

EvoBoids: Incorporating Machine Learning into Artificial Life

Adam Birt

Department of Computer Science
Colby College
Waterville, Maine 04901 USA
birtweiser@yahoo.com

Samuel Shaw

123 Barbourtown Road
Canton, Connecticut 06019
chaosfusion@hotmail.com

Abstract: EvoBoids integrates genetic algorithms and agent learning with Craig Reynolds' "Boids" form of artificial life. "Flocking" is an emergent behavior of Reynold's complex system; EvoBoids seeks to uncover new and interesting emergent properties. By assigning a limited genetic sequence to each boid in the XBoids screensaver system (open source), Reynold's simulated birds can be evolved using Genetic Algorithms. Survival elements, such as feeding requirements, are included as a means of determining boid fitness, and fitnesses are used to bias the selection of parents when new generations are produced. Successfully modeling nature as a combination of complex systems and genetic algorithms provides insight into the underlying organizational principles that govern our environment.




Figure 1: EvoBoids in action

1 Introduction and Background

1.1 Boids

In 1986, Craig Reynolds developed a computer system for modeling synchronized animalian movement, such as that exhibited by flocking birds or schooling fish. He dubbed this system "Boids." Boids are governed in their movements by three simple rules:

1. *Separation*: each boid should steer to avoid crowding its local flockmates;
2. *Alignment*: each boid should steer toward the average heading of local flockmates;
3. *Cohesion*: each boid should steer toward the average position of local flockmates.

"Flocking," Reynolds discovered, is an emergent property of the application of this rule set to each boid. As suggested by the above regulations, any given boid need only consider flockmates within a small radius when calculating the rule set effects. Boids outside this local "neighborhood" are ignored. The neighborhood is defined by an angle, measured from the boid's direction of flight, and a distance, measured from the boid's center (Reynolds 1995).

Observed in nature, flocking is quite awe-inspiring. Perhaps more impressive, however, is its abstraction down to only three governing principles.

1.2 Complex Systems and Artificial Life

Complex systems, e.g. Boids, display intricate behaviors, but obey simplistic rule sets. Such systems have the unique property of short-term predictability. For example, the position of any given boid can be predicted based on its velocity for the coming few seconds (at most). Beyond that, the deterministic aspect of Boids fails. Short-term predictability is in contrast to chaotic systems, which are completely unpredictable, and periodic systems, which are entirely predictable (Reynolds 1995).

The "Boids" complex system is exemplar of artificial life. Artificial life complex systems seek to model natural behaviors demonstrated by living beings. Lindenmayer systems, John Conway's Game of Life and cellular automata are also considered forms of artificial life.

