Biol Philos (2008) 23:625–638
DOI 10.1007/s10539-008-9116-z

The end of the adaptive landscape metaphor?

Jonathan Kaplan

Received: 25 February 2008 / Accepted: 29 April 2008 / Published online: 5 June 2008

Abstract The concepts of adaptive/fitness landscapes and adaptive peaks are a central part of much of contemporary evolutionary biology; the concepts are introduced in introductory texts, developed in more detail in graduate-level treatments, and are used extensively in papers published in the major journals in the field. The appeal of visualizing the process of evolution in terms of the movement of populations on such landscapes is very strong; as one becomes familiar with the metaphor, one often develops the feeling that it is possible to gain deep insights into evolution by thinking about the movement of populations on landscapes consisting of adaptive valleys and peaks. But, since Wright first introduced the metaphor in 1932, the metaphor has been the subject of persistent confusion, from equivocation over just what the features of the landscape are meant to represent to how we ought to expect the landscapes to look. Recent advances—conceptual, empirical, and computational—have pointed towards the inadequacy and indeed incoherence of the landscapes as usually pictured. I argue that attempts to reform the metaphor are misguided; it is time to give up the pictorial metaphor of the landscape entirely and rely instead on the results of formal modeling, however difficult such results are to understand in 'intuitive' terms.

Keywords Adaptive landscape · Fitness landscape · Wright · Gavrilets · Holey landscape · Peak shift · Shifting balance theory · Bateson-Dobzhansky-Muller model · Speciation · Adaptation · Metaphor

---

J. Kaplan (✉)
Philosophy Department, Oregon State University, 208 Hovland Hall, Corvallis, OR 97331-3902, USA
e-mail: Jonathan.Kaplan@oregonstate.edu

Heuristics and history

Wright's metaphor of the fitness[1] landscape, and the associated images and diagrams, date back to his presentation at the Sixth International Congress of Genetics in 1932; according to Michael Ruse, Wright (1931), having been advised to keep his presentation short and non-technical, devised the adaptive landscape metaphor in order to explain his "Shifting Balance Theory" of evolutionary change (first developed in his rather technical 1931 paper) in a non-mathematical form (Ruse 1996). Wright proposed that natural selection would rapidly bring a population to a local optimum, but that in order to sample a larger range of possible adaptations, a mechanism was needed through which populations could sample a larger portion of "genotype space" (Wright 1932, pp. 3–4).

Wright represented this problem (and the associated solution) by first imagining all the possible genotypes of organisms with a particular number of loci and a particular range of alleles. For example, in the case of two loci with two alleles each, there are four possible genotypes, each of which is one "move" (roughly, one mutation) away from two genotypes, and two moves away from one; for five loci with two alleles each, there are 32 different available genotypes, each of which is one move away from 5 genotypes, 2 moves away from 10 genotypes, 3 moves away from another 10 genotypes, etc. (see Fig. 1). For organisms with 1,000 loci and 10 possible alleles at each, there are, Wright notes, 10^{1000} possible genotypes, each of which is one move away from 9,000 distinct genotypes, etc. Within a particular environment, each genotype can be assumed, Wright thought, to have a particular fitness; generally speaking, he suggested, genotypes that are close together (fewer moves apart) will tend to have more similar fitnesses than genotypes chosen at random (fitnesses will tend to be correlated with genotype location).

As Wright realized, if one wishes to represent the number of moves between genotypes as a distance, the dimensionality required increases with the number of loci and the number of alleles; in the above example, to represent the number of moves between any two genotypes as a distance, some 9,000 dimensions would be required. While it is obviously impossible to visualize a 9,000 dimensional genotype-space, Wright suggested that by visualizing an ordinary (three-dimensional) landscape, one could come to understand important aspects of the evolutionary process.

In Fig. 2, genotypic distance (the number of 'moves' between genotypes) is represented by their distance from each other in the two-dimensional picture; the fitness of each genotype is represented by the location of the genotype within the contour lines (the "height" of that genotype) (See Fig. 2). This generates a "fitness landscape" where the fitness of individual genotypes is mapped against their

---

[1] Some authors have suggested distinguishing between *adaptive* landscapes and *fitness* landscapes in the following way: *adaptive* landscapes represent each population as a single *point* on the landscape (representing either allelic or genotype frequencies in the population as a whole), with axes representing *frequencies* of alleles or genotypes, whereas *fitness* landscapes represent each unique genotype as a point on the landscape, with axes representing particular loci. In such fitness landscapes, populations are *collections of points* whose locations represents genetic distance from each other. As will become clear, I think this distinction is becoming moot, but will follow it nonetheless.

