

Critical Interplay Between Density-dependent Predation and Evolution of the Selfish Herd

Randal S. Olson^{1,3}, David B. Knoester^{2,3}, and Christoph Adami^{2,3}

¹Department of Computer Science & Engineering

²Department of Microbiology and Molecular Genetics

³BEACON Center for the Study of Evolution in Action

Michigan State University

East Lansing, Michigan 48824, USA

olsonran@msu.edu, dk@msu.edu, adami@msu.edu

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. 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 the predator attack mode plays a critical role in 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 verifies 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.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*coherence and coordination, multiagent systems*;
I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search

Keywords

group behavior, selfish herd hypothesis, predator-prey co-evolution, density-dependent predation, digital evolutionary model

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'13, July 6–10, 2013, Amsterdam, The Netherlands.
Copyright 2013 ACM 978-1-4503-1963-8/13/07 ...\$15.00.

1. INTRODUCTION

Over the past century, biologists have devoted considerable effort into studying animal grouping behaviors due to their important implications for social intelligence, collective cognition, and potential applications in engineering, artificial intelligence, and robotics [1]. 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 [2]. Western honeybees (*Apis mellifera*) communicate the location of food and nest sites to other bees in their group via a complex dance language [3]. Even relatively simple bacteria exhibit grouping behavior, such as *Escherichia coli* forming biofilms which allow their group to survive in hostile environments [4].

Swarming is one example of grouping behavior, where animals coordinate their movement with conspecifics to maintain a cohesive group while on the move. Although swarm-like groups could arise by chance, e.g., Little Egrets (*Egretta garzetta*) pursuing a common resource in water pools [5], most of the time swarms are maintained via behavioral mechanisms that ensure group cohesion [6]. 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 [7]. With this fact in mind, significant effort has been dedicated to understanding the compensating benefits that grouping behavior provides [8]. Many such benefits of grouping behavior have been proposed, for example, swarming may improve mating success [9], increase foraging efficiency [10], or enable the group to solve problems that would be impossible to solve individually [1]. Furthermore, swarming behaviors are hypothesized to protect group members from predators in several ways. For example, swarming can improve group vigilance [11], reduce the chance of being encountered by predators [11], dilute an individual's risk of being attacked [12], enable an active defense against predators [13], or reduce predator attack efficiency by confusing the predator [14].

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, and make it even more difficult to study the properties of those behaviors as they evolve [14, 15]. In this paper, we use a digital model of predator-prey coevolution to explore Hamilton's

selfish herd hypothesis [12]. 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.

In our model, both predators and prey have the ability to detect and interact with other agents in the environment. We evolve the agents with a genetic algorithm by preferentially selecting predators and prey based on how effective they are at consuming prey and surviving, respectively. Forming a selfish herd is a possible solution for the prey to survive longer, but is not selected for directly. In this study, we first test whether a selfish herd evolves within a two-dimensional virtual environment with different forms of simulated predation. Doing so enables us to experimentally control the effects of specific modes of predation on the evolution of the selfish herd. We found that if predators are able to consistently attack the center of the group of prey, the selfish herd will not evolve. In subsequent experiments, we discovered that density-dependent predation can provide a generalization of Hamilton’s original formulation of “domains of danger.” Following these findings, we coevolve groups of predators and prey in a similar virtual environment to determine if coevolving predators impact the likelihood of the selfish herd to evolve. Finally, this study demonstrates that density-dependent predation provides a sufficient selective advantage for prey to evolve the selfish herd in response to predation by coevolving predators.

2. RELATED WORK

Hamilton’s original formulation of the selfish herd hypothesis introduced the concept of “domains of danger” (DODs), which served as a method to visualize the likelihood of a prey inside a group to be attacked by a predator [12]. 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. Furthermore, 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 [16], investigating the effects of prey vigilance [17], considering the initial spatial positioning of prey when the group is attacked [18], and even confirming Hamilton’s predictions in biological systems [19].

