



How Development May Direct Evolution

JUSTIN GARSON, LINTON WANG and SAHOTRA SARKAR

Section of Integrative Biology, and Department of Philosophy

1 University Station, C3500

University of Texas at Austin

Austin, TX 78712-1180

U.S.A.

E-mail: sarkar@mail.utexas.edu

Abstract. A framework is presented in which the role of developmental rules in phenotypic evolution can be studied for some simple situations. Using two different implicit models of development, characterized by different developmental maps from genotypes to phenotypes, it is shown by simulation that developmental rules and drift can result in directional phenotypic evolution without selection. For both models the simulations show that the critical parameter that drives the final phenotypic distribution is the cardinality of the set of genotypes that map to each phenotype. Details of the developmental map do not matter. If phenotypes are randomly assigned to genotypes, the last result can also be proved analytically.

Key words: adaptive landscape, adaptive walks, developmental evolutionary biology, evolution and development, evolutionary developmental biology, hypercube models

1. Introduction

Direction in phenotypic evolution is generally supposed to be produced by natural selection acting on blind variation or mutation. These mutations are “blind” in the sense that a mutation is not more or less likely to occur if it increases the fitness of the organism. However, it has long been suspected that developmental rules – physical patterns of embryological and subsequent assembly of adult organisms – may, by themselves (that is, without selection), account for some of the directionality of morphological evolution. Haldane (1932), in *The Causes of Evolution*, one of the founding documents of the so-called Modern Synthesis, observed that the steady and systematic degeneration of form seen in the fossil record for some groups such as the ammonites cannot be plausibly attributed entirely to selection. Developmental rules operating on (adaptively) blind variation are the most plausible candidates to explain such cases.¹

In spite of the occasionally spectacular advances of developmental biology during the last few decades, resulting in the many ongoing programs to

synthesize evolution and development (Gilbert et al. 1996; Hall 1998, 2000; Wagner 2000; Arthur 2002), no one knows the exact form or content of any of these developmental rules, or whether they even exist in any non-trivial sense (that is, they carry significant explanatory weight in explanations of features of development). Moreover, at present, it is not possible to model the entire developmental process for a single organism. Even models of single organs of model organisms remain rudimentary.

Nevertheless, Stadler et al. (2001) have recently proposed a promising framework for analyzing the interaction of evolution and development at least in the simplest of cases. Let G be the set of (possible) genotypes of a population; let P be the set of (possible) phenotypes. In this framework development is modeled using the developmental map, or d -map, $\mathbf{d}: G \rightarrow P$.² Obviously this is a gross over-simplification, at first sight appearing to be no more than expressing symbolically the trivial fact that development takes a genotype to a phenotype. Worse, it ignores the critical role of the environment, the fact that it is the history of developmental genotype-environment interactions that leads to the phenotypic features of an adult organism. Differences in these histories typically lead to different phenotypes from the same genotype, that is, to a display of phenotypic plasticity. In general the d -map is many-one, that is, for a single $p \in P$ there are many different $g \in G$ such that $\mathbf{d}: g \mapsto p$. (Phenotypic plasticity can be incorporated into this framework by either making the d -map explicitly dependent on environmental features, that is, $\mathbf{d}: G \times E \rightarrow P$, where E is the set of environmental features; or by making the d -map many-many. The former option has the advantage of including more relevant biological factors besides being more mathematically tractable. For more details of this framework, including a discussion of generalized norms of reaction, see Sarkar and Fuller [2003].)

Nevertheless, as Stadler et al. (2001) suggested, and as this paper will show, considerable insight about the possible role of developmental rules in evolution can be gleaned from the properties of d -maps. Different d -maps are different implicit models of development. They are “implicit” models in the sense that, while different d -maps represent different models of development, the details of those models of development are not explicitly represented in the maps. For instance, a d -map in which many different genotypes are mapped to the same phenotype can be a model of canalization. Nevertheless, it can also be a model of any other developmental process that has the same outcome as canalization. Such differences are not explicitly represented in the d -maps.

