Chapter 3

Invalidation

Every model is an abstraction and a simplification. The goal of a model is to capture the essence of a system in the real world such that the behavior of the model matches the qualitative behavior of interest in the real system. Thus, for any model we may ask if it is a reasonable representation of the real system. Answering this question is the problem of validation. Traditionally we can try to invalidate the model by collecting data from the real system and comparing it to the behavior of the model. However, it should be noted that failure to invalidate can never prove the validity of a model (Caswell 1976).

In artificial life we rarely have the luxury of comparison to the real system. Artificial life models tend to be highly abstract and general because the field is striving to discover general properties of life. This makes experimental testing extremely difficult. The time scale of evolution tends to restrict experiments to observation of the fossil record (Benton 1990, for example) or manipulation of organisms with extremely short life-cycles in simplified environments (Lenski & Travisano 1994; Bull et al. 1997; Krukonis 1996, for example). Similarly, the complexity and size of ecosystems makes ecological experiments cumbersome and difficult to control.

An alternative form of testing can be pursued indirectly through reference to ecological and evolutionary theory. Instead of asking if the model matches the experimental data, we can ask if the model matches our understanding of the dynamics of ecology and evolution. Then, to the extent that the theories of ecology and evolution have been corroborated by experimental observations, we can invalidate the model when it fails to match those theories. These tests must be carried out at the appropriate level of abstraction. Because every model is a simplification, it will fail to match reality at some level of detail. For example, in the model described in Chapter 2, there is no distinction between phenotype and genotype. The fact that this is in blatant conflict with biological reality is peripheral to the issue of whether or not the model captures the essence of diversification. We are only interested in testing the dynamics represented in the model against their analogues in theoretical biology.

What follows is an example of this technique applied to MoD. While the purpose of this model is to explore new theoretical ground in biology, the ecological and evolutionary dynamics in the model have been tested against theories of predation, competition, and adaptation. As a result of these studies, the algorithms for implementing predation had to be elaborated, in accordance with the methodology I laid out in Section 1.4.2. Initially, a predator could automatically find and consume a prey organism in its patch, as long as the predator's search pattern matched the prey's predation resistance chromosome. As we will see below, such simplicity fails to capture the essence of a predator-prey interaction. The implementation of predator-prey interactions had to be fixed, and the final result is described both below and in Section 2.2.2.

The problems of speciation include both ecological and evolutionary dynamics of the model. Thus, it is important to test both the ecological dynamics and the evolutionary dynamics. Ecological theory provides qualitative predictions for predator-prey dynamics as well as competition. Evolutionary theory predicts changes in populations as they adapt to their environments. In the ecological experiments below, the dynamics within a single patch were examined when the mutation rate was 0. This prevented evolutionary change from interfering with the underlying ecological interactions.

3.1 Predator-Prey Oscillations

The Lotka-Volterra equations were an early abstraction of the dynamics of predator and prey populations. The details of the Lotka-Volterra equations are "hopelessly unrealistic" (Crawley 1992, p.68) because they lead to neutrally stable oscillations. Any perturbations to the system should generally result in either a stable fixed point or a stable limit cycle (May & Watts 1992). The introduction of a time lag between changes in the prey population and the response of the predator population tends to cause the two populations to oscillate stably. In our case the model requires three time steps to translate prey matter into predator matter. When predators are rare, prey populations will naturally expand, providing more food for the predators. Once the predator population expands in response to this abundance, they cull back the prey population. Eventually the prey become hard to find and predators begin to starve to death. And so the process repeats.

At a minimum, a model of an ecosystem should exhibit these oscillations. The model with a naive implementation of predation does not. As originally written, a predator could automatically find a prey organism in its patch. This meant that even when prey organisms were rare, the predators could still find them. The prey were driven to extinction and the predator population crashed shortly thereafter. This style of predation, where each predator captures a constant number of prey is called "fixed number predation," and, as we have seen, is unstable when prey are scarce (Crawley 1992).

A *prey-location* probability was introduced to the model to represent the chance that a particular predator might find a particular prey organism in its patch. However, a predator still gets one attempt for every potential prey organism in its patch. This results in stable herbivore-plant oscillations across a wide range of parameters. In particular, the *prey-location* probability can range from about[1] 0.15 to 0.3 and still produce stable oscillations. Higher than that and the herbivores overeat, killing off all the plants. Lower than that and the herbivores cannot find enough food to sustain life.