Additional studies have focused on discovering the movement rules that prey in a selfish herd follow to minimize their DOD [20]. This line of work began by demonstrating that the simple movement rules proposed by Hamilton do indeed reduce predation for prey inside the group [21], then opened some parameters of the movement rules to evolution in an attempt to discover a more biologically plausible set of movement rules [22]. Finally, some work has gone into studying the evolution of predator behavior in response to prey density [23], the coevolution of predator and prey behavior in the presence of the predator confusion effect [24], and elaborating upon the interaction between ecology and the evolution of grouping behavior [25, 26]. This paper builds on this work by studying the effects of coevolving predators and predator attack mode, i.e., how the predator selects a prey in the group to attack, on the evolution of the selfish herd.

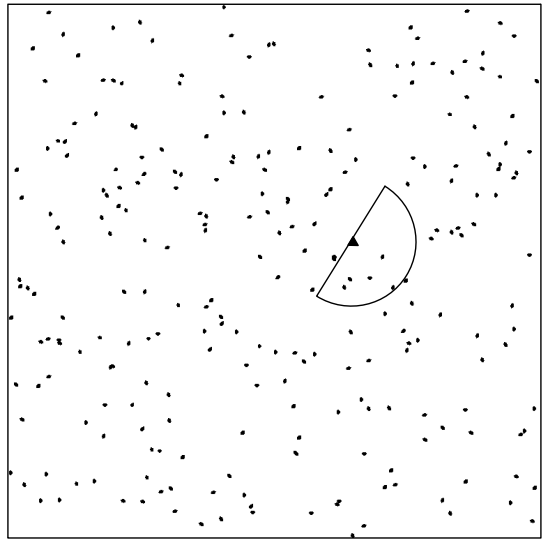


Figure 1: A screen capture of the simulation environment in which the agents interact. Black dots are prey agents, the black triangle is a predator agent, and the lines around the predator agent indicate its field of view. Agents wrap around the edges of the toroidal simulation environment.

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) [27]. PSO applications range from feature selection for classifiers [28], to video processing [29], to open vehicle routing [30]. A related technique within PSO seeks to combine PSO with coevolving “predator” and “prey” solutions to avoid local minima [31]. Researchers have even sought to harness the collective problem solving power of swarming agents to design robust autonomous robotic swarms [32]. Thus, elaborations on the foundations of animal grouping behavior has 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 simulation in which agents interact in a continuous, toroidal virtual environment (736×736 virtual meters), shown in Figure 1. At the beginning of each simulation, we place 250 agents in the environment at 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) [33]. We evolve the agent MNs with a genetic algorithm (GA) [34] 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, 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.




Figure 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 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 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.

The source code¹ and data² for these experiments are available online. 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 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), 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 predator resides within that slice, but not how many. For example, in Figure 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 1. Prey agents turn in 8° increments and move 1 virtual meter each time step.

¹Code: <https://github.com/adamilab/eos-selfish-herd>

²Data: <http://dx.doi.org/10.6084/m9.figshare.663680>

Table 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

In our coevolution experiments, the predator agents detect nearby prey agents and conspecifics using a limited-distance (200 virtual meters), pixelated retina covering its frontal 180° that works just like the prey agent’s retina (Figure 2). Similar 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. Finally, if a predator agent moves within 5 virtual meters of a prey agent that is visible to it, the predator agent makes an attack attempt on the prey agent. If the attack 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 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. Similar to an Artificial Neural Network (ANN), a MN receives input via sensors (e.g., visual retina), performs a computation on inputs and any hidden states (i.e., memory), then places the result of the computation into hidden or output states (e.g., actuators). In contrast to many ANNs, 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.

