Biol Philos (2008) 23:639–657
DOI 10.1007/s10539-008-9127-9

Assessing the fitness landscape revolution

Brett Calcott

Received: 2 May 2008 / Accepted: 16 July 2008 / Published online: 22 October 2008

Abstract According to Pigliucci and Kaplan, there is a revolution underway in how we understand fitness landscapes. Recent models suggest that a perennial problem in these landscapes—how to get from one peak across a fitness valley to another peak—is, in fact, non-existent. In this paper I assess the structure and the extent of Pigliucci and Kaplan's proposed revolution and argue for two points. First, I provide an alternative interpretation of what underwrites this revolution, motivated by some recent work on *model-based* science. Second, I show that the implications of this revolution need to carefully assessed depending on question being asked, for peak-shifting is not central to all evolutionary questions that fitness landscapes have been used to explore.

Keywords Morphospace · Evo-devo · Speciation · Fitness landscapes · Models · Neutral change

1 Introduction

According to Pigliucci and Kaplan, there is a revolution underway in how we understand fitness landscapes (Pigliucci and Kaplan 2006). The story of this revolution is largely based on recent work by Sergey Gavrilets (particularly Gavrilets 2004) and runs something like this:

> Wright (1932) presented us with a simple visual metaphor, where a fitness landscape consisted of peaks of high fitness separated by valleys of low fitness.

B. Calcott (✉)
Philosophy Program, RSSS, Australian National University, Canberra ACT 0200, Australia
e-mail: brett.calcott@gmail.com

B. Calcott
Centre for Macroevolution and Macroecology, Australian National University, Canberra ACT 0200, Australia

> Given such a picture, a central issue has been how a population could shift from one adaptive peak, through a fitness valley, to another adaptive peak. But the visual metaphor has misled us. Mathematical models with realistic assumptions about the dimensions of such landscapes suggest that the problem of peak-shifting is non-existent, as "ridges" connect most, if not all, genotypes with near-equal fitness, obviating the need for any shifts to go through a fitness valley.

If understanding "peak-shifting" has been central to fitness landscapes, and there are, in fact, no such peaks, then a revolution may well be justified.

This paper critically examines some aspects of this proposed revolution, and argues for two main points. First, I suggest a different interpretation of how Gavrilets' work underwrites such a revolution—a different story than the one given above. I then argue that the general impact of Gavrilets' work on fitness landscapes needs to be carefully assessed, for his new insights arise from work on speciation, yet fitness landscapes have been used to understand other questions in evolution. I look at two of these: the evolution of morphological diversity and the evolution of development. I show how their use of fitness landscapes differs, and examine how the dissolution of the peak-shifting problem affects these models. So the arguments in this paper both modify the story of the revolution within work on speciation, and assess the implications of it for some alternative uses of fitness landscapes.

In making these arguments, I shall adopt a *model-based view of fitness landscapes*. Under this view fitness landscapes are (a) a family of related models that share some core defining features, and that (b) these models are used to address a wide variety of related, but different, questions about evolutionary change. This contrasts with Pigliucci and Kaplan's discussion which (a) treats fitness landscapes primarily as metaphors, and (b) surveys a limited use of fitness landscapes. I outline this view and, in addition to the two main points above, show that it is particularly well-suited to fitness landscapes, providing a richer insight into how they do their explanatory work.

It is important to mention some issues I shall *not* be addressing. Pigliucci and Kaplan distinguish between two different types of landscapes, a *fitness* landscape and an *adaptive* landscape. I shall have nothing to say about adaptive landscapes; all of the models in this paper are *fitness* landscapes. They also point out a number of problematic simplifying assumptions that many fitness landscape models make (such as a fixed environment). Although I shall discuss the kinds of assumptions made in these models, for the most part I shall attend to the differences between assumptions made in various models, rather than these more general problems.

I begin the paper with some remarks about metaphors and then describe the "model-based view" of fitness landscapes. I then outline the core features of fitness landscape models. The following three sections each examine a different use of fitness landscape models. The first of these, on speciation, shows how the important insights of Gavrilets' work are underwritten by models, and how my view contrasts with the story given by Pigliucci and Kaplan. The next two sections each demonstrate the use of fitness landscapes for reasons other than exploring speciation. I first look at Niklas' landscape models of morphological change, and argue that they give us reason to limit the generality of some of Gavrilets' conclusions. I then examine

Fontana's RNA fitness landscape models, and show that neutral change in this context plays a quite different role than it does in the speciation models. I conclude with some brief remarks about how these different models inform us about the relationship between macroevolution and microevolution.

2 Metaphors and models

The term "fitness landscape" is used in various ways in the literature. Sometimes it simply refers to the visual representations introduced by Wright (1932). More often it refers to a set of metaphorical ideas associated with these representations: being stuck on peaks; crossing valleys; the roughness and smoothness of a landscape. Most importantly for this paper, fitness landscapes refer to models, described by mathematical formalisms or numerical simulations.[1]

Gavrilets explicitly distinguishes between fitness landscapes *as metaphors* and fitness landscapes *as formal models* (Gavrilets 2004, p. 34). According to Gavrilets, it is the models that form the basis of any claims in his book about how speciation might occur, and throughout the book Gavrilets intends the term "landscapes" to refer to models (Gavrilets 2004, p. 105). Pigliucci and Kaplan (2006) never explicitly mark out the distinction: landscapes sometimes refer to metaphors, and other times to models. Their story of the current revolution in fitness landscapes, however, *relies* on the distinction, as it is the properties of particular multi-dimensional landscape *models* in Gavrilets (2004) that suggest that the low-dimensional visual *metaphors* have misled us.

Despite the distinction between metaphors and models being central to Pigliucci and Kaplan's story, the exact relation between them remains largely unanalysed.[2] Gavrilets does suggest some differences: models are more precise with stricter requirements, and metaphors are "simplified", emphasizing only certain features, whilst neglecting others (Gavrilets 2004, p. 32). But it not clear that Gavrilets' suggestions support Pigliucci and Kaplan's story. Precise models of low-dimensional landscapes can be constructed, and such models could still *mislead* us. So even if lack of precision qualifies something as metaphorical, it is not a *necessary* condition for providing the misleading character to which Pigliucci and Kaplan allude. As for simplicity, a key feature of models themselves is that they emphasize some features of the world, and neglect others. If both models and metaphors are simplified representations of the world, then it is difficult to see what distinguishes the two. It can't be just that one is misleading, while the other is not. Perhaps a principled distinction can be made, but one is not on offer that does the work required by Pigliucci and Kaplan.

