

Sewall Wright's adaptive landscapes: 1932 vs. 1988

Massimo Pigliucci

Received: 13 November 2007 / Accepted: 16 July 2008 / Published online: 6 August 2008
© Springer Science+Business Media B.V. 2008

Abstract Sewall Wright introduced the metaphor of evolution on “adaptive landscapes” in a pair of papers published in 1931 and 1932. The metaphor has been one of the most influential in modern evolutionary biology, although recent theoretical advancements show that it is deeply flawed and may have actually created research questions that are not, in fact, fecund. In this paper I examine in detail what Wright actually said in the 1932 paper, as well as what he thought of the matter at the very end of his career, in 1988. While the metaphor is flawed, some of the problems which Wright was attempting to address are still with us today, and are in the process of being reformulated as part of a forthcoming Extended Evolutionary Synthesis.

Keywords Adaptive landscapes · Evolutionary theory · Genetic drift · Natural selection

The idea of an adaptive landscape is arguably the most influential metaphor in evolutionary biology after Darwin's own parallel between natural and artificial selection. It is commonly presented in textbooks (Hartl and Clark 1989; Futuyma 1998) and has inspired a wealth of (largely theoretical) research (Kauffman and Levin 1987; Conrad 1990; Schluter and Nychka 1994; Whitlock 1995; Coyne et al. 1997; Gavrillets 1997; Svensson and Sinervo 2000; Hadany and Beker 2003; Jones et al. 2004; Kramer and Donohue 2006; Prusinkiewicz et al. 2007). Yet, the argument has been made that taking adaptive landscapes seriously leads to quite different evolutionary scenarios from those classically formulated by Wright (Gavrillets 1999), and in fact to the possibility that the metaphor itself has grossly

M. Pigliucci (✉)
Departments of Ecology & Evolution and of Philosophy,
Stony Brook University, Stony Brook, NY 11794, USA
e-mail: pigliucci@platofootnote.org

mislead evolutionary biologists for several decades, causing them to focus on problems—such as the question of a “peak shift”—that either do not exist or need to be reformulated in a dramatically different fashion (Pigliucci and Kaplan 2006, Chap. 8).

In this essay I intend to go back to the origins of the metaphor, Wright’s 1932 paper, to identify the main themes that run through it; I will then discuss whether and how Wright changed his position in the last paper he published on the subject, in 1988 (the year he died). Finally, I will briefly examine how Wright’s ideas have fared in the recent debate about the conceptual foundations of evolutionary biology and what, if anything, they may contribute to a forthcoming Extended Evolutionary Synthesis. The focus will not be on a detailed historical reconstruction (which is, however, a valuable project in its own right), but rather on the motivations that pushed Wright to introduce the metaphor, and on the impact that it still has, for good and for ill, on evolutionary theory.

Wright 1.0: the introduction of the adaptive landscapes metaphor

The very first image of an adaptive landscape is Fig. 1 of Wright’s (1932) paper on “the roles of mutation, inbreeding, crossbreeding and selection in evolution.” The attempt at a pictorial version of his ideas was inspired by Wright’s mentor, E.M. East, who had organized the Sixth International Congress of Genetics where the paper was presented. According to Wright’s (1988) own account, East was afraid that biologists—notoriously uncomfortable (then and now) with mathematics—might simply not understand the import of Wright’s ideas unless he used pictures. While East’s attitude may have been condescending, history confirmed that there was more than just a grain of truth in it.

East’s own work, which had inspired Wright, contrasted the mostly deleterious effect of mutations with the evolutionary possibilities afforded by recombination. While at the time the number of known mutations in *Drosophila* was a puny 400, and the best guess available for the number of genes in higher organisms was in the “range from 1,000 up” (Wright 1932, p. 356), Wright quickly calculated that reasonable assumptions about the number of allelomorphs possible even with those figures was orders of magnitude higher than the number of particles in the universe. This conclusion brought him to the fundamental question he attempted to address in the 1932 paper: how does evolution gain “access” to different regions of what he called the “field of possible gene combinations”? His back-of-the-envelope calculations showed that, while it is not difficult to account for the fact that there is enough variation so that any given organism has a good chance of having a different genotype from the other members of its species (in outcrossing populations), it also follows that only a tiny fraction of the possible genetic combinations is actually accessible at any given time: “the population is thus confined to an infinitesimal portion of the field of possible gene combinations” (Wright 1932, p. 356). This distinction between the actual and potential genetic combinations anticipates the formal distinction between *standing* variation and

potential variability introduced by Wagner and Altenberg (1996), and which is one of the key ideas in the current study of evolvability (Pigliucci 2008).