How Markov Networks Function

As seen in Figure 2, when we embed an agent into the simulation environment, we provide it sensorial inputs from the retina into its MN every simulation step (labeled “retina” and “Markov Network”, respectively). Once we provide a MN with its inputs, we activate it and allow it to store the resulting 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 3, 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




Figure 3: 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.

same state, the outputs are combined into a single output using the OR logic function. Thus, the MN uses information about the environment and 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 classic logic gates, e.g., AND & OR. A classic 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 provides 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 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 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

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 in the next section.

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 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 $N = \lfloor \frac{\text{byte}}{255/N_{max}} \rfloor$. 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: $\lfloor \frac{\text{byte} \times \# \text{states}}{255} \rfloor - 0.5$, where $\# \text{states}$ is the total number of states in the MN and $\lfloor \bullet \rfloor$ denotes the nearest integer. 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 define 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. 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.




Figure 4: Example circular byte strings encoding the two Markov Gates (MGs) in Figure 3, denoted Gene 1 and Gene 2. The sequence (42, 213) represents the beginning of a new MG (red blocks). The next two bytes encode the number of input and output states used by the MG (green blocks), and the following eight bytes encode which states are used as input (blue blocks) and output (yellow blocks). The remaining bytes in the string encode the probabilities of the MG’s logic table (cyan blocks).

Table 3: 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	1,200
Replicates	100

4. ARTIFICIAL SELECTION

In our first set of experiments, we observe the evolution of prey behavior in response to various forms of artificial selection. This enables us to experimentally control the effects of certain modes of predation 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 MNs. Following this, we evaluate the relative fitness of each prey genome by competing the genomes 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. We assign each prey an individual fitness according to the following equation:

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

where T is the number of time steps the individual prey survived in the simulation environment. Thus, individual prey are rewarded for surviving longer than other group members. Once all of the prey genomes are assigned fitness values, we perform fitness-proportionate selection on the population via a Moran process [36], increment the generation counter, and repeat the evaluation process on the new population until the final generation (1,200) 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 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, we apply artificial selection every 4 simulation time steps to simulate an ideal predator attacking the group. 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




Figure 5: An illustration of the three artificial predator attack modes. A) Random attacks, B) Random walk attacks, C) Outside attacks.




Figure 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. Prey in groups attacked randomly (light grey triangles with a full line) evolved dispersive behavior, whereas prey in groups attacked by a predator that follows a random walk (dark grey circles with a dashed line) or always from the outside of the group (black squares with a dotted line) evolved cohesive swarming behavior. Error bars indicate two standard errors over 100 replicates.

the time it takes to refocus on another prey in the case of an unsuccessful attack attempt.

For each experiment, we characterize the grouping behavior by measuring the *swarm density* of the entire prey population every generation [37]. We 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. 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.

Random Attacks. Our initial study sought to verify Hamilton's selfish herd hypothesis by simulating 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" to the predators by placing as many conspecifics as possible around them [12]. Similar to previous models studying the selfish herd [22], 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 5A.

As seen in Figure 6, if the predators make uniformly random attacks on the prey, the selfish herd did not evolve³ (light grey triangles with a full line). Instead, the evolutionary response to predators performing uniformly random attacks on groups of prey was to move randomly to avoid potential attacks. This finding has important implications, namely that one of the key assumptions of the selfish herd hypothesis—uniformly random predator attacks—does not appear to select for the evolution of the selfish herd. Following this discovery, we hypothesized that the *directionality* of the predators' attacks play a critical role in the evolution

³Video of evolution of prey behavior 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 lifespan of 1,000 simulation time steps. Even when experiencing density-dependent predation, prey in groups attacked randomly (light grey triangles with a full line) evolved dispersive behavior, whereas prey in groups attacked by a persistent artificial predator (dark grey circles with a dashed line) or always from the outside of the group (black squares with a dotted line) evolved cohesive swarming behavior. Error bars indicate two standard errors over 100 replicates.

of the selfish herd. To test this hypothesis, we next explore two different predator attack modes, each with their own distinct directionality of predation.

