1

The concept of the adaptive landscape

The idea of a fitness landscape was introduced by Sewall Wright (1932) and it has become a standard imagination prosthesis for evolutionary theorists. It has proven its worth in literally thousands of applications, including many outside evolutionary theory.

Dennett (1996, p. 190)

What is an adaptive landscape?

An adaptive landscape is a very simple – but powerful – way of visualizing the evolution of life in terms of the geometry of spatial relationships, namely the spatial relationships one finds in a landscape. Consider an imaginary landscape in which you see mountains of high elevation in one region, towering mountains separated by deep valleys with precipitous slopes. In another region these mountains give way to lower elevation rolling hills separated by wide, gently sloping valleys, and that these further give way to broad flat plains in the distance. Now replace the concept of 'elevation' (height above sea level) with 'degree of adaptation' and you have an adaptive landscape. Why is that such a powerful concept? The purpose of this book is to answer that question.

The concept of the adaptive landscape was first proposed by the geneticist Sewall Wright in 1932. Being a geneticist, he thought in terms of genes rather than morphology and Darwinian fitness rather than adaptation, and his original concept is what is termed a fitness landscape today, rather than an adaptive landscape. The two concepts differ only in that the dimensions of a fitness landscape are genetic traits and
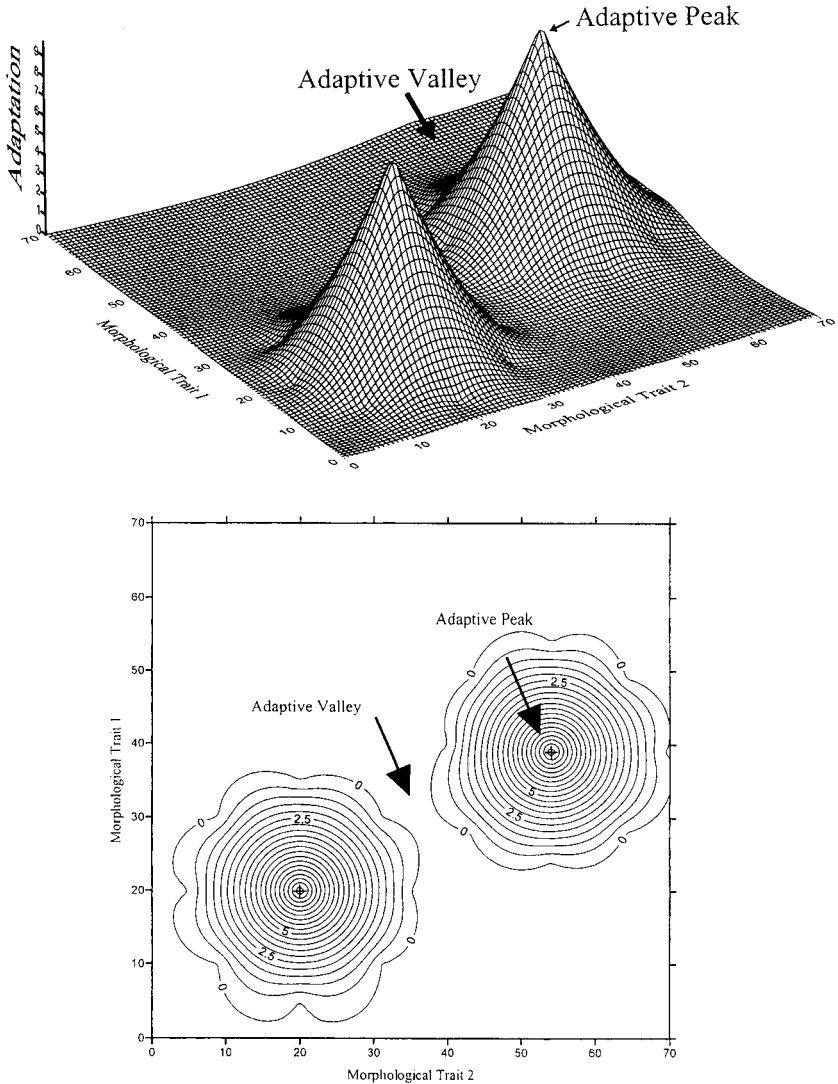
1

