Supplemental Material for Chapter 5, “Neo-Darwinian School”

Evolution, Optimization, and the “No Free Lunch Theorem”

Version 1.0

It is quite common to find species that persist for millions of years in the fossil record. Because these species are so stable, it has to be assumed that they are highly optimized for their environments; otherwise they would be quickly replaced by other species that are. This is an underlying assumption of NDT: species occupy fitness peaks, or what Dennett calls a good area of design space. Now if species sit on fitness peaks, significant physiological changes (or even small ones) will almost invariably result in reduced fitness, moving the species downhill away from its local peak. Thus evolving populations must pass through a region of less fitness, a valley that separates these two peaks. This means that evolution must effectively solve a problem known in mathematics as a *non-convex optimization problem*. Except for the simplest cases, this type of problem is extremely difficult to solve.

Let us see what the problem is, and why it is so difficult. A *convex* optimization problem is easy to understand and easy to solve. A simple example is finding the highest point on a typical hill: you know you are at the top because if you go in any direction, you start descending (see Figure 1). If you are not at the top, you just move in an upward direction and eventually you will get there. The word *convex* is used to describe such a problem because the shape of the hill is convex. ,

Now consider a simple non-convex problem, still involving hills, as shown in Figure 2. Suppose that the object is to find the highest point. Starting from the top of mountain *A*, we need only send out two skiers, one to the right and another to the left, and wait for them to report back. If there is a higher peak, one of them will find it—in this case, mountain *B*. (We assume that they are cross-country skiers and can climb hills! We may assume further that all the mountains are obscured by clouds, so that we cannot see which is the highest peak.) This is considered to be a one-dimensional problem, since movement is permitted only alone one axis (horizontal).

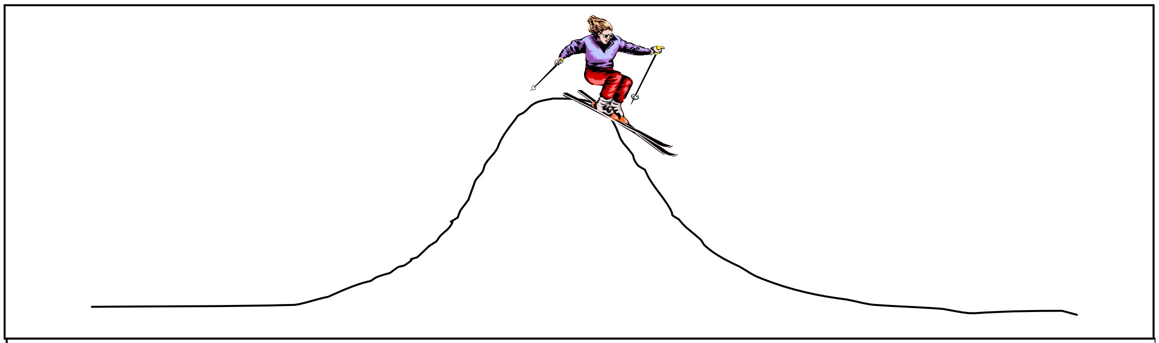

Figure 1. The simple, convex optimization problem. The skier knows she's at the top of the hill because every path leads down.

In this case, the height of the hill corresponds to fitness with higher peaks corresponding to a higher level. To solve such a problem, a population need only generate offspring of two types, some that migrate to the left and some that migrate to the right. After a certain number of generations, organisms originally on peak *A* can migrate to peak *B* if 1) enough variants that migrate to the left are generated and 2) the valley between peaks *A* and *B* is not so deep as to compromise fitness to the extent that the organisms traversing it perish. It is step two that creates problems for Darwinian evolution. Getting from *A* to *B* by incremental changes will always result in an initial decrease in fitness as the valley begins to be traversed.

The only way to explain how to traverse a fitness valley is to explain the valley away, and that is exactly what is done by Neo-Darwinists. The topology of Figure 3 represents an environment in which a population may find itself. In that particular environment, peak *B* represents the fittest genotype and peak *A* represents a fairly fit one although not as fit as peak *B*. If the environment were to change, as environments do, one could imagine that peak *A* could become the fittest peak as it now towers above peak *B* (fig 3(A)). In fact, one might imagine that the valley separating the two peaks might, in some specific environment, become much less step making it much easier to traverse (Fig 3(B)). Although this is possible, *imagining* that the valleys disappear is not the same as *demonstrating* that they do. In addition, there are some valleys that are not driven entirely by the environment (whether it is hot or cold, wet or dry), but rather have to do with internal physical constraints upon the organism. Such constraints do not change with the environment and therefore the valleys remain. For the most part these constraints involve the stability of the system's architecture, the limitations on size posed by an exoskeleton for example, and will be discussed in the next section.

