

Simulated Evolution: Facilitating Students' Understanding of the Multiple Levels of Fitness through Multi-Agent Modeling¹

Uri Wilensky & Damon Centola
 Center for Connected Learning & Computer-Based Modeling
 Northwestern University, Evanston, IL
 Email: uri@northwestern.edu

{complex systems, mean field approx}

I. Introduction

In this paper, we describe a computer-based approach to learning about evolution. In the Simulated Evolution project we investigate students' understanding of evolutionary concepts and design multi-agent computer models for student exploration and analysis of evolutionary scenarios. The Simulated Evolution project is embedded within a set of projects we have undertaken to investigate students understanding of complex systems and to design effective agent-based modeling activities that help students make sense of such systems (e.g., Wilensky, 1999c; Wilensky, Hazzard & Froemke, 1999; Wilensky & Reisman, 1998). In the context of these umbrella projects, we have created units and sets of activities about evolution, among these units on population biology (Wilensky, Hazzard & Longenecker, 2000), classic evolutionary examples (e.g., Wilensky, 1998c) and the evolution of cooperation. (see the EACH project (Centola, Wilensky & McKenzie, 2000)). The agent-based modeling language, NetLogo (Wilensky, 1999a), serves as the computing substrate for building and running the evolutionary models. Recently, we have integrated the various units into a more comprehensive curriculum called BEAGLE (Biological Experiments in Adaptation, Genetics, Learning and Evolution, (Wilensky, Novak & Rand, 2004)). In this paper we report on a small study we conducted in 1998/1999 with students working with the EACH unit to learn about the evolution of cooperation. A more comprehensive description of the Simulated Evolution project and the BEAGLE curriculum can be found in (Wilensky & Novak, in preparation).

The prevailing approaches to teaching evolution in secondary and early post-secondary education are intended to lay the groundwork for students to understand the basic concepts of fitness, natural selection, and evolutionary success. In practice, however, what students retain from this instruction is largely the phrase "survival of the fittest" (Brumby 1984; Centola, Wilensky & McKenzie, 2000; Ferrari & Chi, 1998). Many

¹ This paper was adapted from Centola, D., Wilensky, U., & McKenzie, E. (2000). Survival of the Groupiest: Facilitating students' understanding of the multiple levels of fitness through multi-agent modeling - The EACH Project. *Proceedings of the Fourth International Conference on Complex Systems*, Nashua, NH: New England Complex Systems Institute and *InterJournal Complex Systems*, 377. Adaptations were made to reflect current software and curricula.

students who have taken introductory high school and college courses in evolution come away from their studies understanding natural selection in terms of strong individuals dominating over weak individuals. Even students who participate in curricula that do not explicitly teach this view of natural selection will often come out of instruction with these simplified understandings (Brumby, 1984; Greene, 1990). Research suggests that students' misconceptions about how the evolutionary process works can often be traced beyond the students' coursework to fundamental intuitions about how natural systems behave (Bishop & Anderson, 1990; Ferrari & Chi, 1998).

In over a decade of previous research we have described the difficulties students (and adults) have in making sense of complex systems. In particular, we have documented the difficulties students have with understanding the role of randomness and decentralized control in complex systems and how understanding such systems requires thinking at multiple levels (Resnick & Wilensky, 1993; Wilensky & Resnick, 1999). A large body of research has reported the difficulties students have in making intuitive sense of random variation, probabilistic runs, decentralized control – all sub-component processes of natural selection (Tversky & Kahneman, 1974; Konold, 1991; Resnick, 1996; Wilensky, 1997b; Wilensky & Resnick, 1999). Because most students have not had many opportunities to interact reflectively with complex systems, they have not developed the intuitive knowledge that would help them make sense of such systems. We have designed agent-based modeling environments and activities that enable students to explore complex systems across a wide variety of domains and documented their increased understanding (Abrahamson & Wilensky, 2004; Blikstein & Wilensky, 2004; Wilensky, 1999b; Wilensky & Reisman, 1998). We view evolutionary processes – processes that operate at multiple levels – as fundamental processes of change of complex systems. Without challenging and further developing their intuitions about how natural systems work, students are unable to internalize an understanding of many of the important concepts in evolutionary theory (Greene, 1990). Evolution is therefore a natural domain for the design of agent-based models to further understanding.

One important aspect of evolutionary theory that is difficult for introductory students to understand is the multi-level perspective (Sober & Wilson, 1998; Mitteldorf & Wilson, 2000; Wilensky & Resnick, 1999). While a complex systems approach to evolution encourages students to think about concepts such as fitness and natural selection from a number of perspectives, students given traditional instruction typically try to pin down these concepts into singular meanings. In particular, students' understanding of fitness is often oversimplified by understanding it as an individual's strength in one-on-one competition for resources (Bishop & Anderson, 1990). Because of this narrow conception of fitness, many interesting aspects of evolutionary biology such as the plausibility of altruistic and cooperative behavior being evolutionarily advantageous (Axelrod, 1984; Wilson & Sober, 1998; Mitteldorf & Wilson, 2000) are conceptually intractable to students of evolutionary theory (Jacobson, 1996). The Simulated Evolution project was established to enable students to develop more sophisticated intuitions about evolutionary processes by interacting with multi-agent models of evolutionary systems. The primary goal of the EACH unit is to help students to think about the complex dynamics of evolution that would allow for altruistic and cooperative habits to be

evolutionarily successful. One important goal of this research is to help students to think more critically about their assumptions about fitness, and how these assumptions play into their expectations for evolutionary scenarios. Specifically, we want students to see the significant differences between understanding fitness at the individual-level, at the gene-level, and at the group level, and how understanding fitness at multiple levels can dramatically change their conceptions of other evolutionary processes, such as the evolution of altruism. Before the advent of powerful computation, calculating trait fitness for individual contexts was intractable, so trait fitness was conceived as averaged over all environments. By taking an agent-based modeling perspective on evolution, students are enabled to think about trait fitness from the point of view of a specific individual in a specific environment.