Figure 1.1. A hypothetical adaptive landscape, portrayed as a three-dimensional grid at the top of the figure and a two-dimensional contour map at the bottom. Topographic highs represent adaptive morphologies that function well in natural environments (and therefore are selected for), and topographic lows represent nonadaptive morphologies that function poorly in natural environments (and therefore are selected against). In the contour map portrayal, the top of an adaptive peak is indicated by a plus-sign, following the convention of Sewall Wright (1932).
Source: Modified from McGhee (1980a).

degree of fitness whereas the dimensions of an adaptive landscape are morphological traits and degree of adaptation (Fig. 1.1). A fitness landscape is used by geneticists to visualize evolution, and an adaptive landscape is used by morphologists. As I am a morphologist, a student of the evolution of biological form, this book will concentrate on adaptive landscapes and, beginning in Chapter 4, the very important related concept of the *theoretical morphospace*.

But back to Sewall Wright. His first crucial insight was that it could be possible (at least theoretically) to construct a space of all possible genetic combinations that living organisms might produce, and that one could visualize such a complex space by simply considering the possible combinations of two genes at a time or, in the case of an adaptive landscape, two morphological traits at a time (Fig. 1.1). That is, if genetic trait number one had 10 different variants or alleles, and genetic trait number two had 10 different variants, then the total possible genetic combinations of those two traits would be 100 potential variants.

Wright's second crucial insight was that *the majority of those 100 possible variants probably do not exist in nature*. Perhaps only 10 of the possible variants actually exist as living organisms, the other 90 variants potentially could exist but do not. Why not? Wright proposed that these 90 potential genetic combinations had zero fitness; that is, they represented lethal genetic combinations. The other 10 variants had fitness values greater than zero, some perhaps having higher fitness than others. Wright further proposed that these genetic relationships could be spatially visualized as geometric relationships by simply adding the dimension 'degree of fitness' to the two genetic trait dimensions, producing a three-dimensional grid similar to a landscape (as in the adaptive landscape in Fig. 1.1). If the landscape is portrayed in two-dimensions by using fitness contours to give the fitness dimension, then the result looks very much like a topographic map of a landscape. The 10 existent combinations of genetic traits number one and two would be located on the peaks or slopes of the hills within the landscape (depending upon their degree of fitness), and the 90 possible but nonexistent combinations of genetic traits number one and two would be located in the flat plain of zero fitness. Thus was born Sewall Wright's concept of the fitness landscape.

In adaptive landscapes the *high regions* are called adaptive peaks, and the *low regions* between the peaks are called adaptive valleys (Fig. 1.1). The degree of adaptation of the possible morphological traits

is determined by functional analyses of the potential forms; that is, analyses of how well the potential morphological variants function in nature. The geometric arrangement of the adaptive peaks within the landscape thus represents, in a spatial fashion, the different possible ways of life available to organisms. The spatial distribution of the adaptive valleys and plains represents ecomorphologies that are nonfunctional in nature.

Modelling evolution in adaptive landscapes

Adaptive landscapes are potentially very powerful tools for the analysis of the evolution of life. Life is constantly evolving, and we would like to know why life has evolved the way that it has in the past three and one-half thousand million years of Earth history, and perhaps be able to predict how life might evolve in the future.

Although evolution itself is a fact, an empirical observation, the cause of evolution is theoretical. That is, there exist several different theories to explain how evolution takes place. The most widely subscribed-to theory of how evolution takes place is that of natural selection, first proposed by Charles Darwin. If he had not proposed it, Alfred Wallace would have instead; thus it was clearly an idea whose time had come in the 1800s. What is natural selection? A precise, rather pithy definition is the 'differential change in genotypic frequencies with time, due to the differential reproductive success of their phenotypes' (modified from Wilson and Bossert, 1971). The first part of the definition ('differential change in genotypic frequencies with time') is simply a restatement of evolution itself, in that evolution is genetic change in populations from generation to generation. The real heart of the theory is 'differential reproductive success' of various phenotypes, or morphologies. If certain organisms with certain morphologies in a population reproduce at a higher rate than other organisms with other morphologies, then the next generation will contain more of their genes than the previous one. And that change in gene frequencies, from generation one to generation two, is by definition evolution. Thus natural selection could clearly drive evolution.

