

# A Cautionary Tale of Caterpillars and Selectional Interference

W. Garrett Mitchener

December 14, 2007

## Abstract

Evolutionary theory yields many important insights into why organisms are the way they are. However, for any given problem, there is significant uncertainty about what level of abstraction is appropriate. With such issues in mind, I will discuss an evolutionary scenario concerning the size and development time of a particular species of caterpillar, the tobacco hornworm. A larva of this species grows approximately exponentially; once it reaches a critical size various hormones result in a time delay before pupation during which it continues to grow. Thus, the relationship between genotype and phenotype is non-trivial. In laboratory experiments by Nijhout and in computer simulations, one can selectively breed caterpillars based on size and development time and investigate the resulting evolutionary dynamics. Both laboratory experiments and computer simulations yield striking complexity. Selecting for size and time simultaneously yields unexpected interference, resulting in caterpillars that do not necessarily have the selected-for properties. This interference is due to constraints on the set of possible phenotypes imposed by the caterpillar's fundamental development process. The important lesson here is that when building evolutionary models of even simple phenomena, one must be distrustful of intuition, and consider (1) how the details of the development process of individual organisms affect the ability of a population to explore the space of possible phenotypes, and (2) how selection criteria in combination can interact in unexpected ways.

## 1 Introduction

The basic premise of biological evolution is that there is a population of individuals that live and reproduce in a more-or-less closed environment. Selection is the process by which the number of offspring an individual has is influenced by its interactions with its environment. Mutation is the imperfection in the reproductive process that causes offspring to be similar but not identical to their parents. These principles are used to understand why organisms have the forms that they do: Presumably,

there is (or was) selection in favor of certain features that leads organisms bearing those features to increase in number at the expense of others. The most precise way to reason about selection-mutation processes is to express these assumptions in terms of a formal mathematical or computational model. That precise statement frequently yields a precise solution, and further discussion can focus on whether the assumptions grounding the model are appropriate and how to interpret the solution biologically.

A typical evolutionary model begins by hypothesizing a simplified environment inhabited by organisms playing some sort of abstract game, and the ability of an organism to reproduce depends on its score in this game [11, 23, 24]. Within this general framework, there is plenty of room for variety. The addition of simplifying assumptions often makes the model more mathematically tractable. For instance, if the game can be expressed as a payoff matrix, then mathematical results based on Nash equilibria are available. Alternatively, the model may be more like a simulation, perhaps becoming so detailed that it is mathematically intractable, and the researcher has no alternative but to implement it as a computer program, run it many times, and analyze the results statistically.

Formal evolutionary models can demonstrate complex and unexpected behavior that would go undiscovered if reasoning about evolution were left to imprecise verbal arguments. In particular, models of similar evolutionary scenarios that differ in a few details can have very different outcomes. First, I will summarize two well studied examples that clearly show how details such as spatial structure, environmental limits, and an organism's memory can influence evolution: the prisoner's dilemma and communication games. Both of these focus on the evolution of a single feature of the species in question. Then I will describe in somewhat more detail some new results concerning a different sort of complexity that occurs when two features of a species are evolving simultaneously: Selective breeding experiments performed by Nijhout's lab on the tobacco hornworm show surprising interactions between selection based on size and selection based on development time [31, 32]. These results are the basis for my simulation that shows how limitations imposed by an organism's basic structure and interference between selectional criteria can prevent a species from developing the traits favored by selection.

## 1.1 The prisoner's dilemma

As a first example, the *prisoner's dilemma* [36] is a family of games in which two players each have the choice of two strategies traditionally called *cooperate* and *defect*. In any given round, both players must choose a strategy without knowing what the other will choose. If they both cooperate, they receive a moderate payoff. If one cooperates and one defects, the defector receives a high payoff and the cooperator is punished severely. If both defect, both receive a low payoff. See Figure 1. This game models the tension between acting individualistically versus cooperatively, and there are countless variations in the mathematical and biological literature. Many

	cooperate	defect
cooperate	$\begin{array}{c} \diagdown \\ 3 \\ \diagup \\ 3 \end{array}$	$\begin{array}{c} \diagdown \\ 5 \\ \diagup \\ 0 \end{array}$
defect	$\begin{array}{c} \diagdown \\ 0 \\ \diagup \\ 5 \end{array}$	$\begin{array}{c} \diagdown \\ 1 \\ \diagup \\ 1 \end{array}$

Figure 1: An example of a payoff matrix for one round of the prisoner’s dilemma. Each player selects one of the two strategies, *cooperate* or *defect*. One player selects a row, the other selects a column. The combination specifies one of the four boxes. The row player wins the payoff in the lower left corner. The column player wins the payoff in the upper right corner.

of these variations have astonishingly different behavior, given that they are based on the same fundamental game and the same general principles.

In the simplest version of the prisoner’s dilemma, the game is limited to a one-round contest between rational players with no memory. Both players should always defect because this maximizes their payoff no matter what their opponent does. In an evolutionary setting of this basic version, there are many players in an unstructured population. Each individual plays the game with many others selected at random, and its payoffs from these interactions accumulate. Individuals with greater total payoff are more likely to be selected to reproduce. In general, the simple strategy ALWAYSDEFECT (which, as its name suggests, always chooses to defect) takes over the population. If an individual attempts to play another strategy, it always scores less than an opponent playing ALWAYSDEFECT, so it will die out.

A more realistic alternative is the iterated prisoner’s dilemma. Each pair of individuals plays many rounds of the game consecutively, and their payoffs accumulate. They are allowed to remember the past actions of their opponent during the interaction and take this data into account when deciding on an option for the next round. In this iterated game, players are generally better off risking cooperation. Such strategies include ALWAYSCOOPERATE, which always cooperates; TITFORTAT, in which a player cooperates on the first round and thereafter chooses the action its opponent played in the previous round; and WINSTAYLOSESHIFT, in which a player picks a strategy and uses it repeatedly until it loses, after which it switches to the other strategy, and so on. If the model includes mutation, a population of ALWAYSDEFECT is vulnerable to invasion by mutants that play TITFORTAT because when two such mutants interact they can earn the higher payoff of mutual cooperation, as opposed to the low payoff of mutual defection. So the population tends to change from all ALWAYSDEFECT to all TITFORTAT. However, there is an important effect called neutral evolution. Since ALWAYSCOOPERATE and TITFORTAT are indistinguishable when playing only against each other, a mutation to ALWAYSCOOPERATE is invisible to the selection process in a population of all TITFORTAT. Random

drift in the form of such mutations can allow a population of all TITFORTAT players to be replaced by ALWAYSCOOPERATE, which is then vulnerable to invasion by ALWAYSDEFECT [13].

Another variation is to make the population spatially structured rather than well-mixed. Each individual is located at a fixed position on a grid and interacts only with its neighbors. Reproduction is accomplished by cloning an individual into a neighboring grid point, thereby replacing the former occupant. In this case, ALWAYSCOOPERATE can coexist permanently with ALWAYSDEFECT because cooperators can form clusters that enable them to earn the higher payoff of mutual cooperation [36].

In short, the evolution of a population facing something like the prisoner's dilemma depends greatly on details that are not specifically stated in the basic premises of evolutionary theory.