II. Project Frame

The Simulated Evolution project is one of several curricular interventions undertaken as part of several NSF funded research projects. These projects share the goal of building computer-based tools and curricula to enable students to explore and make sense of complex systems and to study student sense-making using these tools. As part of these projects, we have developed several agent-based modeling languages and environments culminating in the multi-agent modeling environment NetLogo (Wilensky, 1999a). NetLogo is a multi-platform agent-based modeling language and integrated modeling environment. All activities described herein are written in NetLogo or in a previous less developed language, StarLogoT (Wilensky, 1997a). NetLogo was designed explicitly for exploring systems with multiple interacting "agents". NetLogo is a prominent representative of a new class of such multi-agent (AKA agent-based) modeling languages. While most of these languages were designed for researchers, NetLogo was explicitly designed as a tool for both learners and researchers and adheres to the design criterion of "low threshold and high ceiling" (Tisue & Wilensky, 2004). Using NetLogo, students can represent many different types of "agents," such as flashing fireflies, cars in traffic or molecules in a gas. They can then build models of the behavior and interactions of thousands of such individual agents (Wilensky, 2001). Two of the models described herein were originally developed in StarLogoT (Wilensky, 1997) and have subsequently been ported to NetLogo. More recent project models have been developed in NetLogo. Besides its more advanced modeling features, NetLogo also incorporates the HubNet tool (Wilensky & Stroup, 1999), which opens avenues for incorporating interactive agents and multi-user scenarios; these possibilities will be discussed briefly in the concluding section.

The latest version of NetLogo and an associated large collection of sample models (collectively entitled the "NetLogo Models Library (Wilensky, 1999)) are available for download at [.ccl.northwestern.edu/netlogo](http://ccl.northwestern.edu/netlogo). The models are drawn from a wide range of disciplines including physics, biology, mathematics, computer science, chemistry, materials science, ecology, economics, urban studies and linguistics.

III. The Context

As we have mentioned, Simulated Evolution emerged from a succession of projects in which we studied students at many different grade levels as they explored and created

multi-agent models of complex phenomena. This work was conducted at the Center for Connected Learning and Computer-Based Modeling (CCL), first located at Tufts University and then moved to Northwestern University. Students (from middle school to graduate school) came to the CCL and met informally with project staff who mentored them in learning NetLogo. Students explored models from the project library according to their interests. They were encouraged to vary and extend the underlying NetLogo code for the project sample models and, when they felt ready, to create new models from “scratch”. The topic of evolution, in general, is one that has appealed to many of the student modelers. For the high school students and many of the undergraduates as well, evolution was studied largely in “story” form. That is, they learned stories of how the theory of evolution came into being, how it was resisted and how it is supposed to work. Many of the students found these stories intriguing but unsatisfying. The dissatisfaction consisted of feeling that they did not have an adequate methodology for testing the plausibility of particular evolutionary arguments or for evaluating competing evolutionary claims,

The Simulated Evolution project is an on-going endeavor to teach students about complexity in evolution. The EACH unit began with a focus on getting students to explore models wherein altruism, cooperation, and group behaviors could survive under Darwinian selection. The motivation behind this unit was to expand students’ notion of fitness and to get them to take into account the complex field of factors that play into the evolutionary process. The unit concentrates on students’ conceptions of fitness, and is directed at helping them to develop intuitions about evolution as a multi-level process. Additional EACH unit models, not included in this paper, addresses student conceptions of randomness in natural selection, and focuses on fostering students’ intuitions about genetic variation in populations and the mechanism of genetic drift..

The Students

The students participating in the present study were undergraduates at Tufts University. These four volunteers were drawn from the sciences and social sciences departments. Beverly was a junior biology major; she came to the Simulated Evolution project interested in getting a ‘complex systems’ understanding of evolution. David was a junior biology major who came to the Simulated Evolution project because he wanted to understand how altruism could be an evolutionary advantage. Amy was a sophomore sociology major who wanted a better understanding of how the social and group dynamics of populations could evolve. Derek was a freshman chemistry major who came to the Simulated Evolution project because of his desire to combine his interests in evolutionary biology with his interests in mathematics. They had all studied evolutionary theory in either a high school or university class. They were recruited to the study by an advertisement offering the opportunity to interact with computer models and simulations that could result in improved understanding of evolution.

We began the study by meeting with each student individually and asking them to discuss their ideas on evolutionary theory. All of the students expressed common intuitions about individual selection, and at some point in the discussion mentioned “survival of the fittest.” But when asked to flesh out the idea of fitness, the students had difficulty

articulating what ‘fitness’ meant to them beyond the popular notion that fitness was strength: the power of an individual to dominate other individuals. These pre-session interviews gave useful guidance to the research team on how to scaffold student interaction with the models. They also served to set the ground for the project intervention as students started to question their assumptions about the meaning of "fitness" – questions that were elaborated and discussed once they began to interact with the models.

Overview of the Activities

As part of the Simulated Evolution project, we have developed a series of models that enable students to interact with certain evolutionary scenarios in order to discover what outcomes are plausible, how their assumptions about fitness play into their expectations, and how different environmental and social considerations can affect the fitness of different populations. These models serve as seeds for developing intuitions about evolutionary processes. Students run these models exploring their behavior under a variety of conditions. They then modify and extend the models to further refine their thinking. Ultimately, our hope is that most students participating in the Simulated Evolution project will be inspired to develop and discuss their own models of evolutionary scenarios, thereby developing a self-sustaining means for testing their intuitions about evolution, and gaining an increasingly rich understanding of the complexity of the evolutionary process.

The Simulated Evolution project activities are structured around basic sets of “seed” models. For the EACH unit, the goal of these models is to raise basic questions about fitness, altruism, and group dynamics in natural selection that will encourage students to explore these issues further. The activities we describe herein are based on a combination of our “seed” models and a set of models that were developed by students while engaged in the EACH unit. The original “seed” models, the Altruism models (Wilensky, 1997c), were developed during the early stages of the project when the focus of our research was directed more specifically toward helping students to understand how the evolution of altruism is possible. During an early session, we collaborated with a student-learner, who was inspired by the Altruism models to find a more behavioral account of how altruism could evolve, to develop a model called the Cooperation model (Wilensky, 1998b). The Cooperation model gave a more embodied feel to the issues surrounding cost, benefit, and altruistic behavior. The Cooperation model shows how an altruistic behavior that would be selected against under normal conditions, could survive under conditions of high population viscosity. Because the Cooperation model is closer to student experience than the original Altruism model, we decided to adapt a version of it to use in our research on students’ understanding of fitness and how it relates to their understanding of complexity in evolution.