[1] Somewhat less common is to refer to fitness landscapes as things that exist "out there" in the world: "Many authors have assumed that the adaptive landscape is merely a metaphor ... These authors have ignored the possibility that the adaptive landscape—like mutation, drift, and inheritance—is an object for empirical study" (Arnold 2003, pp. 367–368).

[2] It is not clear how these metaphors might relate to any theory of literary metaphor, an area which itself has no univocal view on what metaphor is (Camp and Reimer 2006).

Instead of attempting to work with this distinction, I shall adopt a model-based view of fitness landscapes. Under this view, fitness landscapes are a family of related models that share some core defining features (which I shall outline in the next section). Although sharing a number of common features, particular landscape models may realise or emphasise these features in a variety of ways, making different assumptions in each case. We can make sense of Gavrilets' insights without requiring a distinction between metaphors and models. Instead, what matters is whether or not the particular assumptions made in the model are appropriate for the question being asked. We may be *misled* by the models (rather than metaphors) when we have made inappropriate assumptions relevant to the question we are asking.

To be clear, I'm not claiming that this the only way we might be misled by a model. Gavrilets' suggestion about precision might relevant in some cases—a model may not be sufficiently precise to represent the target of interest. Or we may simply have made a mistake how we chose to represent the model (as simple as a bug in our simulation or a mistake in our algebra). But I shall argue that it is a change in model assumptions that underlies the dissolution of the peak-shifting problem.

What do I mean by "adopting a model-based view"? I am relying on recent work by Michael Weisberg and Peter Godfrey-Smith who, in turn, have been influenced by Giere, Wimsatt, and Levins (Levins 1966; Wimsatt 1987; Giere 1988; Godfrey-Smith 2006; Weisberg 2007). This work describes a particular method of theorising that Godfrey-Smith refers to as the *strategy of model-based science*:

> The modeler's strategy is to gain understanding of a complex real-world system via an understanding of simpler, hypothetical system that resembles it in relevant respects. (Godfrey-Smith 2006)

I will not outline the view in its entirety here (for a good overview see Godfrey-Smith (2006)). What will matter here are two points central to model-based science.

The first is *indirect representation*. Model-based science represents the world in an indirect fashion, and consists of two distinct steps. First, one *specifies a model*—this delineates the properties that the model has. Such a specification or model description may be done in a number of representational media—words, pictures, a computer program, or a mathematical formalism. The second step is to then make some *similarity claim* between the model thus described and the thing in the world of interest—the model's target system. Together, an understanding of the model as specified, and the fact that the model resembles the world in certain ways, is what gives the model power (albeit indirect) to analyze and understand the target system.

Fitness landscape models fit this strategy particularly well. There is no doubt that the models are simplified versions of the target system we want to understand. In addition, as we shall see below, these models are used in a variety of ways both in- and outside of evolutionary biology, suggesting that this second step of resemblance can occur in a very flexible manner.

This highlights the second important feature of model-based science. Because the specification of a model is a separate step from the similarity claims made about it,

assessing the worth of a particular model depends not just on facts about the model, but also how it is being deployed:

> One cannot eliminate the purposes of scientists from the evaluation of any model. (Giere 2002, p. 1060)

This is relevant for fitness landscape models, for as we shall see, they are used to answer a number of different questions. Evaluating "fitness landscapes" requires understanding what questions are being asked in each case, and how the models are being used address these questions.

3 The core features of fitness landscape models

Stadler provides a very concise, mathematical, description of a fitness landscape:

> A fitness landscape is a mapping from a configuration space that is equipped with some notion of adjacency, nearness, distance or accessibility, into the real numbers. (Stadler 2002)

Assuming a biological model (such as those in Gavrilets), we can unpack this description in the following way. The *configuration space* is the set of possible allele combinations; the *mapping* is the assignment of a fitness to each of these; *accessibility or distance* is captured by the number of allele differences between any two genomes, some being *closer* to others. I'll provide more detail below. In passing however, it is worth mentioning why Stadler's definition is couched in such abstract terminology:

> Fitness landscapes are a valuable concept in evolutionary biology, combinatorial optimization, and the physics of disordered systems. (Stadler 2002)

When the abstract features of fitness landscape models are isolated, they provide a structure whose application extends beyond that within evolutionary biology (a fact noted earlier by Kauffman and Levin (1987)). Theoretical advances in understanding landscapes models *in general*, can have implications for the various different ways that they are deployed. Some of the conclusions drawn from models in Gavrilets' book rely on mathematics originally developed to address the percolation of fluids in disordered physical systems (again, see Kauffman and Levin (1987) for some comparisons). This fits well with the view of model-based science endorsed above.

I now describe in more detail the core features of a fitness landscape model as interpreted to explore evolutionary change. The description will be informal and simplified.[3] But it will contain enough detail to demonstrate the variety of choices

[3] In particular, I shall limit the informal description to *topological* configuration spaces—landscape structures whose neighbourhoods are defined by a set of independent dimensions (a euclidean vector space). This excludes *pre-topological* spaces, where neighbourhoods may be defined arbitrarily as a directed graph. Most (but not all) current landscape models in evolutionary biology are topological spaces (all of Gavrilets models, for example). For further details see (Stadler 2002; Stadler et al. 2002; Stadler and Stadler 2006).

available to modellers when they construct a fitness landscape model. What choices are made depend on the question the model is used to address. These different choices will be apparent when we examine some models and the different questions they address in the sections that follow.