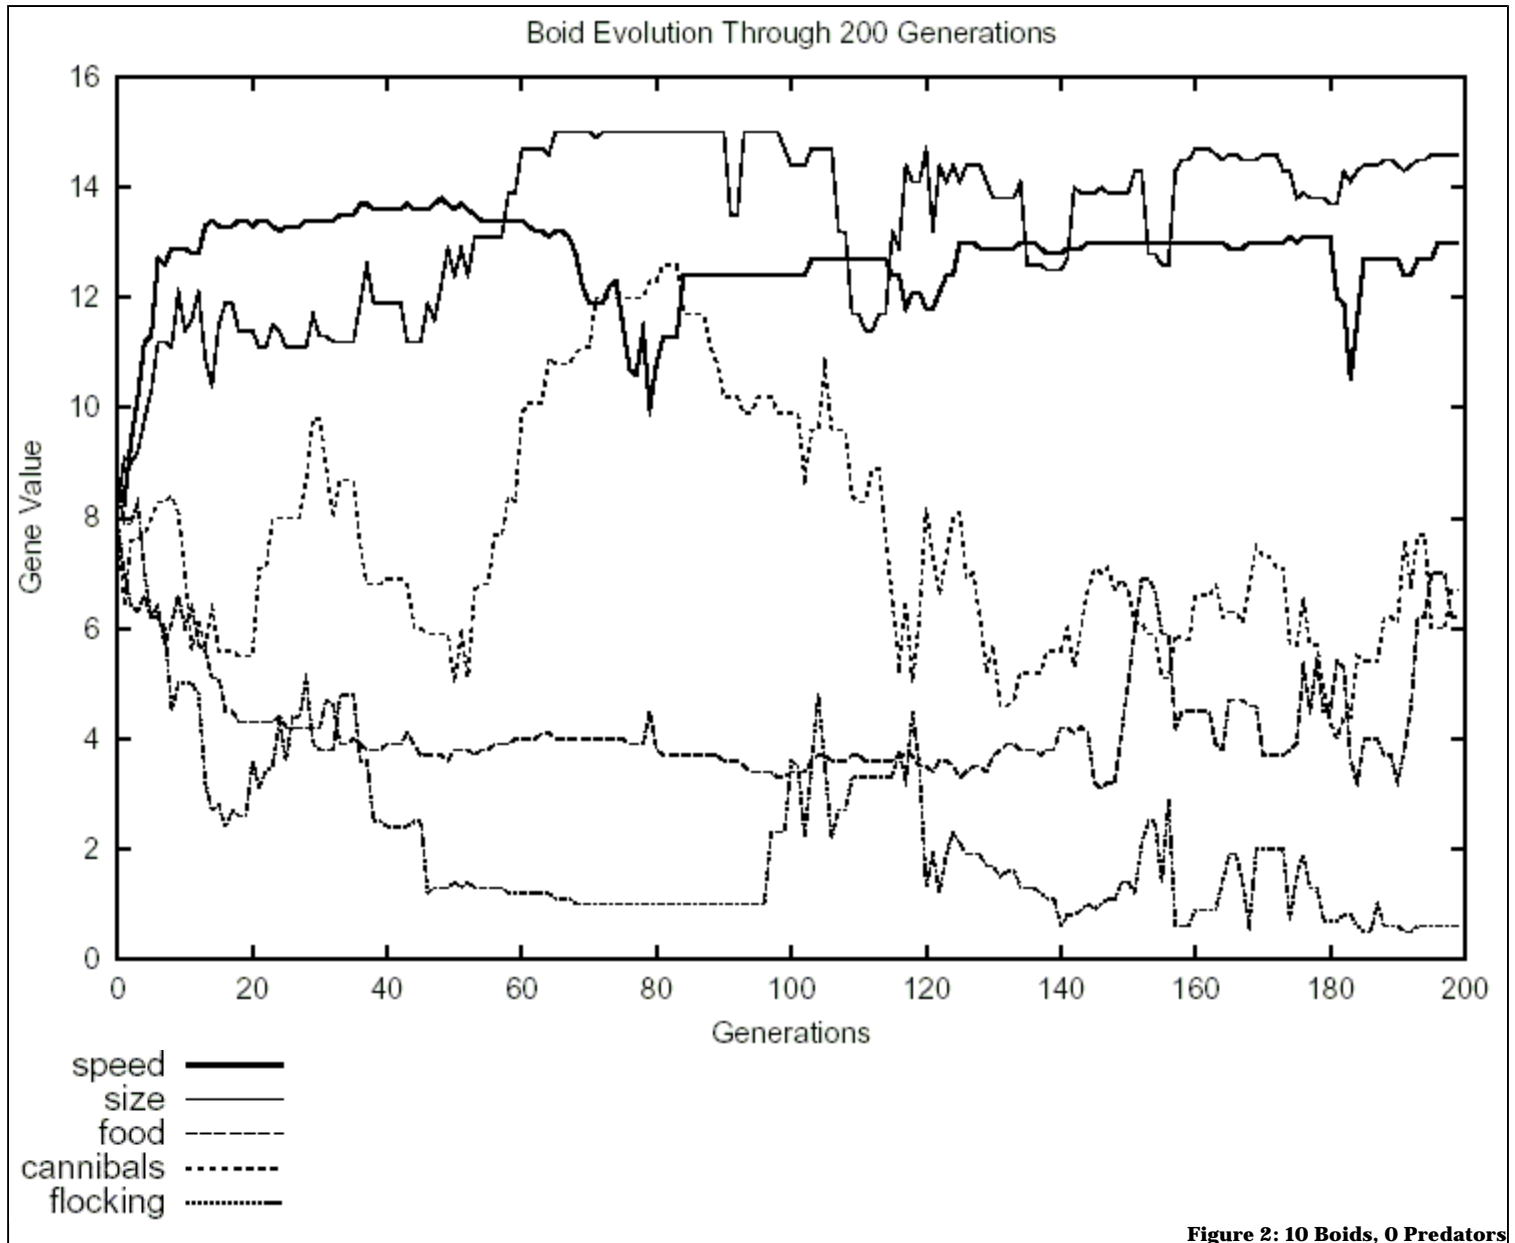


Figure 2: 10 Boids, 0 Predators

1.3 Research Goals

EvoBoids seeks to merge two differing but logical and well-grounded means of modeling nature (complex systems and genetic algorithms) into a single entity contained under the heading "artificial life." In doing so, our aim is to expand conceptions of natural order and to gain a deeper understanding of the interactions that exist between the complex systems and genetic algorithms disciplines.

2 Process and Methodology

2.1 Genesis

Genesis, written by John J. Grefenstette, is a freely-available application for "function optimization using genetic search techniques" (1990). It is highly flexible, providing the user with control over population sizes, generations and trials run, crossover and mutation rates, various facets of each gene in the sequence, and great many other variables.

2.2 XBoids

EvoBoids is based on Conrad Parker's implementation of the boids system (2002). It is written in C, using the X Window system, and available under the GNU General Public License.