## 1.2 Communication games

For a second example that comes from my research, there has been much interest lately in trying to understand the biological origin and genetic history of the language faculty, and in using evolutionary models to understand language change on historical timescales [1–10, 12, 14–16, 18, 20–22, 30, 33–35, 37, 38, 40, 42–45, 48]. One example is the *naming game*, [17, 47] which models a community coming to an agreement on a mapping between words and meanings. Evolutionary game theory determines which such mappings are Nash equilibria or evolutionarily stable. The naming game can be used to precisely describe the insight that when communicating across a noisy channel [39, 41, 46], if there are sufficiently many possible topics of conversation, then the language should switch from using a single symbol per meaning to using a sequence of symbols, that is, a language with syntax.

One can also adapt continuous replicator dynamics [11] to include imperfect learning, with or without genetic variation, resulting in the *language dynamical equation* [19, 25–27, 29]. This equation is a formalization of the idea that language is a strategy for communicating, and the genetically-encoded language acquisition process is a strategy for choosing a language. Typical simplifying assumptions for such models include assuming that the number of relevant languages is finite; that children have a single parent; that they learn from only their parent; that the number of individuals is large enough that a continuous approximation is appropriate, for example by representing the population as a list of what fraction of the population uses each language; that the payoff between players depends on their choices of languages but is otherwise constant; and that each individual interacts equally with all others. Details such as the genetic encoding of language, the syntactic and semantic details of languages, and social structure might be ignored. Mutation might be modeled directly within the resulting differential equations or accounted for as an external perturbation. One case of this model shows how a communication game can lead a population to a coherent state dominated by one language. Another

suggests several initially surprising possibilities. For instance, humans may have once been capable of natively learning many more languages than are now possible; that is, some languages may have been removed from the language faculty because selection has favored individuals that do not consider as many possible languages during acquisition and are therefore less likely to make a mistake that inhibits their communication ability [28].

### 1.3 The dangers of excessive abstraction

The prisoner’s dilemma, naming game, and the language dynamical equation all give important insights into the origins and genetic history of cooperation and language, and they lend precision to any discussion concerning them. However, there is reason to be skeptical. Many details, such as the process by which the brain develops and its methods of representing and manipulating meanings are simply not known and are typically left out of models. Furthermore, these models base each simulated individual’s reproductive success on its ability to play one abstract game and on nothing else. Evolutionary models in general focus on a single aspect of an organism to the exclusion of all else, which means that interactions among the many demands placed on individuals are left out.

These two issues, the development process and the interaction of selectional demands, are the subject of the rest of this paper. A particular example, the tobacco hornworm, illustrates how dangerous it can be to ignore them. The tobacco hornworm *Manduca sexta* is a well studied species of caterpillar. My interest in the species began with a presentation by Fred Nijhout of Duke University [32], in which he observed that specimens from lab strains were distinctly larger than specimens from the wild. Apparently, something about the lab conditions favored larger caterpillars, and they evolved accordingly. That inspired a series of experiments in which his lab bred colonies of these caterpillars, selecting pupae based on various combinations of mass at pupation and time to pupation.

These experiments had very interesting and unexpected results. One might expect that there should be harmony between development time and pupa size: A caterpillar that stays a caterpillar longer before pupation ought to be larger. Selecting simultaneously for longer development time and larger pupa mass eventually yielded caterpillars with the desired combination of traits, but the path through phenotype space taken by the population was erratic, and frequently backtracked. That is, selecting for longer development time and more massive pupae sometimes yielded faster development times or smaller pupae in the next generation. See Figures 2 and 3 for graphs of some of Nijhout’s preliminary data. The populations in these experiments consisted of 221 to 366 caterpillars per generation. Although Figure 2 shows a general upward trend in pupation size and development time, neither feature is increasing monotonically. As Figure 3 shows, the two features also fluctuate independently. Further statistical analysis must be done on this data to determine how significant the trends and fluctuations are, but the overall picture is

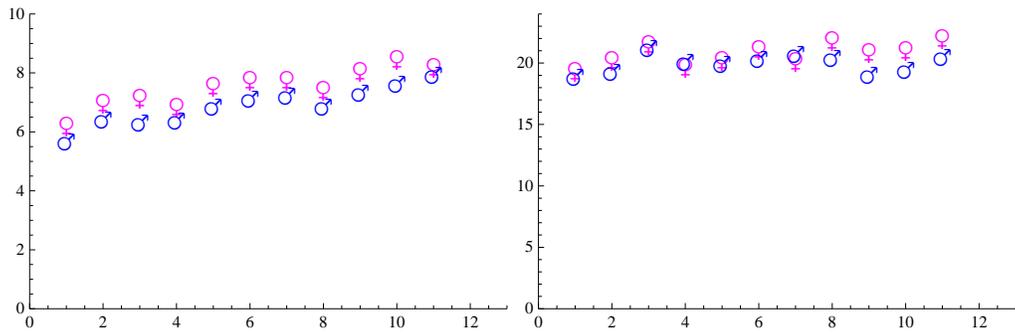


Figure 2: Data from Nijhout's experiments. Left: Mean weight in grams at pupation as a function of generation number. Right: Mean development time in days as a function of generation number. Females: ♀. Males: ♂.

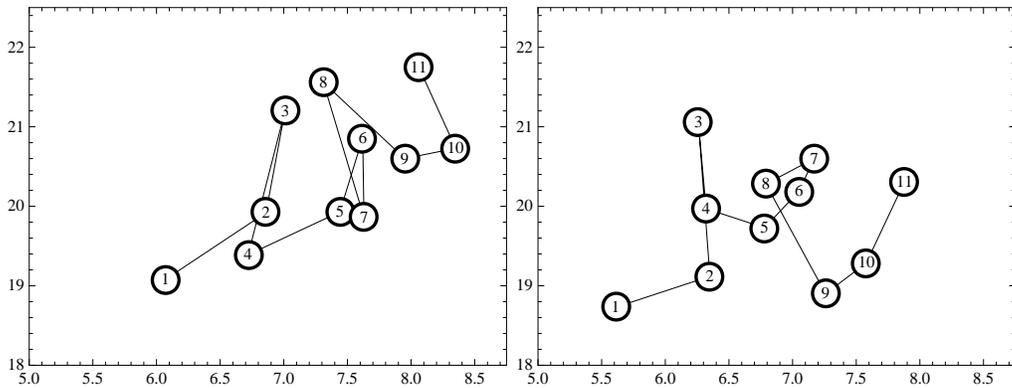


Figure 3: Data from Nijhout's experiments: Mean weight in grams vs. mean time to pupation in days. Numbers indicate the generation. Left: females. Right: males.

that selective breeding for longer development time and larger pupae simultaneously, even in a reasonably large population and over 11 generations, does not immediately lead to the expected results. Selecting for shorter development time and smaller pupae also yielded the desired combination of traits only after meandering through phenotype space. Oddly, selecting non-harmonious combinations (shorter time and larger, longer time and smaller) yielded the same general behavior, in which case the eventual success of the selective breeding was the surprising aspect.

Nijhout explained that the hornworm's development process is a complex sequence of hormonal interactions, resulting in a very jagged mapping from biochemical reaction parameters to the development time and final size. Thus, a small change in the genome leading to a small change in some protein could yield a large change in phenotype.