It is useful to divide the features of a fitness landscape model into two parts: static and dynamic. The static features are those that produce the *structure* of the landscape. The dynamic features are those that control the *movement* across the landscape, given a particular static structure. Stadler's concise description detailed only the static properties of a landscape. He does extend his discussion to dynamic features, suggesting that:

> The challenge for a theory of landscapes is therefore to link these two points of views, for instance by determining how geometric properties influence the dynamical behavior. (Stadler 2002)

I first discuss features that make up the static part of fitness landscapes model, and then briefly discuss the dynamics.

Two components structure the static landscape. The first are *dimensions of variation*. Each of these dimensions describes some property, and the range of values that this property might take. In Gavrilets's work, each dimension of the landscape represents a locus, and each discrete position along that dimension represents an allele that can exist at that locus. But the dimensions need not represent genes; the dimensions of variation may be *phenotypic* characters. So we might use the different axes that construct the dimensions of the landscape to represent such things as limb length, visual acuity, or even the degree that some individual is disposed to behave in a certain way. Together, these dimensions of variation mark out a space of possible individuals (or individual types). An individual is defined by the combination of values drawn from each of these dimensions of variation.[4]

The second component that structures the landscape is the *assignment of fitness*. Each distinct combination of values defining a possible individual is given a fitness value. This can be done in various ways: randomly, manually, or the value might be some function of the values taken from each dimension. In many of Gavrilets's simpler models the fitness is manually assigned to a combination of alleles. In contrast, fitness in Kauffman's NK landscapes fitness are randomly assigned, taking into account epistasis amongst the loci (Kauffman 1993). In Niklas's phenotypic landscapes (which I describe later in this paper) the fitness is calculated from the resulting morphology produced by the various morphological dimensions (Niklas 1994). The fitness values assigned to each combination of values along the dimensions of variation mark out the (multi-dimensional) surface of the fitness landscape.

These two structural features shape the surface of the fitness landscape, which provides the basis for constraining movement. The dimensions of variation

[4] There is already an important difference between genotypic and phenotypic models in how change might occur in a particular dimension. Typically, any allele might change into any other allele, but changes along a phenotypic dimension are often themselves constrained in some way to small increments in (for example) size, shape, or angle.

constrain movement across the landscape as movement typically proceeds by small changes in one dimension at a time. Additionally, in the case of phenotypic landscapes, the values that lie close together along any particular dimension are assumed to be more accessible to one another. The assignment of fitness at each of these points constrains movement as a step is more likely in the upward direction than the downward direction. Roughly speaking, given some initial position, lower fitness points that are far away are the hardest to get to, and higher fitness points that are nearby are easier to get to.

These two features in a fitness landscape model deployed in evolutionary biology map to assumptions about the supply and selection of variation. Constraining movement to small increments in the dimensions of variation captures the assumption that evolution predominately proceeds by small incremental change, not large leaps.[5] Constraining movement largely to upward rather than downward movement captures the assumption that natural selection will act to maintain fitter individuals.

The static structure of the landscape dictates, in large part, how movement across the landscape might take place. But there are varying ways that the dynamics of movement can be modelled. This provides yet another place for choice in the construction of fitness landscape model. Movement across a surface may be guided by something as simple as random walk biased by the fitness slope, or as complex as tracking a "cloud" of sexually reproducing individuals diffusing across the surface. Given the same static landscape structure, the choice of dynamics may produce different results. For example, assuming a sexual versus asexual population dynamics produces different conclusions about the likelihood of speciation on similarly structured landscapes (Gavrilets 2004, p. 76).

Although fitness landscape models share a number of core features, exactly how these features are realised or implemented may differ. What dimensions are chosen and what they represent can vary; how fitness is assigned might be manual, random, or some function of type; and the dynamics on the landscapes can be modelled in different ways. As we shall see in the following sections, the different features used to construct a fitness landscape models take on more or less importance depending on the aims that the modeller has. The varying ways that these features are realised will also reflect different assumptions that the modeller is prepared to make.

Lastly, notice that it is the structural features of these models that makes them amenable to visual representation. The two constraints on movement are naturally seen as (a) limiting the size of each step (b) limiting each step to be largely upward. Such a visual representation is limited to low dimensions. But it is important to realise that terms such as "walk", "step", "roughness", and "distance" can be given precise mathematical interpretations that continue to make sense in high dimensional models.

[5] In general, fitness landscape models assume something like point mutations. This is, of course, a huge simplification. Some models do consider the affects of other types of mutation, such as unequal crossover (Shpak and Wagner 2000). The upshot is complicated, to say the least.

4 Speciation and fitness landscapes

> ... the problem of how a population crosses an adaptive valley in its way from one adaptive peak to another ... may be non-existent. (Gavrilets 1997, p. 307)

Gavrilets's dissolution of the peak-shifting problem forms the basis of Pigliucci and Kaplan's revolution. In this section I look more closely at how Gavrilets's work supports this conclusion. I argue that the story is more complex than that told by Pigliucci and Kaplan. In particular, there is a more indirect role for high dimensionality than they assume. The key factor introduced in Gavrilets' models of speciation is actually *neutral change*, rather than high-dimensions. These two things are linked, but not the same. Indeed, it is essential for many of Gavrilets' speciation models that one can construct a model that includes neutral change but *does not* have very high (biologically realistic) dimensions.

Gavrilets' deals with the "peak-shifting" problem for speciation in two steps, deploying a different model for each step. Gavrilets first establishes the likelihood of neutral ridges in fitness landscapes by examining the *static* structure of a high dimensional landscape. He then constructs relatively low-dimensional models that assume a neutral structure to study the *dynamics* of speciation on these landscapes.

Gavrilets begins his analysis of the "peak-shifting" problem in Chap. 4 of his book (Gavrilets 2004). He presents a number of models demonstrating that, as dimensionality increases, it becomes more and more likely that genotypes of nearly-equal fitness become connected via a neutral ridge—a series of intermediate steps which neither descends into a fitness valley nor ascends to an isolated peak. These neutrally-connected genotypes form larger and larger *neutral networks*, which grow in size as dimensionality increases. Eventually, a single neutral network of high fitness genotypes "percolates" throughout the entire landscape.