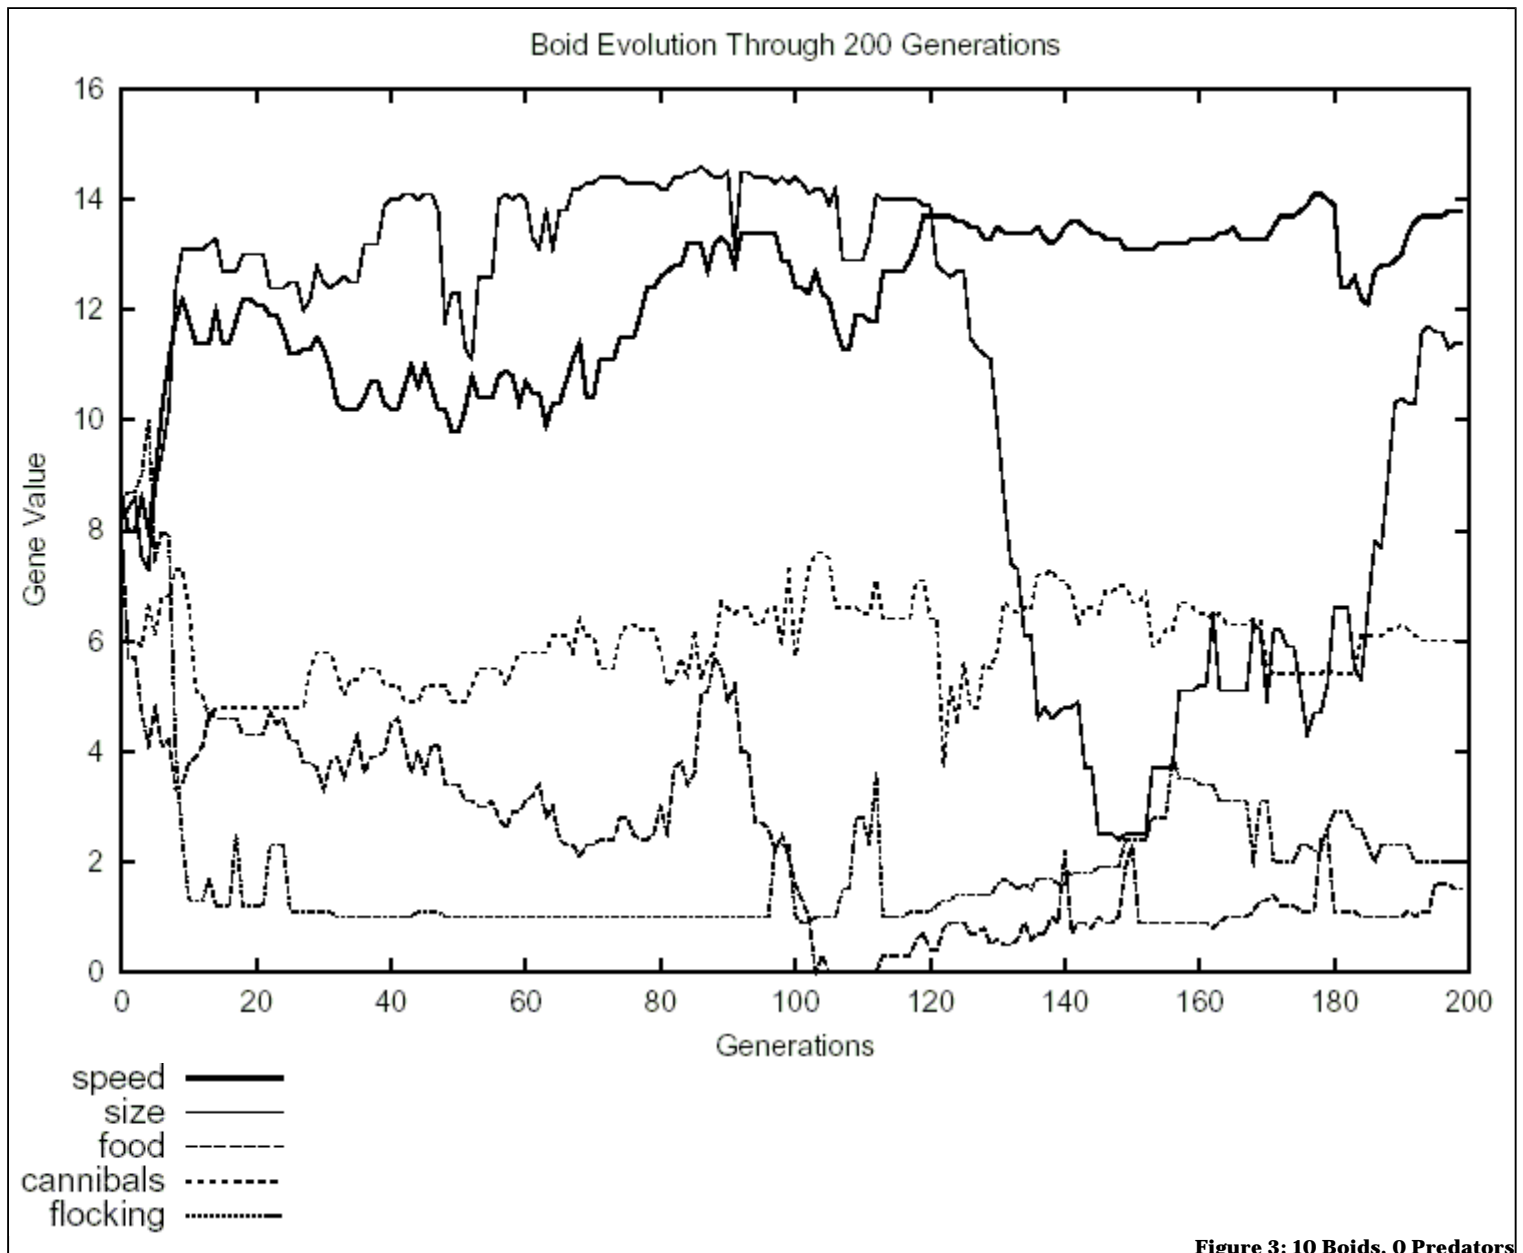


Figure 3: 10 Boids, 0 Predators

2.3 Integrating Genesis and XBoids

We first altered both the Genesis and XBoids systems to allow for their direct communication. Genesis creates two-dimensional arrays containing five genes per EvoBoid (gene values are randomized for the initial EvoBoids population) for a set number of agents (we experimented with two different population sizes). The XBoids program then accesses these arrays for the purpose of adding EvoBoids to its simulation.