IV. Activities

In the research described in this paper, we used two sets of models, the Altruism set and the Cooperation set, as complementary ways of exploring the issues of fitness and complexity in evolution. The activities are organized in two parts. We first presented the Altruism models in order, letting students become familiar with the first model, and its

unsurprising consequences, and then gave them the second, more problematic model. After some discussion, we gave them the first Cooperation model, and let them explore the parameter space. Once they became comfortable with this model, we asked them to apply the lessons from the Altruism model to the case of the Cooperation model. Over the course of their involvement with the EACH unit, students made surprising advances in their understanding of group behavior and in their conceptions of the differences between individual, group, and gene-level fitness. They came away from the activities with a more critical sense of how distinctions between individual and gene-level fitness determine how they understand group-level phenomena such as the evolution of altruism. Most importantly, they began to see that group behaviors are phenomena that need to be considered from multiple levels of interpretation in order to understand how they can function in evolution.

The First Altruism Model²

The first model used in the EACH unit was a conservative model of evolutionary behavior based on evolutionary biology theory (Hamilton, 1964; Wilson, Pollock & Dugatkin, 1992). This first model (written in StarLogoT) has two types of agents: selfish and altruistic. The premise of the model is that each agent “looks around” to its neighbors and sees whether they are altruistic or selfish³. For each altruistic neighbor, the agent increases its fitness by a fixed value. Thus, each agent with altruistic neighbors would have a higher fitness in the reproductive lottery. Each agent then calculates its fitness with one of the following simple equations:

If I am a selfish agent, my fitness is $1 + \text{my benefit from my altruistic neighbors } (B_n)$.
If I am an altruistic agent, my fitness is $1 - \text{the cost of being an altruist } (C) + B_n$

The values of altruistic benefit (B) and cost (C) are determined by slider-variables. We initialized these variables to values such that the cost of being an altruist significantly weakens the fitness value of altruists, while the benefit from the neighboring altruists is only significant when there are two or more of them.

Each agent calculates B_n as $B * \text{proportion of ones neighbors (including oneself) that are altruists}$.

To illustrate the fitness updating, we present a quick example:

Let $B = .5$ and $C = .2$.

Suppose you are an altruistic agent with fitness 1 and that you are surrounded by two altruistic neighbors and two selfish neighbors. Then the proportion of your altruistic

² The two Altruism models were combined into one model in the StarLogoT models library. That model can be downloaded at <http://ccl.northwestern.edu/cm/models/altruism/>. To get the behavior of the Altruism 1 model, set both the harshness and disease sliders to zero. The Altruism model was converted to NetLogo and can be downloaded at (Wilensky, 1998a, <http://ccl.northwestern.edu/netlogo/models/altruism>).

³ In agent-based modeling, it is common to take the point of view of the agent and talk about it as if it is active. But, in fact, this model does not assume that the agents are active, the benefit conferred on the central agent can be purely passive.

neighbors is $3/5$. You would therefore update your fitness as $1 - .2 + .6 * .5$ and your new fitness would be 1.1 (See Figure 1).

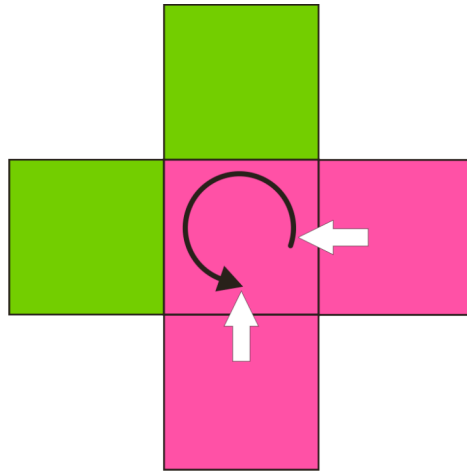


Figure 1. The altruistic agent (pink) at the center recalculates its fitness based on the proportion ($3/5$) of altruistic agents in its neighborhood.

Each agent is at the center of a five-agent neighborhood. To enter the reproductive lottery, each agent looks at its neighbors (in the four cardinal directions) and assesses their fitnesses. All of the altruistic agents' fitnesses are summed (including the central agent, if it is an altruistic agent), and all of the selfish agents' fitnesses are summed (including the central agent, if it is a selfish agent). These sums constitute the weight of altruistic 'seeding,' and the weight of selfish 'seeding', respectively, for the center spot in the reproductive lottery. To complete a generational cycle, a random number is picked between 0 and the total of the weights for a spot. The spot is then given to the type of agent whose weighted chance fell on the side of the random number. The higher an agent-type's weight, the greater the chance that that type of agent will win the spot.

We included control parameters in the model, so that students could set the relative density of the altruistic and selfish populations. The clear outcome of this initial model is that the fitness of the selfish agents is higher than the fitness of the altruistic agents. Independent of the initial population densities, after on average two-hundred generations, all of the altruistic agents are extinct; the upper bound on this time, starting with almost no selfish agents, is about five-hundred generations before the altruists are extinct. The outcome of this model fits well with standard adaptationist accounts of evolutionary theory (Hamilton, 1964; Williams, 1992).

Figure 2a below is the control interface for the StarLogoT Altruism-Model 1. Figure 2b shows the model when it starts running. The pink agents are altruistic and the green are selfish. Figure 2c shows the model after it has run for about 100 generations (about thirty seconds running on a typical project computer at the time, but about one second on a typical current personal computer).