The main purpose of this paper is to introduce a rigorous framework for the construction of d -maps in which models of development incorporating different philosophical assumptions and conceptual resources can be

constructed. Two different haploid models are studied here. Genotypes are modeled as strings of n loci with 2 alleles at each locus. Many-one d -maps define a phenotype for each genotype. Evolution occurs through one-mutation steps. There is no selection in these models. Genotypic changes occur because of random mutation and drift across the one-mutation evolutionary landscape. In the models studied here, phenotypic change is tracked as genotypes evolve through single mutations. The number of such mutations or steps provides a measure of evolutionary time. This framework may be interpreted in two ways: (i) each string representing a genotype may be regarded as a population in which that genotype has become fixed by drift. Thus the sequence of phenotypic changes as a result of mutations would represent evolutionary change over a relatively long time-scale; or (ii), alternatively, the initial set of genotypes subject to mutation forms a population with a genotypic and a phenotypic distribution. Evolution takes place through mutation and is sequentially tracked for the entire population after each individual genotype has, on the average, undergone a mutation. Under this interpretation, the evolution of the phenotypic distribution is followed over a relatively shorter time-scale. In this paper the second interpretation will be used because it is more consonant with the usual modeling practices of theoretical population genetics. The central question asked in this paper is whether directional evolution, leading to systematic phenotypic change, can occur in the absence of selection because of the nature of the d -rules of development.

Such one-mutation evolutionary landscapes have the structure of an n -dimensional hypercube which may be represented as an undirected graph with 2^n vertices, each with degree n . They have been extensively studied in the past, though only in the context of adaptive walks (Kauffman and Levin 1987; Macken and Perelson 1987; Weinberger 1988; Sarkar 1990; Aggarwal et al. 1992; Aggarwal and Sarkar 1992). In those models, different genotypes were typically assigned different fitness ranks; the distribution of the lengths of adaptive walks and the degree of optimization of fitness that can be achieved were typically the parameters of interest. The assignment of fitness ranks and the assumption that adaptive walks go uphill in the fitness landscape formally convert the underlying graph structure of these models into that of a directed graph with sinks representing local maxima of fitness (Sarkar 1990). These models were mostly constructed to study adaptation through natural selection in haploid models or affinity maturation through clonal selection during the vertebrate immune response. Unfortunately, the present context is sufficiently different that very little of the formalism developed in these earlier models can be co-opted for the type of study reported here.

2. Framework

The general framework used here is essentially the same as that of Stadler et al. (2001) stripped of some unnecessary complexity. Let $G' \subseteq G$ be a sub-set of the set of (possible) genotypes of a population and let P be the set of phenotypes. If, for a particular phenotype, $p \in P$, $\forall g \in G'$, $\mathbf{d}: g \mapsto p$, that is each genotype in G' leads to the same phenotype $p \in P$, then G' will be called the *developmentally equivalent genotype set* (or *deg-set*) for that $p \in P$, and designated $D(p)$.³ The cardinality of $D(p)$, $|D(p)|$ turns out to be an evolutionarily important parameter in the results reported in this paper. In models with selection it also measures the number of ways it is possible to mutate a genotype without loss of fitness because the phenotype, p , remains unchanged. Thus there are many more different options to serve as the starting point for mutations that make the same selective difference. Thus $|D(p)|$ is one measure of evolvability. (For more on such selectionist models, see Stadler et al. [2001].)

For any phenotype $p \in P$, the boundary of the *deg-set* for p , designated $\partial D(p)$ will consist of those genotypes $g' \in G$ that are one-mutation neighbors of the genotypes belonging to $D(p)$. Let $p_i, p_j \in P$ and let $D(p_i)$ and $D(p_j)$ be their respective *deg-sets*. Then the cardinality of $\partial D(p_i) \cap D(p_j)$, that is, $|\partial D(p_i) \cap D(p_j)|$, is a measure of the likelihood of a one-mutation transition from p_i to p_j . The one-mutation accessibility likelihood from p_i to p_j , $\mathbf{L}^1(p_i, p_j)$ is defined by

$$\mathbf{L}^1(p_i, p_j) = \begin{cases} |\partial D(p_i) \cap D(p_j)| & i \neq j \\ \text{undefined} & i = j. \end{cases} \quad (1)$$

$\mathbf{L}^1(p_i, p_j)$ is undefined when $i = j$ because any phenotype is “infinitely” accessible from itself. More importantly, the one-mutation relative accessibility of p_j from p_i is defined by:

$$\mathbf{R}^1(p_i, p_j) = \frac{|\partial D(p_i) \cap D(p_j)|}{|\partial D(p_i)|} \quad (2)$$