## 2 A simulation

In the rest of this chapter, I will describe a computational simulation of Nijhout's caterpillar experiments. It allows the experimenter to directly vary the details of the development process and to adjust the combination of selection pressures placed on the population. The simulation can run on very large populations for many more generations than would be possible in the lab.

### 2.1 Details of the computer program

The computer program simulates the growth and evolution of tobacco hornworms under a variety of selection pressures. It includes some details of the caterpillars' development that turn out to be crucial to understanding Nijhout's results. Each simulated caterpillar is characterized by three genetically specified numbers,  $k$ ,  $\rho$ , and  $\delta$ . From an initial size of 0.1 at time 0, it grows exponentially at rate  $k$  until it reaches  $\rho$  times its original size. Living caterpillars are kept in their larval state by the presence of a juvenile hormone, which begins to break down once they reach this critical size. They do not pupate until the juvenile hormone's concentration decreases below a certain threshold. In the simulation, the breakdown of this hormone is represented by a time delay  $\delta$  during which the simulated caterpillars continue to grow. As a further complication, living caterpillars have a daily hormonal cycle that only allows pupation to begin at a certain time of day. Once the juvenile hormone is gone, the caterpillar must wait (perhaps most of another day) to pupate, during which time it continues to grow. In the simulation, this constraint is represented by requiring that the simulated caterpillars grow beyond the critical size and time delay until time reaches the next whole number.

Mathematically, the time to maturation is

$$t = \left\lceil \frac{\log \rho}{k} + \delta \right\rceil \quad (1)$$

where  $\lceil x \rceil = \text{least integer } \geq x$

and the pupa size is

$$m = m_0 e^{kt} \quad (2)$$

where  $m_0 = \text{initial size}$

To produce the data in this chapter, the simulation maintains a population of 400 virtual insects. Each has a 90 bit genome that encodes three numbers. Each number is encoded by 30 consecutive bits that are interpreted as an integer between 0 and  $2^{30} - 1$ , then divided by  $2^{30}$  to yield a real number between 0 and 1. These numbers,  $a$ ,  $b$ , and  $c$ , determine  $k$ ,  $\rho$ , and  $\delta$  as follows.

$$\begin{aligned} \rho &= 5 + 100a \\ k &= \frac{\log 2}{0.5 + 4b} \\ \delta &= c \end{aligned} \quad (3)$$

Thus, the critical size ratio varies from 5 to 105, the doubling time for the exponential growth (that is,  $\log 2/k$ ) varies from 0.5 to 4.5 days, and the time delay representing the breakdown of juvenile hormone varies from 0 to 1. (These ranges were selected to give reasonable results and are not explicitly based on any experimentally derived estimates.)

Each generation is constructed from the previous by selecting part of the population for reproduction (details to follow). From the reproductive subset, the simulation picks a pair of genomes uniformly at random, and recombines them at a crossover point selected uniformly at random. After crossover, one of the resulting genomes then undergoes mutation, in which each bit is toggled with probability 0.05. This genome is added to the next generation and the other is discarded. Four hundred repetitions of this process produce the next generation.

In living sexual organisms, crossover occurs during the formation of sperm and eggs cells and involves parallel chromosomes within a single individual. Strictly speaking, the step in the simulation that performs crossover between individuals is biologically unrealistic. It is used here because it enables the simulation to make use of gene recombination without the added complexity of diploid genomes. The program has an option to use simple asexual reproduction instead, but the results are essentially the same either way.

The selection process is very flexible. The four basic criteria are to take the largest or slowest half, or the half with the fastest or slowest development. The simulation can be configured to combine several of these. For example, it can select

the fastest half of the largest half, thereby choosing one quarter of the original population for reproduction. Alternatively it can select the largest half of the fastest half, which results in a slightly different quarter of the original population.

## 2.2 Results of the simulation

Out of all the possible variations that this program can compute, we will consider the following selection criteria:

- no selection at all
- the fastest half
- the slowest half
- the largest half
- the smallest half
- the fastest half of the largest half
- the largest half of the fastest half
- the slowest half of the largest half
- the largest half of the slowest half
- the fastest half of the smallest half
- the smallest half of the fastest half
- the slowest half of the smallest half
- the smallest half of the slowest half

The composite criteria are similar to those used in breeding experiments on living caterpillars performed by Nijhout's lab. The results of these simulations are shown in an appendix in Figures 6 through 18. The simulated population generally settles to a steady state after ten to twenty generations. The summary diagram in Figure 4 depicts the 100th generation of the simulated population under these selection criteria.

A few striking properties are immediately apparent. First, the four basic (non-composite) selection criteria are clearly effective. One hundred generations are ample time to drive the population to the fastest, slowest, and smallest extremes. In the final population for each of these selection criteria, three quarters of the population is quite close to the extreme of the phenotypic range, with some outliers. The one oddity is selection for the largest. The median settles near the upper extreme of possible sizes, and the larger half of the population is concentrated between sizes

80 and 100, but the population remains very diverse, with around a quarter of the individuals spread out between sizes 30 to 80. This much spread is not present in the results of any of the other basic selection criteria.

Among the composite selection criteria, many of the combinations give reasonable results. However, the combinations involving selection in favor of larger pupae show some unexpected behavior. When selecting for larger size and slower development time, the caterpillars evolve slower development time but do not become much larger than they would be without selection. They are concentrated around size 15, which is nowhere near the maximum of 100 or so that occurs when they are selected for large size alone. To understand this phenomenon, it is useful to plot the set of all possible phenotypes, as in Figure 5. This graph shows that the development process confines the set of phenotypes such that caterpillars of size near 100 and development time near 35 are simply not possible. Most of the population ends up around the upper left tip of the set of possible phenotypes. Intuitively, a caterpillar can only develop slowly if  $k$  is near the low end of its range. The growth is exponential, and the additional development time turns out not compensate for the smaller rate constant.

Another unexpected result is the fact that selecting the largest half of the fastest half results in caterpillars that develop quickly and are nearly as large as those that result from selecting based on size alone. This success is possible because faster development requires a large value of  $k$ . The additional time spent waiting for the juvenile hormone to degrade and the daily cycle allow them to grow to the maximum possible size.

Oddly, reversing the order of the selection criteria leads to failure. When we select the largest half of the fastest developing half, the caterpillars evolve to be a bit larger than with no selection at all, but the bulk of the population is concentrated around size 22, which is nowhere near the maximum. Referring back to Figure 5, the region of possible phenotypes has very narrow spikes for the fastest and largest caterpillars. Selecting the fastest-developing caterpillars first favors the phenotypes just below these spikes, and the largest of these are relatively small. Selecting the largest caterpillars first favors individuals within the spikes, for which the development time is short anyway. It is therefore not surprising that when selecting short development time first, the caterpillars do evolve a very short development time around 4 at the cost of remaining small, but when selecting larger pupae first, the caterpillars take a little extra time to develop but grow much larger. This particular development process happens to exaggerate the size difference more greatly than the time difference.

If we select for smaller pupae, then as seen in Figure 5, the possible range of development times is much greater, and there is less interference between selection for small size and selection for faster or slower development. Selection for small size and fast development is quite successful. Even though small size and longer development time seem to be contradictory criteria, the combination is within the

range of possible phenotypes and is found by the evolutionary process. However, the smallest pupae can only be obtained with development times less than 25 or so because of how the set of possible phenotypes curves away from the vertical axis in Figure 5. Thus, their development time is not nearly as long as it could be without selection for small size.