The models presented in this chapter are, for the most part, about the understanding the *static structure* of such landscapes. There is some brief discussion of dynamics on these surfaces, but none is about speciation. Gavrilets also provides additional reasons to think that neutral networks might be a prevalent phenomenon, citing work on RNA- and Protein-folding, and some suggestive empirical evidence from such phenomena as ring species (Gavrilets 2004, pp. 97–100). The central message of this chapter is to argue that we have good reasons to expect that fitness landscapes contain nearly neutral networks.

Gavrilets also introduces a new concept that is important for his models of speciation: the *holey* landscape ("holey" as in swiss cheese, *not* "holy" as in the shroud of Turin). Although there may be *a* neutral path between two genotypes, it is not the case that *every* path between these genotypes will be neutral. Straying off a particular neutral path may well lead to falling into a fitness "hole". This is important for Gavrilets' speciation models. Although drift might incrementally shift one population along a neutral path, compatibility between genotypes separated by such a neutral path is not ensured. A combination of alleles from the two populations might not lie on the neutral path, instead falling into a "hole"—a genetic incompatibility that leads to the reproductive isolation of populations occupying different regions of the neutral network. Gavrilets suggests that these

holey landscapes form the right kind of structure in which to explore the problem of speciation.

Gavrilets's next chapter contains models showing how populations might move and speciate on such landscapes. Given the importance of high dimensions in the previous chapter, it is surprising that the model Gavrilets uses to introduce speciation on holey landscapes is one with very few dimensions. The model studied in this chapter is a two allele, two locus model—dimensions low enough that it *can* be visualised in a simple 3D rendering. Despite this low-dimensionality, the model manages to represent a holey landscape possessing a ridge of nearly neutral genotypes. This is done by manually assigning fitnesses to different allele combinations such that this low-dimensional landscape effectively captures the holey structure expected to emerge in biologically realistic dimensions (as shown in the previous chapter). Gavrilets draws a number of important conclusions about speciation from these simple models (Gavrilets 2004, pp. 144–145). In particular, these models show that speciation is likely, in contrast with rugged landscape models of similar low dimensions in earlier chapters, where speciation was highly unlikely. The difference between these models is the structure of the landscape—one includes neutral change, the other rugged peaks and valleys. An assumption of neutral change, rather than increasing the number of dimensions, is sufficient to make a difference in these models of speciation.

Gavrilets does go on to study models with higher dimensions. But the exploration of these higher dimensional models "... lead to conclusions similar to those reached on the basis of the two-locus model" (Gavrilets 2004, p. 194). In addition, though many of the models in the book contain *higher* dimensions, they are still not *biologically realistic* dimensions (see the individual simulations in Chap. 7, for example). They do, however, assume that networks of neutral change exist, and this is sufficient to produce significant new results about our understanding of speciation.

I am not suggesting taking into account realistic biological dimensions is unimportant. Rather, it is that the models with high dimensions play an indirect role in supporting the conclusions from Gavrilets's (largely low-dimensional) speciation models.

The key results about speciation are best interpreted as involving two steps, each deploying a different model. Gavrilets' simpler dynamic models show how speciation occurs *given certain assumptions about the structure of the fitness landscapes*. The high-dimensional static models give us justification for the assumptions in these dynamic models. More precisely, the static models tell us that, when simplifying our models (by decreasing the dimensionality) the essential property to maintain is *networks of nearly-neutral change*. Notice that simplifying the dynamic models is essential, as much of the mathematical analysis and numerical simulations would simply not be possible if the models really did include realistic biological dimensions.

The two-step analysis I provide above is important, for it tells us something about the *practice* of model-building. Levins (1966) first emphasised that building models of complex phenomena often requires making certain trade-offs between realism, precision, and generality (see Matthewson 2008 for an in-depth evaluation of the

relationship between these different desiderata). In order to overcome this, Levins suggested one option is to build a number of inter-related models, each supporting the other (Weisberg 2006).[6] Reading Pigliucci and Kaplan, however, one gets the impression that, given fresh insights about landscapes with biologically realistic dimensions, the right move as a modeller would be to build high-dimensional models.[7] But the actual strategy that Gavrilets pursues appears to be one much closer to that suggested by Levins. A number of interrelated models are constructed, and the structure of a model with realistic biological dimensions supports the assumptions made in simpler dynamic model of speciation.

My analysis of Gavrilets' work in this section suggests a different interpretation of the fitness landscape revolution, one without reference to visual metaphors at all. The story of the revised revolution, in contrast to that given in the introduction, runs like this:

> Our simplified mathematical models of fitness landscapes have often assumed that different gene combinations will have different fitnesses. Even when we move beyond dimensions that can easily be visualised[8] such assumptions produce rugged landscapes, with peaks and valleys. But simple static models with much higher dimensions suggest that a better assumption is that many gene combinations have the same fitness. Making such an assumption in our simpler models changes their dynamics dramatically.

This is a subtly different picture than that presented by Pigliucci and Kaplan. It is not that our models lacked enough dimensions, rather it is that they lacked the right assumptions. The incorporation of neutral change into these models removes the problem of peak-shifting *even in models which have low dimensions*. Under this interpretation, we avoid having to articulate a principled distinction between metaphor and model.

5 Morphospace and fitness landscapes

> ... I suspect that roughening [of the fitness landscape] played a major role in driving both diversity and disparity in the Cambrian "explosion" of animals. (Marshall 2006, p. 376)

[6] Weisberg has pursued this idea, which Levins called *robustness analysis* (Levins 1966; Weisberg and Reisman 2008). However, it is not clear that what is going on here is exactly the same. Robustness analysis looks at how models with different idealizations produce the same results, with the intent of showing that the results are not simply artifacts of particular idealizations. In this case, however, one model provides reasons for believing the assumptions of another. This suggests that robustness is not the only relation of support that might obtain between multiple models.

[7] Pigliucci and Kaplan' discussion may lead one to believe there is a single fitness landscape model that Gavrilets analyses (Pigliucci and Kaplan 2006, pp. 193–194). But they mention both a model from "percolation theory" and the "Bateson–Dobzhansky–Muller model". These are distinct models.

[8] For example, many of Kauffman's rugged landscapes have up to 24-dimensions (Kauffman 1993). Well beyond anything I can visualise!