2.4 The Gene Sequence

Each EvoBoid has a genetic sequence consisting of five genes:

1. *Size*: the size of this EvoBoid.
2. *Maximum Speed*: the EvoBoid's top flight speed.
3. *Food Capacity*: how much food this EvoBoid can hold.
4. *Tendency to Cannibalize*: how frequently this EvoBoid will eat other EvoBoids to survive.
5. *Tendency to Flock*: if this EvoBoid will actually flock (turns on or off Reynold's flocking rules).

These genes were selected empirically, based on several underlying tenets. We wanted genes that were biologically sensical. We wanted genes that we could envision producing interesting results. We wanted to limit the size of our gene sequence initially, to a length that would not be overwhelming and unmanageable and upon which it would not be impossible to speculate concerning cause and effect. Size, speed, food capacity, the tendency to cannibalize, and the tendency to flock fit this list of requirements perfectly.

Genes affect one another. That is, the effects of one gene may interact or overlap with the effects of another gene. For example, an EvoBoid that is large has a high food capacity but a lower maximum speed. The values for each gene are set by Genesis. As mentioned, the first generation is assigned pseudo-random values, and successive generations have values derived through Genesis' evolutionary process, which employs crossover and mutation. The EvoBoid C structure also has other variables indicating certain properties of the boid's state:

1. *Food*: the amount of food an EvoBoid has on board at a given time. This value is necessarily less than or equal to the boid's food capacity. EvoBoids lose food by flying and gain food by feeding.
2. *Hungry (a boolean)*: an EvoBoid becomes hungry when its food level falls to or below 1. When hungry, it will tend to fly toward the ground, where it can perch and feed. If an EvoBoid is cannibalistic and hungry, it targets the nearest other EvoBoid. If that EvoBoid is closer than the ground (where the hungry EvoBoid can perch and eat), the hungry boid will chase after its nearest potential prey. If the hungry agent is fast enough, it will overtake the other. If this occurs (implemented in our program as a computation of distance; that is, if the hungry EvoBoid comes within a relatively small distance of its target), then the EvoBoid being chased is killed and messily devoured. The cannibal is changed from hungry to not hungry, its food is increased, and its starvation counter is reset.
3. *Starve*: each EvoBoid has a starvation counter. If an EvoBoid remains hungry for more than a specified period of time, it starves to death.

Though an EvoBoid has values for each of its genes, these values are not used exactly when the boid's behavior is calculated in its environment. Gene values are generally set in expressions, effectively scaling their potency up or down. This is done in an effort to achieve a balance of behaviors by eliminating tendencies toward extremes. For example, EvoBoids are set to cannibalize at a frequency of approximately one eighth their genetic value. Previously, cannibal EvoBoids had been set to eat other EvoBoids at approximately one half their genetic value, but this proved too frequently: agents were dying much too quickly to allow for adequate observations of behavior.

2.5 Environmental Additions

As described previously, survival provides the EvoBoids system with its fitness evaluation function. The

longer an EvoBoid survives, the greater the fitness its gene sequence is awarded. The addition of certain aspects to the system creates an environment more suitable for making use of "survival" in this fashion. The feeding and perching requirements to prevent starvation are one such environmental factor affecting an EvoBoid's survival. The inclusion of predators is another.