3.2 Trophic Cascades

While the introduction of the *prey-location* probability stabilized the two-trophic level dynamics, the difficulties multiply when we move to three trophic levels by the addition of a carnivore species. Unfortunately, there is only a knife-edge of parameter space that can sustain both a two and a three trophic level ecosystem. Since herbivores are rare, relative to plants, the *prey-location* probability must be high to allow the carnivores to find the herbivores. However, if it is too high, the carnivores will kill off all of the herbivores. All three trophic levels could only be sustained by lowering the *energy-conversion* parameter to unrealistic levels, requiring the consumption of only two prey organisms to produce a new predator organism.

[1]Tweaking other parameters like *energy-conversion* or the amount of energy given to the plants can increase this range of stability.

The problem comes when the predator population has expanded and culled back the prey population to low levels. For a prey organism to survive during a time step, it must evade location by *every* predator in the patch. Even a low *prey-location* probability cannot save a prey organism from a large predator population.

The spatial structure of the real world tends to soften intense predation. A prey organism is generally not directly threatened by all the predators in the patch, because some of the predators are busy stalking other prey or pursuing other interests. Predators sometimes patrol a fixed territory from which they exclude one another. Furthermore, there are many forms of refuge in which prey can find safety. Reasonable forms of refuge that might be captured in the modeled system include hiding places, microhabitat selection by the prey, and aggregation of the predators leaving some areas unpatrolled. These are but a small subset of the forms of refuges found in biology (Crawley 1992). However, the important characteristic of all of them is that there is variation across prey individuals in the probability they will be attacked by predators.

The model was elaborated with two mechanisms, effectively simulating the addition of interference between the predators, or possibly the presence of prey refuges. First of all a predator *search-radius* was added, restricting the amount of the patch that a predator can patrol in one time step. Each predator starts searching in a randomly chosen location in the patch's organism array, and then looks for prey in the *search-radius* consecutive following locations[2] in the array. Second, a prey organism that survives an encounter with a predator, because of the *prey-location* probability[3], becomes more difficult to find by other predators. Specifically, the *prey-location* probability is multiplied by $e^{-(\text{number of encounters}\times \textit{predation-distribution})}$. When the *predation-distribution* is 0, there is effectively no spatial structure, and the model behaves as it did before. However, when the *predation-distribution* is positive, a form of spatial structure is imposed on the patch, and a small prey population has a better chance of surviving intense predation. This introduces interference competition between the predators.

An ecological model with multiple trophic levels, in this case plants, herbivores, and carnivores, should demonstrate a trophic cascade in the plant density. A trophic cascade is a

[2]The array was formed into a "ring" so that the last location in the array is adjacent to the first location.
[3]If the predator search pattern mismatches the prey's predation resistance chromosome in m bits, then an additional e^{-m} is factored into the probability of the predator consuming the prey.

change in plant or autotroph biomass in response to predation. In the absence of herbivores, plant biomass ought to be relatively high. In contrast, the introduction of herbivores tends to dramatically reduce plant biomass. However, the further introduction of carnivores to prey upon the herbivores tends to restore plant biomass to some intermediate level. Does the model demonstrate a trophic cascade?

3.2.1 Results

The model was run for 500 time steps with a single patch under varying parameters. As long as the carnivores were able to survive, a trophic cascade was found. In most cases, the model was run at least 40 times for each parameter setting, although in some cases it was run as many as 100 times.

The addition of the *search-radius* had little effect. A small *search-radius* actually helps to support a larger herbivore population by preventing over predation in a two trophic level system. However, if we want carnivores to survive, they must find the herbivores scattered amongst the plants, and so require a large *search-radius* to survive. In a patch that could hold 2048 organisms (with a maximum of 1024 plants), carnivores required a *search-radius* of more than 128, and tended to be stable with a radius of 512 as shown in Figure 3-1. When the search range is small, the carnivores cannot find the herbivores scattered about amongst the plants. The carnivore population crashes and the herbivores are left unmolested in their consumption of the plants. However, when the search range is greater than 128, the carnivores sometimes manage to hang on, reducing the herbivore population and allowing the plants to fill their environment to the carrying capacity of 1024 plants. With a search radius of 256, the model settled into one of two states, represented by the two clusters of points at that setting. The curves, representing the averages of the data at each parameter setting, lie between the two clusters. Data for single trophic level ecosystems are not shown because the plants always fill the environment to its carrying capacity.