### 3 Conclusion

The first lesson to be learned from these experiments is that the devil is in the details. The development of these caterpillars has some complexities that restrict the space of possible phenotypes. Selective breeding can evolve larger or smaller caterpillars, or lengthen or shorten their time to pupation. However, when selection is applied to both characteristics at the same time, there is interference between them, yielding unexpected results. Combinations of traits that at first glance appear to be harmonious, such as larger pupae and slower development, are impossible. Seemingly disharmonious combinations, such as larger pupae and faster development, evolve surprisingly easily. Furthermore, the details of the selection process, such as the order in which the selection criteria are applied, are sometimes highly significant.

Genomes in this particular simulation are limited to 90 bits that control three biochemical parameters that lead deterministically to two phenotypic numbers. The living caterpillars that inspired this simulation have much larger genomes and far more potential for variation, so at first it may seem unjustified to claim that certain phenotypes are impossible. However, these caterpillars are very complex biochemical machines. Consequently, many possible variations are fatal. Since the evolutionary path through genome space must always be from one viable genome to another, these fatal variations form barriers that confine the species to a certain region of viable genomes. A long jump from one region to another is always possible since enough simultaneous mutations could happen, but the small probability of such events makes them unlikely to occur within the one hundred generation range of this study. Thus from any initial population and with a fixed time scale, any selection-mutation process will be limited in how much of genome space it can be expected to explore, and it is reasonable to speak of certain genotypes and phenotypes as being impossible. Any species therefore has a limited set of possible phenotypes that it can achieve in the foreseeable future, and the shape of that set can generate the kind of interference between selection criteria seen in the caterpillar simulation.

These details about possible phenotypes are typically ignored when setting up a model based on evolutionary game theory. For example, when studying the prisoner's dilemma, it seems at first perfectly reasonable to suppose that tit-for-tat is an available phenotype. However, implementing tit-for-tat in a living social organism requires an episodic memory and a reliable means of identifying each other individual in the community. In mammals, such abilities are common. However, in bacteria or insects, they are almost inconceivable. Common sense does not necessarily cor-

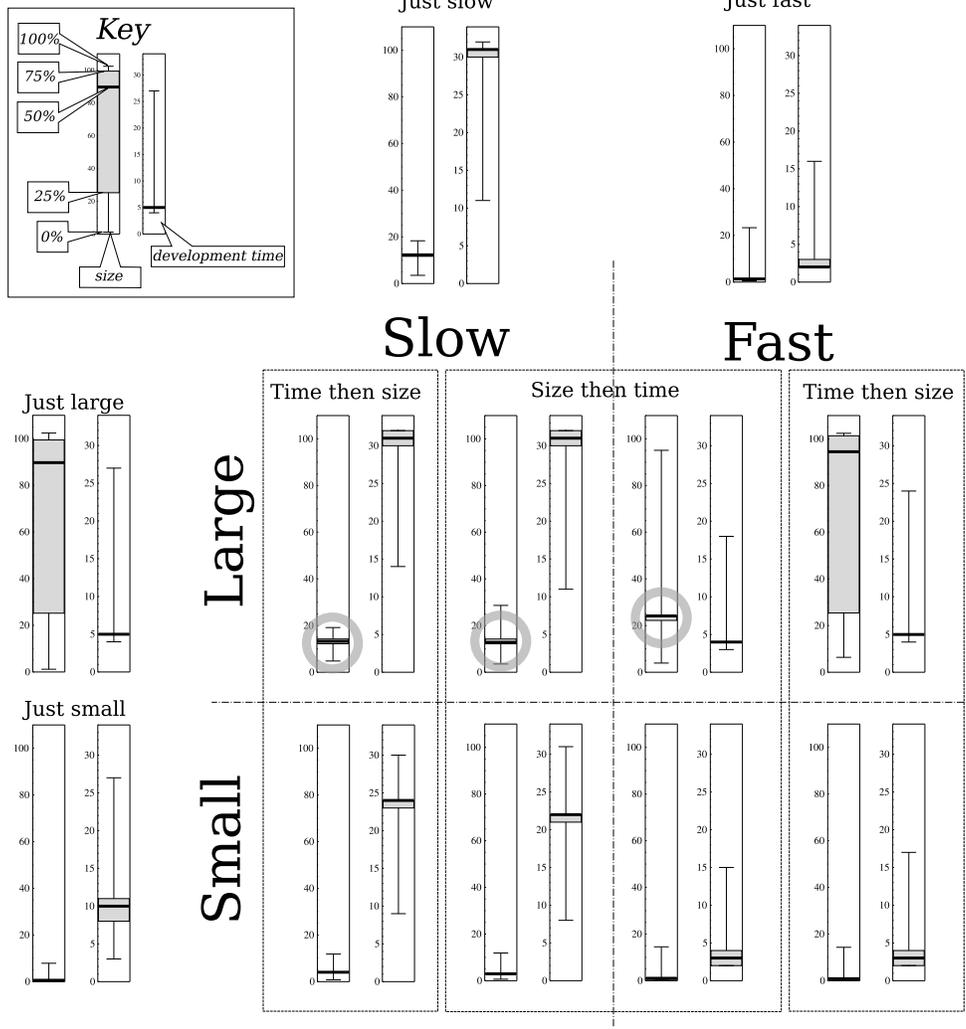


Figure 4: A summary of the results. The box-and-whisker plots for the 100th generation under each selection criterion are shown. Instances where the selection criteria do not lead to the specified results are circled.

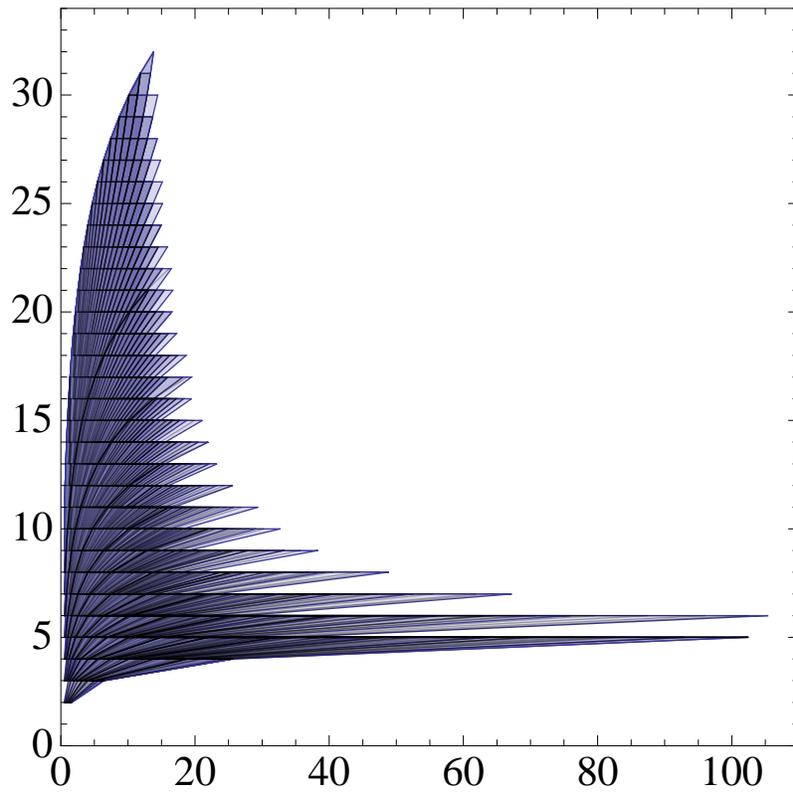


Figure 5: The set of all possible phenotypes. The horizontal axis represents size. The vertical axis represents development time. This figure was created by plotting the phenotype parametrically as a function of  $\rho$ ,  $k$ , and  $\delta$  across their respective ranges. The darker areas are phenotypes that happen to be covered by more than one value of  $(\rho, k, \delta)$ .

rectly identify which phenotypes are possible and which are not. In the caterpillar example, it was genuinely surprising to Nijhout and to me that the combination of slow development and large size is so difficult to achieve, and this fact only makes sense in light of detailed knowledge about the caterpillar's development process.