Predators exist in two forms. Cannibalistic EvoBoids are one form. "Strict" predators are the other. These are modeled graphically exactly as EvoBoids are, but with different colors (in the current system, EvoBoids are blue and yellow, predators are green and purple). Swapping color palettes was simply an efficient means of allowing a user easy visual differentiation; more complicated or distinct polygonal models would have required more time to implement and would slow the system down, which negatively affects the user's ability to conduct experiments (with large population sizes, for example). Predators behave more simplistically than EvoBoids. They do not evolve. Instead, they merely fly around, seeking out the nearest EvoBoid and attempting to capture it for prey.

2.6 The Fitness Function

XBoids creates an array of fitnesses based on the period for which a given EvoBoid survives. An EvoBoid that survives for the entire duration of the system run in which its genes are evaluated receives the highest fitness.




Figure 4: 10 Boids, 2 Predators

Likewise, EvoBoids who die quickly are assigned low fitnesses. The fitness array is delivered to Genesis, which copies it into its own evaluation function for use in producing the next generation of agent gene sequences. Thus, those EvoBoids whose genes are evolutionarily most viable are favored in the reproduction algorithm and more likely to survive through successive generations.

2.7 System Configuration and Experiments

Experiments were conducted using a crossover rate of 0.9 and a mutation rate of 0.005, both of which were kept steady throughout. Instead of altering these variables between experiments, we elected to change the EvoBoid population size and vary the presence or absence of "strict" predators in the environment. Each experiment was allowed to run for two-hundred generations, and each generation was set to run for two-hundred fifty clock ticks. The particular number and duration of generations per experiment were selected to maximize the time available to gauge any given gene sequence against the rubric of "survival of the fittest" while also permitting us to meet our situational time constraints. The breakdown of experiments is as follows:

- 5 experiments were conducted with EvoBoid population size 10 and without any "strict" predators present in the environment.
- 5 experiments were conducted with EvoBoid population size 10 and with 2 "strict" predators present in the environment.
- 5 experiments were conducted with EvoBoid population size 50 and without any "strict" predators present in the environment.
- 5 experiments were conducted with EvoBoid population size 50 and with 10 "strict" predators present in the environment.

The proportion of EvoBoids to strict predators in a given environment was purposefully kept constant, as a control. Considering the number of influential factors already existent in the EvoBoids system, we concluded that other, secondary or outside influential factors would prove more of a hindrance than a benefit. That is, when speculating on cause and effect, we preferred a carefully limited set of causes upon which to pin the set of effects.

3 Results

Upon conducting our experiments, we noticed several differences in the various evolutionary paths taken by our EvoBoid populations. For smaller populations with no predators, speed and size were consistently the genes that evolve the highest values. EvoBoids evolved to grow larger and faster. On a couple occasions, however, large variations in certain genes occurred. For instance, in Figure 2, cannibalism evolves to be very prevalent (78 percent of the population) around generation 60. Around generation 120, it drops back down to its average (37 percent). This appeared to be a random occurrence: we could not find any patterns in the data to explain it.

Another example of such a surge occurs in Figure 3, where size (which normally averaged 10-15 (63-94 percent) dipped around generation 120 to an average of about 4 (25 percent).

Food capacity and flocking frequently evolve to have the lowest values in the system. The average food capacity is 0.5 to 6, and the average for flocking appeared to hover between 0 and 5 (0 to 31 percent). There were a couple of correlating spikes in values of pairs of genes. For example, one experiment showed food capacity and cannibalism spiking around generation 150 to a value of about 5. We do not believe any link exists between these two genes, however, because this pattern failed to reappear. Similarly, between generations 150 and 180 of one experiment, food capacity and flocking both spiked, though food capacity spiked to a greater magnitude of 8, while flocking spiked at 4. We do not believe there is a correlation here, either, because this is also a pattern which does not recur.

Cannibalism for this set of data seemed to take a middle-of-the-road approach for all generations. Its values averaged from 4 to 8. This gene seemed to vary widely, though it usually settled around its average relatively

quickly. Its median is stable around 6.

