attack mode of the predator that

critically determines the evolvability of swarming behavior. Second, we show that density-dependent predation is sufficient for the selfish herd to evolve as long as the predators cannot consistently attack prey in the middle of the group. Finally, we show 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 consider not only the behavior of the prey in the group, but the attack mode of the predators as well. Following these experiments, we 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.

Of course, the evolved prey behavior shown in the videos accompanying this article may not closely resemble the anti-predator behavior of many species of group-living prey that we observe in nature. We provide these videos to demonstrate that grouping-like behavior has indeed evolved—and to confirm that the swarm density count metric accurately captures when the prey evolve grouping-like behavior—but we do not seek to claim that we have evolved a particular behavioral phenotype that would match the grouping behavior we observe in nature. Presumably such grouping behavior has been selected for by a variety of environmental factors that are not completely captured in this model. Including those factors would make a fascinating venue of research in the future.

While this work shows one method by which the the evolution of grouping behavior can be studied, there remain many different hypotheses explaining the evolution of grouping behavior [30]. Our future work in this area will focus on directly exploring these hypotheses in similar digital evolutionary models, as detailed in [42].

Acknowledgments

We thank the three anonymous reviewers for their many insightful comments and suggestions. This research has been supported in part by the National Science Foundation (NSF) BEACON Center under Cooperative Agreement DBI-0939454, by NSF grant OCI-1122617, and by Michigan State University through computational resources provided by the Institute for Cyber-Enabled Research. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

References

1. Angeline, P. J., & Pollack, J. B. (1993). Competitive environments evolve better solutions for complex tasks. In S. Forrest (Ed.), *Proceedings of the 5th International Conference on Genetic Algorithms* (pp. 264–270). San Francisco: Morgan Kaufmann.
2. Ballerini, M., Cabibbo, N., Chandelier, R., Cavagna, A., Cisbani, E., Giardina, I., Lecomte, V., Orlandi, A., Parisi, G., Procaccini, A., Vale, M., & Zdravkovic, V. (2008). Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences of the U.S.A.*, *105*, 1232–1237.
3. Beauchamp, G. (2004). Reduced flocking by birds on islands with relaxed predation. *Proceedings of the Royal Society B*, *271*, 1039–1042.
4. Beauchamp, G. (2007). Vigilance in a selfish herd. *Animal Behaviour*, *73*, 445–451.
5. Couzin, I. (2009). Collective cognition in animal groups. *Trends in Cognitive Sciences*, *13*, 36–43.
6. Şahin, E. (2004). Swarm robotics: From sources of inspiration to domains of application. In E. Şahin & W. M. Spears (Eds.), *Proceedings of the International Conference on Swarm Robotics (SAB)* (pp. 10–20). Berlin, Heidelberg: Springer-Verlag.
7. Demšar, J., Hemelrijk, C. K., Hildenbrandt, H., & Bajec, I. L. (2015). Simulating predator attacks on schools: Evolving composite tactics. *Ecological Modelling*, *304*, 22–33.
8. Diabaté, A., Yaro, A., Dao, A., Diallo, M., Huestis, D., & Lehmann, T. (2011). Spatial distribution and male mating success of *Anopheles gambiae* swarms. *BMC Evolutionary Biology*, *11*, 184.
9. Dyer, F. (2002). The biology of the dance language. *Annual Review of Entomology*, *47*, 917–949.