Evolutionary game theory models also tend to focus on variation in a single trait. In the caterpillar simulation, selection for size or development time alone yields nothing out of the ordinary: The population becomes larger or smaller or slower or faster in agreement with a single selection criterion. Yet selecting for various combinations of these traits results in very interesting outcomes. The simulation suggests that the two traits interfere with each other primarily because the set of possible phenotypes is oddly shaped. Surely this same condition must affect other combinations of selectional pressures in other organisms.

Returning to my own primary line of research, it is clear that these complexities also affect the evolution of language. The growth of a brain capable of symbolic processing from a single cell is at least as complicated as the cartoon of caterpillar growth simulated here. The human communication system clearly grants our species a tremendous survival benefit, but it competes with selection in favor of less powerful but more energy efficient brains, faster growth from a helpless baby, and safer swallowing, for example. The set of possible linguistic phenotypes on the time scale of human evolution is almost completely unknown. The set of all selection criteria and their relative strengths in combination is likewise unknown. In light of the results of the caterpillar experiment, the problem of understanding language evolution seems almost impossible. However, the simulation does provide some suggestions on how to proceed.

First, given some knowledge of the biochemical workings of the tobacco hornworm, the paradoxical results of Nijhout's selective breeding experiments do begin to make sense. Not all biochemical details are immediately relevant, however, the most important ones such as the juvenile hormone are obviously key once their function is identified. In trying to understand evolution of more complex organisms, we must therefore understand as much as possible about the relevant biological machinery if we are to have any hope of correctly identifying the possible phenotypes. When studying evolution of language, for instance, it is imperative for researchers to know as much as possible about how the brain works and how language is learned and used, and to use that knowledge in formulating evolutionary models.

Second, there is a clear need for improved evolutionary models that can handle selectional interference. This does not mean that standard frameworks such as the replicator equation are obsolete, rather, the part of the modeling process in which biological observations are translated into abstract mathematical games must be improved. Instead of assuming that all strategies are possible or focusing on just one or two, a model could hypothesize a restricted continuum of strategies like Figure 4.

Finally, the results of evolutionary simulations and modeling with constrained

phenotypes and selectional interference can also aid the understanding of particular organisms. For example, by comparing computer simulations to selective breeding experiments, it might be possible to sketch out the space of possible phenotypes, thereby inferring the actual range of decay rates for variations of hormones.

I would like to thank the NSF for supporting my research (grant number 0734783). I would also like to thank Fred Nijhout and his colleagues for introducing me to this problem and for sharing their data.

## **A All simulation results**

See Figures 6 through 18. In each figure, two graphs are shown, one for pupa size and one for development time. Every tenth generation out of a total of 100 is depicted as a box-and-whisker plot. The heavy line represents the median. The gray box spans the 25th to 75th percentiles, and the whiskers span the 0th to 100th percentiles. Sizes span roughly 0 to 100, and development times span roughly 0 to 35.

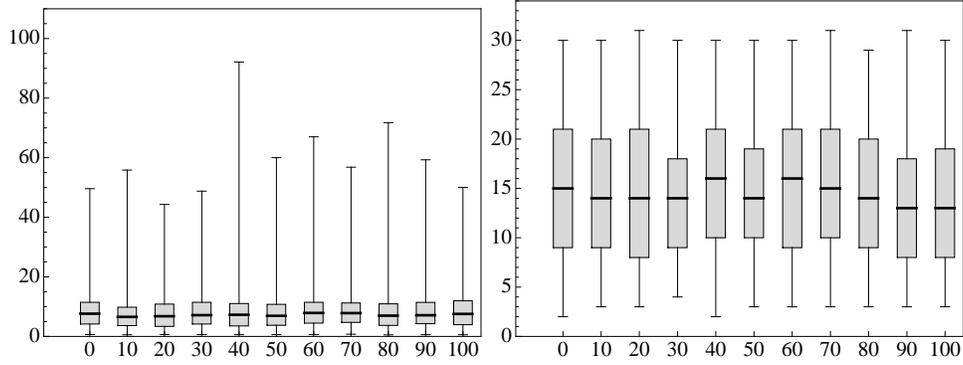


Figure 6: Result of simulation with no selection. Left: Size as a function of time. Right: Development time as a function of time.

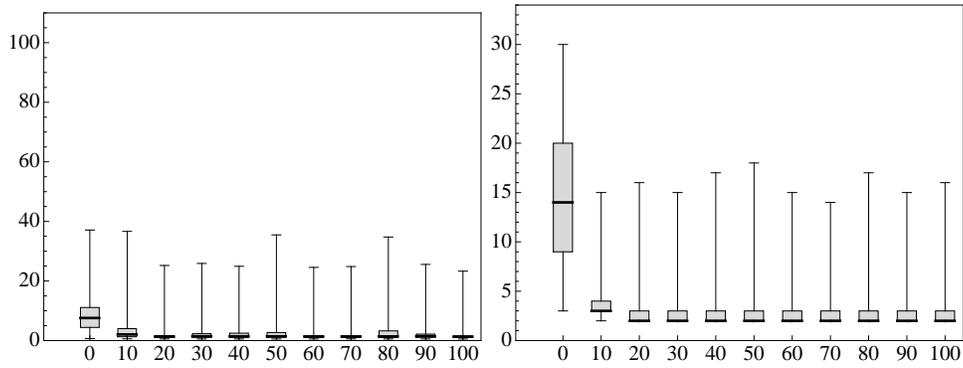


Figure 7: Result of simulation selecting for fast development. Left: Size as a function of time. Right: Development time as a function of time.

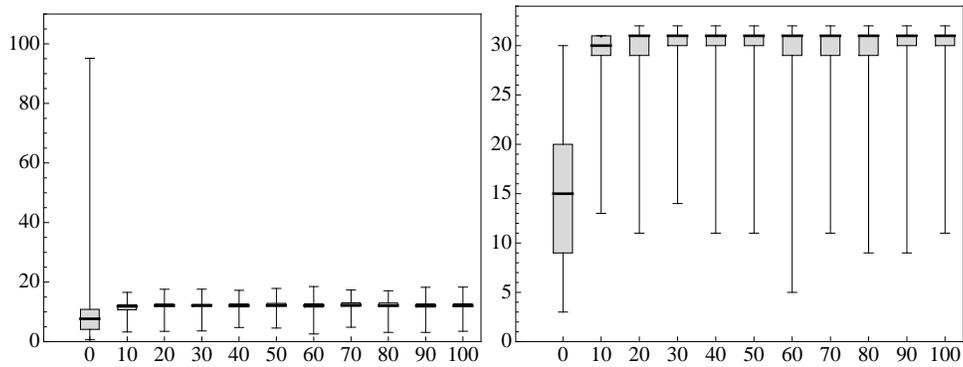


Figure 8: Result of simulation selecting for slow development. Left: Size as a function of time. Right: Development time as a function of time.