Keeping a small population size but adding two predators had a large mitigating effect on the degree to which gene values varied widely. After adding predators to the system, we saw fewer spikes, and EvoBoids seemed to quickly settle on, and stick to, values for the genes. Cannibalism became the most important gene in the system. It surged almost immediately for all but the last trial. Its average value hovered around 15 (again, excepting the last trial). We speculate that the reason for this newfound importance of cannibalism results directly from introduction of strict predators to the system. It becomes evolutionarily more viable to cannibalize, to attack and devour flockmates (often in mid-flight), when "normal" feeding requires perching. Perched EvoBoids are much easier prey for strict predators than moving EvoBoids. Similarly, speed was also a highly valued gene, with averages around 13. Higher speed allows flying EvoBoids to more easily avoid both strict predators and cannibalistic flockmates.




Figure 5: 50 Boids, 0 Predators

This is not to say that there were no spikes or variations in the genes following the addition of predators. The spikes that did occur, infrequently though they appeared, are significant. In Figure 4, size drops from 15 around generation 100 to near zero by generation 120, and stays low for the remainder of the experiment. Similarly, in one experiment, flocking surged at the outset of the experiments, and then dropped around generation 50, staying low for the remainder of the experiments. Flocking at generation 100 plummets in another experiment, falling to a minimum near the end of the experiment. It went from a value of 15 to nearly zero. A huge spike occurred with food capacity in a separate experiment. At generation 150, the value jumped from 2 to 13, and remained this high for the rest of the experiment.

With a larger population size (50 EvoBoids) but no predators, we observed fewer large spikes in the gene data. Speed and size were once again the most highly valued genes, hovering around 11. Cannibalism averaged around 5 and 8, remaining fairly constant throughout the experiments. In Figure 5, cannibalism rises as size falls. However, shortly thereafter size resumes its previous levels, while cannibalism remains the same. The opposite effect, with cannibalism falling and size increasing mildly, occurred in another experiment. Food and flocking once again hovered in the lower value ranges. They both maintained levels between 1 and 3.

When we added predators to this larger population, it failed to have the mitigating effect we saw with fewer EvoBoids. Once again, speed and size evolved to have the highest values. In some experiments, flocking takes on a much higher average value (up to approximately 10, from its original non-predators value of 3). In others, cannibalism seems less valued, and it falls to an average of about 5. Without predators, it stayed around 8.

Once again, food capacity was not highly valued in these experiments. Interestingly, however, one experiment with 50 EvoBoids and 10 predators shows a spike around generation 180, where food capacity, cannibalism, and flocking all shot up. Size and speed were increasing at this point, as well, so the spike may be an anomaly, or it may be linked somehow to those genes. Another coincidental spiking occurred in a similar experiment. There is a small but significant spike in food capacity around generation 150, at which point flocking, speed, and size also rise briefly.

4 Conclusions

It is clear that this system needs more balancing. Food capacity should have more of an effect on the system; we believed initially that the addition of starvation to our system would kill off EvoBoids with low food capacity. Either our EvoBoids simply didn't starve, or they evolved to be cannibalistic so that they had a quick and ready food supply. However, there is little evidence of this in our data. As for flocking, we naturally assumed that flocking would become prevalent, as we observe this behavior in real birds. However, our EvoBoids largely moved away from flocking. This was perhaps due, in part, to the cannibalism factor. EvoBoids did not avoid cannibals, because they were unaware of the fact that their neighbor might have cannibalistic tendencies. This shows a correlation between flocking and cannibalism: the closer EvoBoids came to one another, the easier they were for cannibals to prey upon.

We also expected to observe a correlation between food capacity and size. Larger EvoBoids should have a great food capacity, and larger EvoBoids also provide more food for cannibals. However, we saw no correlation between cannibalism and EvoBoid size. This indicates that our system needs to be tweaked. We need to refine the connections between genes and make them more realistic. Perhaps such overlaps in gene effects could also be included as genes (see section 5).