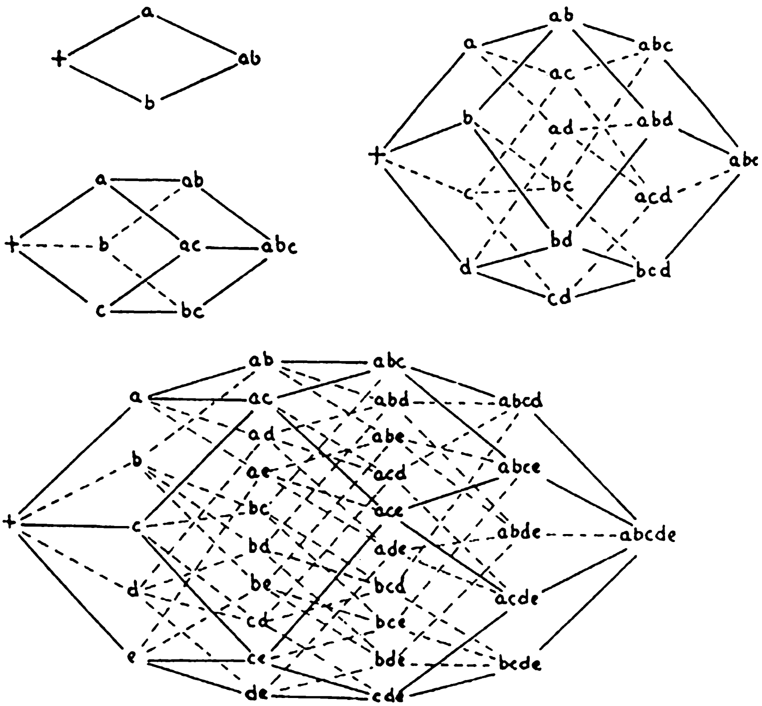

Fig. 1 Wright's fitness landscape becomes increasingly complex as one considers more alleles at one locus and/or more loci. It is obvious that the visualization of the "landscapes" soon becomes unmanageable for any attempt to model a realistic dimensionality of gene combinations. (From Wright 1932 Fig. 1)

distance from other possible genotypes. Wright argued that adaptive "peaks" (collections of relatively close genotypes that are of high-fitness) would be separated from each other by "valleys" of genotypes that are of relatively low fitness. In order to "move" from one adaptive peak to another, populations would have to "traverse" adaptive valleys of low fitness. So in order to explore any sizable portion of genotype space, a population would have to cross valleys of low fitness. While natural selection could not drive such a process, Wright argued that drift could do so, at least in small populations (see Fig. 3). From this, Wright argued, it followed that small populations experiencing drift were of vital importance to adaptive evolution; without such drifting populations, only a tiny portion of genotype space would ever be accessible to a population. This was, in essence, the aspect of Wright's Shifting Balance Theory that was most controversial.

There is an ongoing argument about the heuristic value of Wright's landscape diagrams. Provine (Wright's biographer) suggested that (one version of) the landscape diagram Wright presented was actually *incoherent*; the diagram attempts to represent the distance between discrete genotypes as a continuous surface, but there is no continuous variable associated with discrete genotypes (Provine 1986). That is, there is no way to get a surface out of a collection of genotypes and their one-step neighbors. Such a collection, visualized as Wright originally suggested, would yield a very

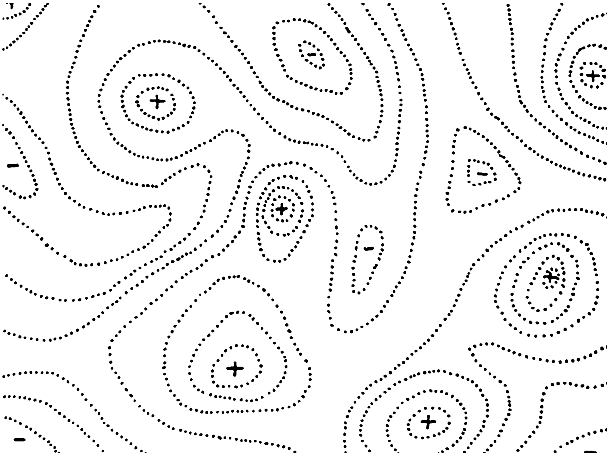

Fig. 2 Wright's visual representation of the fitness landscape idea, a field of n-dimensional gene combinations (represented here in two dimensions) and the relevant adaptive topography; note that the topography is not based on empirical evidence or computational models, but rather represents what Wright thought such landscapes likely resembled. Note too that there are no labels on the axes. (From Wright 1932 Fig. 2)

complex version of the pictures in Fig. 2, none of which suggest a continuous surface. For this reason, Provine suggested that the interpretation of the axes of landscape diagrams as either allelic or genotype frequencies was to be preferred. This produces an *adaptive* landscape diagram, in which each point represents a population with a particular mean fitness. While the allelic frequency interpretation of adaptive landscapes is more often represented pictorially in descriptions of landscapes, it is deeply problematic for several reasons. First, in general mean population fitness does not map onto allelic frequencies in a population—it matters greatly how the alleles are distributed in the population[2]. Second, it is broadly unexplanatory—it is unclear why envisioning populations as particular points is supposed to be helpful in understanding the nature of adaptive evolution.