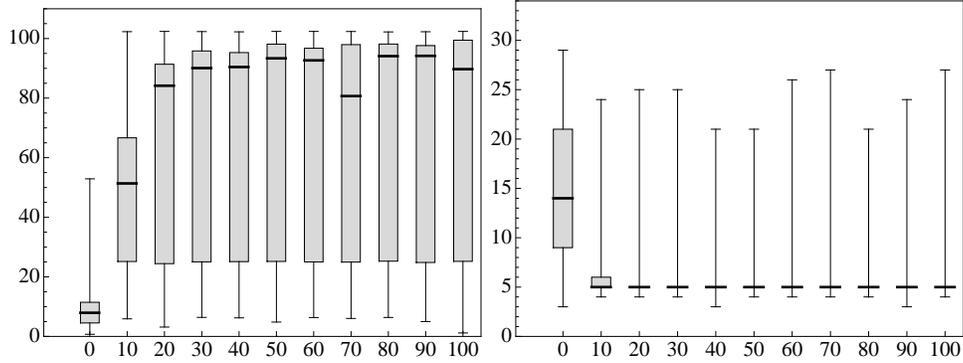


Figure 9: Result of simulation selecting for large pupae. Left: Size as a function of time. Right: Development time as a function of time.

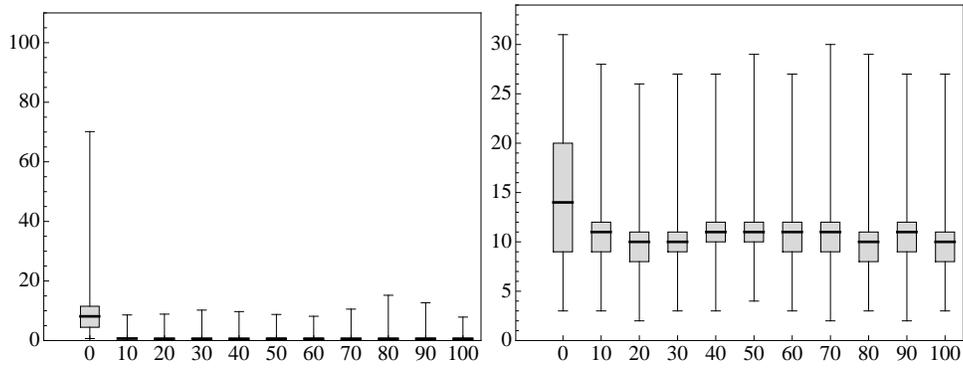


Figure 10: Result of simulation selecting for small pupae. Left: Size as a function of time. Right: Development time as a function of time.

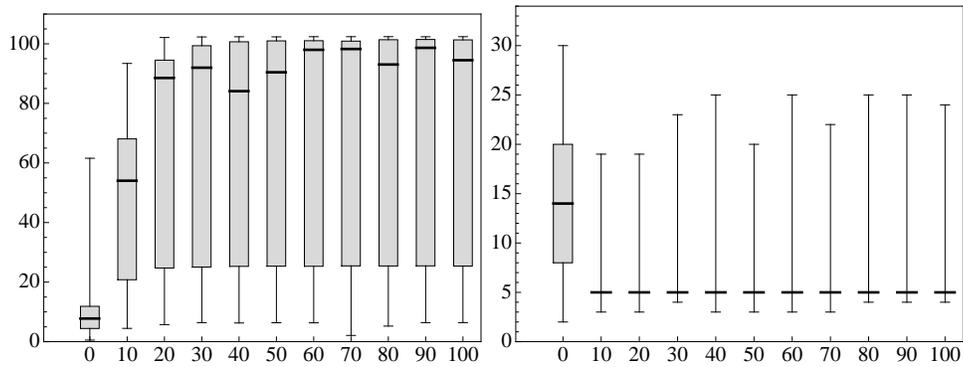


Figure 11: Result of simulation selecting for fast development, then large pupae. Left: Size as a function of time. Right: Development time as a function of time.

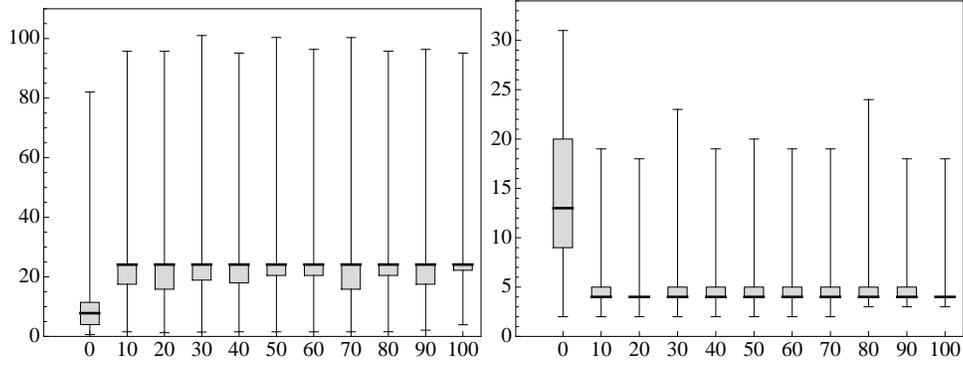


Figure 12: Result of simulation selecting for large pupae, then fast development. Left: Size as a function of time. Right: Development time as a function of time.

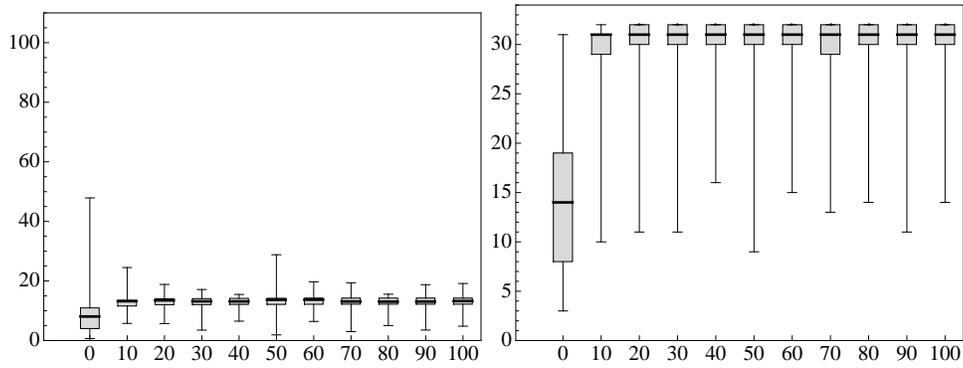


Figure 13: Result of simulation selecting for slow development, then large pupae. Left: Size as a function of time. Right: Development time as a function of time.