Wright's first attempt at visualizing the problem is the above-mentioned Fig. 1 of the 1932 paper, a series of polyhedrons of increasing complexity, with vertices indicating combinations of allelomorphs for a growing number of loci. The most complex of these diagrams shows 32 homozygous combinations, and Wright makes the point that each of these is one step removed from five others, and two steps removed from ten others—thereby introducing the idea of step-wise movement across a genetic landscape that was later formalized by Kauffman and Levin (1987), and which represents an important metric in Gavrilets' (1997) model of “holey” adaptive landscapes. Wright (1932, p. 357) admits that it would take five dimensions to represent his simplified field of gene combinations, adding a sixth one “to represent level of adaptive value”—the first reference to the idea that the landscape is, in fact, an adaptive one.

Crucially, Wright also immediately acknowledges that even with his very conservative figures concerning the number of loci and alleles present in a species, it would actually require about 9,000 dimensions to truly represent the field. Nonetheless, he then introduces Fig. 2 of the paper (Wright 1932, p. 358), a two-dimensional representation (with unlabeled axes), of the genetic field, where for the first time fitness peaks and valleys are visualized by means of a contour plot. Interestingly, the landscape looks continuous, which may induce a modern reader to think that it is analogous to most modern representations, where the non-fitness axes are gene frequencies. But in this case Wright is visualizing (as he makes explicit in the 1988 paper) “genotypes ... packed, side by side ... in such a way that each is surrounded by genotypes that differ by only one gene replacement” (Wright 1988, p. 116). But this makes the landscape in Fig. 2 confusing at best, as there is no metric that allows one to “pack” genotypes side-by-side in that fashion, let alone any intelligible way of interpreting what a “peak” or “valley” would represent, since the contour lines are not, in fact, contour lines at all (they cross a dense but fractioned, not continuous, space).

Nonetheless, Wright states that in such landscapes “it may be taken for certain that there will be an enormous number of widely separated harmonious combinations” (Wright 1932, p. 358), i.e., adaptive peaks in modern parlance. This is where Wright then sets up what he conceives of as the central point he is addressing, which is worth citing in full:

In a rugged field of this character, selection will easily carry the species to the nearest peak, but there may be innumerable other peaks which are higher but which are separated by ‘valleys.’ *The problem of evolution as I see it* [my emphasis] it is that of a mechanism by which the species may continually find its way from lower to higher peaks in such a field. In order for this to occur, there must be some trial and error mechanism on a grand scale by which the species may explore the region surrounding the small portion of the field which it occupies. To evolve, the species must not be under strict control of natural selection. Is there such a trial and error mechanism? (Wright 1932, pp. 358–359)

But why on earth should this be *the* problem of evolution, or for that matter, a problem at all? There are two reading keys that I will entertain here. One is the possibility that Wright somehow espoused a belief in the inevitability of progress in evolution; the other that he in a sense anticipated one of the crucial questions of current theoretical research in evolutionary biology, the evolution of the ability to evolve itself. Neither reading, I suspect, is entirely satisfactory, and yet a brief discussion of both may be enlightening not just about Wright's own thought, but more importantly about the recent and future direction of evolutionary research.

The idea that there is something inevitable in evolution, some sense in which evolution makes "progress" is a persistent one throughout the history of evolutionary biology, and it does not have to imply intelligent agency. At the turn of the 20th century, during what some have referred to as the (temporary) "eclipse" of Darwinism (Bowler 1983), many paleontologists in particular held to an "orthogenetic" view of evolution, which de-emphasized natural selection in favor of the role played by internally generated forces. More recently, discussions of whether and in what sense evolution may yield an increase in complexity (McShea 1991; Gould 1996), or in adaptation (Van Valen 1973) still occasionally rage in the literature. It is difficult, of course, to make reasonable inferences on exactly why Wright thought that what later became known as the problem of peak shift is so central to evolution, but certainly a good number of empirical and especially theoretical biologists followed suit, providing a variety of possible solutions to the problem (Phillips 1993; Price et al. 1993; Wagner et al 1994; Whitlock 1995; Whitlock 1997; Hadany and Beker 2003), besides or alongside Wright's own shifting balance theory.