The definition of natural selection does not specify what causes differential reproductive success; it simply holds that if it does occur, evolution will result. The next question is obvious: what determines the differential reproductive success of differing phenotypes, or morphologies, such that

different animals and plants reproduce at different rates? It is here that the concept of adaptation enters the equation. Organisms must function in their environments, and they must interact with other organisms. If some organisms possess morphologies and behaviours (aspects of their phenotypes) that allow them to function well in their ecological setting then they are described as well adapted. Well adapted organisms are healthy, well fed and potentially able to devote more time and energy to reproduction. If other organisms possess morphologies that do not allow them to function as well − say, they cannot run as fast due to the different structure of their legs, or cannot find their prey or other food as quickly due to the different structure of their eyes or ears (their visual and auditory systems) − then they are described as poorly adapted. Poorly adapted organisms must spend more time simply trying to escape predators and to find food, are generally less healthy and spend less time and energy in reproduction.

Wright's concept of a fitness, or adaptive, landscape is firmly rooted in the theory of natural selection (we shall see in Chapter 4 that the concept of the theoretical morphospace *is not*). In the previous section we have seen that an adaptive landscape is an actual spatial map of the different possible ecomorphologies that are available to organisms, and of other possible ecomorphologies that are nonfunctional and thus not available to organisms. What would happen now if we place a population of actual organisms within the adaptive landscape, say half-way up the side of an adaptive peak, and observe the evolution of that population with time? A basic rule of modelling evolution in adaptive landscapes is that *natural selection will operate to move a population up the slope of an adaptive peak, from lower degrees of adaptation to higher degrees of adaptation*. That seemingly simple rule has some intriguingly complicated consequences, however, as we shall see in the next chapter.

2

Modelling natural selection in adaptive landscapes

'Wedges in the economy of nature' wrote Darwin in his diary, leaving us with a glimpse of his own first glimpse of natural selection ... Later biologists, by the fourth decade of the twentieth century, would invent the image of an adaptive landscape whose peaks represent the highly fit forms, and see evolution as the struggle of populations of organisms driven by mutation, recombination, and selection, to climb toward those high peaks. Life is a high-country adventure.

Kauffman (1995, p. 149)

Visualizing natural selection

We have seen in the last chapter that an adaptive landscape is a way of visualizing the evolution of life in terms of the geometry of the spatial relationships one finds in a landscape, where the landscape consists of adaptive hills and valleys. If we use the theory of natural selection to model evolution within an adaptive landscape, we saw that natural selection will operate to move a population up the slope of an adaptive peak, from lower degrees of adaptation to higher degrees of adaptation.

What happens, however, when an evolving population reaches the top of an adaptive peak? Or what happens if an evolving population encounters two peaks in an adaptive landscape, rather than one? Clearly natural selection will operate in different ways at different times in the evolution of any group of organisms, depending upon the environmental and ecological context within which that group of organisms is evolving. The adaptive landscape concept allows

6

us to visualize the possible effects of natural selection through simple spatial relationships, rather than complicated modelling of changing environmental or ecological conditions.

Modelling directional selection

Let us consider again the situation where a population of animals of plants is positioned half-way up the slope of an adaptive peak. In this situation, natural selection will operate to move the population up the slope of an adaptive peak, from lower degrees of adaptation to higher degrees of adaptation. But how does natural selection actually accomplish this?