In a sense $\mathbf{L}^1(p_i, p_j)$ gives an absolute measure of the accessibility of p_j from p_i . In contrast $\mathbf{R}^1(p_i, p_j)$ measures the relative accessibility of p_j from p_i when compared to the accessibility of some other $p_k \in P$ from p_i ; in the analysis below it plays a critically important role. Similarly, the k -mutation accessibility likelihood (that is accessibility after k mutations) of p_j from p_i is designated as $\mathbf{L}^k(p_i, p_j)$; the corresponding relative accessibility is $\mathbf{R}^k(p_i, p_j)$. For all k , the $\mathbf{R}^k(p_i, p_j)$ are probabilities (that is, they obey the axioms of the probability calculus – see Sarkar, Garson and Wang [in preparation]); in contrast the $\mathbf{L}^k(p_i, p_j)$ are not. This partly explains why the $\mathbf{R}^k(p_i, p_j)$ are

more useful than the $L^k(p_i, p_j)$ for analytic purposes even though the latter have, in a sense, a more natural interpretation.

For any $p \in P$, let the relative weight of its *deg*-set be defined by $\rho(p) = \frac{|D(p)|}{|G|}$, where $|G|$ is the cardinality of G . (Trivially, these are also probabilities.) This parameter provides a measure of the relative dominance of a phenotype, p , with respect to other phenotypes, among the genotypes. As the results below will show, in certain types of d -maps this turns out to be the critical parameter for phenotypic change.

3. Models

Models are constructed for a haploid genome with n loci and 2 alleles (represented by “0” and “1”). As noted in Section 1, this gives the one-mutation evolutionary landscape the structure of an n -dimensional hypercube which can be represented as an undirected graph with 2^n vertices, each with degree n . Thus each genotype has n one-mutant neighbors. In the simulations reported below, $n = 10$. Thus the number of genotypes, that is, the cardinality of G , $|G|$ is $2^{10} = 1024$. The set G will be given an order, called here the “lexical order”; this is necessary to characterize explicitly the different d -maps (see below). This order is given as follows: $G = \langle 1111111111, 1111111110, 1111111101, 1111111011, \dots, 1111111100, 1111111010, \dots, 0000000000 \rangle$.

The set of phenotypes consists of 11 strings, $P = \langle 1111111111, 1111111110, 1111111100, 1111111000, \dots, 1100000000, 1000000000, 0000000000 \rangle$. Once again, the order will be important. The fact that the phenotypes are being represented as these strings will not be explicitly used in any of the analyses of this paper (they could just have been symbolically distinguished). However, in Model II below, it is assumed that there is a linear response of the phenotype to the genotype: adding all the “1”s in a phenotype establishes its lexical location. (The assumption is the usual one made in quantitative genetics models with no interaction between alleles, that is, no epistasis and also, in diploid and polyploid models, no dominance. Moreover, this representation is useful because it immediately suggests that a continued transition along either direction of the [ordered] set P constitutes a directional phenotypic change.)

Different models are distinguished by having different d -maps:

Model I: The d -map, d_1 , assigns 512 genotypes at random to the first phenotype, p_1 ; 256 to the second phenotype, p_2 ; 128 to the third phenotype, p_3 ; and so on. In this way p_{10} and p_{11} , each get assigned a single random genotype. For this model there is a simple formula for the relative weights of a phenotype’s *deg*-set: $\forall i, i = 1 \dots 10, \rho(p_i) = \frac{1}{2^i}$; $\rho(p_{11}) = \frac{1}{2^{10}}$.

Model II: d -map, d_{II} , now assigns the first 512 genotypes by lexical order, from g_1 to g_{512} to p_1 ; the next 256 genotypes, from g_{513} to g_{768} to p_2 ; the next 128 genotypes, from g_{769} to g_{896} to p_3 ; and so on. Finally g_{1023} gets assigned to p_{10} and g_{1024} gets assigned to p_{11} . Once again, $\forall i, i = 1 \dots 10$, $\rho(p_i) = \frac{1}{2^i}$; $\rho(p_{11}) = \frac{1}{2^{10}}$.

4. Analytic results

For Model I, in which there is a random though unequal assignment of phenotypes to genotypes, two relatively trivial but, nevertheless, interesting analytic results can be obtained:

PROPOSITION 1. For all $p_i, p_j \in P$, the expected value of $\mathbf{R}^1(p_i, p_j)$, $\mathbf{E}_i[\mathbf{R}^1(p_i, p_j)] = \rho(p_j)$.