10. Edlund, J., Chaumont, N., Hintze, A., Koch, C., Tononi, G., & Adami, C. (2011). Integrated information increases with fitness in the evolution of animats. *PLoS Computational Biology*, 7, e1002236.
11. Eiben, A., & Smith, J. (2003). *Introduction to evolutionary computing*. Heidelberg, New York: Springer-Verlag.
12. Feare, C. (1984). *The starling*. New York: Oxford University Press.
13. Foster, W., & Treherne, J. (1981). Evidence for the dilution effect in the selfish herd from fish predation on a marine insect. *Nature*, 293, 466–467.
14. Goldberg, D. (1989). *Genetic algorithms in search, optimization and machine learning* (1st ed.). Boston: Addison-Wesley.
15. Goldbogen, J., Calambokidis, J., Oleson, E., Potvin, J., Pyenson, N., Schorr, G., & Shadwick, R. (2011). Mechanics, hydrodynamics and energetics of blue whale lunge feeding: Efficiency dependence on krill density. *Journal of Experimental Biology*, 214, 131–146.
16. Haley, P., Olson, R., Dyer, F., & Adami, C. (2014). Exploring conditions that select for the evolution of cooperative group foraging. In H. Sayama, J. Rieffel, S. Risi, R. Doursat, & H. Lipson (Eds.), *Proceedings of the 14th International Conference on the Simulation and Synthesis of Living Systems (ALIFE)* (pp. 310–311). Cambridge, MA: MIT Press.
17. Hall-Stoodley, L., Costerton, J., & Stoodley, P. (2004). Bacterial biofilms: From the natural environment to infectious diseases. *Nature Reviews Microbiology*, 2, 95–108.
18. Hamilton, W. (1971). Geometry for the selfish herd. *Journal of Theoretical Biology*, 31, 295–311.
19. Hemelrijk, C., & Hildenbrandt, H. (2011). Some causes of the variable shape of flocks of birds. *PLoS ONE*, 6, e22479.
20. Hemelrijk, C. K., & Kunz, H. (2005). Density distribution and size sorting in fish schools: An individual-based model. *Behavioral Ecology*, 16, 178–187.
21. Huepe, C., & Aldana, M. (2008). New tools for characterizing swarming systems: A comparison of minimal models. *Physica A*, 387, 2809–2822.
22. Inman, A., & Krebs, J. (1987). Predation and group living. *Trends in Ecology & Evolution*, 2, 31–32.
23. Ioannou, C., Guttal, V., & Couzin, I. (2012). Predatory fish select for coordinated collective motion in virtual prey. *Science*, 337, 1212–1215.
24. Ioannou, C., Tosh, C., Neville, L., & Krause, J. (2008). The confusion effect—from neural networks to reduced predation risk. *Behavioral Ecology*, 19, 126–130.
25. James, R., Bennett, P., & Krause, J. (2004). Geometry for mutualistic and selfish herds: The limited domain of danger. *Journal of Theoretical Biology*, 228, 107–113.
26. Jeschke, J., & Tollrian, R. (2007). Prey swarming: Which predators become confused and why? *Animal Behaviour*, 74, 387–393.
27. Kenward, R. (1978). Hawks and doves: Factors affecting success and selection in goshawk attacks on woodpigeons. *Journal of Animal Ecology*, 47, 449–460.
28. Kersten, M., Britton, R., Dugan, P., & Hafner, H. (1991). Flock feeding and food intake in little egrets: The effects of prey distribution and behaviour. *Journal of Animal Ecology*, 60, 241–252.
29. Krakauer, D. (1995). Groups confuse predators by exploiting perceptual bottlenecks: A connectionist model of the confusion effect. *Behavioral Ecology and Sociobiology*, 36, 421–429.
30. Krause, J., & Ruxton, G. (2002). *Living in groups*. New York: Oxford University Press.
31. Kunz, H., & Hemelrijk, C. K. (2003). Artificial fish schools: Collective effects of school size, body size, and body form. *Artificial Life*, 9, 237–253.
32. Kunz, H., Züblin, T., & Hemelrijk, C. (2006). On prey grouping and predator confusion in artificial fish schools. In L. M. Rocha, L. S. Yaeger, M. A. Bedau, D. Floreano, R. L. Goldstone, & A. Vespignani (Eds.), *Proceedings of the 10th International Conference on the Simulation and Synthesis of Living Systems (ALIFE)* (pp. 365–371). Cambridge, MA: MIT Press.
33. Lenski, R., Ofria, C., Pennock, R., & Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423, 139–144.