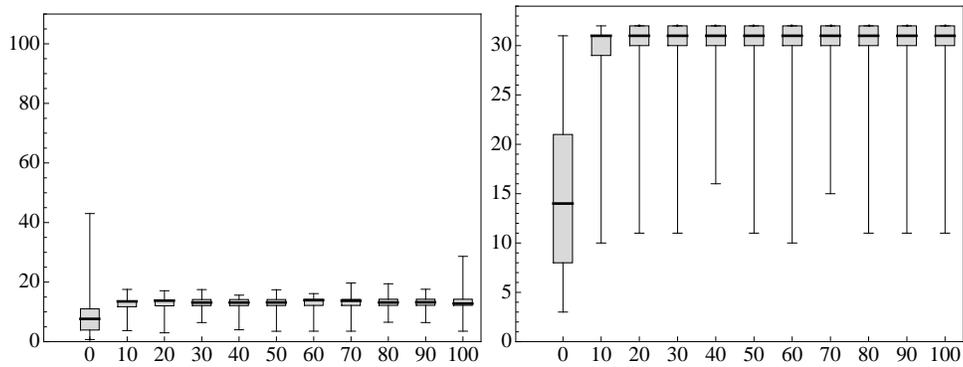


Figure 14: Result of simulation selecting for large pupae, then slow development. Left: Size as a function of time. Right: Development time as a function of time.

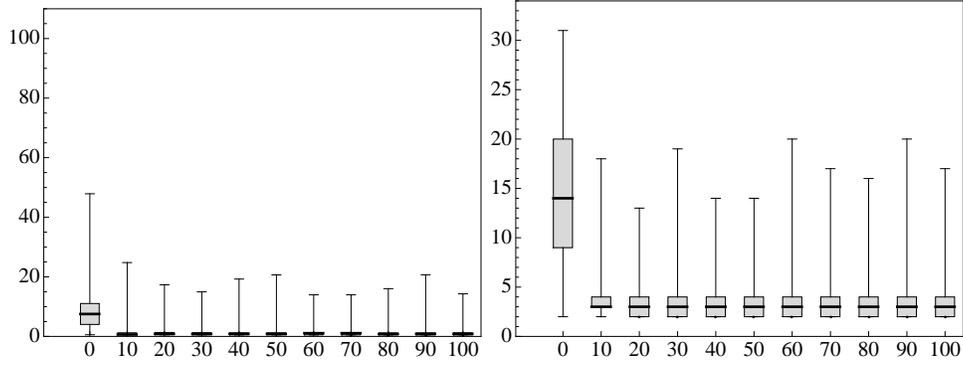


Figure 15: Result of simulation selecting for fast development, then small pupae. Left: Size as a function of time. Right: Development time as a function of time.

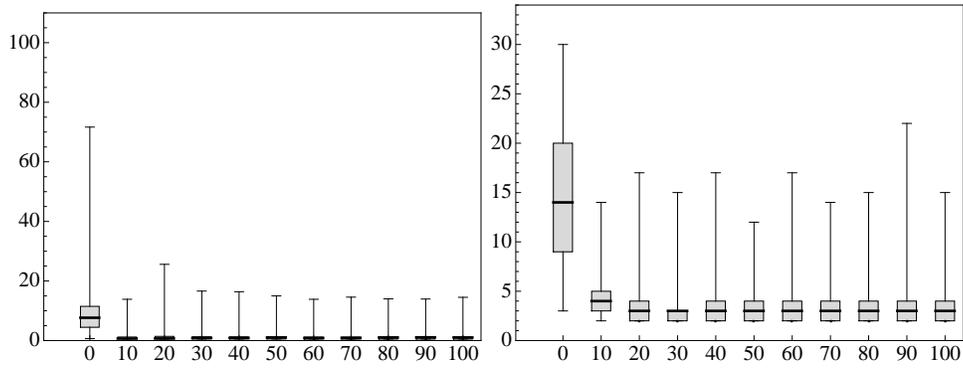


Figure 16: Result of simulation selecting for small pupae, then fast development. Left: Size as a function of time. Right: Development time as a function of time.

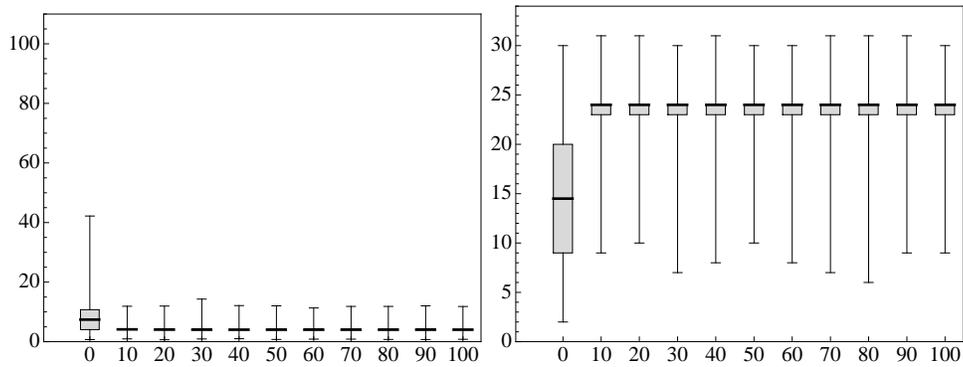


Figure 17: Result of simulation selecting for slow development, then small pupae. Left: Size as a function of time. Right: Development time as a function of time.

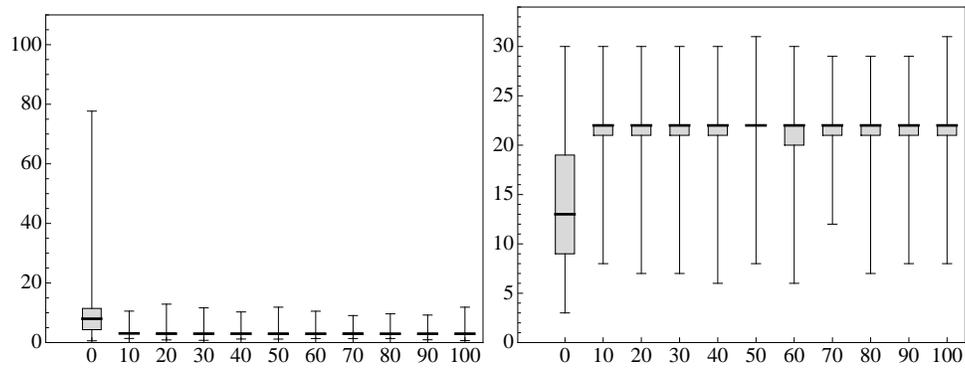


Figure 18: Result of simulation selecting for small pupae, then slow development.  
 Left: Size as a function of time. Right: Development time as a function of time.