Natural selection operates on variation in nature. If there were no variation in nature, natural selection would cease (note that evolution itself may not cease, however, because evolution may be driven by more than natural selection − we shall consider this possibility in more detail later). That is, if a population of animals is composed of individuals that are all identical in the state of their adaptive morphologies − for example, if they are all clones inhabiting the same environment − then they should all function equally as well in a given environment. Natural selection would have no differences in adaptive morphology to 'select' and all the individuals should reproduce at more or less the same rate, with some random variation. Such a hypothetical situation is very rare in nature, however, where variation is the normal natural condition.

There are two main sources of variation for natural selection to operate with. One is genetic recombination, the other is genetic mutation. Genetic recombination is the constant reshuffling of genes that occurs from generation to generation in sexually reproducing organisms. Imagine all the genes present within a species, its genome, to be represented by a deck of 52 playing cards, where each card is a gene. Imagine further that the morphology of any individual animal is determined by four cards (genes), then you can divide your deck of 52 cards into 13 individual animals in generation number one. Each time the animals reproduce to produce a new generation you reshuffle all the cards again and draw another 13 sets of four cards each. In this simple exercise (actual genetic recombination is much more complicated) you can easily see how much variation is produced from generation to generation by merely reshuffling the same genes over and over again. Genetic mutation, on the other hand, is the appearance of a new genetic coding − a new card in the deck that was not present there previously.

8   *Modelling natural selection in adaptive landscapes*

Consider that our hypothetical population is composed of two major variants, animals with A-type morphologies and animals with B-type morphologies, and that in generation number one the population is equally divided in numbers of individuals with A-type and B-type morphologies (Fig. 2.1). However, let us further imagine that animals with A-type morphologies function a bit better in the environment than animals with B-types; that is, A-types have a somewhat higher degree
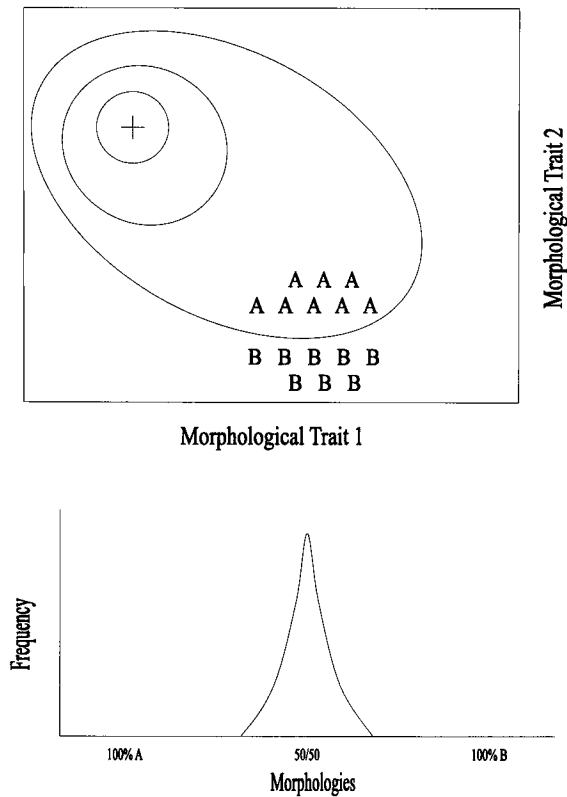

Figure 2.1. Modelling directional selection, part one. The spatial positions of individuals of a hypothetical species population, composed of organisms with morphological variants A and B, are depicted within an adaptive landscape in the top figure. Individuals with morphological variants A have a higher adaptive value than B (these variants are shown upslope from B) and thus, under the expectations of the theory of natural selection, organisms with A-type morphologies should reproduce at a higher rate than those with B-type morphologies. The initial frequency of organisms with morphological variants A and B within the species is depicted as roughly equal in the graph given in the bottom figure.

of adaptation than B. We can represent this selective difference in an
adaptive landscape by bisecting the population with an adaptive contour,
where individuals with A-type morphologies are on the upslope side
of the contour, and individuals with B-type morphologies are on the
downslope side (Fig. 2.1).