Proof: For all $p_j \in P$, the probability of finding any genotype in $D(p_j)$ is given by the ratio $\frac{|D(p_j)|}{|G|}$, that is, $\rho(p_j)$. Therefore, because of the random assignments of phenotypes to genotypes by $d_i: G \rightarrow P$, the expected number of genotypes that are in $\partial D(p_i)$ and also in $D(p_j)$ is given by $|\partial D(p_i)| \times \frac{|D(p_j)|}{|G|}$. Thus:

$$\begin{aligned} \mathbf{E}_i[\mathbf{R}^1(p_i, p_j)] &= \mathbf{E}_i\left[\frac{|\partial D(p_i) \cap D(p_j)|}{|\partial D(p_i)|}\right] \\ &= |\partial D(p_i)| \times \frac{|D(p_j)|}{|G|} \times \frac{1}{|\partial D(p_i)|} \\ &= \rho(p_j). \end{aligned}$$

PROPOSITION 2. For all $p_i, p_j \in P$, and for all positive integers, k , then the expected value of $\mathbf{R}^k(p_i, p_j)$, $\mathbf{E}_i[\mathbf{R}^k(p_i, p_j)] = \rho(p_j)$.

Proof: For all $p_i, p_j \in P$,

$$\begin{aligned} \mathbf{E}_i[\mathbf{R}^k(p_i, p_j)] &= \mathbf{E}_i\left[\sum_{l=1}^{|P|} \mathbf{R}^{k-1}(p_i, p_l)\right] \mathbf{E}_i[\mathbf{R}^1(p_i, p_j)] \\ &= \rho(p_j) \mathbf{E}_i\left[\sum_{l=1}^{|P|} \mathbf{R}^{k-1}(p_i, p_l)\right] \quad (\text{by Proposition 1}) \\ &= \rho(p_j) \times 1 \quad (\text{because the } \mathbf{R}^{k-1} \text{ are probabilities}) \\ &= \rho(p_j). \end{aligned}$$

Together these two results underscore the remark made in Section 2 that the cardinality of the *deg*-set of a phenotype is critical to determining how prevalent it becomes during evolution in the absence of selection. The random nature of the map, $d_1: G \rightarrow P$, is critical to the proof of Proposition 1. Consequently this proof does not carry over to Model II.

5. Simulation results

The computer program used for the simulations was written in Microsoft Visual C++; random numbers were generated using its internal “rand” function. The program allows the user to specify the number of loci n , to be used in a simulation. Here, $n = 10$. Each simulation consisted of 10 000 runs, each starting from an initial genotype that is mutated at each step. Thus, there is an initial genotypic distribution of 10 000 genotypes and a corresponding distribution of phenotypes in the population. The resulting distribution of phenotypes and other parameters are tracked for differing numbers of mutations (m). Figures 1–2 represent the results obtained from Model I; Figures 3–4 represent the results from Model II. Both models were run using two different initial distributions of genotypes; in order to facilitate model comparison, exactly the same two initial genotypic distributions were used for the two models.

Figure 1a shows the distribution of phenotypes when 10 000 initial genotypes were selected at random from the 1024 possibilities using the d -map of Model I (d_1) which assigns phenotypes at random to each genotype. Each of these genotypes is the starting-point of a walk on the one-mutation landscape. This is the first initial distribution of phenotypes explored with this model. Figure 1b shows the distribution of the resultant phenotypes after 10 mutations.

Figure 2a shows the distribution of phenotypes when 10 000 initial genotypes were chosen from what are lexically the bottom 125 genotypes (using the lexical order defined in Section 3), once again using d_1 . This constitutes the second initial distribution of phenotypes studied with Model I. Figures 2b, c, d, e, and f show the distribution of phenotypes attained when $m = 1, 2, 5, 10,$ and 20 , respectively.

The results represented by Figure 1a are as expected. The initial phenotypic distribution mimics the distribution of phenotypes across the genotypes reflecting the cardinalities of the *deg*-sets. The differences between Figures 1a and 1b do not appear to be significant. There has been very little distance traveled in phenotypic space. In Figure 2a, that there should be no phenotype, $p_7, p_8, p_9,$ and p_{11} is expected, but that there are 82 of p_{10} is worth notice. It appears to be a fluctuation. Figures 2b–2f show a convergence to the distri-