One way to think of the issue is to phrase it as a question of optimizing vs. "satisficing" processes. It is commonplace in evolutionary theory to use optimization models to predict which phenotypes will evolve under a given set of conditions and tradeoffs (Maynard Smith 1982; Charlesworth 1990; Williams 1992; Abrams et al. 1993; Seger and Stubblefield 1996), although of course such models are very different from Wright's approach in the crucial respect that the latter is a genetic model, while optimization theory deals with phenotypes (then again, evolutionary biologists since Simpson have moved between genotypic and phenotypic adaptive landscapes rather too easily, as Wright himself acknowledged: Wright 1988, p. 120). If natural selection is an optimizing process, then it makes sense for Wright and followers to ask how on earth could selection get a population unstuck from a local adaptive peak, which comes to be seen as a sub-optimal situation.

But there is a different and, I think, more biologically sound, way to look at the "problem," which actually makes the problem itself disappear, or at least changes the terms of the discussion significantly. Some authors have used the concept of "satisficing," imported from foraging theory in ecology (Ward 1992), to suggest that we should think not of natural selection as "the survival (and reproduction) of the fittest," but rather more modestly as the "non-survival (and non-reproduction) of the non-fit" (den Boer 1999). In other words, if it is good enough for survival (and reproduction), it is as far as natural selection will push it, even though the

(phenotypic) solution may be sub-optimal when considered with respect to all *possible* solutions. There is some intriguing empirical evidence that natural selection does not, in fact, improve the fit of populations to their environments even over long periods of time, but that it rather works at maintaining whatever local adaptive peak the population has achieved (usually simply inherited from the previous generation, a phenomenon known as phylogenetic niche conservatism). Van Valen's (1973) paleontological data are consistent with this idea, termed "Red Queen hypothesis" in reference to the homonymous character in Lewis Carroll's *Through the Looking Glass*, who was forced to keep running just to maintain her place. Analogously, Van Valen suggested, species are under constant natural selection to keep tracking environmental changes, and if they fall behind (for instance because of lack of sufficient appropriate genetic variation) they slide into extinction. There is no question of moving from one peak to a higher one if the realities of everyday biological evolution are such that it is difficult enough to simply hold on onto one's current position.

The above notwithstanding, there is a much more modern interpretation of Wright's insistence that peak shifting is "the problem of evolution": he may have been concerned, at least in some fashion, with what today is known as evolvability (Conrad 1990; Wagner and Altenberg 1996; Pigliucci 2008), i.e. with what makes it possible for evolutionary mechanisms themselves to evolve over time. This is certainly not the place for even a short review of the large literature on evolvability, but I think it important to note a little appreciated paper by Conrad (1990), which actually offers a solid conceptual bridge between Wright's original work and more modern treatments of adaptive landscapes (Gavrilets 1997, 1999), within the context of the concept of evolvability. Conrad begins by comparing Wright-type landscapes with molecular landscapes defined by nucleotide variation. The advantage of the latter is that one can actually do experimental evolution with them, and that the Genotype \Rightarrow Phenotype mapping function (Alberch 1991), where the phenotype is the structure and behavior of a protein, is actually computable. This computability in turn makes the transition from genetic to phenotypic landscapes tractable, and it is this transition that Conrad exploits in his conceptualization of phenotypic spaces as phase spaces with basins of attraction. From there he then tackles the possibility that the high dimensionality of adaptive landscapes may create "extra-dimensional" bypasses (a concept taken up in a more formal fashion by Gavrilets 1997, 1999) as alternative solutions to (in fact, reformulations of) the peak shift problem.

Considering how Wright continues his original paper, however, it seems unlikely that his primary concern was the evolution of evolvability in the long run. At most, his goal was to provide an account of the current level of local evolvability of a given population, as when he refers to his conviction that there must be "some trial and error mechanism on a grand scale" (Wright 1932, p. 359) to solve the problem of evolution on rugged adaptive landscapes. Indeed, the next crucial figure of the 1932 paper presents a complex classification of mechanisms to explain what happens to a population when it sits on an adaptive peak (Wright 1932, Fig. 4, p. 361). This is a remarkable figure, in that it both summarizes much of what was then very recent theoretical advances in population genetics, and anticipates several

of the major lines of research that unfolded in the field during the following several decades.