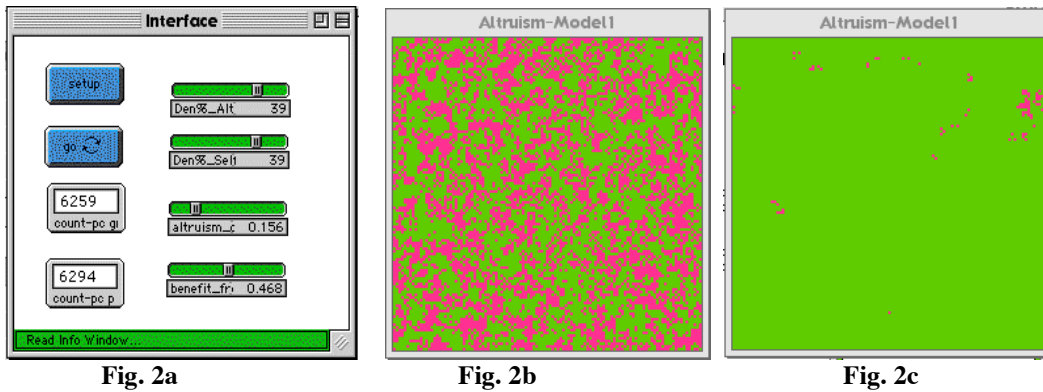


Figure 2. The Altruism Model 1

The results of running this model were consistent with the students' expectations for evolutionary theory. All of the students predicted that the selfish agents would win. Clearly, the students noted, the selfish agents were the stronger agents, and clearly they had a higher individual fitness. After running this model a few times, students felt comfortable with the parameters, and were confident that they understood how the evolutionary scenario was playing out. By enabling the students to develop a feel for the modeling environment by running a model with a familiar outcome, we hoped to encourage them to develop explanations about how and why the scenario was playing out as it was.

When we asked her why the selfish agents had won, one student, Derek, a freshman chemistry major, said, "The selfish agents are stronger because they have a higher fitness value." But, when we asked Derek to expand on the idea of having a higher fitness value, he said "Having a higher fitness is being strong enough to get more food, which makes it more likely to reproduce." This was a typical response; the students viewed fitness as an individual's strength in competition for resources. Another student, David, a junior in biology, said, "fitness is the strength of a genotype to out-compete other individuals." David's notion of fitness demonstrated a conflict between his understanding that there are gene-level phenomena, and his intuition that fitness must be a property of individuals.

The Second Altruism Model

After getting the students' reactions to the first model, we introduced them to the second model. The second model, based on recent work in the evolutionary biology of cooperation and altruistic behavior (Mitteldorf & Wilson, 2000), introduces a new element to the model: adversity. The second model adds a slider-variable called 'Harshness,' and a slider variable called 'Disease.' The harshness variable gives each empty "patch" (a unit of the screen that is not occupied by a selfish agent or an altruistic agent) a chance of staying empty (resisting population) each turn. Envisioning the model-world as an environment in which individuals need to occupy a space on the grid in order to live, the harshness variable limits population growth by, at each clock "tick", making some of the spaces uninhabitable. The disease variable is incorporated into the reproductive lottery for each spot. The value of this variable corresponds to the chance that an agent that occupies a spot on the grid will fail to reproduce, and thus that the spot will become empty. The introduction of these elements into the model alters the

relationship between individuals and their environment. The new threats to the well-being of individuals reframe the importance of group behavior on individual success. This second model has a much richer parameter space -- students can explore the effects of the various harshness and disease values on the stability of the altruistic and selfish populations.

Figure 3a below is the control interface for the Altruism-Model 2. Figure 3b shows the model after thirty generations. Figure 3c shows the model after two-hundred-fifty generations.

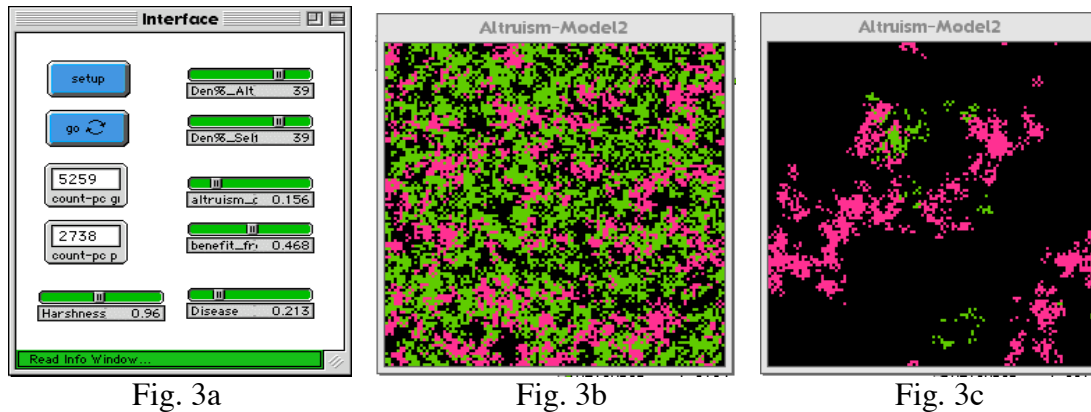


Figure 3. The Altruism Model 2

Students found that when the harshness value is set around .96, and the disease value is set around .2, the model shows the surprising result that the altruistic population fares far better than the selfish population. Why should this be? This counter-intuitive result comes from the fact that when the natural conditions are harsh enough, single agents cannot survive against nature. That is to say, even though the selfish agents would fare better in competition against the altruistic agents, the altruistic agents would fare better against nature than the selfish agents. The reason for this is that altruistic agents survive in groups of altruists. Each altruist adds to the fitness of the other altruists. While a lone altruistic agent can survive no better than a lone selfish agent, because the altruists contribute to the well-being of the agents around them, the community can survive, with its combined altruistic benefit, in the face of harsh conditions⁴.

The students who participated in the project were all very surprised by this result. After running the model a few times, students began to ask questions about how the altruists were winning. It was clear from the outset that the selfish agents were the stronger, or more fit, agents, and yet under harsh conditions the altruists were running the selfish agents to extinction. We asked students to describe what was happening on the screen, and how it was different from the first model. They all recognized that, unlike the first model, the altruists were grouping together.

⁴ While the Altruism 2 model presented here relies on harsh conditions to enable altruism to survive, Mitteldorf and Wilson's result does not require harsh conditions – large regions of parameter space can enable altruism to survive as long as the density of the total population is not equal to 1 – so if there are vacant areas or multi-populated areas then altruism can survive.