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. Shown in Figure 5B, after each attack made by this predator, it is then moved to a random location within 50 virtual meters. This models a predator that persistently feeds on a group of prey, rather than ambushing prey.

Figure 6 shows that swarming did evolve when the prey were attacked by a predator following a random walk⁴ (dark grey circles with a dashed line). Therefore, this experiment demonstrated that the selfish herd can indeed evolve if the predators follow a persistent hunting pattern. However, the selection pressure exerted by such predators appears to be weak since it took roughly 800 generations for prey to evolve swarming behavior. Thus, the random walk predator attack mode appears to capture an important aspect of predation that selects for the selfish herd, but the strength of selection for the selfish herd is relatively weak.

Outside Attacks. Finally, we simulate a predator that always approaches from the outside of the group and attacks the prey nearest to it. This predator attack mode effectively has the predators consistently attacking prey on the outer edges of the group. As shown in Figure 5C, we 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, we convert the angle into a location on the edge of simulation space and attack the prey nearest to that location.

As shown in Figure 6, 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 with a dotted line). Taken together, the results of these artificial selection experiments demonstrate another discovery of this work: the more consistently predators attack prey on the outside of the group, the more likely the selfish herd is to 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 [12, 16, 21], an alternative hypothesis is that of density-dependent predation.

Density-Dependent Predation. To study the impact of density-dependent predation on the evolution of the selfish herd, we 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 [12, 16, 21], 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:

$$P_{\text{capture}} = \frac{1}{A_{\text{density}}}$$

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. As a consequence of this mechanism, the prey experience density-dependent predation.

Figure 7 demonstrates the effect of density-dependent predation on the previous artificial selection 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 with a full line), the selfish herd did not evolve. Thus, even if the prey experience density-dependent predation, the evolutionary pressure for prey to disperse will outweigh the evolutionary pressure for prey to swarm if predators can consistently attack prey in the center of the group. In contrast, when the predators followed a random walk (dark grey circles with a dashed line) or always attacked from the outside of the group (black squares with a dotted line), the prey experiencing density-dependent predation evolved the selfish herd. The most noticeable effect of density-dependent predation is on the random walk treatment, where the mean nearby prey at generation 1,200 increased from 14.39 ± 1.97 (mean \pm two standard errors) to 19.11 ± 2.32 , indicating stronger selection for swarming.

In summary, the artificial selection experiments provided us with two important pieces of information regarding the evolution of the selfish herd: (1) the predator cannot consistently attack prey on the inside of the group and (2) prey in less dense areas, such as those on the outside of the group, must experience a higher predation rate than in areas of dense prey, such as found in the center of the group.

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

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

5. PREDATOR-PREY COEVOLUTION

Building upon the artificial selection 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 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 and simulation environment area 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:

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

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 proportionate 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 most-fit prey MN in the final population until we reached the randomly-generated ancestral MN with which the starting population was seeded (see [38] for an introduction to the concept of a LOD in the context of digital evolution). We again characterized the prey grouping behavior by measuring the swarm density of the entire prey population every generation.

Figure 8 depicts the prey behavior measurements for the coevolution experiments with density-dependent predation⁶ (black circles with a dashed line; mean swarm density at generation $1,200 \pm$ two standard errors: 26.2 ± 2.3) and without density-dependent predation (light grey triangles with a full line; 3.9 ± 0.8). Without density-dependent predation, the prey evolved purely dispersive behavior as a mechanism to escape the predators. However, with density-dependent predation, the prey evolved cohesive swarming behavior in response to attacks from the predators.

This result highlights the final discovery of this work: density-dependent predation provides a sufficient selective advantage for prey to evolve the selfish herd in response to predation by coevolving predators. Accordingly, these results uphold Hamilton's hypothesis that grouping behavior could evolve in animals purely due to selfish reasons, with-

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




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 lifespan of 1,000 simulation time steps. Prey in groups experiencing density-dependent predation (black circles with a dashed line) evolved cohesive swarming behavior, whereas prey in groups not experiencing density-dependent predation (light grey triangles with a full line) evolved dispersive behavior. Error bars indicate two standard errors over 100 replicates.

out the need for an explanation that involves the benefits to grouping [12]. Moreover, the discoveries in this work refine the selfish herd hypothesis by clarifying which modes of predation would allow for the evolution of the selfish herd.