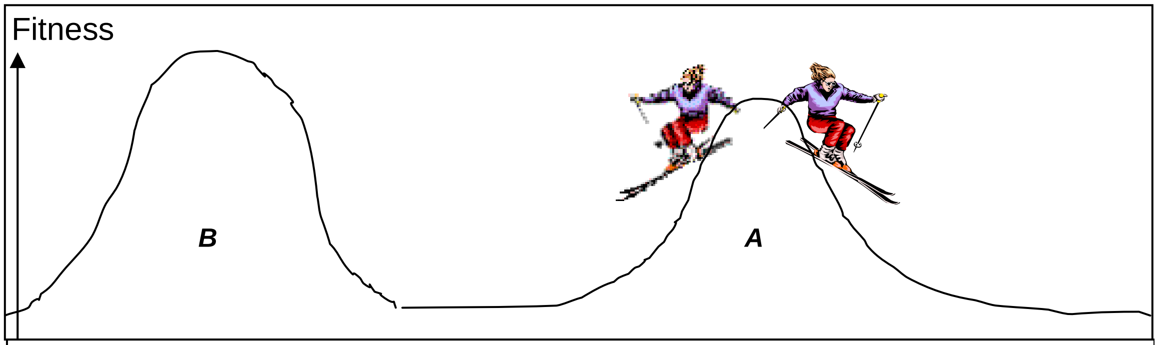

Figure 2. The simplest non-convex optimization problem. There are only two possibilities, and one skier will find a higher peak if it exists.

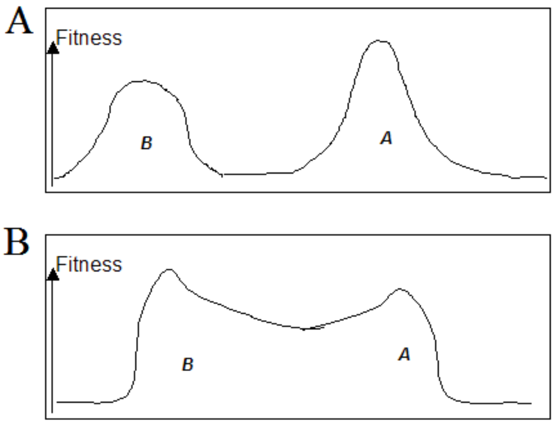

Figure 3. Hypothetical examples of a shifting fitness landscape under varying environmental conditions. Notice how the valley becomes less deep in environment B, facilitating organism change.

Now let us return to non-convex optimization problems and turn the heat up a notch. So far we have looked at simple one-dimensional problems, now we move on to a more realistic (and more difficult) case, a two-dimensional problem, as illustrated in Figure 4. In this case, movement is permitted in two directions, and the number of possible paths to be explored is significantly greater. In biological terms, this means that two different characteristics must be simultaneously changed in order to explore all possibilities, e.g., both color and size. We see that most of the explorers will fail to reach any mountain, and even those headed in the right direction may be killed enroute, if the journey is a long one and traverses a large valley. We may quantify the new situation if we assume that the distant mountains subtend about 1 degree each. Then if 360 explorers are sent out, only four will reach a new mountain, for a probability of success of about 1%. (Again making the favorable assumption that they don't die in the valley) This contrasts sharply with the previous one-dimensional case, where the probability of success was 50%.

If we go to yet another dimension—and advance to a still more realistic problem—the situation, as the reader may have surmised, is considerably more difficult. Biologically, this would correspond to changing three characteristics simultaneously, such as height, color, and limb length. The number of 1 degree square patches in the sky visible from a point in space, which corresponds to the number of explorers to be sent out, is no longer 360 but 360 x 360 = 129,600. The probability of success will depend, clearly, on how many "mountains" we assume are out there, but the relatively small number of actual species, as compared to the number of possible species, indicates that the number is small. In any case, the number of explorers needed is going up at an alarmingly rapid rate, exponential in fact, and the problem just gets worse with higher dimensions. We may note that 360^N becomes very large even for modest values of N (N corresponds to the number of dimensions minus 1), say 10, on the order of 10^{25}. The brute force approach of sending out explorers in every direction becomes cost-prohibitive very quickly, even for the massive parallelism available to biological systems.