In this section, I look at fitness landscape models by Karl Niklas and a further suggestion for their application by Charles Marshall (Niklas 1994, 2004; Marshall 2006). Both Niklas and Marshall are interested in how morphological diversity evolves, so the dimensions of these models are phenotypic, rather than genotypic. The results from these models suggest that diversity might be driven by the increased "roughening" of landscapes over time. This roughening is bought about by increasingly complex environmental demands on individuals. I briefly present the model, then examine the different assumptions required in contrast to those models addressing speciation. I then suggest that Niklas's models provide a reason to limit the generality of Gavrilets' static models of percolation in neutral networks.

Niklas constructs phenotypic fitness landscapes to explore the evolution of morphological disparity in plants (Niklas 1994; Niklas 2004). The target of explanation is the rapid diversification of plant form in the Devonian (417–354 MYA). Niklas outlines a multidimensional morphospace, where the dimensions of variation are such things as the branching angle, the rotation angle, and the probability of an axis branching. These parameters are used to "grow" different plant forms. Niklas makes these dimensions discrete by dividing them up into small evenly spaced steps—so the branching angle varies by 1° increments. Despite the limited number of parameters, they are sufficient to produce a wide variety of forms, representative of early plant forms (Niklas 1994, p. 6774).

These forms are then evaluated for how well they perform a number of tasks: light interception, maintaining mechanical stability, dispersing spores, and conserving water. The evaluation is done using some simple physics and engineering principles. For example: "the ability of each variant to conserve water is gauged to be a simple linear function of total plant surface area" (Niklas 2004, p. 53). Fitness is assigned to each plant form according to how well it performs *all* of the assigned tasks. These dimensions of phenotypic variation and the resulting fitnesses are used construct the fitness landscape.

Niklas then uses a simple adaptive walk to navigate the resulting landscapes. The walk begins in a part of morphospace representative of early plant life, and then steps to a neighbouring point if it is of either equal or higher fitness. The walk continues until a peak is reached. In some of the simulations, the fitness function is changed as the walk progresses. This is done by first assessing fitness with respect to only one task and then increasingly adding new tasks over time. Niklas also experiments with developmental constraints, by limiting the independence of the different dimensions of variation (Niklas 2004, p. 60).

The results reveal that the structure of the landscapes is very sensitive to the number of tasks that individuals are required to perform. As the number of tasks is increased, the landscape roughens producing more peaks, and adaptive walks on these landscapes produce a wider variety of forms. This occurs because increase in number of tasks requires *trade-offs*—a plant needs to spread its branches to increase light interception, but horizontal extension of the branches reduces mechanical stability. A single plant form cannot simultaneously optimise the different functions to perform. Different trade-offs amongst the optima produce the multiple peaks on the landscape. As the number of tasks goes up, the number of peaks increases.

At the same time the relative fitness of the optima decreases as it becomes more difficult to simultaneously achieve each task.

Niklas's simulations also demonstrate the historical contingency of particular solutions evolving. When the differing selection regimes are changed (varying the number of tasks), the peaks subsequently reached can depend on the ordering of tasks that are introduced, even though the final set of tasks is the same.

Marshall, inspired by Niklas' simulations of plant evolution, has suggested something similar to explain the diversification of animal forms in the Cambrian (Marshall 2006). He argues that the increasing ecological complexity, largely bought about demands of competition, raised the number of demands on individuals. As in Niklas' model, this created a fitness landscape which roughened over time, as individual demands on morphology increased. Marshall proposes that this same process drove the morphological diversity that arose in the Cambrian. If Niklas and Marshall are right, then these simple models tell us something important about phenotypic change at macroevolutionary scales.

As Niklas and Marshall are both aware, modelling phenotypic landscapes such as this requires making different assumptions than when modelling landscapes with genotypic dimensions. The first important assumption is a *smooth genotype-fitness mapping*. The small increments along the dimensions of variation are assumed to be the result of mutations. For this to be the case there must be a smooth genotype to phenotype map: neighbours in genotype space, separable by small mutational changes, will map to phenotypes that are neighbours in morphospace. This does not seem an unreasonable assumption—"realistic" genotypic models are often assumed to have correlated landscapes, where the fitness values of neighbouring genotypes are similar[9]. A second, less obvious assumption is *change-appropriate representation*. Here it is assumed the parameters used to define the morphospace capture both the range and direction of morphological possibility, and does so in way that it can represent how it might change. After all, we can parameterise a particular morphological shape with various coordinate systems. Not all of these will map to a way that captures how morphological change might differ under mutational change. For example, the morphospace of shells and their coiling is represented using polar coordinates (see Stone (1996) for an overview). Representing the morphospace via Cartesian coordinates would imply a very different (and probably quite useless) notion of accessibility in shell morphospace.

Niklas and Marshall have very different aims than Gavrilets when they employ their fitness landscape models. Speciation is not central to their arguments. Rather, it is the arise of morphological disparity that they wish to explain. Though we expect there to be some relation between morphological disparity and speciation, it will not be straightforward. Morphological diversity can occur both over time (change along a lineage) and at a time (phenotypic plasticity) without speciation.[10] Likewise, speciation can occur with little morphological change—there are many cryptic

[9] I take it that the assumption is that the *fitness* of neighbours is highly correlated depends on the *phenotypes* of neighbours being highly correlated.

[10] Speciation in Gavrilets' sense of genetic divergence causing reproductive isolation (Gavrilets 2004, p. 9).

species which look the same but cannot interbreed. The complex relationship between the two is a problem in conservation biology, where simple species counts are often inadequate guides to what people wish to preserve (Maclaurin and Sterelny 2008).

One can frame the different aims in these two cases as the search for difference-makers or enabling conditions. Enumerating the various conditions that *enable speciation* is, to a good approximation, the central focus of Gavrilets book. But these are not the same conditions that *enabled the rapid increase in diversity* in both the Devonian or the Cambrian that Niklas and Marshall wish to understand.

Despite the different aims in the two projects, it seems reasonable to ask whether the revolution described in the previous section is relevant here. After all, central to that argument was that biologically realistic fitness landscapes contain large interconnected neutral networks, rather than peaks. Yet the conclusions drawn from Niklas's model depend on peaks being present.