Briefly: (a) either increased mutation or reduced selection will cause the area occupied by the population on a peak to expand, essentially reducing average fitness; (b) the opposite, either decreased mutation or increased (stabilizing) selection will shrink the adaptive cloud, increasing adaptation to the local environment; (c) a change in the environment, combined with Fisherian mass selection (i.e., acting on large effective population sizes) will move the population to a new peak, but at a very slow pace; (d) inbreeding in small populations will cause a random walk off the adaptive peak, and eventually extinction; (e) moderate levels of inbreeding in medium-sized populations also lead to wandering off the current peak, but with a low likelihood of reaching a new one; (f) the population fractures into local demes (“races,” in Wright’s language) characterized by relatively small numbers of individuals. It is the latter case, which he later termed “shifting balance,” that Wright introduces as the trial and error mechanism likely to solve “the problem of evolution.”

According to Wright (1932, p. 363), this process is characterized by continuous nonadaptive shifts, which leads to “good chances” that at least one population will come under the influence of a new adaptive peak. In his view, the process of intergroup selection among the various demes will be more efficient than standard intragroup selection, because it depends not on the mutation rate, but on the much higher migration rate. The question of how likely the shifting balance actually is, however, has not been settled more than seventy years after Wright’s original suggestion that it was a likely and efficient process (see, for instance, Coyne et al 1997 for a contrary view, and Peck et al. 1998 for a favorable one; very recently, work by Schemske and Bierzychudek (2007) seems to undermine one of the best-known empirical examples of shifting balance, the evolution of blue- and white-colored morphs of the desert snow plant of the Mojave desert).

Yet, Wright appears to be confident already in the 1932 paper that his shifting balance works, partly because of what he considered highly suggestive empirical evidence. One of the standard classes of examples he adduces is the observation that the subspecific and interspecific differences often used by systematists to separate species appear to be nonadaptive, implying that whatever mechanisms lead to speciation are likely not driven by natural selection. Although this is a *prima facie* weak argument, it is an early direct link between the shifting balance and speciation theory, and it is interesting to note that still today much of the discussion about *adaptive* landscapes actually refers to allegedly non-adaptive speciation events. Most interesting among the empirical evidence that Wright invokes in his favor is an astoundingly early recognition of the fact that the fossil record is often compatible with a process “in which evolution has proceeded by extensive differentiation of local races, intergroup selection, and crossbreeding” (p. 365). This sounds a lot like Eldredge and Gould’s (1972) theory of punctuated equilibria, which in fact they did connect to episodes of selection among small demes which, if successful, then expand their geographical distribution (see also Eldredge et al. 2005 for an updated treatment of evolutionary stasis in light of Gavrilets’ ideas of highly dimensional landscapes).

Wright 56 years later

It is not often the case that we have the opportunity to compare what a scientist thought of a specific topic when he first introduced his ideas and at the end of his career. In this case, however, one of Wright's last papers was published in *American Naturalist* in 1988, and it is appropriately titled "Surfaces of selective value revisited."¹ The amazing thing about the '88 paper is that it does not differ much from the 1932 version, a fact that may reflect Wright's own remarkable consistency as much as the lack of substantial progress on the question of adaptive landscapes since the concept had been proposed more than half a century earlier. Indeed, Wright (1988) reproduces the figures of the 1932 paper, providing a figure-by-figure commentary on what they mean, but in fact adding little to what he had written originally.

One rather small exception to this consistency is the fact that Wright points out that he had omitted one possible mechanism from his original scenariogram (Wright 1932, 1988, Fig. 4), although it was included in the text of the 1932 paper: the appearance of entirely new, favorable, mutations in the population. As he writes in the 1988 paper (p. 118), a graphical representation of this phenomenon would have entailed adding more contour lines to the adaptive landscape, to reflect the fact that the mutation had altered the landscape itself. However, Wright originally as well as later dismissed this possibility as "an exceedingly slow process," because such mutations are presumably very rare.

What is perhaps most interesting in the 1988 paper is Wright's discussion of paleontology (p. 120). He explicitly makes the connection I have highlighted earlier between his shifting balance theory and Eldredge and Gould's punctuated equilibria. Moreover, Wright—correctly, it seems to me—links both to Simpson's (1944) idea that evolution can proceed at different speeds in different lineages, with "tachytely" (fast evolution) yielding a punctuated equilibrium pattern. According to Wright, tachytely is what one would expect if the shifting balance were taking place. Interestingly, Wright acknowledges that Simpson took the idea of adaptive landscapes—defined in genotypic space—and translated it with little fanfare at the phenotypic level. This is actually a bold and highly questionable move (Pigliucci and Kaplan 2006, Chap. 8), and yet Wright comments that "the choice ... is practically irrelevant in connection with pictorial representation of changes that occur in populations" (p. 120). How can it be irrelevant? One representation portrays a non-continuous surface of genotypes packed in an ill-defined conceptual space, the other outlines a continuous space of metric phenotypic characters. The first one can be described—albeit approximately—by mathematical population genetic formalisms, while the latter admits only of a posteriori statistical descriptions à la Lande and Arnold (1983). To move between the two one would