To see why adaptive landscapes (those in which populations are modeled as points, and the axes represent allelic frequencies) are unexplanatory, consider for a moment why fitness landscapes, in which populations are envisioned as clouds of points, might be considered explanatory. If there is a mapping function in a particular environment between genotype-phenotype and phenotype-fitness, then each individual (each point) will contribute a number of offspring to the next generation with some probability. The "cloud" of points will tend to grow or to shrink, and/or to move across the landscape, based on the performance of the individual organisms that make up that population. The dynamics of the population are provided by the expected fitnesses of the members, the population structure, and the like. This is just what we would expect, and the landscape, which here represents the genotype-fitness mapping of the individuals, is (at least part of) what explains the population dynamics, including speciation events.

---

[2] This of course does not hold if every locus is at Hardy–Weinberg equilibrium; however, complete H–W equilibrium is not generally expected in real populations.

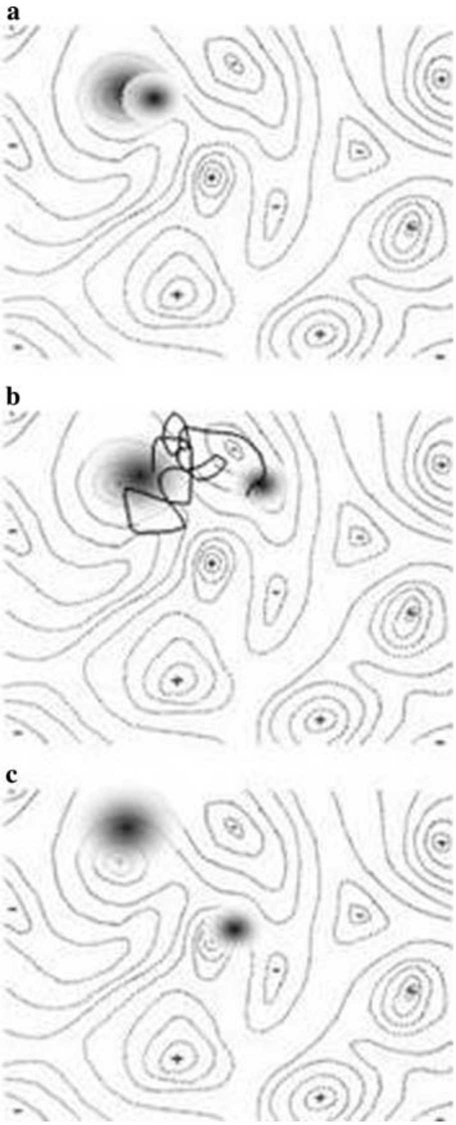

Fig. 3 Wright's "Shifting Balance" Theory. Loosely adapted from Wright (1932), Fig. 4e. (**a**) A small part of a population becomes isolated. (**b**) The small populations "wanders" around the landscape via "drift." (**c**) The small population establishes itself at a new fitness peak

But this kind of interpretation won't do at all for an adaptive landscape, in which allelic frequencies are mapped onto average population fitness. Here, the shape of the landscape is simply a description of the population dynamics, *not* an explanation of those dynamics. If one asks why a population moves in a particular way, the only 'answer' is in terms of the shape of the landscape, which again is simply a representation of the (likely) paths that populations will take. It is interesting to note that the interpretation of adaptive landscapes in terms of genotype frequency does not suffer from exactly this problem. When the mean fitness of the population is determined by summing over each extant genotype, and the landscape is shaped by all possible combinations of genotypes/population sizes, movement on the landscape is determined (again) by the individual genotype-fitness mapping function. But note that the value of this descriptions depending critically on the individual genotype interpretation; unless we can determine the individual genotype-fitness mapping functions, population structure, etc., we cannot determine the genotype-frequency landscape. Note as well that this style of landscape requires an absurdly high dimensionality, and is likely computationally untractable for any plausible populations in any event.

In any event, Ruse (1996) argued that Provine misunderstood what the Wright's fitness landscape diagram was trying to represent—the surface is not *really* continuous, but because of the huge number of genotypes surrounding each possible genotype (in very high dimensional space), it can be visualized as continuous when those dimensions are compressed and represented in two dimensions; this explains, for example, why Wright drew populations as *clouds* of points (Wright 1932, p. 6). The problem, of course, is that this compression misrepresents the distances between most of the genotypes—accurate representations of distance cannot survive the packing of many dimensions into a few. But Wright's intuition was that this misrepresentation wouldn't matter—that the basic structure of the landscape would be similar enough that the picture could serve a crutch for the math phobic despite the fact that such a picture "cannot accurately represent relations that are multidimensional" and hence the pictures being "useless for mathematical purposes" (Wright 1988, p. 117).