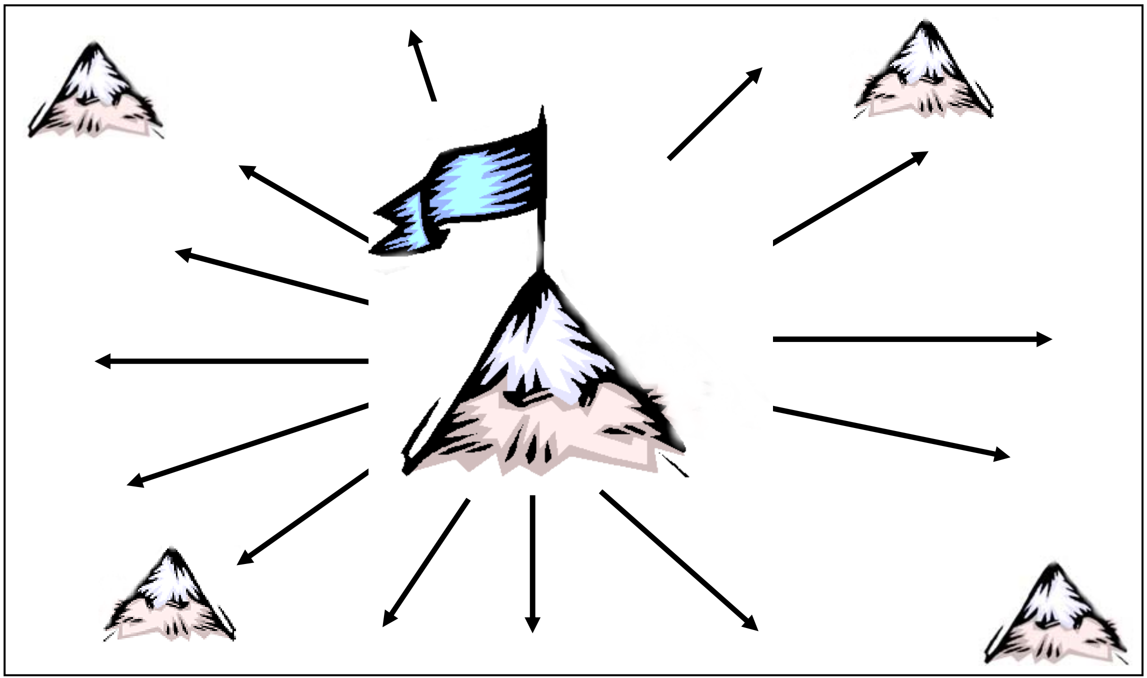

Figure 4. A simple two-dimensional non-convex optimization problem. The number of possible exploration routes is much larger, and the probability of success much lower.

This is the problem faced by evolution in its quest to create improved species: how to explore all of the possibilities, when only finite resources are available, most of which must be devoted to maintaining the status quo. There is no indication of which direction is most favorable, how far it is to get to the nearest fitness peak (even if it is lower than the current one), or even if it is possible to get there at all using small steps to traverse a valley of low fitness (which may be a function of environmental conditions). There is, in addition, the problem of demonstrating that the architecture of a given species can be modified sufficiently to enable organisms to remain viable at all over large changes in one or more characteristics, regardless of environmental constraints.

The "No Free Lunch" Theorem

So what mathematical methods are available to solve such problems, and what, concretely, does evolution do to solve them? The methods are many, including hill-climbing (mentioned above), and hill-descending. A general discussion of these solution methods to non-convex optimization problems is beyond the scope of this book, but we may note that nature utilizes mutation and crossover as techniques. Nature's own methods have been mimicked by an algorithmic technique known as "Genetic Algorithms" or GAs. The question, then, becomes just how good these alternative techniques can be at solving the problem. That question was examined recently by two researchers from the Santa Fe Institute, David Wolpert and William Macready. They proved a most surprising and unexpected result, since known as the "No Free Lunch" (NFL) Theorem:

> …all algorithms that search for an extremum of a cost function perform exactly the same, according to any performance measure, when averaged over all possible cost functions. In particular, if algorithm A outperforms algorithm B on some cost functions, then loosely speaking there must exist exactly as many other functions where B outperforms A.[1]