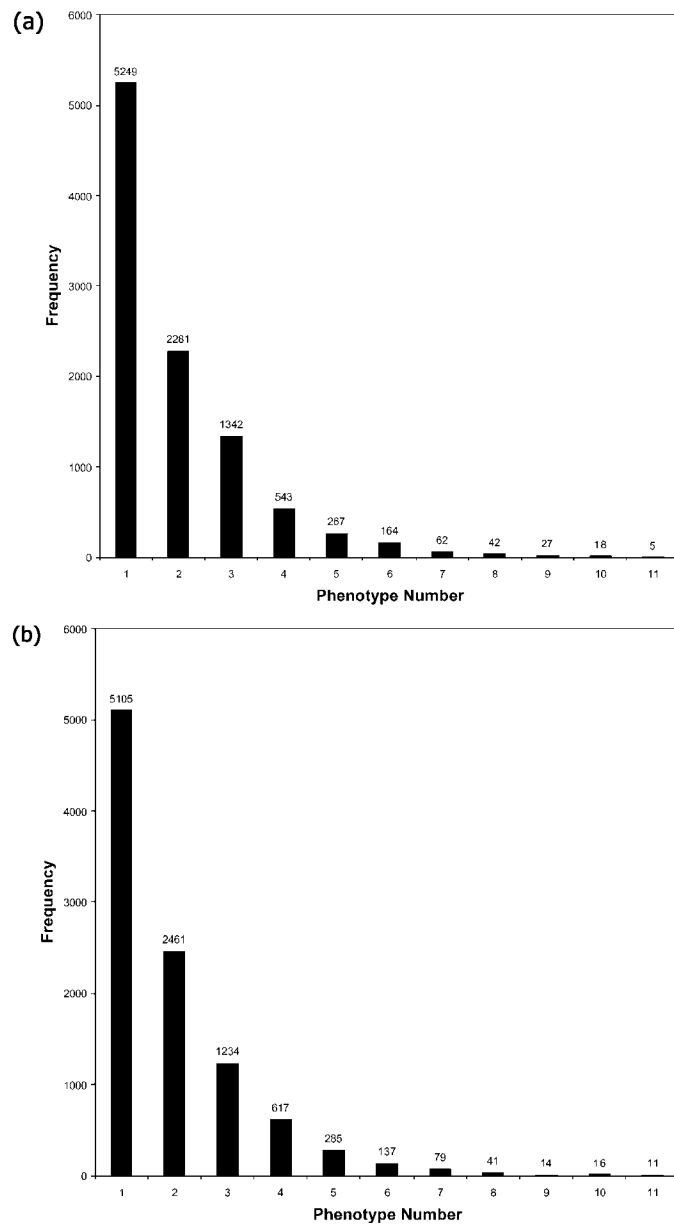


Figure 1. Phenotypic evolution with random assignment of genotypes sequentially to phenotypes and random initial genotypic distribution (**Model I**): (a) initial distribution of phenotypes; (b) distribution of phenotypes after 10 mutations.

bution of phenotypes across the genotypes reflecting the cardinalities of the *deg*-sets. These figures, along with Figure 1b, are consistent with the analytic results obtained in Section 4.

Figures 3a–b correspond to Figures 1a–b, except that the *d*-map used is that of Model II (d_{II}) which uses the lexical order of Section 3 to assign phenotypes to genotypes. Similarly, Figures 4a–f correspond to Figures 2a–f for Model II. Once again, Figures 3a–b are as expected.

Figures 4a–f are much more interesting. There is a gradual convergence to the distribution of phenotypes across the genotypes reflecting the cardinalities of the *deg*-sets. However, the convergence does not appear to be monotonous; this will merit further quantitative analysis. Comparing Figures 2a–f to Figures 4a–f also shows that the differences between d_I and d_{II} make a difference as to how many mutations it takes before there is convergence to the final phenotypic distribution; the former clearly requires many fewer mutations than the latter though there has yet been no quantitative analysis of the difference.

6. Discussion

Figures 4a–f, which depict the most interesting results obtained from Model II, give clear evidence of the possibility of directional phenotypic change because of the form of the *d*-map with natural selection in the sense of the reproductive advantage *of a genotype* playing no role. The initial genotypic distribution that was used was skewed towards genotypes with lower-ranked phenotypes (with the number of “1”s in the phenotypic string measuring rank). With each mutation, the phenotypic distribution shifts to phenotypes of higher ranks. By about 20 mutations the phenotypic distribution appears to be converging to the distribution of phenotypes across the genotypes reflecting the cardinalities of the *deg*-sets. We expect that this will be the “equilibrium” distribution (with “equilibrium” defined using a statistical test – see Sarkar, Garson and Wang [in preparation]) though the simulations done so far do not establish that result definitively. The change in the phenotypic distribution mimics what is typically achieved through directional selection. Thus there appears to be a sense in which developmental rules, as modeled with *d*-maps of the sort used here, can achieve the same results as selection. Presumably more complex rules, modeled by complex *d*-maps, can produce even more interesting results. There is much that remains to be studied in the future.