¹ The reference in the title is to a paper he published in 1967, "Surfaces of selective value," in the *Proceedings of the National Academy of Science*, USA. However, Wright actually discusses his 1932 contribution, as well as a longer paper he published in 1931 in *Genetics*, entitled "Evolution in Mendelian populations."

have to have access to some sort of Genotype \Rightarrow Phenotype mapping function, one of the most obstinate problems in modern evolutionary biology (Pigliucci 2008).

Wright (1988, p. 120) comes close a second time to tackling a serious problem with the adaptive landscape metaphor, only to once again let it go much too prematurely. He writes that he became recently (about a decade earlier than the 1988 paper) dissatisfied with the 2D representation of landscapes, given that their actual dimensionality is much higher. Nevertheless, he doesn't seem to be bothered by the inaccuracy, writing that:

I did not arrive at any changes that should be made in the diagrams as presented. Irrespective of the complexity of the true pattern, the likelihood of a shift from the currently controlling selective peak to a higher one depends only on the number of replacements of currently predominant alleles by ones that are fairly abundant and capable of contributing significantly toward reaching the saddle leading to the most available higher peak.

Less than a decade later, Gavrilets (1997, 1999), borrowing percolation analysis from physics, showed that the dimensionality of the landscapes matters indeed, as high dimensional landscapes are characterized by very different evolutionary dynamics, so that the entire "peak shift" problem may in fact disappear altogether (Pigliucci and Kaplan 2006, Chap. 8).

Wright concludes the '88 paper with another interesting clarification: he suggests (p. 121) that his original intentions had been much misunderstood, as many people read his 1932 work as a contribution to the then ongoing debate concerning the degree of adaptedness (or lack thereof) of interspecific differences. He mentions that Julian Huxley took Wright's insistence on the role of genetic drift as confirmation of his own view that interspecific differences were largely non-adaptive, and that Fisher also adopted the same interpretation, in his case leading however to disapproval because of Fisher's contention that natural selection is the main agent explaining biological diversity. Wright, on the other hand, apparently saw genetic drift as the generator of random variation that fueled inter-group selection, just as random mutation fuels standard natural selection. Indeed, he states explicitly that "my chief interest has been not evolution in general, but *adaptive* evolution" (Wright, 1988, p. 121, original emphasis). Wright concludes in a rather ecumenical and conciliatory mood: "it is to be noted that the mathematical theories developed by Kimura, Fisher, Haldane, and myself dealt with four very different situations. ... All four are valid." So much for seven decades of what everyone else perceived as a raging controversy (Wade and Goodnight 1998; Crow 2002).

Is there a future for adaptive landscapes?

Given the serious conceptual issues surrounding Wright's metaphor of adaptive landscapes, one could reasonably ask whether it is not time to simply drop the metaphor altogether. As mentioned above, Kaplan and I (Pigliucci and Kaplan 2006, Chap. 8), following an in-depth analysis of Gavrilets's (1997, 1999) work on the properties of high-dimensionality landscapes, come close to strongly endorsing this

suggestion. At the very least, biologists ought to be very specific about what sort of “landscape” they are referring to (individual genotypes, gene frequencies, or phenotypes), as well as constantly be reminded that they simply cannot assume that the intuitive properties of low-dimensional landscapes are a reliable guide for the real thing.

That being said, the crucial problem Wright was attempting to address remains: as Darwin (1859) noted, the theory of evolution does require an account of the intermediate steps taken by the evolutionary process in the production of new structures and functions. Obviously, Darwin was not talking about genotypic landscapes, but he was aware that one needs to explain how movement in what is today called phenotypic space actually occurs. That problem is still with us, though we are finally making some significant progress towards its solution. A brief discussion of three recent examples will give some indication of how the field is moving.