In practice, this means that no algorithm, on average, can solve optimization problems faster than the brute force approach. Since that approach is extremely limited in what it can do, this is a very serious problem for biological evolution. Still worse is to come, however. We have been assuming, up to this point, that the search directions in space can be selected independently—equivalent to saying that organism physiological traits can be varied independently, for example height can be varied without varying weight. In fact, this is not always possible, as many traits are controlled or influenced by a single gene. This constraint will severely restrict the range of searches that can be done, and was discussed in connection with the question of pleiotropy[*] in the book, and thus significantly reduce the probability of a successful search. Researchers have recognized a fundamental inconsistency between the variations in genes observed and those required by the theory:

> The results of the last 20 years of research on the genetic basis of adaptation has led us to a great Darwinian paradox. *Those* [genes] *that are obviously variable within natural populations do not seem to lie at the basis of many major adaptive changes, while those* [genes] *that seemingly do constitute the foundation of many, if not most, major adaptive changes apparently are not variable within natural populations*.[2]

This is a very serious problem for any theory which relies on incremental change to create new species and higher taxa. Along these lines we may also discuss the related issue of competitive displacement, which is a prediction of the NDT, based as it is on the notion of superior fitness leading to larger numbers of offspring—essentially the basis for the type of optimization predicted by any type of Darwinian theory. What should be observed is that replacement occurs fairly rapidly, and that evolutionary success is correlated closely with ecological dominance. This is not what is observed, a point conceded by evolutionists:

> …we have considered two reasons why one higher taxon may replace another, and how to test which theory has operated. First, competitive displacement has often

---

[*] Pleiotropy refers to the case of a single gene controlling multiple phenotypic traits. The fact that a single gene controls more than one trait makes changes to any of those traits much more difficult, as improvements to one phenotypic characteristics usually result in degradation for others.

> been put forward as a hypothesis, but it is difficult to support with evidence. Second, the rises and falls of higher taxa suggest that replacements often happen after the earlier group has gone extinct; the second group then radiates into the empty ecological space.[3]

But if competitive displacement cannot be supported with evidence, and replacements only occur after a group has become extinct, then there is a fundamental problem with the whole Darwinian paradigm of survival of the fittest, based as it is on this very principle. Other investigators have observed similar behavior. Searching for large-scale patterns in evolution, one group came to a rather startling conclusion:

> [Our] observations demonstrate that evolutionary success and ecological dominance can be decoupled and profoundly different, even over tens of millions of years.[4]

The investigators looked at the interaction between two branches of bryozoans,[†] the *Cheilostomata* and the *Cyclostomata*. New arrivals (*Cheilostomata*) came on the scene, but were unable to drive their rivals (*Cyclostomata*) to extinction. What is of particular interest in this case is the way in which the data are interpreted to fit into the Neo-Darwinian paradigm, when an equally obvious interpretation is that they contradict it.

Implications of the "No Free Lunch" Theorem

For now, however, let us concentrate on the rather startling biological implications of the theorem, which have been clearly delineated by Wolpert and Macready, who are not members of any anti-evolution camp:

> …the empirical evidence of the biological world does not indicate in any sense that natural selection [i.e., natural selection + mutation + crossover, etc.] is an effective search strategy. It does not even indicate that natural selection is an effective search strategy in the biological world. We simply have not had a chance to observe the behavior of alternative strategies. According to the NFL theorem, for all we know, the strategy of breeding only the *least fit* members of the population may have done a better job at finding the extrema of the cost function [i.e., fitness] faced by biological organisms. (This is exactly analogous to the fact that hill-descending can beat hill-climbing at finding fitness maxima.) The breed-the-worst strategy will in general result in worse recent *generations*, but simply the fact that you are using that strategy implies nothing about the quality of the *populations* over the long term.[5]

We may note that the authors are here assuming something which has not been proved, namely, that some form of naturalistic process *can* account for the observed characteristics and history of the world's flora and fauna; so they cannot be accused of bias against the NDT. But their own work does cast some doubt on this assumption, at least in terms of its presumed efficacy at generating new, improved species.