The most striking result of the simulations is that, for both models, the critical parameter that can be used to predict the final phenotypic distributions is the cardinality of the *deg*-set. For Model I, because of the analytic results presented in Section 4, this is not surprising. Moreover, there is an analogy

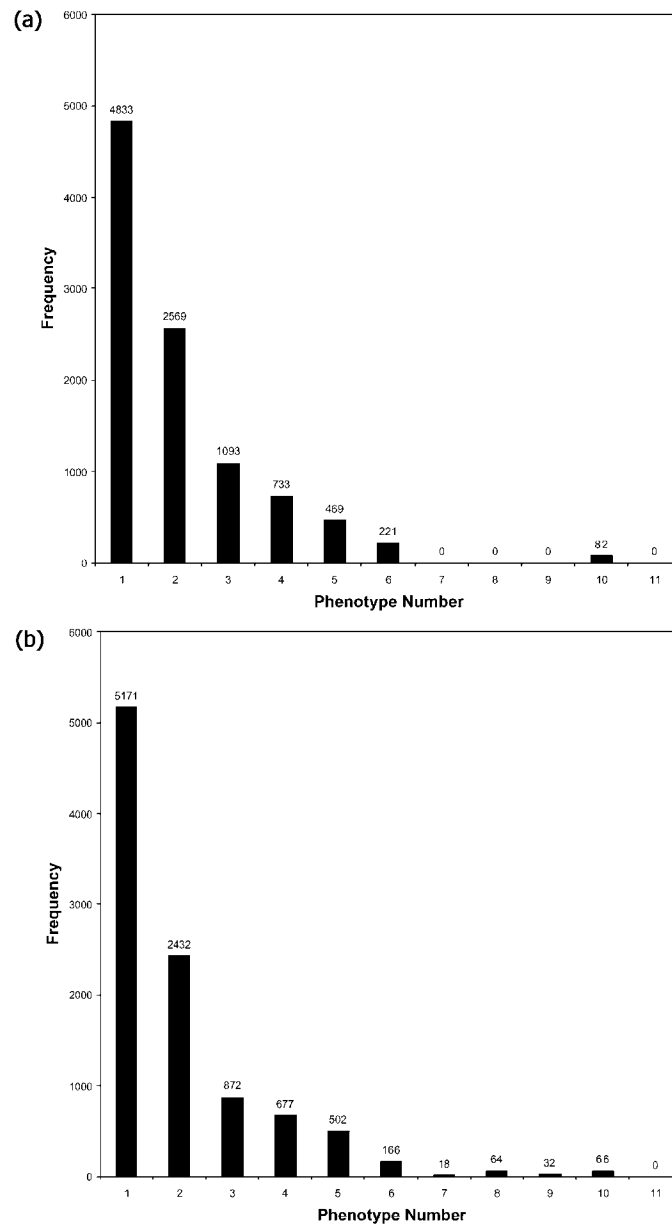


Figure 2. Phenotypic evolution with random assignment of genotypes sequentially to phenotypes and lower-rank initial genotypic distribution (**Model I**): **(a)** initial distribution of phenotypes; **(b)** distribution of phenotypes after 1 mutation, $m = 1$; **(c)** distribution of phenotypes after 2 mutations, $m = 2$; **(d)** distribution of phenotypes after 5 mutations, $m = 5$; **(e)** distribution of phenotypes after 10 mutations, $m = 10$; **(f)** distribution of phenotypes after 20 mutations, $m = 20$.





Figure 2. Continued.





Figure 2. Continued.