Each data point indicates an independent run of the model, and the curves connect the means for data points at each parameter setting. A trophic cascade can be seen when the plant biomass in a two trophic is significantly less than the plant biomass with three trophic levels. This is illustrated by the solid curve diverging from the dashed curve in the graph of plant populations.

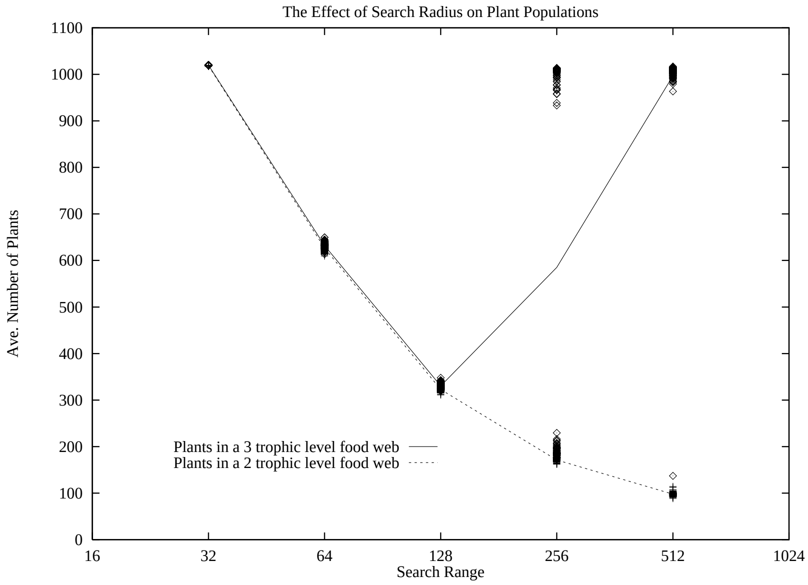

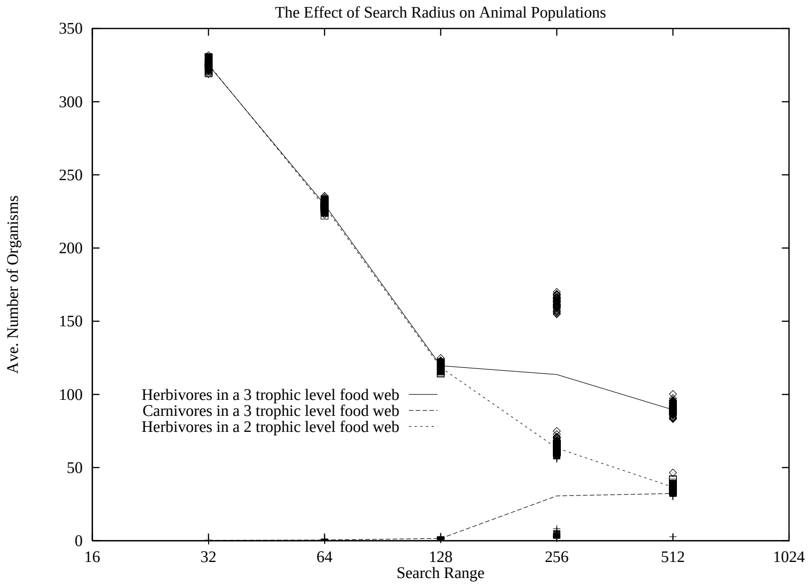

Figure 3-1: Each data point represents the average number of plants in a run, while the lines represent the average across the 100+ runs for each parameter setting. The "search-range" is the number of organisms that may potentially be encountered in a time step by a predator (both herbivores and carnivores). The upper graph shows the effects of the search range on the plant populations. The lower graph shows the effects of search range on the herbivores, in both a two and three trophic level food web, as well as the carnivores in the three trophic level food web. It is interesting to note that when the search range is small the herbivore population thrives relative to a moderate search range, even in the absence of carnivores.