This problem is sometimes broached in qualitative terms, and then handwaved away:

[†] Bryozoans are small aquatic animals from the phylum Bryozoa which reproduce by budding, and which form mosslike or branching colonies that attach to stones or seaweed. The two branches investigated are from the Cretaceous period

> An organ must be advantageous to its bearer at all stages in its evolution if it is to be produced by natural selection. Some adaptations, it is said, although undoubtedly advantageous in their final form, could not have been when in a rudimentary form: "What is the use of half a wing?" is a familiar example. The anatomist St. George Jackson Mivart particularly stressed this argument in his work *The Genesis of Species* (1871). The Darwinian reply has been to suggest ways in which the character could have been advantageous in rudimentary form. In the case of the wing, partial wings might have broken the force of a fall from a tree, or proto-winged birds might have glided from cliff tops or between trees—as many animals, such as flying foxes, do now. These early stages would not have required the muscular back-up of a full, final wing. The concept of preadaptation…provides another solution to the problem."[6]

But this just papers over the difficult optimization problem. No partial-winged birds have been found, and no evidence produced that such wings would actually work as described. To even break the force of a fall, such a wing would have to be rather far along in development. Undoubtedly such a structure, before it became useful, would be a hindrance to the animal in its movements, thus having negative selective value. And how did the birds survive in the trees before they could fly anyway? Preadaptation doesn't solve the general problem either.[‡] First, the structures which are pre-adaptive had themselves to evolve. So this only pushes the problem back one stage. Moreover, there are certainly going to be hard limits to the number of structures and systems which can function as preadaptions—surely not all of evolution can be predicated on such fortuitous cases. But no attempt has been made to quantify these limits or even delineate the types of preadaptations which are possible. Worse, if preadaptation were widespread, it could readily be used as evidence for design.

So where does all of this leave us? Unfortunately—one might say inexplicably—still somewhat in the dark. It can reasonably be assumed that brute force searching of higher-dimensional spaces will not work, unless a number of arbitrary and *ad hoc* assumptions about those spaces are made, and in that case, one might as well concede design. The key research on this subject, namely, how close the nearest stable (high fitness) peaks are, what the dimensionality of the relevant space is, and how fatal the intervening valleys typically are, has not even begun. All we have is some indirect knowledge based on the Random Matrix Eigenvalue Problem, the implications of which do not bode well for biological evolution based on the paradigm of random mutation and natural selection.[§] Nor has any investigation of the efficacy of alternative search strategies even begun, as the authors of the NFL theorem have pointed out. For some reason, such important questioning of the fundamental principles of evolutionary biology is not done, in contrast to other branches of science. Perhaps this is because of their assumed "obviousness" or logical inevitability—qualities which, as we have seen in other essays, have been replaced by often counter-intuitive results and complex mathematics.

[‡] A preadaptation is a characteristic that evolved in an ancestral species or population that then serves an adaptive though different function in a descendant species or population.

[§] The Random Matrix Eigenvalue problem addresses the probability that a randomly connected system is stable. The goal is to figure out what percentage of all possible systems of any dimension are stable, and thus indirectly gain insight into the nature of the spaces to be explored in optimization problems such as those faced by biological systems.

The NFL theorem does not definitively *prove* that postulated biological processes do not or cannot work to produce new species; we simply don't know enough yet to draw such a conclusion on the basis of the NFL theorem alone.[**] But it, along with the general difficulties of the non-convex optimization problem, do cast considerable doubt on the efficacy of those processes—doubt that needs to be dispelled by proof that they *can* work. This indeed is an ideal case for further theoretical and empirical investigation, one that all schools can support, because it is an area where definitive results can contribute directly to resolution of the evolution controversy. (Of course, this assumes that the schools *want* to have the controversy resolved!)

---

[1] Wolpert, David H., Macready, William G., "No Free Lunch Theorems for Search," Santa Fe, NM: Santa Fe Institute Technical Report SFI-TR-95-02-010, 1996, p. 1.

[2] McDonald, John, "The Molecular Basis of Adaptation," *Annual Review of Ecology and Systematics* **14**:93 (1983).

[3] Ridley, Mark, *Evolution.*

[4] McKinney, Frank K., Lidgard, Scott, Sepkoski, J. John Jr., Taylor, Paul D., "Decoupled Temporal Patterns of Evolution and Ecology in Two Post-Paleozoic Clades," *Science* **281**:807-809 (1998).

[5] *Ibid.*, p. 29.

[6] Ridley, Mark, *Evolution*, *op. cit.*, p. 345.

[**] Stuart Kaufmann, realizing the difficulties this poses, argues that organisms evolve in such a way as to "create" optimization problems that natural selection can solve. This is discussed in the Meta-Darwinian chapter.