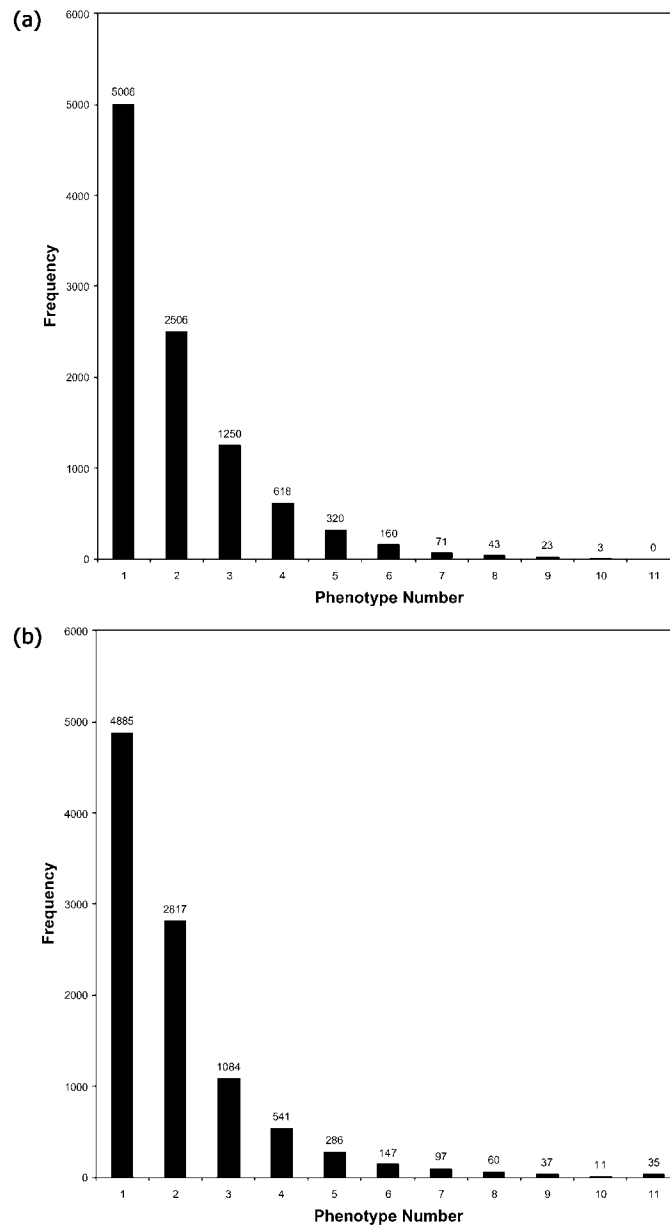


Figure 3. Phenotypic evolution with lexical assignment of genotypes sequentially to phenotypes and random initial genotypic distribution (**Model II**): (a) initial distribution of phenotypes; (b) distribution of phenotypes after 10 mutations, $m = 10$.