After discussing this for a few minutes, one student, Beverly, a junior biology major said, “Well, it means that all of the altruistic benefit is going only to altruists; which means that under harsh conditions, either you are an altruist, or you’re on your own.” She suggested that altruists were in fact more fit in the second model, but only as a group. We discussed her earlier comments about individual fitness, and asked her how she could make sense of the fact that this group advantage was at odds with her previous focus on individual strength. After some group discussion about how it was possible to have multiple levels of fitness, Beverly resolved this apparent conflict, “So the altruistic gene is more fit under harsh conditions, even though the individuals have a low fitness, because the harshness allows the altruists to clump together, and benefit only one another.”

Most of the students made similar inferences about the relevance of grouping to gene-level fitness. Many of them remarked that the selfish agents were still more ‘fit’, even though the altruists could survive better. But, they were puzzled about this apparent equivocation in the term ‘fitness.’ We encouraged them to address this confusion by concluding this part of the sessions with a discussion of the idea of the “survival of the fittest”, and what that meant. Amy, a sophomore sociology major, had earlier been confident that this catch-phrase summed up natural selection as a process that favored the strongest individuals, but was now unsure of how to think of ‘fitness’: “The altruistic group has a higher fitness, as a group, than the selfish group, but I don’t think that they are any stronger as individuals. They [the altruists] are stronger now [under harsh conditions] because being in a group makes gene-level selection matter.” Amy was beginning to see that thinking about fitness at the gene-level could change her perspective on group-level selection. Ultimately, she suggested that the phrase “survival of the fittest” was misleading. Instead of “survival of the fittest”, Amy and the other students preferred to talk about how a population could survive, or ‘be more fit’ under certain circumstances, due to the benefits that a trait gave to the individuals that carried it. Indeed, most students were talking about traits as having fitness, and individuals as ‘carrying’ this fitness. Another project participant, Derek, suggested that in this case, “survival of the groupiest” was a better phrase to describe the selection process.

The Cooperation (Behavior) Models⁵

After the students had explored the Altruism models, and had reflected on how the complexity of environmental and social factors affect an agent’s fitness, we gave them the first Cooperation model. The Cooperation model consists of two parts: cows and grass. The grass was designed to grow in such a way that above a certain height, it would have a high percentage chance of growing back to its full height each turn, but below a certain height, it would grow back very slowly. The high grass is thus considered to be

⁵ In the NetLogo models library, the two Cooperation models were combined into one NetLogo Cooperation Model. The VISCOSITY slider was replaced by a STRIDE-LENGTH slider – increased viscosity corresponds to decreasing the STRIDE-LENGTH. To get the behavior of Cooperation Model 1, set the STRIDE-LENGTH slider to its maximum value. The NetLogo Cooperation model can be downloaded from <http://ccl.northwestern.edu/netlogo/models/Cooperation> .

healthy, or sustainable grass, and the low grass is considered to be unhealthy grass. In terms of the model, these variables can be expressed in terms of the following parameter settings: the maximum grass height is 10; the threshold for healthy grass is 5; the chance for growing back for healthy grass is 75%; and, the chance for unhealthy grass growing back is 30%. These parameters are the initial values for slider variables in the Cooperation model that can be modified by students.

There are two types of cows. The first type of cow, the “greedy”, eats the grass as far down as possible. The second type of cow, the “cooperative”, only eats the grass if it is above the fast-growth threshold. All the cows have the same metabolism, and require a fixed amount of food to live. Each turn they lose a percentage of their energy, and if their energy runs out, they die. Eating grass restores a cow’s energy by a fixed amount. Finally, if a cow’s energy reaches a certain threshold, it reproduces. In terms of the model, these variables can be expressed in terms of the following parameter settings: all the cows get 51 energy units for eating from a patch of grass; they all lose 10 energy units each turn for moving (regardless of how far they move); they all reproduce when their energy level reaches 101; and, they all lose 40 energy units for reproducing. These parameters, too, are initial values for the sliders in the Cooperation model.

Thus, all things being equal, the only difference between the greedy cows and the cooperative cows is that the former eat all the grass and over-harvest the land, while the latter will go hungry rather than damage the well-being of the food supply. Supposing that these habits are not intentional states, but manifestations of different genetic traits, we asked students to discuss the scenario, and predict the outcomes of the models.

All of the students said that the evolutionary advantage goes to the cows who can get as much energy as quickly as possible. David said, “The greedy cows will have the clear advantage because they will be able to eat more food.” He suggested that the higher an agent’s energy got, the more it would reproduce, and the more it would reproduce the more resources it would use. Another student, Beverly, said, “once the greedy cows begin to reproduce, they will spread out and eat the grass below the level where the cooperative cows can eat.” The cooperative cows would be keeping the grass high, so that they and the other cooperative cows would have plenty of food; however, the greedy cows would eat the grass below the healthy threshold, and the cooperative cows would have no food. Amy said, “the cooperative cows will not be able to eat once the greedy cows start to eat the grass; and they [the cooperative cows] will go to extinction.”

Below is a picture of the NetLogo Cooperation model (Wilensky, 1998b). Figure 4a depicts the initial state of the model. The sliders determine the values of the following variables: metabolism, energy from grass, reproduction- cost, level of energy needed to reproduce, the maximum grass height, the healthy grass threshold, the percentage chance to grow back for healthy grass, the percentage chance to grow back for unhealthy grass, and the numbers of the cooperative and greedy populations. Students explored changing these variables, but found the result to be the same in most every scenario where there was a substantial environmental difference between healthy and unhealthy grass. Figure 4b shows the Cooperation model after 30 generations. The greedy cows out-grazed and

out-reproduced the cooperative cows; eventually, they ran the cooperative population to extinction.