Work by Kavanagh et al. (2007) on the evolution of mammalian teeth has produced a developmental model that accounts for the actual variation found in a simple two-dimensional phenotypic space defined by the ratios of sizes between the second and first molar on one axis and the third and first molar on the second axis. The model is based on the observation that the three molar teeth develop sequentially from front to back, and the assumption that the size of each tooth depends on the ratio of an activator to inhibitor molecules. Other than this, there seems to be a constraint such that regardless of other considerations the second molar is always about one third the size of the other two molars combined. Polly (2007) extended the application of this simple model to 35 species of extant and extinct mammals, and found that it is consistent with the large majority of the data, with only three species of bear, the horse, and three species of raccoons failing to fit the pattern. The model is based on experimental data, although the precise identity of the molecules involved is at the moment only suspected. What makes this work important, as Kavanagh et al. (2007) point out, is that it provides “an example of ecologically driven evolution along a developmentally favored trajectory.” Of course, we do not know why molars have to develop in a particular sequence, nor why the second one is constrained to a particular ratio of sizes with the other two; moreover, we do not have a good explanation of the exceptions to the model mentioned above. Still, this is precisely the sort of work that goes a long way toward satisfying Darwin's requirement that we have to have an account of how evolution in phenotypic space proceeds and why.

Of course, the work of Kavanagh et al. (2007) refers to phenotypic landscapes, not genotypic ones. Although Wright seemed to be strangely unconcerned about the difference between the two (Wright 1988, p. 120), the two landscapes are not at all easy to relate to each other, with a few exceptions. As mentioned above, the connection between genotypic and phenotypic landscapes can be articulated using the idea of a “genotype-phenotype mapping function” ($G \Rightarrow P$), as proposed by Alberch (1991). The problem, of course, is that we rarely have any idea of how to construct such a “map” (which is, obviously, yet another metaphor, and therefore should carry all the usual warnings about the conceptual traps hidden in the use of metaphors). One of the exceptions is the study of the evolution of protein function, as in that case the $G \geq P$ map is relatively simple, depending on a small number of

relatively well-understood molecular steps. Poelwijk et al. (2007) review work on the evolution of enzymes using adaptive landscapes as a conceptual framework. Unlike most other phenotypes, we can now artificially engineer variants of known

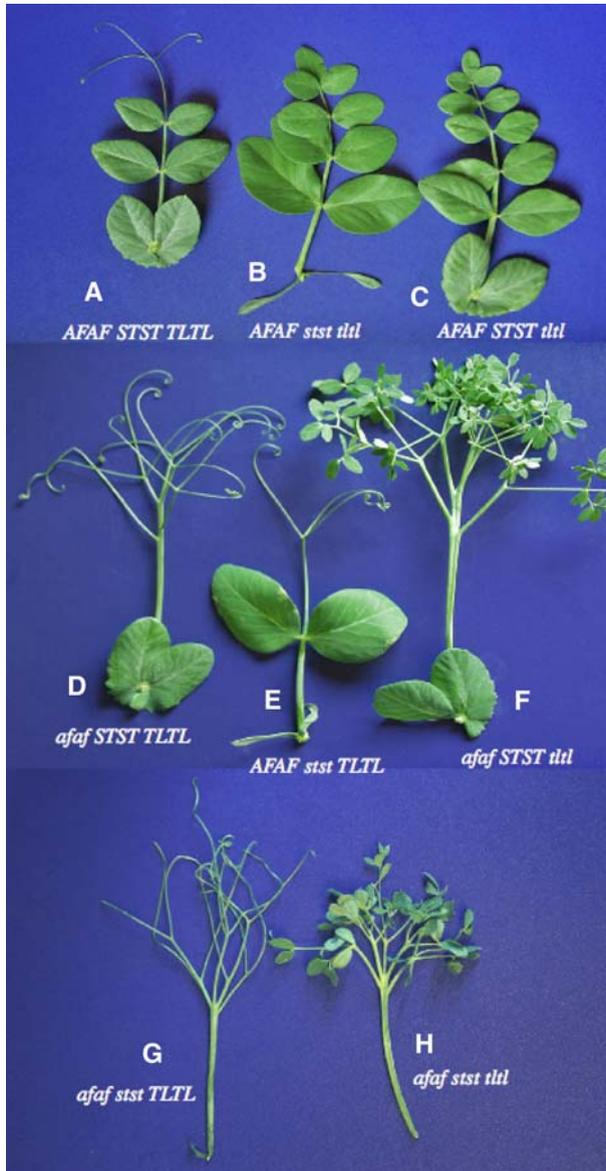


Fig. 1 Phenotypes of leaves of pea plants differing in their genotypic constitution at three loci, *af* (afila), *st* (stipules reduced), and *tl* (tendrill-less). Although we do not know the mechanistic details of how the allelic changes produce the final phenotypes, this is one of few available examples of an outline of the $G \Rightarrow P$ map for high-level phenotypes in eukaryotes. (Photo courtesy of Karl Niklas, Cornell University)

enzymes, like the isocitrate dehydrogenase and the *lac* repressor and operator in the bacterium *Escherichia coli*, and see how a variety of one-step mutations from the wild type (or, even better, from hypothetical ancestors) influence protein function.