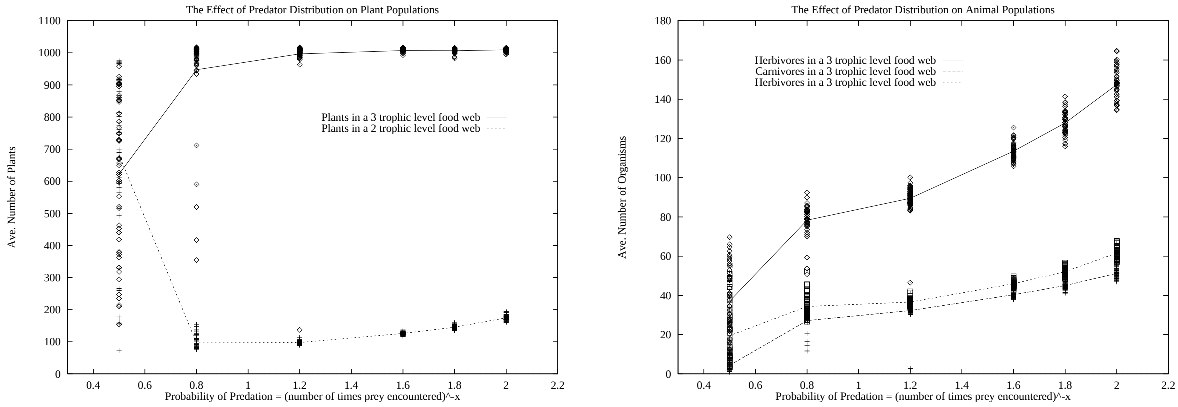

Figure 3-2: The effects of enforcing more spatial structure and so adding prey refuges with the *predation-distribution* parameter on the biomass of the three species. The left graph shows the plant biomass in the two and three trophic level ecosystems. The right graph shows biomass of the herbivores and carnivores in the two and three trophic level ecosystems.

The *predation-distribution* parameter had a much more dramatic effect. Both herbivore and carnivore populations were stable with a *predation-distribution* ≥ 0.8. In fact, the average biomass of both herbivores and carnivores increased with the *predation-distribution*, up to the maximum value that was examined (2.0), as seen in the right-hand graph of Figure 3-2. The predator populations expand as the *predation-distribution* parameter makes it harder for them to capture prey. A similar thing could be said of the herbivores in Figure 3-1. As the search radius decreases and the herbivores are restricted to feeding on fewer plants, herbivore biomass actually rises.

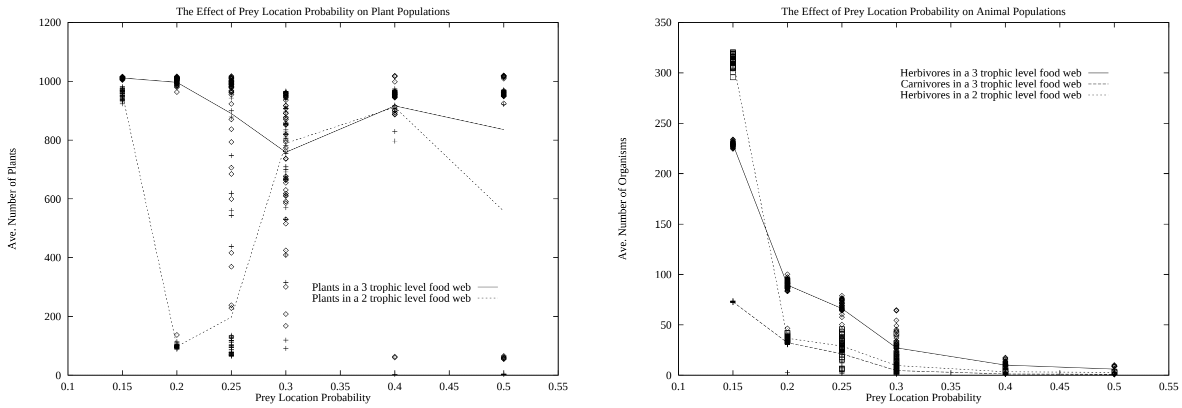

Figure 3-3: The effect of the prey location probability on the population sizes of plants (left) and animals (right).

A similarly counter-intuitive result can be seen if we look at the effect of the *prey-location* probability on the populations of herbivores and carnivores as shown in Figure 3-3.