6. 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 which critically determines whether the selfish herd evolves. 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 center 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 not only consider the behavior of the prey in the group, but the attack mode of the predators as well. While this work shows one method by which the evolution of grouping behavior can be studied, there remain many different hypotheses explaining the evolution of grouping behavior [8] that have yet to be studied using digital evolutionary models. As such, future work in this area should focus on directly testing these hypotheses in digital evolutionary models, such as the model described here.

7. ACKNOWLEDGMENTS

This research has been supported in part by the National Science Foundation (NSF) BEACON Center under Cooperative Agreement DBI-0939454, and NSF grant OCI-1122617. 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. We wish to acknowledge the support of the Michigan State University High Performance Computing Center and the Institute for Cyber Enabled Research (iCER).

8. REFERENCES

- [1] I.D. Couzin. Collective cognition in animal groups. *Trends Cogn. Sci.*, 13:36–43, 2009.
- [2] C.K. Hemelrijk and H. Hildenbrandt. Some causes of the variable shape of flocks of birds. *PLoS ONE*, 6:e22479, 2011.
- [3] F.C. Dyer. The biology of the dance language. *Annu. Rev. of Entomol.*, 47:917–49, 2002.
- [4] L. Hall-Stoodley, J.W. Costerton, and P. Stoodley. Bacterial biofilms: from the natural environment to infectious diseases. *Nat. Rev. Micro.*, 2:95–108, 2004.
- [5] M. Kersten, R.H. Britton, P.J. Dugan, and H. Hafner. Flock feeding and food intake in Little Egrets: the effects of prey distribution and behaviour. *J. Anim. Ecol.*, 60:241–252, 1991.
- [6] 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*, 105:1232–7, 2008.
- [7] J.K. Parrish. Complexity, pattern, and evolutionary trade-offs in animal aggregation. *Science*, 284:99, 1999.
- [8] J. Krause and G.D. Ruxton. *Living in groups*. Oxford University Press, USA, 2002.
- [9] B. Yuval, J.W. Wekesa, and R.K. Washino. Effect of body size on swarming behavior and mating success of male *Anopheles freeborni* (Diptera: Culicidae). *J. Insect Behav.*, 6:333–342, 1993.
- [10] R.H. Pulliam and T. Caraco. Living in groups: is there an optimal group size? In J. R. Krebs and N. B. Davies, editors, *Behavioral Ecology: An Evolutionary Approach*. Sinauer, Mass., 2nd edition, 1984.
- [11] M. Treisman. Predation and the evolution of gregariousness. I. Models for concealment and evasion. *Anim. Behav.*, 23:779–800, 1975.
- [12] W.D. Hamilton. Geometry for the selfish herd. *J. Theor. Bio.*, 31:295–311, 1971.
- [13] B.R.C. Bertram. In J. R. Krebs and N.B. Davies, editors, *Behavioral Ecology: An Evolutionary Approach*, pages 64–96. Blackwell, Oxford, 1978.
- [14] J.M. Jeschke and R. Tollrian. Prey swarming: which predators become confused and why? *Anim. Behav.*, 74:387–393, 2007.
- [15] G. Beauchamp. Reduced flocking by birds on islands with relaxed predation. *P. Roy. Soc. B*, 271:1039–42, 2004.
- [16] R. James, P.G. Bennett, and J. Krause. Geometry for mutualistic and selfish herds: the limited domain of danger. *J. Theoret. Biol.*, 228:107–13, 2004.