More recently, Skipper (2004) has argued that both Provine and Ruse misunderstood the value of the landscape metaphor; Skipper argues that by focusing on the technical question of what was being visualized in the concrete drawings of it, both authors miss the real problems with Wright's interpretation, in part because they fail to treat the landscape metaphor as a tool for *theory evaluation*, focusing instead on the metaphors role as a heuristic crutch for the math-phobic (pp. 1180–1181). The value of the metaphor as a tool for evaluating the plausibility of particular theories depends not on the broad outlines of what the drawing represents, but rather on the details of the landscape represented; it is here, Skipper suggests, that Wright went wrong, and any analysis of the value of the metaphor will depend on understanding those details (pp. 1184–1186).

Diagrams, metaphors, and models

In Skipper's view, Ruse's defense of the value of adaptive landscapes is inadequate, because Ruse focused on just one heuristic role that adaptive landscapes played, namely that of *illustrating* particular theories; Skipper argues, in contrast, that "the landscape heuristic" plays "a central role in *theory evaluation*" (p. 1184, emphasis in original). But what might be meant by the "landscape heuristic," or for that matter, by "illustrating" as opposed to "evaluating" theories?

In *Inventing Temperature*, Chang distinguishes between *theories*, *ideal models*, *partially-concrete models*, and what we might, for lack of a better word, call the world (2004, pp. 206–209). *Theories*, on this view, are abstract entities—strictly speaking, they are not 'true' or 'false' of anything[3]. In order to *apply* theories to the world, one must generate an *image* or a *model* of the theory; at first, this image or model may be highly *idealized*, but, if it is to be useful, it must be turned into a *partially concrete model* or a *partially concrete image*. It is only these latter models or images that can be compared to the world, and hence used to "test" the particular applications of the theories in question. In the case of quantitative genetics, the "theories" are sets of equations; while these are usually interpreted to be about allele frequencies, mutation and migration rates, and the like, they *need not be*. "Solving" these equations for particular values of the variables does not demand any particular assumptions about how the variables are to be interpreted. A more-or-less idealized image of the theory determines what the theory applies to, that is, what the equations are supposed to mean and what the variables are supposed to "stand for". So for example, we note that in the equations of quantitative genetics, particular variables stand for relative fitness values, heritability estimates, etc. Finally, these images or models can be made more or less concrete by relating them to particular kinds of physical structures in the world, in this case, by claiming for example that the behavior of our model is similar in some key respects to that of real populations.

But what roles are played by the "landscape heuristic"? It is not obvious that this question has an unambiguous answer. There are a number of roles that might be played by particular *diagrams* of adaptive landscapes, a number of roles that might be played by the *metaphor* of adaptive landscapes more generally, and a number of roles that might be played by the models generated from the abstract equations themselves. If by the "landscape heuristic" we mean merely these semi-concrete models based on the mathematical equations, then the usefulness of the heuristic will depend on, for example, the degree to which such models generate results that prove to be useful for testing predictions, generating and improving hypotheses, and the like. Here, though, one might legitimately question whether it makes sense to call such a semi-concrete model a "landscape" heuristic—after all, there isn't anything in the mathematics that demands any sort of topological interpretation (though there is of course nothing that forbids such an interpretation, either). But if the "landscape heuristic" refers to the uses to which the physical *pictures* or *diagrams* of landscapes are put, then the usefulness of the heuristic depends on for example the degree to which such pictures provide meaningful insights into the systems they are supposed to represent. Here, there is an additional ambiguity, in that we might wish for insights into the more or less concrete models/images produced from the theories, or we might wish for insights into the world. In the latter case, we might for example wish to compare aspects of the pictures produced using a particular (semi-concrete) model to some aspect of the world.

In the next section, Sergey Gavrilets' metaphor of the "holey" fitness landscape will be used to explore some of these different possible uses of the "landscape heuristic" and to argue that, if Gavrilets' understanding of the semi-concrete model underlying his work on speciation is correct, *both* the metaphorical understanding of fitness (and adaptive) landscapes *and* the physical diagrams and images produced have outlived their usefulness. The semi-concrete models must stand or fall on their

---

[3] In this, he broadly follows Cartwright's position as expressed in *The Dappled World* (1999).

own merits, and these merits are unrelated to either the pictures produced or even to the metaphorical understanding of the models in terms of "landscapes" that such pictures inspire.