The most dramatic trophic cascade appears when the *prey-location* probability is 0.2 (the setting used for future experiments). It is interesting to note that all three trophic levels thrive when the probability is at a minimum 0.15. Once again, as the parameter values make predation more difficult, the average predator populations expand. This is probably due to a more efficient allocation of resources to the predators. The average maximum and minimum population sizes of the predators are shown in Figure 3-4 along with the mean values. When predation is easy, the predator populations can expand rapidly, reaching higher values than the maximum population sizes attained when predation is hard. However, the success of the predators has a dramatic effect on the prey, and so the predator populations collapse, reaching lower minima as compared to the minima in the runs where predation is more difficult. Ease of predation tends to destabilize the populations, causing larger fluctuations and resulting in lower average predator populations over time. Crawley (1992) argues, from the distribution models, that prey refuges will stabilize predator-prey dynamics as long as the refuges support a large enough prey population to preserve the predator population in lean times.

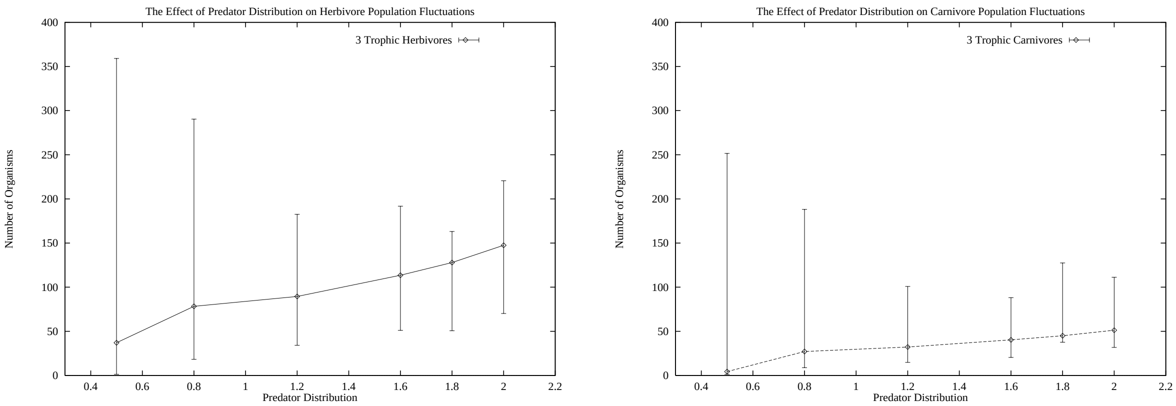

Figure 3-4: The average maximum and average minimum population sizes have been added to the herbivore and carnivore curves from the lower graph of Figure 3-2.

The effects of two further parameters were investigated, *energy-input*, as shown in Figure 3-5, and *energy-conversion*, as shown in Figure 3-6. When the *energy-input* is less than sufficient to give each plant a single energy unit in a time step, the plant populations do not replenish themselves fast enough to support many animals. Once the *energy-input* rises above 1, herbivores begin to thrive. Note the difference between the herbivores with no predation and the herbivores that suffer predation. The presence of carnivores seems to reduce competition amongst the herbivores, when the *energy-input* is 1, and consequently allows the herbivores to expand to greater population sizes than they could maintain without the carnivores. Yet, when the *energy-input* is greater than 1, the presence of the carnivores prevents the herbivores from responding to the greater plant abundance. A trophic cascade can be observed in the left graph of Figure 3-5 when the *energy-input* is 1 and 1.25.