Now let us consider the state of the population after several genera-
tions of reproduction under the influence of natural selection. We would
predict that the somewhat better adapted animals with A-type morphol-
ogies should reproduce at a somewhat higher success rate than the less
well-adapted animals with B-type morphologies, and that the number of
individuals with A-type morphologies now comprise a greater percentage
of the total population numbers than individuals of B-type (Fig. 2.2).
That is, we now have more individuals in the population on the upslope
side of the adaptive gradient than on the downslope side (Fig. 2.2).
The population is moving uphill.

Sooner or later, however, the uphill movement of the population will
cease when all of the individuals in the population have A-type morphol-
ogies, and all have the same degree of adaptation. Now we need to
introduce the second source of variation into the equation: genetic
mutation. Let us introduce three new morphological variants into the
scenario at random, as genetic mutation is random. One new variant, X,
has a morphology that is further upslope than the parent population;
another new variant, Y, has a morphology that is further downslope
than the parent population; the last variant, B, is a backmutation to
a previously existent morphology (Fig. 2.3). Under the expectations of
the theory of natural selection, variants Y and B will be selected against —
they will have an even lower success rate of reproduction than the
individuals in the parent population that produced them. On the other
hand, however, variant X will be selected for, and would have a higher
success rate of reproduction than the individuals in the parent popu-
lation itself. With time, individuals with X-type morphologies should
become more and more numerous, and individuals with A-type morphol-
ogies less numerous; thus the population continues to move uphill.

The effect of natural selection in this particular scenario is termed
directional selection. We can summarize the effects of directional selec-
tion in Figure 2.4 with a series of vectors that indicate that the effect of
natural selection will always be to select genetic mutational morphologies
that possess higher degrees of adaptation, and that the net result of
natural selection is evolution that always proceeds in the uphill direction
in an adaptive landscape.
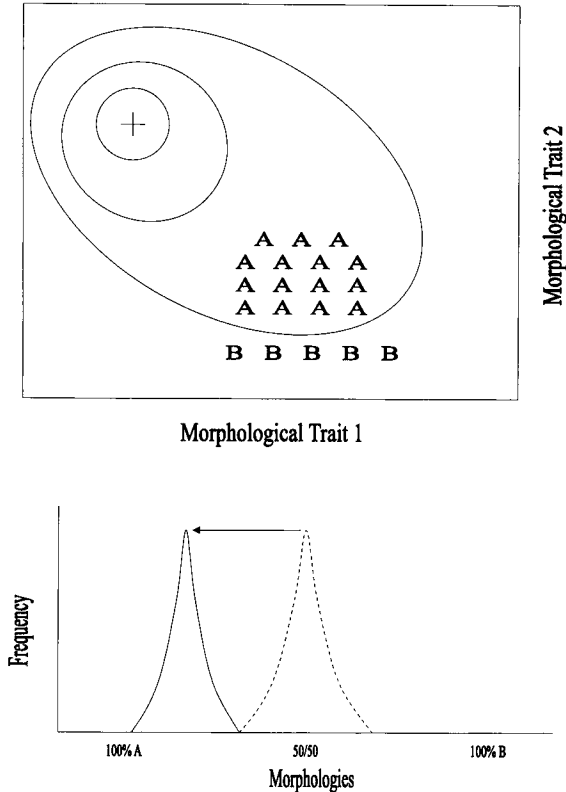
Figure 2.2. Modelling directional selection, part two. The spatial positions of individuals of the hypothetical species population within the adaptive landscape after several generations of natural selection (top figure). The number of organisms possessing the higher-adaptive A-type morphologies has increased within the species (top figure), and the frequency distribution of morphologies within the species has shifted to the left in the graph given in the bottom figure.

Modelling stabilizing selection

As directional selection operates to produce evolution in the uphill direction, sooner or later the evolving animals or plants will reach the adaptive peak, the local point of maximum degree of adaptation within the adaptive landscape. What happens then?

Once at the peak, any major new source of variation will always be in the downslope direction, and thus will be selected against. Consider a hypothetical population sitting on top of an adaptive peak, and