34. Marinakis, Y., & Marinaki, M. (2011). A honey bees mating optimization algorithm for the open vehicle routing problem. In N. Krasnogor (Ed.), *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation* (pp. 101–108). New York: ACM Press.
35. Martin, G. (1986). The eye of a passeriform bird, the European starling (*Sturnus vulgaris*): Eye movement amplitude, visual fields and schematic optics. *Journal of Comparative Physiology A*, *159*, 545–557.
36. Moran, P. (1962). *The statistical processes of evolutionary theory*. Oxford, UK: Clarendon Press.
37. Mori, M., & Butterworth, D. (2006). A first step towards modelling the krill-predator dynamics of the Antarctic ecosystem. *CCAMLR Science*, *13*, 217–277.
38. Morrell, L., Ruxton, G., & James, R. (2010). Spatial positioning in the selfish herd. *Behavioral Ecology*, *22*, 16–22.
39. Morton, T., Haefner, J., Nugala, V., Decino, R., & Mendes, L. (1994). The selfish herd revisited: Do simple movement rules reduce relative predation risk? *Journal of Theoretical Biology*, *167*, 73–79.
40. Olson, R., Hintze, A., Dyer, F., Knoester, D., & Adami, C. (2013). Predator confusion is sufficient to evolve swarming behaviour. *Journal of the Royal Society Interface*, *10*, 20130305.
41. Olson, R., Knoester, D., & Adami, C. (2013). Critical interplay between density-dependent predation and evolution of the selfish herd. In C. Blum (Ed.), *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation* (pp. 247–254). New York: ACM Press.
42. Olson, R. S. (2015). *Elucidating the evolutionary origins of collective animal behavior*. Ph.D. thesis, Michigan State University.
43. Olson, R. S., Haley, P. B., Dyer, F. C., & Adami, C. (2015). Exploring the evolution of a trade-off between vigilance and foraging in group-living organisms. *Royal Society Open Science*, *2*, 150135.
44. Parrish, J. (1999). Complexity, pattern, and evolutionary trade-offs in animal aggregation. *Science*, *284*, 99.
45. Poli, R. (2008). Analysis of the publications on the applications of particle swarm optimisation. *Journal of Artificial Evolution and Applications*, *2008*, 4:1–4:10.
46. Pulliam, H. (1973). On the advantages of flocking. *Journal of Theoretical Biology*, *38*, 419–422.
47. Pulliam, R., & Caraco, T. (1984). Living in groups: Is there an optimal group size? In J. Krebs & N. Davies (Eds.), *Behavioral ecology: An evolutionary approach* (2nd ed.) (pp. 122–147). Oxford, UK: Blackwell Scientific.
48. Quinn, J., & Cresswell, W. (2006). Testing domains of danger in the selfish herd: Sparrowhawks target widely spaced redshanks in flocks. *Proceedings of the Royal Society B*, *273*, 2521–2526.
49. Reluga, T., & Viscido, S. (2005). Simulated evolution of selfish herd behavior. *Journal of Theoretical Biology*, *234*, 213–225.
50. Reynolds, C. (1987). Flocks, herds and schools: A distributed behavioral model. In M. C. Stone (Ed.), *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (pp. 25–34). New York: ACM Press.
51. Reynolds, C. W. (1993). An evolved, vision-based behavioral model of coordinated group motion. In J.-A. Meyer, H. L. Roitblat, & S. W. Wilson (Eds.), *From animals to animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 384–392). Cambridge, MA: MIT Press.
52. Rudell, R. (1986). *Multiple-valued logic minimization for PLA synthesis* (Tech. Rep. UCB/ERL M86/65). EECS Department, University of California, Berkeley.
53. Silva, A., Neves, A., & Costa, E. (2002). An empirical comparison of particle swarm and predator prey optimisation. In M. O’Neill, R. F. E. Sutcliffe, C. Ryan, M. Eaton, & N. Griffith (Eds.), *Proceedings of the 13th Irish Conference on Artificial Intelligence and Cognitive Science (AICS)* (pp. 103–110). London: Springer-Verlag.
54. Sims, K. (1994). Evolving 3D morphology and behavior by competition. *Artificial Life*, *1*, 353–372.
55. Spector, L., Klein, J., Perry, C., & Feinstein, M. (2003). Emergence of collective behavior in evolving populations of flying agents. In E. Cantu-Paz, J. Foster, K. Deb, L. Davis, R. Roy, U.-M. O’Reilly, H.-G. Beyer, R. Standish, G. Kendall, S. Wilson, M. Harman, J. Wegener, D. Dasgupta, M. Potter, A. Schultz, K. Dowsland, N. Jonoska, & J. Miller (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)* (pp. 61–73). Berlin, Heidelberg, New York: Springer-Verlag.

56. Tosh, C. (2011). Which conditions promote negative density dependent selection on prey aggregations? *Journal of Theoretical Biology*, 281, 24–30.
57. Treherne, J., & Foster, W. (1981). Group transmission of predator avoidance behaviour in a marine insect: The Trafalgar effect. *Animal Behaviour*, 29, 911–917.
58. Treherne, J., & Foster, W. (1982). Group size and anti-predator strategies in a marine insect. *Animal Behaviour*, 30, 536–542.
59. Treisman, M. (1975). Predation and the evolution of gregariousness. I. Models for concealment and evasion. *Animal Behaviour*, 23, 779–800.
60. Vellasques, E., Sabourin, R., & Granger, E. (2012). Gaussian mixture modeling for dynamic particle swarm optimization of recurrent problems. In T. Soule (Ed.), *Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation* (pp. 73–80). New York: ACM Press.
61. Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., & Shochet, O. (1995). Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75, 1226–1229.
62. Viscido, S. V., Millery, M., & Wethey, D. S. (2002). The dilemma of the selfish herd: The search for a realistic movement rule. *Journal of Theoretical Biology*, 217, 183–194.
63. Viscido, S. V., Miller, M., & Wethey, D. S. (2001). The response of a selfish herd to an attack from outside the group perimeter. *Journal of Theoretical Biology*, 208, 315–328.
64. Ward, C., Gobet, F., & Kendall, G. (2001). Evolving collective behavior in an artificial ecology. *Artificial Life*, 7, 191–209.
65. Wood, A., & Ackland, G. (2007). Evolving the selfish herd: Emergence of distinct aggregating strategies in an individual-based model. *Proceedings of the Royal Society B*, 274, 1637–1642.
66. Wood, A. J. (2010). Strategy selection under predation; evolutionary analysis of the emergence of cohesive aggregations. *Journal of Theoretical Biology*, 264, 1102–1110.
67. Xue, B., Zhang, M., & Browne, W. (2012). Multi-objective particle swarm optimisation (PSO) for feature selection. In T. Soule (Ed.), *Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation* (pp. 81–88). New York: ACM Press.
68. Yuval, B., Wekesa, J., & Washino, R. (1993). Effect of body size on swarming behavior and mating success of male *Anopheles freeborni* (Diptera: Culicidae). *Journal of Insect Behavior*, 6, 333–342.