We saw fewer spikes with 50 EvoBoids because agent deaths had a lesser impact the gene pool. They tend to average out the fitnesses, because a single predator cannot greatly impact the system, as it can with only 10 EvoBoids. When we had 10 EvoBoids and we added 2 predators, we observed a mitigating effect on the rapid fluctuation of gene values. However, we do not see this in larger populations with EvoBoids.

We theorize that this is due to the fact that a small population will only cover a limited number of values for the EvoBoids' genes. If there are a small number of EvoBoids with large speeds, and the slower ones are quickly killed off by predators, quick EvoBoids will evolve and be more prevalent in future generations. There are fewer ranges of speeds, due to the decreased number of EvoBoids. With 50 EvoBoids, there is a large enough range of genes so that a predator cannot kill off an entire subset of genes. That is, a predator cannot kill off every slow EvoBoid; there will be some very slow EvoBoids, and some slightly faster EvoBoids, that it could not kill.

5 Future Work

Our future work on the EvoBoids combination system will include a number of courses.

The gene sequence requires expansion and revision. The genes for each agent EvoBoid could also include other characteristics an actual bird might have. The possibilities in this direction are quite vast, as the actual genetic sequence for a living, breathing organism is vast. The system could grow to incorporate qualities such as:

- Physical features: the phenotype of a given organism, the outward manifestation of that organism's genotype, obviously affects its ability to survive. EvoBoids should be given genes representing physical characteristics beyond simple "size." Some birds have talons, for example, with which they hunt and defend themselves. EvoBoids could have the option of evolving this trait. Other physical traits affect a bird's ability to perceive its environment, such as eyesight, or to otherwise prevent death, such as camouflage body coloring.
- Predispositions toward certain behaviors: predatory EvoBoids could evolve a tendency to hunt either solitary prey or prey that flock more frequently, for example.

Certain characteristics of an EvoBoid might also require further development of the environment in which they dwell. Perhaps an agent EvoBoid could have the option of being amphibious (not particularly realistic, but cannibalistic birds are a far-fetched concept as well). This would require the inclusion of some form of water into our system. This introduces the potential for new types of prey. Taken to its extreme, this concept could metamorphose into a full-blown attempt at a simulated ecosystem.

As more complications arise, however, the system moves increasingly far from under the heading "complex system," where simple rule sets rule. Maintaining a balance between the simplicity inherent in the coordinating laws of a complex system and the intricacies of implementing such a simulated ecosystem requires a great deal of care. However, as emergent behaviors, such as flocking, are characteristic of complex systems, any relatively simple addition to the EvoBoids/Genesis system that does generate new emergent behaviors would not seem to have the effect of moving the system away from the definition of "complex."

Other future work could include the evolution of predators. Different species competing with one another could lead to what is referred to (in the complex systems arena) as the "evolutionary arms race," where different species are continuously evolving to become more fit than others to survive.

Still further future work could include more runs on system variations already considered, and runs on those variations not considered. Differing population sizes, crossover and mutation rates, etc, are all subject to further experimentation.

Acknowledgements

We would like to especially thank Clare Bates "Den Mother" Congdon, our advisor on this project, for her knowledge, wisdom, general helpfulness, and handy candy stash. Also, our colleagues, Marc "Asleep Till Noon" Attiyeh, Kevin "C-Master" Septor, Jon "C-Virgin" Nickerson, and David "Concussion" Cohen, for their feedback and assistance.

Bibliography

Grefenstette, J. J. (1984). GENESIS: a system for using genetic search procedures. Proc. Conf. Intelligent Systems and Machines, Rochester MI, 161-165. Software available at:
<ftp://www.aic.nrl.navy.mil/pub/galist/src/genesis.tar.Z>

Parker, C. XBoids C source code (2002). Available at: <http://www.vergenet.net/~conrad/boids/download/>

Reynolds, C. W. (1987) Flocks, Herds, and Schools: A Distributed Behavioral Model, in Computer Graphics, 21(4) (SIGGRAPH '87 Conference Proceedings) pages 25-34.