Rejecting the "Landscape" image(s)

As Skipper points out, on one plausible interpretation, the problem with Wright's landscapes was not that they were incoherent, but that Wright was *wrong* about the basic structure of the landscapes (Skipper 2004, pp. 1180–1181). Wright assumed that the topography of a many thousand dimensional landscape would not be significantly different than that of "ordinary" two or three dimensional landscapes; Wright assumed, for example that there would be "peaks" of high fitness separated from each other by "valleys" of low fitness. This assumption, built into both verbal descriptions of fitness landscapes and the pictures drawn of them, framed what came to be known as the "peak-shift" problem—the problem of how a population could move from a lower adaptive peak to a higher one, given that these peaks were separated by a valley; this was, after all, the problem that Wright's "shifting balance theory" was meant to solve. But for very high dimensional systems, our intuitions about the "shape" of the landscape might be very poor, and the actual topology quite different. There might not, for example, be anything like "peaks" separated by "valleys" at all. Indeed, this is the conclusion that Gavrilets claims his computer simulations of very high-dimensional "landscapes" support.

Gavrilets starts from the "Bateson-Dobzhansky-Muller" (BDM) model of speciation. This model starts from the assumption that some gene combinations have relatively high fitness and some have very low fitness. In such a case, some allelic substitutions will be "neutral" or "nearly-neutral" with respect to fitness, but some won't be. Imagine, for example, a two-loci, two-allele, diploid model in which the 'b' allele is incompatible with the 'A' allele such that organisms with a 'b' allele and an 'A' allele will have very low fitness, but other combinations have relatively high fitness (see Fig. 4). In such a system, organisms of genotype aabb will be connected to organisms of genotype AABB by a series of one-step neutral substitutions, all of which are high-fitness. However, organisms of type AABB will be effectively reproductively isolated from those of type aabb, as organisms with any A/b combination will have very low fitness.

Gavrilets extends the idea behind the BDM model by noting that in the sort of one-step representation envisioned by Wright (a fitness landscape, see Fig. 1), each "point" (here representing a genotype) will be one-step removed from a number of points representing alternative genotypes and that this number increases with the number of loci and the number of alternate alleles[4]. In models of organisms with thousands or tens-of-thousands of loci and tens or hundreds of alternate alleles at each, then, the number of genotypes one-step away from a given genotype will be

---

[4] If L is the number of loci, and A is the number of alleles, then for haploid organisms, the number of one-step neighbors for each genotype is equal to $L(A - 1)$; for diploid organisms, the number of one-step neighbors is $2L(A - 1)$.

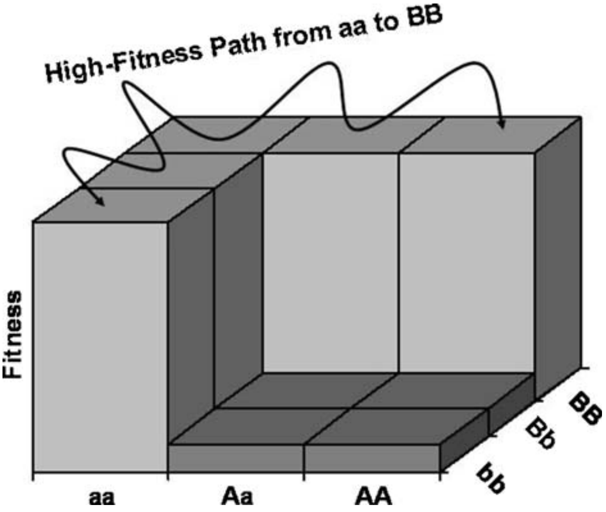

Fig. 4 The "Bateson-Dobzhansky-Muller" (BDM) Model of Speciation. Note that it is possible to move from aaBB to aabb and to AABB without passing through a zone of low fitness, and yet that organisms of aabb are effectively reproductively isolated from organisms of type AABB. Redrawn from Gavrillets (2004)

enormous—on the order of tens or hundreds of thousands. The dimensionality of the system—the Wright-style "genotype space"—will be similarly huge. One of the implications of this is that, under certain plausible assumptions, it is very likely that a high-fitness genotype will be one-step removed from at least one (and likely several) other high-fitness genotypes. Further, if the chance that a genotype chosen at random is high-fitness is above a certain threshold (the so-called "percolation" threshold), it is likely that the vast majority of high-fitness genotypes will be connected to each other through a series of one-step moves all of which are to another high-fitness genotype; hence, most of the high-fitness genotypes will be part of a "nearly-neutral network" of high-fitness genotypes. Since the percolation threshold *decreases* with increases in dimensionality, for the very high dimensional systems considered here, the chances that a random genotype is high-fitness can be very small and still be above the percolation threshold[5]. The upshot of this is that on such a landscape there will be no "fitness peaks" separated by "fitness valleys"—rather, there will be an extensive network (in very high-dimensional space) of high-fitness genotypes connected to each other through one-step high-fitness substitutions.