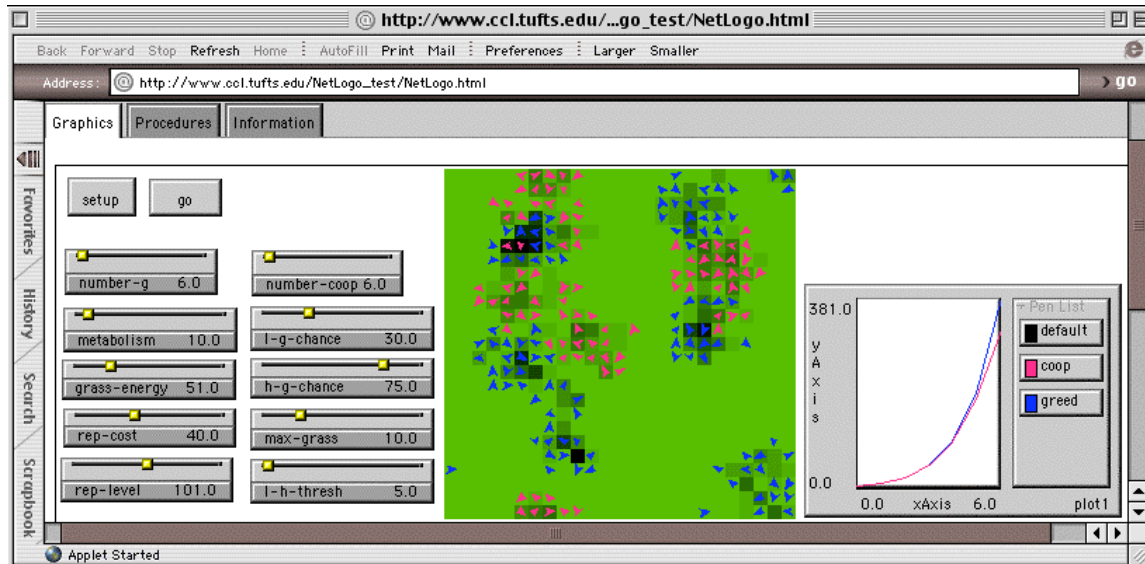


Fig. 4a

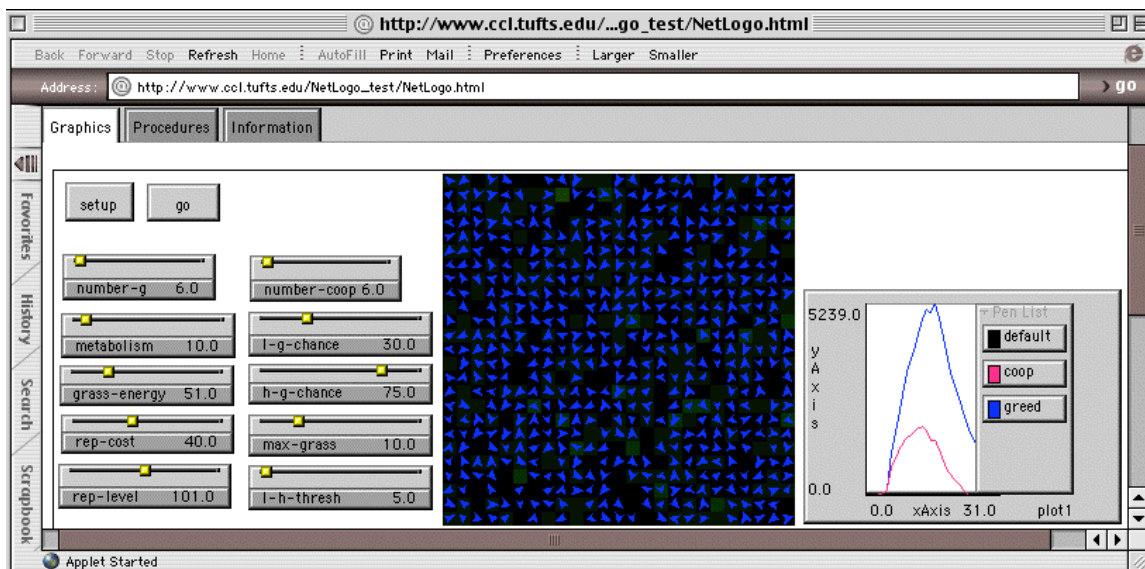


Fig. 4b

Figure 4. The Cooperation Model 1

In discussing the outcome of this model with the students, we asked them to explore ways in which the cooperative population could possibly survive. Reflecting on their work with the Altruism models, a number of the students suggested that if we could get the cooperative population to group together, then they might be able to survive. We encouraged this line of thought, and asked students to consider how this change would affect the fitness of cooperative individuals. A few students immediately responded: Derek said, “if we got the cooperators to group together, they would all have higher grass

to eat, and would all benefit from their cooperative behavior...The cooperative gene would benefit only the cooperative gene, and it would be more fit.”

For the second Cooperation model, we introduced the variable of population viscosity into the model. We explained population viscosity as the limitation in a cow’s range of movement due to environmental barriers, and then asked students, individually, to predict what would happen when different values were entered in the Viscosity slider. Many of them suggested that if the Viscosity value was high enough, that the greedy population couldn’t reach the cooperative population, and that the cooperative population would therefore survive. Figure 5, below, depicts the second Cooperation model with the Viscosity set at eleven (which effectively limited the cows to one-eleventh their normal travelling distance) after 300 generations.