First, notice that Niklas and Marshall are not interested in peak-*shifting*—at no stage do they consider how a population might move from one peak to another. The dynamics of movement in their model always start at a low point and climb up. A simulation ends when a peak is reached; no further peak-shifting is necessary. Second, notice that at least some of Niklas' models include environmental changes (the introduction of new tasks). So these models sometimes include dynamic landscapes, where the peaks are appearing, disappearing, and changing over time. It is not clear how these dynamic landscapes should be treated when comparing them with Gavrilets' static landscapes.

We might still question the general *structure* of the landscapes from which Niklas draws his speculative conclusions. The morphological landscape he uses have few dimensions, and contain peaks separated by valleys. According to Gavrilets, biologically realistic landscapes are not like this. Rather, they contain fitness values of roughly equal value to be connected by a neutral network.[11] Let us consider that the peaks on a Niklas-type landscape are roughly the same in their fitness—so it is possible to make trade-offs in several ways that provide roughly equal fitnesses.[12] If Gavrilets model applies in this situation, then there should be some series of neutral changes in the genotype that take us from one peak to another without going through a valley of low fitness. Again, let us assume that such a route exists. What would this entail about the corresponding path of phenotypic change? Given that any deviation from either morphologies present on the peaks will reduce fitness, there seem to be only two ways such a route could exist. First, somewhere along this neutral path a single genetic change produces a flip from one phenotype to the other—an instantaneous shift between morphologies. Second, the morpho-space is missing some dimensions of variation which, if available, would describe some intermediate phenotypes that do not fall into a valley. Notice that these two solutions exactly violate the two important assumptions described above.

Perhaps we should expect the assumption of *smooth genotype-fitness mapping* to be violated. After all, we know that small changes in genes can *sometimes* lead to

[11] See the "Generalized Russian Roulette Model" (Gavrilets 2004, p. 89).

[12] Niklas' results suggest this is not an unreasonable assumption (Niklas 2004, p. 54).

large morphological change. We might also think that a phenotypic description could never be an entirely *change-appropriate representation*. The range of possible underlying developmental changes is simply to complex to be captured by a few high level morphological parameters. If we accept that Gavrilets insights are universally applicable to any fitness landscape model, then these objections, or something close to them, must be correct.

But it seems equally plausible that, once when we take into account actual physiological constraints on the development and performance of phenotypes (the very thing that is assumed away in Gavrilets' models), then there simply may not be a neutral path between to genotypes that produce very different morphologies that share the same fitness. Or perhaps we might think that path is so long and narrow, and thus incredibly unlikely to be traversed, that for practical purposes, it is simply irrelevant.

I think that something like this second position is more likely, and for this reason: Gavrilets' models were constructed to describe how genetic divergence and the resultant reproductive isolation could happen without going through a fitness valley. As such, the focus is on a single population diverging *just enough* to be reproductively isolated, rather than on *vastly* morphologically distinct populations after millennia of isolation.

Gavrilets's assumptions about pervasive neutral networks were established with a particular goal in mind—exploring how a population speciates. It is dubious whether these assumptions continue to hold when the models are being used to explore a very different question—the factors that underly morphological diversity.

6 Evo-devo and fitness landscapes

> The topological properties of configuration and phenotype landscapes can only be understood through insights into developmental mechanisms. (Wagner 2007, p. 150)

This section examines another use for fitness landscape models: exploring how development evolves. Again, the central question is not peak-shifting. But nor is it (as it was in the last section) understanding the evolution of morphological diversity. Instead, the models explore how genetic change maps to phenotypic change. The question these models are used to address is "... what kinds of systems can evolve successfully by random variation and selection for fitter variants" (Kauffman 1993, p. 19). This is a far from trivial question, as these systems must meet two simultaneous constraints: they must both work *at a time*, and yet be able to change *over time* (Calcott forthcoming). The study of such systems often centers around a number of putatively *general* properties that such systems have, such as robustness, plasticity, evolvability, and modularity. These kinds of properties are entirely different to those that are central in Gavrilets's work, whose focus was on such properties as "expected time to speciation".

To explore these properties, static landscape models are constructed by assuming some kind of mapping from genotype to phenotype, and then the general (statistical)

properties of these landscapes are examined.[13] Although this approach was pioneered by Kauffman with his work on *rugged* landscapes (Kauffman 1993), much recent work has focused on demonstrating the importance of *neutral* change in these landscapes. As with the landscape models for speciation, including neutral change turns out to have some important implications. In this case, however, the effect of including neutral change is interpreted in a very different manner: it affects some of the general properties mentioned above, notably robustness and evolvability.

I'll focus here on a paper by Fontana: "Modelling 'Evo-Devo' With RNA" (Fontana 2002).[14] As the title suggests, Fontana uses this paper to articulate a particular "modelling strategy":

> In a narrow sense, the relation between RNA sequences and their shapes is treated as a problem in biophysics. Yet, in a wider sense, RNA folding can be regarded as a minimal model of a genotype–phenotype relation ...
> ... the RNA folding map transparently implements concepts like epistasis and phenotypic plasticity, thus enabling the study of constraints to variation, canalization, modularity, phenotypic robustness and evolvability. (Fontana 2002, pp. 1164–1165)

Much as Godfrey-Smith (2006) suggests, Fontana's aim is to understand a more complex system by studying a simpler one. In this case the model system is itself a biological one.[15] The project is not to understand how RNA folding works, but to use what we know about the RNA folding as a simple model for the mapping function from genotype to phenotype, and to examine the properties of the resulting landscapes that are generated.

The process of RNA folding maps the *primary* structure of RNA—a sequence consisting of one of four nucleotides—to the *secondary* structure of RNA—a two dimensional structure where some consecutive nucleotides bind together to form *runs* and others remain unpaired forming *loops*.[16] Which particular folded structure (a combination of runs and loops) is most likely depends on the sequence of nucleotides in the sequence of RNA and the temperature at which it folds. Although

[13] In a similar manner to how Gavrilets' first established the holey properties of landscapes—what I have called the first step of his argument.