The "holey landscape" name emerges from Gavrilets' attempt to represent those networks in three dimensions; here, Gavrilets resorts to visualizing the multi-dimensional nearly-neutral network as a (nearly) flat plane peppered with many "holes"—areas of very low fitness (see Fig. 5). In Gavrilets' interpretation,

---

[5] But note that the probability that an arbitrarily chosen genotype will be of high-fitness is in part an empirical question, and in part depends on difficult-to-interpret background assumptions regarding development and the structure of "genotype space"; more on this topic below.

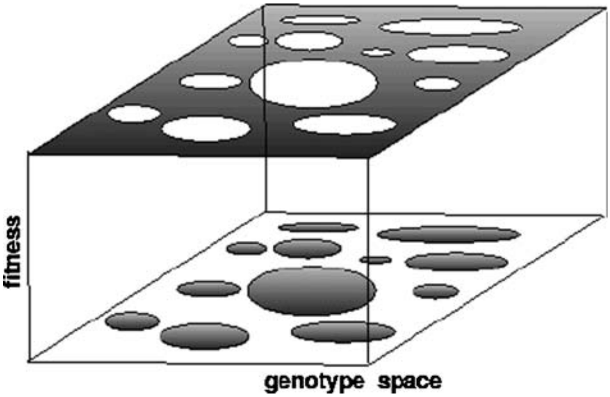

Fig. 5 Gavrilets' "Holey" Landscape Image. More recent images by Gavrilets' depict the surface as "nearly flat" rather than flat, but are otherwise similar

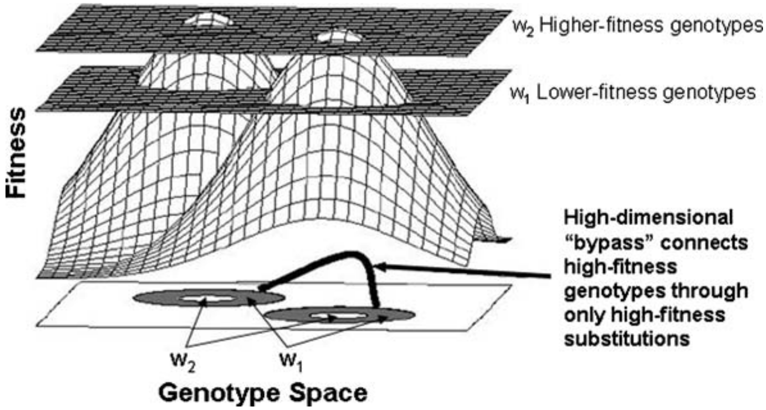

Fig. 6 Gavrilets' attempt to visually represent the relationship between speciation on nearly-neutral networks and "ordinary" adaptation. Adaptation brings populations into the w1–w2 fitness network, speciation can result when populations traverse the nearly-neutral network. Note that a population can "move" between the two "peaks" without traversing a "valley" as those peaks are connected via a "ridge" in high-dimensions (represented by the "extra-dimensional bypass"). Modified from Gavrilets (2004)

speciation takes place when a population splits, and, through a series of changes that are (usually) largely neutral with respect to fitness, end up on opposite sides of a 'hole'. "Ordinary" adaptation occurs when a population is not "on" the high-fitness network, but is on a network of (at least) viable genotypes *connected* to the network of high-fitness genotypes via a series of one-step substitutions some of which increase fitness (and the rest of which are neutral or nearly neutral); on this larger network, the population tends towards moves that increase fitness, until the population is once again part of the high-fitness network (see Fig. 6). The two processes—adaptation and speciation—*can* but *need not* take place simultaneously, and indeed, Gavrilets suggests that there are reasons to think that a decoupling of the two kinds of processes will be the more common scenario (e.g. 2004, pp. 103–104).

If this interpretation of fitness "landscapes" is correct, then the "problem" of peak-shifting simply vanishes; there are, on this view, neither local "peaks" to climb nor "valleys" to cross. Speciation is not a matter of populations arriving at different fitness peaks, but rather of populations becoming genetically incompatible by ending up in (different) zones of broad compatibility separated by zones of incompatibility (the "holes" in the holey landscapes).

It should be immediately obvious that the "holey" landscape pictured in Fig. 5 is inadequate. For example, the "surfaces" are not surfaces at all, but rather areas of *relatively* dense networks of high-fitness genotypes. Just as in Wright's original pictures, there is no continuous surface that corresponds to an area in genotype-space. Rather, what is being pictured as a surface here are "hyper-volumes" in genotype space in which there are many high-fitness genotypes and/or relatively many paths between high-fitness genotypes. The "holes" are "hyper-volumes" in which there are relatively few high-fitness genotypes and/or relatively few paths between those high-fitness genotypes.