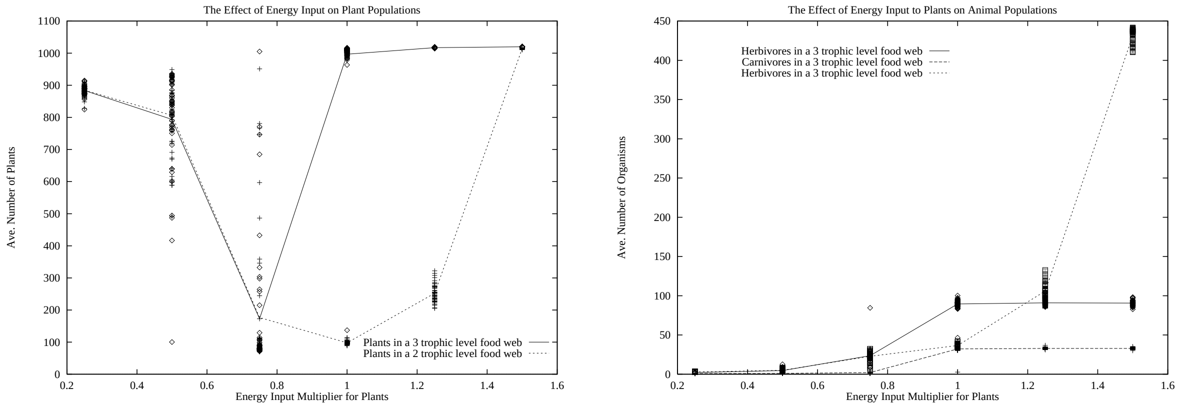

Figure 3-5: The effect of *energy-input* on population sizes. The *energy-input* is the number of energy units a plant can absorb in one time step if it is perfectly matched to its environment (as all the plants are in these experiments).

Finally, the right graph of Figure 3-6 shows that the animal populations decline as the *energy-conversion* parameter forces them to eat more prey organism before they can reproduce. When the *energy-conversion* parameter is high, the predators have difficulty gathering enough food to survive and their populations die off. This accounts for why the plant populations are high when the energy conversion is 8 : 1. Biologically realistic conversion values range from 3 : 1 (for plant to animal matter as in the above) to 10 : 1 (for abiotic to plant matter), though distribution models often use 1 : 1 (Schmitz & Booth 1997). In all other runs, the energy conservation parameter was set to 3. A trophic cascade can be seen in the left graph of Figure 3-6 when the *energy-conversion* parameter was set to 3 or 4.

3.3 Competitive Exclusion

Competition between two species in an ecosystem often results the competitive exclusion (extinction) of one species by the species better adapted to the environment. While relative "adaptation" is often difficult to quantify independently of competitive exclusion experiments, it is easily identifiable in the model. A series of competitive exclusion experiments

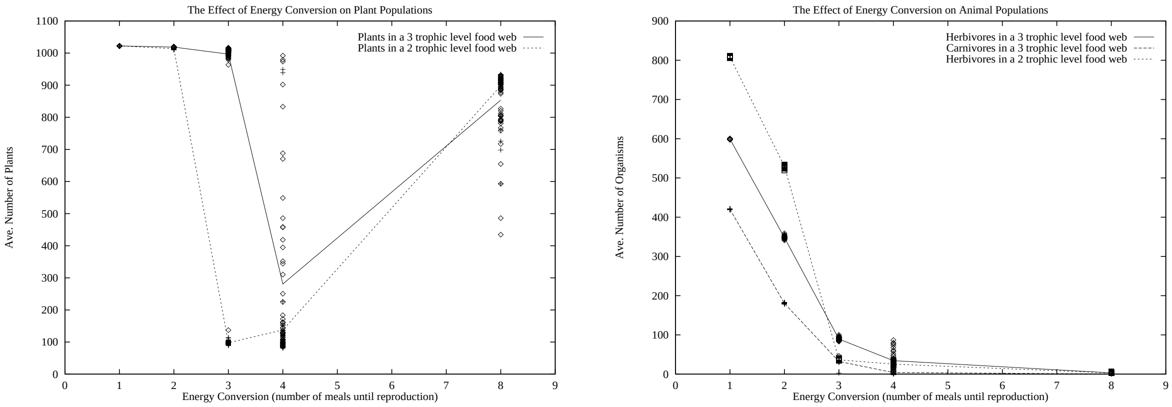

Figure 3-6: The effect of *energy-conversion* on population sizes. The *energy-conversion* parameter is the number of meals an organism must accumulate in order to produce a single offspring. This is the scaling factor by which prey matter is transformed into predator matter.

were designed wherein the environment was seeded with equal numbers of two species of plant. The "climate" of the environment had all its bits set to 0 and the prey template chromosomes of the plant species differed from each other by only one bit. That is, one of the two species was exactly one bit better adapted to the environment than the other. This should result in a slightly better energy absorption and a consequent slightly higher reproduction rate in the species with fewer bits set to 1. A species was considered to have excluded the other species when it had expanded to 90% of the carrying capacity of the environment. 100 trials were run with 1, 4, 8, 16, 24, and 31 bits set to 1 in the prey template chromosome of the inferior species. Mutation was turned off for these experiments. The superior species excluded the inferior in all the trials.