Studying the $G \geq P$ map, and therefore the transition between genotypic and phenotypic adaptive landscapes is much more difficult when it comes to high-level phenotypes expressed by complex organisms. Yet, even here there are intriguing examples. Work by Karl Niklas, Darleen DeMason and Edward Cobb at Cornell (unpublished material) has compared the development of known mutants of peas for which the allelic composition at three developmentally crucial loci is known (Fig. 1). It is clear that one can, in this case, make a direct connection between the allelic make up (genetic architecture) at those loci and the resulting phenotypes, despite the fact that the myriad epigenetic effects that must fill the gap between the G and the P level are, largely, unknown. Note that all these phenotypes are viable, and that some resemble phenotypes present in other species (especially the transformation of leaves into tendrils); when this sort of data will be available within a phylogenetic context we will have some idea of how historically the $G \Rightarrow P$ map makes possible (or limits) the exploration of ecologically useful phenotypes, given a particular genetic architecture.

The Modern Synthesis (MS) in evolutionary biology took place during the 1930s and '40s. A number of researchers now think that it is time for an Extended Evolutionary Synthesis (EES: Pigliucci 2007), the foundations of which will be provided by a series of key concepts with the aim of addressing several questions left unanswered by the MS. Evolvability (Pigliucci 2008) will certainly be one of the conceptual pillars of the EES, and how genetic architecture, developmental constraints and natural selection interact to make possible or limit the evolvability of lineages will be one of the crucial questions. It is, essentially, the question that motivated Wright and worried Darwin, although we are now actually in a position to see the outline of an answer.

Acknowledgments I wish to thank Jonathan Kaplan for inviting me to write this paper, as well as for countless insightful discussions of evolutionary biological theory. Thanks also to Sergey Gavrillets for his patient explanations of his ideas about holey adaptive landscapes. I regret that I never got to meet Sewall Wright, he died (at the age of 99) just before attending a genetic conference where I gave one of my first papers. He was scheduled to talk about adaptive landscapes, naturally.

References

- Abrams PA, Harada Y, Matsuda H (1993) On the relationship between quantitative genetic and ESS models. *Evolution* 47:982–985
- Alberch P (1991) From genes to phenotype: dynamical systems and evolvability. *Genetica* 84:5–11
- Bowler PJ (1983) *The eclipse of Darwinism*. Johns Hopkins University Press, Baltimore
- Charlesworth B (1990) Optimization models, quantitative genetics, and mutation. *Evolution* 44:520–538
- Conrad M (1990) The geometry of evolution. *BioSystems* 24:61–81
- Coyne JA, Barton NH, Turelli M (1997) A critique of Sewall Wright's shifting balance theory of evolution. *Evolution* 51:643–671
- Crow JF (2002) Here's to Fisher, additive genetic variance, and the fundamental theorem of natural selection. *Evolution* 56:1313–1316