Further, the collapse of dimensionality of these networks hides the irreversibility of movements in these areas. Because each high-fitness genotype is one-move away from a very large number of genotypes, only some of which are high-fitness, the existence of a path comprised of only high-fitness genotypes between any two high-fitness genotypes implies nothing about the *likelihood* of such paths being taken; indeed, any *particular* path is very unlikely. 'Retracing one's steps' is therefore nearly impossible in such a system. But the image of a holey landscape hides this fact – the "surfaces" appear to permit free and equiprobable movement throughout the areas of high-fitness.

Neither of these problems is fatal for the image—no (semi-concrete) image of a model can be expected to capture every aspect of it.[6] But the picture also misrepresents elements in ways that encourage confused thinking about the issues—as a means of heuristically testing theories (theory evaluation), such pictures fail by falsely suggesting particular tests and comparisons.

Gavrilets', for example, considers the existence of "ring species" to support the holey landscape model over other topographic possibilities (see Gavrilets 1997, Box 3). But insofar as this is true, the prediction that emerges from Gavrilets' model is not exactly what the image suggest; modeling the likely changes in populations in terms of areas of relatively dense complex high-fitness networks and areas where the networks are less-densely packed, one does indeed find that "ring species" are possible, but this has less to do with the "holes" in the image of the "holey" landscape and more to do with the details of the structure of the networks as represented in the computer models. In some sense, this shouldn't be a surprise—as noted above, the "holes" are entirely metaphorical, and merely represent less-dense networks of high-fitness genotypes within the hyper-volume of genotype space. Similarly, the uniform "density" of the "surface" in the holey image is misleading;

---

[6] Nor should an image attempt to capture every aspect of the model; this is related to the uselessness of a one-to-one full-scale "map".

one of the results of the computer simulations is that populations will tend to move from areas of less-dense nearly neutral high-fitness networks to areas of increased density. The image of populations ‘wandering’ around ‘holes’ misses this.

Other attempts at producing a low-dimensional representation might be better at representing some of these aspects of the model but at the cost of mis-representing other aspects of the model. Should we picture the high-fitness network as relatively dense (to encourage thinking about the internal structure of populations) or as relatively sparse (to encourage thinking about the non-reversibility of the moves that populations make)? For complex networks in very high dimensional space, there is no image that is fully adequate; further, every image that is adequate for representing some aspect of the network will not only fail to model other aspects, but will be misleading in important ways. One might, then, choose to use a variety of different images, and to specify what each kind of image was good at highlighting, and where each image failed. But then, the usefulness of the images as theory evaluation tools, or even as heuristic aids to understanding particular theories, would be very much in question.

Indeed, the value of the “holey landscape” image is cast into some doubt by its very *insensitivity* to the values of the main parameters in the models. Gavrilets recognizes this, and argues that in some sense, this is an advantage of the image—the details of the models are unlikely to change the basic idea behind the holey landscape images (see 1997, p. 309 and 2004, e.g. pp. 100–105). For example, for many different values of the percentage of genotypes that will be of high-fitness, a nearly-neutral network of high-fitness genotypes will result if the dimensionality is high enough. As noted above, however, the number of high-fitness genotypes that any particular high-fitness genotype will be connected to via one-step substitutions remains, essentially, a guess, as does the dimensionality of the system. Note for example that estimates of the dimensionality rely on particular estimates of the number of loci and the number of alleles. But not only are these numbers not well established for most kinds of organisms, it isn’t even clear that the problem of individuating loci and alleles is particularly well-formed! Determining how many loci there are would seem, for example, to presuppose the ability to individuate “genes” (at least conceptually)—and the ability to do that in general is still very much in doubt (see e.g. Stotz and Griffiths 2004; Pearson 2006). None of this is to deny the importance of Gavrilets’ results; again, for many of his results, he suggests that detailed answers to these questions are unnecessary.

But while even a rough estimate of the number of loci and a rough estimate of the percentage of viable genotypes may not be important to the existence of a connected network of high-fitness genotypes, it *is* important to other features of the models. For example, given particular assumptions, high-dimensionality BDM models (models in which genotypes are conceived of in terms of their genetic distance from each other and are assigned either high or low fitness values, see above) are able to make particular predictions about such phenomena as how geography influences speciation patterns (including the average number of species likely to form in a particular environment), the relationship between rates of extinction and the size of the clade, the relationship between diversity (roughly, the number of reproductively isolated populations) and disparity (roughly, how different those populations are) within clades (Gavrilets 2004, Chapter 7), the role of adaptive evolution in influencing reproductive isolation (2004, pp. 187–188), and the relationship between speciation and non-random mating (e.g. 2004, Chapter 9). However, these kinds of predictions are generated from the computational models given assumptions about the *particular* values of the key parameters; different values result in different predictions. While the extent to which these predictions have been, or even can be, tested varies, surely the ability to generate predictions for a variety of different possible values of the parameters is one of the key strengths of the models developed. But if that is so, then the *metaphor* of the holey landscape, as well as the pictures of it produced, do the models a serious injustice, as they are insensitive to variations in the very assumptions that make the model testable!