[14] Note that Gavrilets refers to the work on RNA landscapes in his book on speciation (and Pigliucci and Kaplan subsequently refer to it in their analysis of Gavrilets). Notably, it is used as additional justification for the presence of neutrality (step 1 of his argument, as I described it in "Speciation and fitness landscapes"). Despite their mention by Gavrilets and Pigliucci and Kaplan, however, these papers on RNA folding have little to do with speciation.

[15] There are some unresolved issues regarding the relationship between models and model descriptions here. Although the putative model system is RNA folding, the mathematical representations examined by Fontana already simplify what we know about actual RNA folding (For example, they ignore tertiary structure). So the is the model being used to understand evo-devo RNA folding, or a simplified model of RNA folding? I shall not attempt to address these issues here.

[16] Fontana points out that, strictly speaking, it is not the structure that is treated as the phenotype in this model, but the topology of pairings. This topology can be represented in various ways, one of which will show the shape of the folded RNA (Fontana 2002, Figure 1, p1165). For the discussion here, this distinction will not matter.

there may be a single mostly likely structure (the minimum free energy or "mfe" structure), given a particular sequence and the temperature, the actual structure taken by the RNA sequence can fluctuate amongst a number of different low-energy forms (something that can be described by another type of landscape: an "energy" landscape). The set of structures that a sequence may fold into, the likelihood of these structures, and how these change under both environmental (temperature) and mutational change provides a rich model for representing a number of important relationships between genotype and phenotype.

If we consider only the most likely structure as the phenotype of a particular RNA sequence, then neutral change can arise when neighbouring sequences, differing by only a single nucleotide, fold into the same structure. A neutral network will consist of all the neighbouring sequences that fold into the same structure.[17] Study of RNA folding suggests that these neutral networks are both pervasive and large. A deeper understanding of the neutral networks is relevant to a number of biologically relevant properties.

> The possibility of changing the genotype while preserving its phenotype is both a manifestation of phenotypic robustness to genetic mutations and a key factor underlying evolvability. (Fontana 2002, p. 1169)

A system is robust if it continues to function under perturbations, both internal and external. Biological systems, despite being vastly complex, are often very robust (especially in contrast to many engineered systems) (Kitano 2004; Wagner 2005). What kinds of properties underly this robustness, how they operate, and how they might have evolved is something in need of explanation. The fact that change can occur in the sequence without affecting the resultant folded structure means that this structure is robust to underlying changes in the sequence. Andreas Wagner, in his book-long study of robustness in living systems concludes that "... the concept of a neutral space can provide a unified view of robustness in many different systems ..." (Wagner 2005). Wagner's neutral spaces are equivalent to fitness landscapes models. These neutral spaces also underly (one conception of) evolvability.

> Neutral networks enable phenotypic innovation by permitting the gradual accumulation of silent mutation from these sequences .... (Fontana 2002, p. 1170)

To illustrate the importance of neutral change to evolvability, consider the moving around a department store via its escalators (instead of moving around a rugged landscape). In my experience, ascending the floors is never straightforward task, for the top of one escalator is *never* next to bottom of the next one (no doubt some devious attempt to increase the chances of an unanticipated purchase). So, in a popular department store in Melbourne (Australia), arriving at level 1 from the ground floor, I find I must wander aimlessly past glass cabinets full of perfume and

[17] RNA strands with the same secondary structure may not act the same, as they have different bases at different locations. So it is an assumption in these models is that the same secondary structure is sufficient to cause the same fitness.

lipstick, in order to find the escalator that takes me up to level 2. Once I'm on level 2, in order to get to escalator that takes me level 3, I am required to walk across some plastic grass covered with garden furniture. Now contrast being stuck on a particular floor looking for the next escalator with being stuck on peak of a rugged landscape, staring over at a higher peak. In the department store, some random wandering can eventually get me to the bottom of the next escalator *without ever going down*. Not so in the case of a rugged landscape, descending into a valley is required in order to mount that higher peak. We can interpret Fontana's statement as something like this: increasing the size of department store floor (the neutral space), increases the probability that it will intersect with an escalator at some point, providing a ride to the next floor (and being better adapted).

So the introduction of neutral change to these models paints a very different picture of the constraints on how adaptive change takes place over time. Getting stuck on a peak is no longer a central problem. Instead, what matters is understanding how the genotype-phenotype map shapes accessibility between different phenotypic states via neutral pathways.

Although neutral change is important in both Gavrilets's and Fontana's work, the aim of introducing it is quite different. As I showed in the previous section, a simple way to see the different aims is to contrast the kinds of difference-makers or enabling conditions that are important in each case. Rather than enumerating the conditions that *enable speciation*, Fontana's work on landscapes examines the conditions that *enable adaptive systems to be robust and evolvable*.

The assumptions made in these models are also very different. Rather than manually or randomly assigning fitness to genotypes, how a particular genotype produces a phenotype (and its concomitant fitness) is the focus of the exploration. The details of this process matter, for it affects the accessibility between various phenotypes on the landscape. So, in the case of the RNA folding, we want to understand *how* neutrality comes about, rather than simply showing that it is inevitable given large dimensions (as Gavrilets did). For example, it turns out that the distribution of RNA structures that appear is skewed: many genotypes form the same structure, and a few form rare structures. More importantly, the structure of the boundaries between bordering neutral networks can result in asymmetric accessibility at the phenotypic level: the probability of mutating from shape A to shape B can be significantly different from mutating in the other direction. Fontana interprets this asymmetry as a *developmental constraint*. None of these detailed assumptions are required to understand how genetic divergence leads to reproductive isolation and speciation.

7 Conclusion

The numerous uses to which fitness landscape models can be put suggest that their general structure captures some important aspects of evolutionary change. Günter Wagner even suggests that the ability for fitness landscapes to represent both adaptive constraints (via hill climbing) and developmental constraints (via the

asymmetric accessibility between neutral networks) may be a way to establish some rapprochement between population genetics and evo-devo (Wagner 2007, p. 150).