- Darwin C (1859/1910) *The origin of species by means of natural selection: or, the preservation of favored races in the struggle for life*. A.L. Burt, New York
- den Boer PJ (1999) Natural selection, or the non-survival of the non-fit. *Acta Biotheor* 47:83–97
- Eldredge N, Gould SJ (1972) Punctuated equilibria: an alternative to phyletic gradualism. In: Schopf TJM (ed) *Models in paleobiology*. Freeman, Cooper and Co, San Francisco, pp 82–115
- Eldredge N, Thompson JN, Brakefield PM, Gavrillets S, Jablonski D, Jackson JBC, Lenski RE, Lieberman BS, McPeck MA, Miller W (2005) The dynamics of evolutionary stasis. *Paleobiology* 31(2):133–145
- Futuyma D (1998) *Evolutionary biology*. Sinauer, Sunderland
- Gavrillets S (1997) Evolution and speciation on holey adaptive landscapes. *Trends Ecol Evol* 12:307–312
- Gavrillets S (1999) A dynamical theory of speciation on holey adaptive landscapes. *Am Nat* 154:1–22
- Gould SJ (1996) *Full House: the spread of excellence from Plato to Darwin*. Harmony, New York
- Hadany L, Beker T (2003) Fitness-associated recombination on rugged adaptive landscapes. *J Evol Biol* 16:862–870
- Hartl DL, Clark AG (1989) *Principles of population genetics*. Sinauer, Sunderland
- Jones AG, Arnold SJ, Burger R (2004) Evolution and stability of the G-matrix on a landscape with a moving optimum. *Evolution* 58:1639–1654
- Kauffman SA, Levin S (1987) Towards a general theory of adaptive walks on rugged landscapes. *J Theor Biol* 128:11–45
- Kavanagh KD, Evans AR, Jernvall J (2007) Predicting evolutionary patterns of mammalian teeth from development. *Nature* 449:427–433
- Kramer EM, Donohue K (2006) Traversing the adaptive landscape in snapdragons. *Science* 313:924–925
- Maynard-Smith J (1982) *Evolution and the theory of games*. Cambridge University Press, Cambridge
- McShea DW (1991) Complexity and evolution: what everybody knows. *Biol Philos* 6:303–324
- Peck SL, Ellner SP, Gould F (1998) A spatially explicit stochastic model demonstrates the feasibility of Wright's shifting balance theory. *Evolution* 52:1834–1839
- Phillips PC (1993) Peak shifts and polymorphism during phase three of Wright's shifting balance process. *Evolution* 47:1733–1743
- Pigliucci M (2007) Do we need an extended evolutionary synthesis? *Evolution* 61(12):2743–2749
- Pigliucci M (2008) Is evolvability evolvable? *Nat Rev Genet* 9:75–82
- Pigliucci M, Kaplan J (2006) *Making sense of evolution: toward a coherent picture of evolutionary theory*. Chicago University Press, Chicago
- Poelwijk FJ, Kiviet DJ, Weinreich DM, Tans SJ (2007) Empirical fitness landscapes reveal accessible evolutionary paths. *Nature* 445:383–386
- Polly PD (2007) Development with a bite. *Nature* 449:413–415
- Price T, Turelli M, Slatkin M (1993) Peak shifts produced by correlated response to selection. *Evolution* 47:280–290
- Prusinkiewicz P, Erasmus Y, Lane B, Harder LD, Coen E (2007) Evolution and development of inflorescence architectures. *Science* 316:1452–1456
- Schemske DW, Bierzychudek P (2007) Spatial differentiation for flower color in the desert annual *Linanthus parryae*: was Wright right? *Evolution* 61(11):2528–2543
- Schluter D, Nychka D (1994) Exploring fitness surfaces. *Am Nat* 143:597–616
- Seger J, Stubblefield JW (1996) Optimization and adaptation. In: Rose MR, Lauder GV (eds) *Adaptation*. Academic Press, San Diego, pp 93–123
- Simpson GG (1944) *Tempo and mode in evolution*. Columbia University Press, New York
- Svensson E, Sinervo B (2000) Experimental excursions on adaptive landscapes: density-dependent selection on egg size. *Evolution* 54:1396–1403
- Van Valen L (1973) A new evolutionary law. *Evol Theory* 1:1–30
- Wade MJ, Goodnight CJ (1998) The theories of Fisher and Wright in the context of metapopulations: when nature does many small experiments. *Evolution* 52:1537–1553
- Wagner GP, Altenberg L (1996) Complex adaptations and the evolution of evolvability. *Evolution* 50:967–976
- Wagner A, Wagner GP, Similion P (1994) Epistasis can facilitate the evolution of reproductive isolation by peak shifts: a two-locus two-allele model. *Genetics* 138:533–545
- Ward D (1992) The role of satisficing in foraging theory. *Oikos* 63:312–317
- Whitlock MC (1995) Variance-induced peak shifts. *Evolution* 49:252–259
- Whitlock MC (1997) Founder effects and peak shifts without genetic drift: adaptive peak shifts occur easily when environments fluctuate slightly. *Evolution* 51:1044–1048

-
- Williams GC (1992) Optimization and related concepts. In: Natural selection: domains, levels and challenges. Cambridge University Press, New York, pp 56–71
- Wright S (1932) The roles of mutation, inbreeding, crossbreeding and selection in evolution. In: Proceedings of the sixth international congress of genetics, pp 356–366
- Wright S (1988) Surfaces of selective value revisited. *Am Nat* 131(1):115–123