But these are not the only problems, or perhaps even the main problems, with fitness landscapes as normally represented in the literature. For example, the models of such landscapes assume a genotype-fitness mapping that is unlikely to be maintained through evolutionary—or even developmental!—time scales. It is unclear that assumptions about, for example, how many independent loci there are, or about how often particular loci can be expected to interact (and how), even make sense without strong assumptions about ontogeny being built into the models. But these models are forced to ignore the complexities of ontogeny, and hence it is unclear just what the assumptions they make about development actually are, and, *a fortiori*, unclear if the assumptions are reasonable.

Similarly, even contemporary computer models tend to assume a “fixed” landscape when it is obvious that the topography will change in response to the actions of the population itself as well as to “external” factors (see e.g. Lewontin 1978). Put another way, these models tend to ignore *ecology*. But surely the ecological context will have important implications for both adaptation and speciation. While there is some interest in finding ways to integrate evolutionary biology and ecology (see for example Johnson and Stinchcombe 2007), evo-eco is still in a very early phase and it is unclear how (or even if) any insights that emerge from it can be incorporated into the high-dimensional landscape models.

It is possible that, with increases in computational power and more sophisticated understandings of the complex interactive roles of ecology and development in evolution, such problems could be rigorously addressed. And even without explicitly addressing such problems, the computational models generated by researchers such as Gavrilets can suggest hypotheses to be tested and be used to evaluate particular models. But the visual images generated do none of these things; they are, by necessity, too over-simplified and too misleading.

If this is correct, then it could well be time to give up generating visual images of the models representing fitness “landscapes” (or whatever term we prefer when referring to networks in very high dimensional space). In any event, if such images *are* generated, the areas in which they are misleading must be very carefully evaluated and prominently featured. Indeed, our inability to generate visual images that accurately portray more than a very few significant aspects of the models, when combined with the other fundamental problems in the semi-concrete images of the models themselves, make even the *metaphor* of the fitness landscape problematic. Finding a better way to talk about the models and the results generated from those models – a way that does not encourage the use of misleading images—may prove to be difficult, but that should not stop us from trying.

Acknowledgments This paper is based in part on a book chapter co-written with Massimo Pigliucci (SUNY Stonybrook). Sergey Gavrilets was generous with his time in providing assistance in understanding his models and approaches. Massimo provided helpful feedback on earlier versions of this material, as did Sharyn Clough, Anya Plutynski, members of the UBC Philosophy Department, Department of Botany, and Department of Zoology, members of the University of Houston's Department of Biology and Biochemistry, and participants in the 2006 Pacific APA mini-conference "Scientific Images."

References

Cartwright N (1999) The dappled world: a study of the boundaries of science. Cambridge University Press, Cambridge

Chang H (2004) Inventing temperature: measurement and scientific progress. Oxford University Press, New York

Gavrilets S (1997) Evolution and speciation on holey adaptive landscapes. Trends Ecol Evol 12(8):307–213

Gavrilets S (2004) Fitness landscapes and the origin of species. Princeton University Press, Princeton

Johnson M, Stinchcombe J (2007) An emerging synthesis between community ecology and evolutionary biology. Trends Ecol Evol 22(5):250–257

Lewontin RC (1978) Adaptation. Sci Am 239:213–230

Pearson H (2006) Genetics: what is a gene? Nature 441:398–401

Provine WB (1986) Sewall Wright and evolutionary biology. University of Chicago Press, Chicago

Ruse M (1996) Are pictures really necessary? The case of Sewall Wright's 'adaptive landscapes'. In: Baigrie BS (ed) Picturing knowledge: historical and philosophical programs concerning the use of art in science. University of Toronto Press, Toronto, Canada, pp 303–337

Skipper RA (2004) The heuristic role of Sewall Wright's 1932 adaptive landscape diagram. Philos Sci 71:1176–1188

Stotz K, Griffiths P (2004) Genes: philosophical analyses put to the test. Hist Philos Life Sci 26:5–28

Wright S (1988) Surfaces of selective value revisited. Am Nat 131:115–123

Wright S (1931) Evolution in mendelian populations. Genetics 16:97–159

Wright S (1932) The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In: The sixth international congress of genetics, vol I, pp 356–366