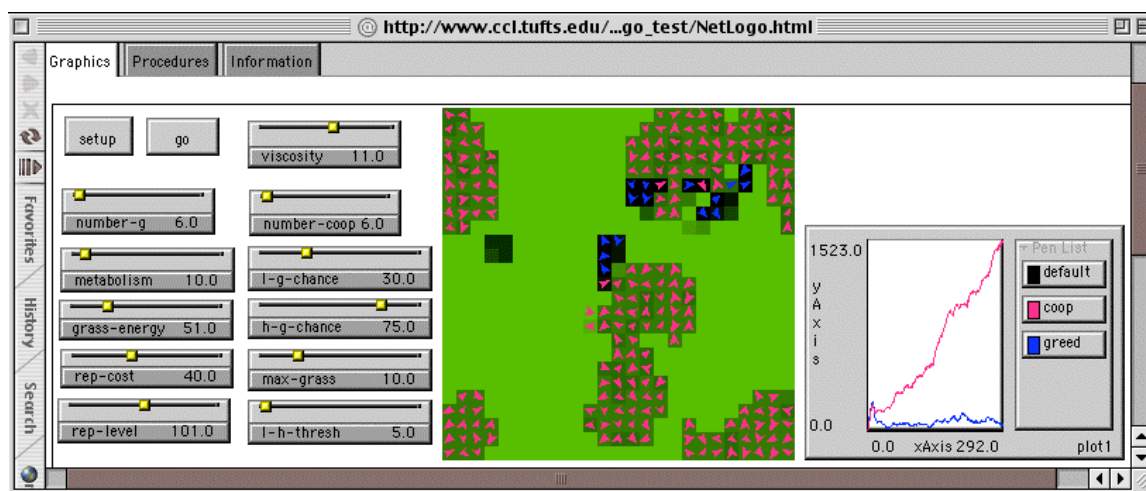


Figure 5. The Cooperation Model 2

After running the model, students found that, given the above values for the other variables, at a Viscosity value of around about eleven the greedy population was kept to about a twentieth of the cooperating population. The students were surprised to find that at higher values of Viscosity, the greedy population was so localized that it over-grazed and killed itself off. The students found a number of interesting metaphors for explaining the success of the cooperative population under high population viscosity. But, more interestingly, they all understood that the localization of the populations was key to the success of cooperative behavior because the benefit of cooperation fell on the cooperators alone. Beverly said, “the greedy cows were stronger when they could spread out because it takes a larger portion of land to support them. They are more fit when there is no population viscosity. But, the cooperative cows are more fit when there is no competition with greedy cows, because all of their cooperative benefit goes to other cooperative cows.” After a few more minutes of discussion, Derek said, “So, the cooperators succeed because of gene-level fitness; the cooperation gene is benefited by itself.” Amy concluded, “When the cooperative gene helps itself, there is no way to beat it.”

After their work with the cooperative models, many of the students began, in discussion, to explore important issues in the evolution of altruism and cooperation. The level at

which they had envisioned agents being altruistic or cooperative had implicitly assumed individual fitness as the standard for cost and benefit. Their new conceptions of multi-level fitness raised many questions about whether the altruistic habits were really altruistic, or whether they were only genes helping themselves. Of course, these are precisely the questions that occupy contemporary thought in the field, and we were encouraged to see students beginning to understand the different levels at which these problems can be addressed: what is altruistic at the individual level, is selfishness at the gene level.

V. Conclusion

The goal of the EACH unit was to enable students to explore their understanding of fitness and its relation to their conceptions of natural selection and the evolution of altruism. All the students who participated in this study initially believed that altruism could not survive because it was not a benefit to individual fitness. After students interacted with the models, they began to think of fitness as a multi-level phenomenon, and to think of the possible ways in which group and environmental factors can make altruism an advantageous trait. Through experimenting with the models, students came to a new understanding of fitness, and of the role that conceptions of fitness play in their understanding of evolutionarily plausible scenarios.

We encouraged students participating in the project to explore the NetLogo environment, and to build models that would explore their ideas of how different evolutionary scenarios would work. We have also started to involve students in participating in on-line interactive simulations of evolutionary systems. Such participatory simulations are a new and emerging technology in which individual students act out the role of individual agents in a simulation. As this technology advances, we expect that participatory simulations will be a powerful new tool for exploring the roles of individual and group-level phenomena in evolutionary systems. We have begun to develop on-going evolutionary scenarios in which each user is an agent that develops a survival strategy that is either passed on, if successful, or dies out, if unsuccessful. We anticipate that enabling users to participate in these simulations as first-person agents will add new dimensions of understanding to their conceptions of evolutionary processes and the multiple levels of fitness.

Acknowledgements

The preparation of this paper was supported by the National Science Foundation (Grants RED-9552950, REC-9632612). The ideas expressed here do not necessarily reflect the positions of the supporting agency. We are indebted to Josh Mitteldorf for the ideas in the Altruism model and for reviewing project models. We thank Dan Dennett for stimulating discussions and critique.

References

- Axelrod, Robert M. (1984). *The Evolution of Cooperation*. New York : Basic Books.
- Bishop, B. A., & Anderson, C. W. (1990). Student conceptions of natural selection and its role in evolution. *Journal of Research in Science Teaching*, 27(5), 415-427.
- Brumby, M.N (1984), Misconceptions about the Concept of Natural Selection by Medical Biology Students. *Science Education*, 68(4), 493-503.

- Abrahamson, D., & Wilensky, U. (2004). SAMPLER: Collaborative interactive computer-based statistics learning environment. Paper presented at the *10th International Congress on Mathematical Education*, Copenhagen, July 4 - 11, 2004.
- Blikstein, P. & Wilensky, U. (2004). MaterialSim: an agent-based simulation toolkit for Materials Science learning. Paper presented at the *International Conference on Engineering Education*, Gainesville, FL.
- Centola, D., Wilensky, U. & McKenzie, E. (2000). A Hands-on Modeling Approach to Evolution: Learning about the Evolution of Cooperation and Altruism Through Multi-Agent Modeling - The EACH Project: In Fishman, B. (Ed.) *Proceedings of the Fourth International Conference on the Learning Sciences (ICLS 2000)*, pp 164-173.
- Ferrari, M., & Chi, M. T. H. (1998). The nature of naive explanations of natural selection. *International Journal of Science Education*. Volume 20, Issue 10 , pp. 1231 - 1256
- Greene, E. D. (1990). The logic of university students' misunderstanding of natural selection. *Journal of Research in Science Teaching*, 27(9), 875-885.
- Hamilton, W.D. (1964). The Genetic Theory of Social Behavior (I&II). *Journal of Theoretical Biology*. 7, 1-16, 17-32.
- Jacobson, M. J., A. Sugimoto, et al. (1996). Evolution, hypermedia learning environments, and conceptual change: A preliminary report. International Conference on the Learning Sciences, 1996: *Proceedings of the ICLS 96*. D. C. Edelson and E. A. Domeshek. Charlottesville, VA, Association for the Advancement of Computing in Education: 151-158.
- Konold, C.: 1991, 'Understanding Students' beliefs about Probability', Von Glaserfeld, E. (Ed.), *Radical Constructivism in Mathematics Education*, Kluwer Academic Press, Netherlands, pp. 139-156.
- Mitteldorf, J. & Wilson, D.S. (2000). Population Viscosity and the Evolution of Altruism. *Journal of Theoretical Biology*, 204, 481-496.
- Papert, S. (1980). *Mindstorms: Children, Computers, and Powerful Ideas*. New York: Basic Books.
- Resnick, M. (1996). Beyond the Centralized Mindset. *Journal of the Learning Sciences*, 5 (1), 1-22.
- Resnick, M., & Wilensky, U. (1993). *Beyond the deterministic, centralized mindsets: A new thinking for new science*. Paper presented at the annual meeting of the American Educational Research Association, Atlanta, GA.
- Samarapungayan, A., & Wiers, R. W. (1997). Children's thoughts on the origin of species: A study of explanatory coherence. *Cognitive Science*, 21(2), 147-177.
- Settlage, J. (1994). Conceptions of natural selection: A snapshot of the sense-making process. *Journal of Research in Science Teaching*, 31(5), 448-457.
- Tamir, P., & Zohar, A. (1991). Anthropomorphism and teleology in reasoning about biological phenomena. *Journal of Biological Education*, 25, 57-67. *Conference on the Learning Sciences, 1996: Proceedings of ICLS 96* (pp. 151-158). Charlottesville, Virginia.
- Tisue, S., & Wilensky, U. (2004). *NetLogo: Design and implementation of a multi-agent modeling environment*. Paper presented at the Agent 2004 conference, Chicago, IL, October 2004.