- [17] G. Beauchamp. Vigilance in a selfish herd. *Anim. Behav.*, 73:445–451, 2007.
- [18] L.J. Morrell, G.D. Ruxton, and R. James. Spatial positioning in the selfish herd. *Behav. Ecol.*, 22:16–22, 2010.
- [19] J.L. Quinn and W. Cresswell. Testing domains of danger in the selfish herd: Sparrowhawks target widely spaced redshanks in flocks. *P. Roy. Soc. B*, 273:2521–6, 2006.
- [20] S.V. Viscido, M. Millery, and D.S. Wettheywz. The dilemma of the selfish herd: the search for a realistic movement rule. *J. Theor. Bio.*, 217:183–194, 2002.
- [21] T.L. Morton, J.W. Haefner, V. Nugala, R.D. Decino, and L. Mendes. The selfish herd revisited: do simple movement rules reduce relative predation risk? *J. Theor. Bio.*, 167:73–79, 1994.
- [22] A.J. Wood and G.J. Ackland. Evolving the selfish herd: emergence of distinct aggregating strategies in an individual-based model. *P. Roy. Soc. B*, 274:1637–42, 2007.
- [23] C.R. Tosh. Which conditions promote negative density dependent selection on prey aggregations? *J. Theor. Bio.*, 281:24–30, 2011.
- [24] R.S. Olson et al. Predator confusion is sufficient to evolve swarming behavior. arXiv e-print. <http://arxiv.org/abs/1209.3330>.
- [25] L. Spector, J. Klein, C. Perry, and M. Feinstein. Emergence of collective behavior in evolving populations of flying agents. In *Proceedings GECCO '03*, pages 61–73, Berlin, 2003. Springer.
- [26] C.R. Ward, F. Gobet, and G. Kendall. Evolving collective behavior in an artificial ecology. In *Artificial Life*, volume 7, pages 191–209, 2001.
- [27] R. Poli. Analysis of the publications on the applications of particle swarm optimisation. *J. Artif. Evol. Appl.*, 2008:4:1–4:10, 2008.
- [28] B. Xue, M. Zhang, and W.N. Browne. Multi-objective particle swarm optimisation (PSO) for feature selection. In *Proceedings GECCO '12*, pages 81–88, New York, NY, USA, 2012. ACM.
- [29] E. Vellasques, R. Sabourin, and E. Granger. Gaussian mixture modeling for dynamic particle swarm optimization of recurrent problems. In *Proceedings GECCO '12*, pages 73–80, New York, NY, 2012. ACM.
- [30] Y. Marinakis and M. Marinaki. A honey bees mating optimization algorithm for the open vehicle routing problem. In *Proceedings GECCO '11*, pages 101–108, New York, NY, USA, 2011. ACM.
- [31] A. Silva, A. Neves, and E. Costa. An empirical comparison of particle swarm and predator prey optimisation. In *Proceedings AICS '02*, pages 103–110, London, UK, 2002. Springer-Verlag.
- [32] E. Şahin. Swarm robotics: from sources of inspiration to domains of application. In *Proceedings SAB '04*, pages 10–20, Berlin, Heidelberg, 2005. Springer-Verlag.
- [33] J.A. 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.*, 7:e1002236, 2011.
- [34] D.E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Boston, MA, 1st edition, 1989.
- [35] G.R. Martin. The eye of a passeriform bird, the European starling (*Sturnus vulgaris*): eye movement amplitude, visual fields and schematic optics. *J. Comp. Physiol. A*, 159:545–557, 1986.
- [36] P.A.P. Moran. *The statistical processes of evolutionary theory*. Clarendon Press, 1962.
- [37] C. Huepe and M. Aldana. New tools for characterizing swarming systems: a comparison of minimal models. *Physica A*, 387:2809–2822, 2008.
- [38] R.E. Lenski, C. Ofria, R.T. Pennock, and C. Adami. The evolutionary origin of complex features. *Nature*, 423:139–44, 2003.