One might ask how long does it take for the superior species to take over 90% of the carrying capacity? If competitive exclusion tends to occur quickly, then we can infer that the slope of the fitness landscape in that region of the prey template chromosome is relatively steep. Figure 3-7 shows the inverse of the average number of time steps required for the exclusion of the inferior plant along with a linear sketch of the fitness landscape. When there is only 1 bit different from the optimum, the left graph indicates that the slope of the fitness landscape is very steep. The landscape has a single optimum surrounded by a steep slope that becomes more gentle by the time that the prey template chromosome is 4 bits away from the optimum. This is sometimes called a "Mount Fuji" fitness landscape because there is only one fitness peak in the landscape (there are no local optima besides the global

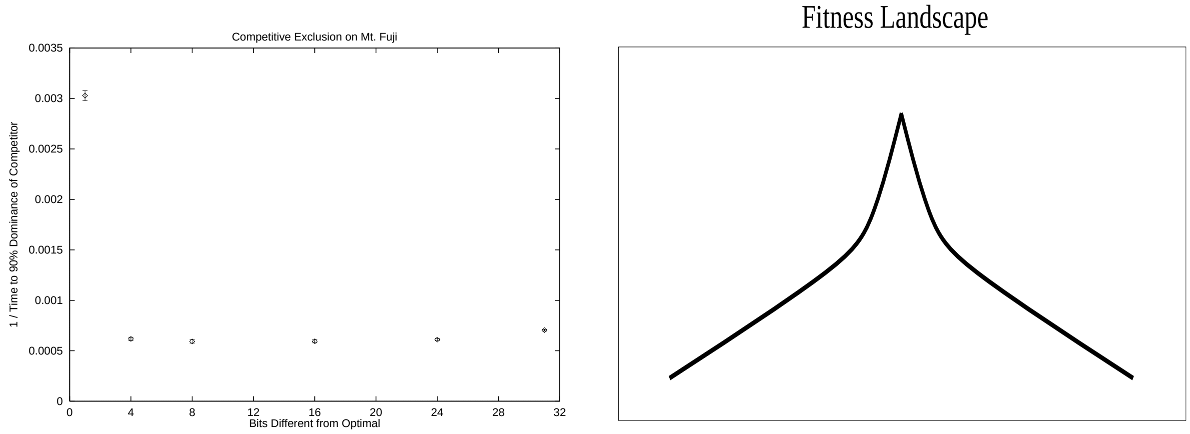

Figure 3-7: The fitness landscape for plants. The left graph shows the inverse of the average number of time steps, with standard error bars, required for the superior plant species to take over 90% of the carrying capacity in a patch. The intensity of selection is inversely proportional to the amount of time it takes for one species to exclude the other. Thus, the slope of the fitness landscape is inversely proportional to the amount of time before competitive exclusion occurs. From this, we can infer the slope of the fitness landscape at the various distances from the optimum, depicted in the right graph.

one) and all other points in genotype space rest on a slope that leads to that peak.

3.4 Adaptation

One of Darwin's (1859) key insights was that biological species are not immutable Platonic entities but rather they are dynamic conglomerations of lineages flowing over a fitness landscape. His theory of natural selection is in part a prediction that a species will adapt to its environment. That is, populations should tend to climb the peaks in the fitness landscape.

A model of evolution should at least demonstrate that, all other things being constant, organisms evolve to be better adapted to their environment. This can be tested in the model when mutation is introduced into the dynamics. Consider a plant species evolving in the conditions of the competitive exclusion experiments. The optimal genotype matches the climate bit pattern of all 0's. This is achieved by having 0's (or wild-cards) in all the loci of the organism's search pattern. If the prey template and generalism chromosomes were evolving neutrally, with no natural selection, then we would expect one in four bits in the plant's prey template to have a 1 that is not masked by the generalism chromosome. We can thus detect natural selection in any significant reduction in that proportion of unmasked

1's below 25%.