## References

- [1] E. J. Briscoe. Grammatical acquisition: Inductive bias and coevolution of language and the language acquisition device. *Language*, 76(2):245–296, 2000.
- [2] E. J. Briscoe, editor. *Linguistic Evolution through Language Acquisition: Formal and Computational Models*. Cambridge University Press, 2002.
- [3] E. J. Briscoe. Grammatical acquisition and linguistic selection. In *Linguistic Evolution through Language Acquisition: Formal and Computational Models* Briscoe [2]. URL <http://www.cl.cam.ac.uk/users/ejb/creo-evol.ps.gz>.
- [4] Angelo Cangelosi and Dominico Parisi, editors. *Simulating the Evolution of Language*. Springer-Verlag, 2002.
- [5] Wolfgang Enard, Molly Przeworski, Simon E. Fisher, Cecelia S. L. Lai, Victor Wiebe, Takashi Kitano, Anthony P. Monaco, and Svante Pääbo. Molecular evolution of FOXP2, a gene involved in speech and language. *Nature*, 418(6900):869–872, August 2002.
- [6] W. Tecumseh Fitch. The evolution of speech: a comparative review. *Trends in Cognitive Sciences*, 4:258–267, 2000.
- [7] M. Gopnik and M. B. Crago. Familial aggregation of a developmental language disorder. *Cognition*, 39(1):1–50, April 1991.
- [8] Marc D. Hauser. *The Evolution of Communication*. Harvard University Press, Cambridge, MA, 1996.
- [9] Marc D. Hauser, Noam Chomsky, and W. Tecumseh Fitch. The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298(5598):1569–1579, November 2002.
- [10] John A. Hawkins and Murray Gell-Mann. *The Evolution of Human Languages*. Addison-Wesley, Reading, MA, 1992.
- [11] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.
- [12] James R. Hurford, Michael Studdert-Kennedy, and Chris Knight, editors. *Approaches to the Evolution of Language*. Cambridge University Press, 1998.
- [13] Lorens A. Imhof, Drew Fudenberg, and Martin A. Nowak. Tit-for-tat or win-stay, lose-shift? *Journal of Theoretical Biology*, 247:574–580, 2007.
- [14] Ray Jackendoff. Possible stages in the evolution of the language capacity. *Trends in Cognitive Sciences*, 3:272–279, 1999.

- [15] Simon Kirby. Spontaneous evolution of linguistic structure: an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, 5(2):102–110, 2001.
- [16] Simon Kirby and James R. Hurford. The emergence of structure: An overview of the iterated learning model. In Cangelosi and Parisi [4], chapter 6.
- [17] Natalia L. Komarova and Martin A. Nowak. The evolutionary dynamics of the lexical matrix. *Bulletin of Mathematical Biology*, 63(3):451–485, 2001.
- [18] Natalia L. Komarova and Martin A. Nowak. Natural selection of the critical period for language acquisition. *Proceedings of the Royal Society of London, Series B*, 268:1189–1196, 2001.
- [19] Natalia L. Komarova, Partha Niyogi, and Martin A. Nowak. The evolutionary dynamics of grammar acquisition. *Journal of Theoretical Biology*, 209(1):43–59, 2001.
- [20] Cecelia S. L. Lai, Simon E. Fisher, Jane A. Hurst, Faraneh Vargha-Khadem, and Anthony P. Monaco. A forkhead-domain gene is mutated in a severe speech and language disorder. *Nature*, 413(6855):519–523, October 2001.
- [21] Philip Lieberman. *The Biology and Evolution of Language*. Harvard University Press, Cambridge, MA, 1984.
- [22] David Lightfoot. *The Development of Language: Acquisition, Changes and Evolution*. Blackwell Publishers, 1999.
- [23] John Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.
- [24] John Maynard Smith and Eors Szathmary. *The Major Transitions in Evolution*. W. H. Freeman, New York, 1995.
- [25] W. Garrett Mitchener. Bifurcation analysis of the fully symmetric language dynamical equation. *Journal of Mathematical Biology*, 46:265–285, March 2003.
- [26] W. Garrett Mitchener. Game dynamics with learning and evolution of universal grammar. *Bulletin of Mathematical Biology*, 69(3):1093–1118, April 2007. DOI 10.1007/s11538-006-9165-x.
- [27] W. Garrett Mitchener. *A Mathematical Model of Human Languages: The interaction of game dynamics and learning processes*. PhD thesis, Princeton University, 2003.
- [28] W. Garrett Mitchener and Martin A. Nowak. Competitive exclusion and coexistence of universal grammars. *Bulletin of Mathematical Biology*, 65(1):67–93, January 2003.

- [29] W. Garrett Mitchener and Martin A. Nowak. Chaos and language. *Proceedings of the Royal Society of London, Biological Sciences*, 271(1540):701–704, April 2004. DOI 10.1098/rspb.2003.2643.
- [30] Frederick J. Newmayer. Functional explanation in linguistics and the origin of language. *Language and Communication*, 11:3–96, 1991.
- [31] H. F. Nijhout. The control of body size in insects. *Developmental Biology*, 261(1):1–9, September 2003. DOI: 10.1016/S0012-1606(03)00276-8.
- [32] H. Frederik Nijhout. The systems biology of body size regulation: A multidimensional mechanism. Presentation at the conference “Applications of Analysis to Mathematical Biology” at Duke University, May 2007.
- [33] Partha Niyogi. *The Informational Complexity of Learning*. Kluwer Academic Publishers, Boston, 1998.
- [34] Partha Niyogi. *The Computational Nature of Language Learning and Evolution*. MIT Press, Boston, 2006.
- [35] Partha Niyogi and Robert C. Berwick. Evolutionary consequences of language learning. *Linguistics and Philosophy*, 20:697–719, 1997.
- [36] Martin A. Nowak. *Evolutionary Dynamics: Exploring the equations of life*. Harvard University Press, 2006.
- [37] Martin A. Nowak and Natalia L. Komarova. Towards an evolutionary theory of language. *Trends in Cognitive Sciences*, 5(7):288–295, July 2001.
- [38] Martin A. Nowak and David C. Krakauer. The evolution of language. *Proceedings of the National Academy of Sciences, USA*, 96:8028–8033, 1999.
- [39] Martin A. Nowak, D. C. Krakauer, and A. Dress. An error limit for the evolution of language. *Proceedings of the Royal Society of London, Series B*, 266:2131–2136, 1999.
- [40] Martin A. Nowak, Joshua Plotkin, and David C. Krakauer. The evolutionary language game. *Journal of Theoretical Biology*, 200:147–162, 1999.
- [41] Martin A. Nowak, Joshua Plotkin, and V. A. A. Jansen. Evolution of syntactic communication. *Nature*, 404(6777):495–498, 2000.
- [42] Martin A. Nowak, Natalia L. Komarova, and Partha Niyogi. Evolution of universal grammar. *Science*, 291(5501):114–118, 2001.
- [43] Martin A. Nowak, Natalia L. Komarova, and Partha Niyogi. Computational and evolutionary aspects of language. *Nature*, 417(6889):611–617, June 2002.

- [44] Steven Pinker. *The Language Instinct*. W. Morrow and Company, New York, 1990.
- [45] Steven Pinker and Paul Bloom. Natural language and natural selection. *Behavioral and Brain Sciences*, 13:707–784, 1990.
- [46] Joshua Plotkin and Martin A. Nowak. Language evolution and information theory. *Journal of Theoretical Biology*, 205:147–159, 2000.
- [47] Peter E. Trapa and Martin A. Nowak. Nash equilibria for an evolutionary language game. *Journal of Mathematical Biology*, 41:172–188, 2000.
- [48] Charles D. Yang. *Knowledge and Learning in Natural Language*. Oxford University Press, Oxford, 2002.