But the variety of uses they serve also makes it difficult to make general claims about their success or failure. One aim of this paper was to assess how well Pigliucci and Kaplan's claim of a revolution held up when a wider range of models was surveyed. For the most part the revolution appears limited to issues of speciation. Here it is really is revolutionary—the dissolution of peak-shifting make a huge difference to the models of speciation. But the effect it has on fitness landscapes that model large-scale morphological change seems minimal in contrast. The differences between the peaks do not represent a speciation event, and peak-shifting is not a problem to be solved. In the case of the evolution of development, the issue is less clear cut. Like speciation, the introduction of neutral change into the landscapes comes from looking at higher dimensional landscapes, and contrasts with earlier work that focussed on rugged landscapes. But the upshot of introducing the neutral change, as I pointed out in the last section, has a very different effect.

I've also argued that the model-based approach provides a clearer view of the current revolution in fitness landscapes. For it takes into account the modelling practices, and these demonstrate strongly pluralistic use of models even when assessing a single question, rather than suggesting that increasing biological realism is sufficient to construct better models. This had the effect of rewriting the story of the revolution, without the need for talk of metaphors.

I want to conclude by briefly showing how a wider survey of the use of fitness landscapes might affect our view of macroevolution. Pigliucci and Kaplan suggest that Gavrilets's work on fitness landscapes provides a distinction between macroevolutionary and microevolutionary processes. Microevolution occurs as short non-neutral climbs on slopes, whereas macroevolution is a large-scale wandering on the neutral networks (Pigliucci and Kaplan 2006, pp. 196–197). But this assessment is based on an interpretation of fitness landscapes (and macroevolution) tied closely to speciation. A different picture arises from the other two uses of fitness landscapes outlined here. The work on RNA landscapes suggests that phenotypic innovation can be influenced by the availability of neutral "wandering". But this suggests that the periods of phenotypic change would be those in which a new "escalator" was found, and that periods of morphological stasis would occur whilst wandering the neutral networks. If we focus on morphological change, rather than speciation, we get the opposite picture than that of Pigliucci and Kaplan—macroevolutionary (morphological) change occurs on the slopes, not on the plains. Looking at Niklas's work, we get a different picture again. Here, microevolution occurs as movement across the landscape, whilst the macroevolutionary processes change the shape of the landscape itself. In each of these cases, the different questions that fitness landscape models are used to address provide different ways of viewing the relation between microevolution and macroevolution.

Fitness landscapes are clearly a powerful tool for simplifying and providing insight into evolutionary processes. But what they tell us is influenced by assumptions that mould the tool for different purposes.

References

Arnold S (2003) Performance surfaces and adaptive landscapes. Integr Comp Biol 43:367–375

Calcott B Lineage explanations: explaining how biological mechanisms change. Br J Philos Sci (forthcoming)

Camp E, Reimer M (2006) Metaphor. In: Lepore E, Smith BC (eds) The Oxford handbook of philosophy of language. Oxford University Press, New York, pp 845–863

Fontana W (2002) Modelling 'evo-evo' with RNA. BioEssays: News Rev Mol Cell Dev Biol 24(12):1164–1177

Gavrilets S (1997) Evolution and speciation on holey landscapes. Trends Ecol Evol 12(8):307–312

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

Giere R (1988) Explaining science: a cognitive approach. University of Chicago Press, Chicago

Giere R (2002) The nature and function of models. Behav Brain Sci 24(06):1060–1060

Godfrey-Smith P (2006) The strategy of model-based science. Biol Philos 21:725–740

Kauffman SA (1993) The origins of order: self-organization and selection in evolution. Oxford University Press, New York

Kauffman S, Levin S (1987) Towards a general theory of adaptive walks on rugged landscapes. J Theor Biol 128(1):11–45

Kitano H (2004) Biological robustness. Nat Rev Genet 5(11):826–837

Levins R (1966) The strategy of model building in population biology. Am Sci 54(4):421–431

Maclaurin J, Sterelny K (2008) What is biodiversity? University Of Chicago Press, Chicago

Marshall C (2006) Explaining the cambrian "explosion" of animals. Annu Rev Earth Planet Sci 34: 355–384

Matthewson J, Weisberg M (2008) The structure of tradeoffs in model building. Synthese. doi: 10.1007/s11229-008-9366-y

Niklas K (1994) Morphological evolution through complex domains of fitness. Proc Natl Acad Sci USA 91(15):6772–6779

Niklas K (2004) Computer models of early land plant evolution. Annu Rev Earth Planet Sci 32:187–213

Pigliucci M, Kaplan J (2006) Making sense of evolution: the conceptual foundations of evolutionary biology. University of Chicago Press, Chicago

Shpak M, Wagner G (2000) Asymmetry of configuration space induced by unequal crossover: implications for a mathematical theory of evolutionary innovation. Artif Life 6(1):25–43

Stadler P (2002) Fitness landscapes. In: Lassig M, Valleriani A (eds) Biological evolution and statistical physics. Springer-Verlag, Berlin, pp 187–207

Stadler P, Stadler B (2006) Genotype–phenotype maps. Biol Theory 1:268–279

Stadler B, Stadler P, Shpak M, Wagner G (2002) Recombination spaces, metrics, and pretopologies. Z Phys Chem 216:217–234

Stone JR (1996) The evolution of ideas: a phylogeny of shell models. Am Nat 148(5):904–929

Wagner A (2005) Robustness and evolvability in living systems. Princeton University Press, Princeton

Wagner G (2007) How wide and how deep is the divide between population genetics and developmental evolution? Biol Philos 22(1):145–153

Weisberg M (2006) Forty years of 'the strategy': Levins on model building and idealization. Biol Philos 21(5):623–645

Weisberg M (2007) Who is a modeler? Br J Philos Sci 58(2):207–233

Weisberg M, Reisman K (2008) The robust volterra principle. Philos Sci 75(1):106–131

Wimsatt W (1987) False models as a means to truer theories. In: Nitecki M, Hoffmann A (eds) Neutral models in biology. Oxford University Press, Oxford, pp 23–55

Wright S (1932) The roles of mutation, inbreeding, crossbreeding and selection in evolution. Proc 6th Int Cong Genet 1:356–366