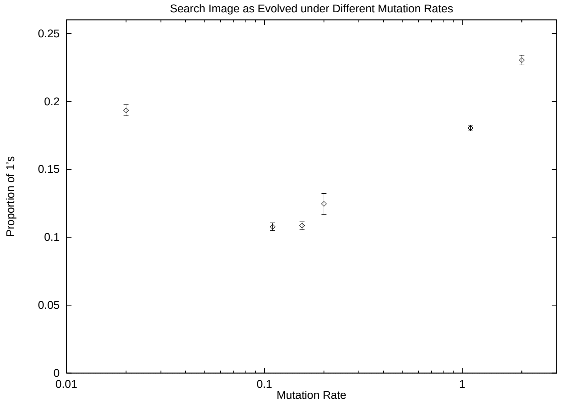

Figure 3-8: Adaptation in plants. The normalized distance (with standard errors) of plant populations from the optimal genotype under different mutation rates. In all cases the null hypothesis of neutral evolution can be rejected.

The model was initialized with a plant species with exactly 8 out of 32 positions in its prey template set to 1 and no wild-card positions in its generalism chromosome. After 5000 time steps the proportion of unmasked 1's was counted and averaged. The data for at least 32 runs in each setting of the *mutation-rate* parameter is shown in Figure 3-8. For all mutation rates, the resulting data is significantly below a mean of 25% unmasked 1's as tested by a one-sided t-test ($p < 0.001$ in all cases). At low mutation rates, the populations have not yet had time to reduce the number of 1's in their search patterns, and so the proportion of 1's is still high after 5000 time steps. At high mutation rates, mutation keeps adding new 1's to the search pattern and so the proportion of 1's stays high. At intermediate mutation rates, selection has enough time to weed out the deleterious mutations while preserving the beneficial ones. There can be little doubt that the plant species have adapted to their environment.

3.5 Conclusion

The initial failure of the model and subsequent elaboration of the predation algorithms illustrates an important benefit of attempts to invalidate a model. Failure helps to sharpen our understanding of the essential features of the system and so helps to guide the further development of the model. It is also important to note that even when there exists a paucity of theory that can be used to test the central results of a model, there generally exists a rich field of theory that impacts upon the dynamics of the model. In this case, a model of species diversification was designed for the very reason that there is a lack of established theory covering that topic. However, such a model must include both ecological and evolutionary dynamics and so I have tested it against both ecological and evolutionary theory. Because the dynamic of diversification intimately depends upon ecological and evolutionary interactions, these studies help to corroborate the end results of the model.

The fact that a model passes a series of tests against theory does not guarantee the quality of its results. The model described in Chapter 2 is only one instantiation of the requirements for modeling species diversity, and so stands as a single data point for theory. This perspective on models is particularly important to artificial life where seemingly trivial implementation details often manifest in artefactual results. In the end, the model does not provide Truth but rather a hypothesis that should be experimentally tested in the laboratory and field, similar to any other theoretical result. It must be granted that such experimental corroboration is difficult. Meanwhile, by testing our models against theory, we may at least make progress toward significant theoretical insight.

3.6 Parameters for the Following Chapters

Species	**Chromosome**	**Loci**
Plant	reproduction	0000000000000000000000000000000011111111111111111111111111111111
	predation resistance	00000000000000000000000000000000
	prey template	00000000000000000000000000000000
	generalism	00000000000000000000000000000000
Herbivore	reproduction	0000000000000000000000000000000000000000000000000000000000000000
	predation resistance	10101010101010101010101010101010
	prey template	00000000000000000000000000000000
	generalism	00000000000000000000000000000000
Carnivore	reproduction	1111111111111111111111111111111111111111111111111111111111111111
	predation resistance	01010101010101010101010101010101
	prey template	10101010101010101010101010101010
	generalism	00000000000000000000000000000000

Table 3.1: The genotypes of the three species used to seed the model. All the species are extreme specialists on their respective food sources. Furthermore, the reproductive genotypes of the three species are different from each other in at least 32 loci.

For the experiments in the following chapters, I chose to set the parameters that were not being examined in the experiments to maximize biomass in both the plants and animals.

Yet, I have kept to the parameter ranges that demonstrated realistic ecological behavior. Specifically, I used a *search-radius* of 512, a *predation-distribution* of 2.0, a *prey-location* of 0.15, an *energy-input* of 1.0, an *energy-conversion* of 3, and a *mutation-rate* of 0.02.

In every case, the model was initialized with 500 plants and 200 animals. The animals were roughly composed of 150 herbivores and 50 carnivores. The chromosomes for the three species are shown in Table 3.6. Initially, all members of a species had identical chromosomes.

And now for something completely different...