- Tversky, A. & Kahneman, D. (1974). Judgment Under Uncertainty: Heuristics and Biases, *Science*, 185, pp. 1124–1131.
- Wilensky, U. (1995). Paradox, Programming and Learning Probability. *Journal of Mathematical Behavior*. Vol. 14, No. 2. p 231-280
- Wilensky, U. (1996). “Modeling Rugby: Kick First, Generalize Later?” *International Journal of Computers for Mathematical Learning* 1(1): 125-131.
- Wilensky, U. (1997b). What is Normal Anyway? Therapy for Epistemological Anxiety. *Educational Studies in Mathematics. Special Edition on Computational Environments in Mathematics Education*. Noss, R. (Ed.) 33 (2), 171-202.
- Wilensky, U. (1999c). GasLab: An extensible modeling toolkit for exploring micro-and-macro-views of gases. In N. Roberts, W. Feurzeig & B. Hunter (Eds.), *Computer Modeling and Simulation in Science Education* (pp. 151-178). Berlin: Springer Verlag.
- Wilensky, U. (2001). *Modeling nature's emergent patterns with multi-agent languages*. Paper presented at the Eurologo 2001 Conference, Linz, Austria.
- Wilensky, U., Hazzard, E., & Longenecker, S. (2000). *A Bale of Turtles: A Case Study of a middle school science class studying complexity using StarLogoT*. Paper presented at the meeting of the Spencer Foundation, New York, New York, October 11-13, 2000.
- Wilensky, U., Hazzard, E., & Froemke, R. (1999). *GasLab - An extensible modeling toolkit for exploring statistical mechanics*. Paper presented at the Seventh European Logo Conference (EUROLOGO '99), Sofia, Bulgaria.
- Wilensky, U., & Novak, M. (in preparation). Understanding Evolution as an Emergent Process: Learning with Agent-Based Models of Evolutionary Dynamics. In R.S. Taylor & M. Ferrari (Eds.), *Epistemology and Science Education: Understanding the Evolution vs. Intelligent Design Controversy*. New York: Routledge.
- Wilensky, U. & Resnick, M. (1999). Thinking in Levels: A Dynamic Systems Perspective to Making Sense of the World. *Journal of Science Education and Technology*. Vol. 8 No. 1.
- Wilensky, U. & Reisman, K. (1998). ConnectedScience: Learning Biology through Constructing and Testing Computational Theories--An Embodied Modeling Approach. *InterJournal of Complex Systems*, M. 234, pp. 1 – 12.
- Wilensky, U. & Stroup, W. (1999b). Participatory Simulations: Network-based Design for Systems Learning in Classrooms. *Proceedings of the Conference on Computer-Supported Collaborative Learning, CSCL '99*, Stanford University
- Williams, George C. (1992). *Natural Selection: domains, levels, and challenges*. New York: Oxford University Press.
- Wilson, D. S. & Sober, E. (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, Mass. : Harvard University Press.
- Wilson, D. S., Pollock, G. B. & Dugatkin, L. A. (1992). Can altruism evolve in purely viscous populations? *Evol. Ecol.* 6, 331-341.

Software and Models

- Wilensky, U. (1997a). StarLogoT. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University. <http://ccl.northwestern.edu/cm/StarLogoT/>.
- Wilensky, U. (1999a). NetLogo. Evanston, IL: Center for Connected Learning and Computer-

- Based Modeling. <http://ccl.northwestern.edu/netlogo>.
- Wilensky, U. (1999b). NetLogo Models Library. Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL.
<http://ccl.northwestern.edu/netlogo/models>.
- Wilensky, U. (1997c). StarLogoT Altruism model.
<http://ccl.northwestern.edu/cm/models/altruism> .Center for Connected Learning and Computer-Based Modeling. Northwestern University. Evanston, IL.
- Wilensky, U. (1998a). NetLogo Altruism model.
<http://ccl.northwestern.edu/netlogo/models/Altruism>. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL
- Wilensky, U. (1998b). NetLogo Cooperation model.
<http://ccl.northwestern.edu/netlogo/models/Cooperation>. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.
- Wilensky, U. (1998c). NetLogo Peppered Moths model.
<http://ccl.northwestern.edu/netlogo/models/PepperedMoths>. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.
- Wilensky, U., & Stroup, W. (1999a). HubNet. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University.
<http://ccl.northwestern.edu/netlogo/hubnet>.
- Wilensky, U., Novak, M., & Rand, W. (2004). BEAGLE curriculum. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University.
<http://ccl.northwestern.edu/curriculum/simevolution/beagle.shtml>.