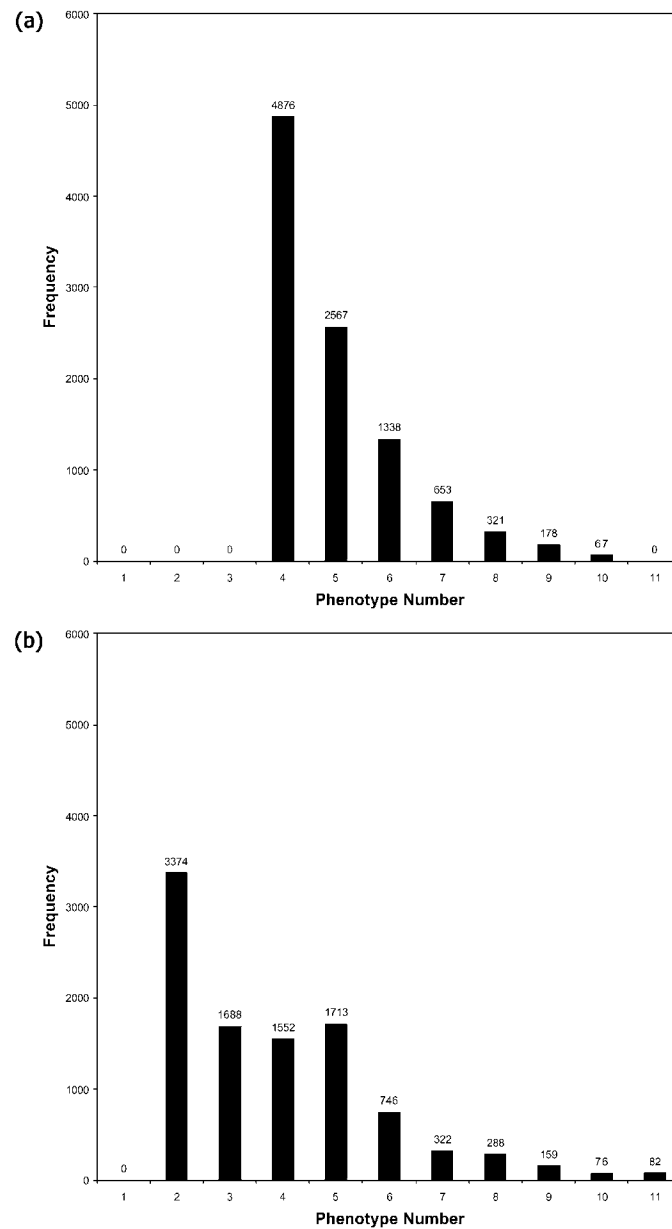


Figure 4. Phenotypic evolution with lexical assignment of genotypes sequentially to phenotypes and lower-rank initial genotypic distribution (**Model II**): (a) initial distribution of phenotypes; (b) distribution of phenotypes after 1 mutation, $m = 1$; (c) distribution of phenotypes after 2 mutations, $m = 2$; (d) distribution of phenotypes after 5 mutations, $m = 5$; (e) distribution of phenotypes after 10 mutations, $m = 10$; (f) distribution of phenotypes after 20 mutations, $m = 20$.





Figure 4. Continued.





Figure 4. Continued.

between the situation here and the evolution of neutral alleles: at least in one-locus models, the probability of fixation of an allele in a population is proportional to its initial frequency (Kimura 1983). However, the same result seems to hold for Model II for which we have no analytic results. An interesting question is how long it takes (measured by the number of mutations, m) to achieve this convergence to the final distribution. Even if it turns out that eventual convergence to the phenotypic distribution that reflects the cardinalities of the *deg*-sets holds for all, or at least most, d -maps, it is very likely that the time to convergence critically depends on the nature of the d -map (besides, obviously, the initial genotypic distribution). This is already some evidence of that because of the differences between Models I and II. Convergence is much more rapid for Model I than Model II.

In a sense, the results obtained here are not entirely unexpected. As philosophers have often emphasized though only in the context of selection (see Sober 1984), all mechanisms of evolutionary change act directly on the phenotype, not genotype. (This is exactly why we await a theory of developmental evolution where evolutionary trajectories are followed in phenotypic, rather than genotypic, space.) From the point of view of those mechanisms, it is irrelevant whether a distribution of phenotypes, so long as it is stably inherited (as is ensured through genetic transmission in the absence of phenotypic plasticity), arises because of selection or, as in this case, a certain sort of mutational bias incorporated into the d -maps explored here. The cardinality of the *deg*-set is the formal analog of fitness in our models. As is often the case in theoretical biology, it took simulations to reveal what turns out to be an unknown result that can subsequently be easily reconciled with conventional ideas. More troubling results will probably arise with more complex models, particularly when phenotypic plasticity is modeled.

Acknowledgements

This work was supported by NSF Grant No. SES-0090036 to SS. These results were presented to the Second Annual Workshop on Philosophy and Developmental Biology in Austin, TX in Spring, 2002. Comments by participants of that workshop were very helpful.

Notes

¹ For more on Haldane and the role of *The Causes of Evolution* in the “synthesis”, see Sarkar (2003).

² Stadler et al. (2001) call it the genotype-phenotype map. The more general terminology adopted here is designed to include environmental considerations (see below).

³ Stadler et al. (2001) call these sets “neutral sets”; that terminology is unfortunate in an evolutionary context because of the controversies, irrelevant for the present purpose, surrounding the neutral theory of molecular evolution.

References

- Aggarwal, A., Ma, W., Sandri, G. vH. and Sarkar, S.: 1992, ‘Adaptive Graph Computations with a Connection Machine’, in V.K. Prasanna and L.H. Canter (eds), *Proceedings of the Sixth International Parallel Processing Symposium*, IEEE Computer Society Press, Washington, pp. 18–21.
- Aggarwal, A. and Sarkar, S. 1992. ‘On the Evolution of Complex Genomes: Adaptive Graph Computations with a CM-2 Connection Machine’, in H. Berghel, E. Deaton, G. Hedrick, D. Roach and R. Wainwright (eds), *Applied Computing: Technological Challenges of the 1990s*, ACM Press, New York, pp. 798–807.
- Arthur, W.: 2002, ‘The Emerging Conceptual Framework of Evolutionary Developmental Biology’, *Nature* **415**, 757–763.
- Gilbert, S.F., Opitz, J.M. and Raff, R.A.: 1996, ‘Resynthesizing Evolutionary and Developmental Biology’, *Developmental Biology* **173**, 357–372.
- Haldane, J.B.S.: 1932, *The Causes of Evolution*, Harper and Brothers, London.
- Hall, B.K.: 1998, *Evolutionary Developmental Biology* (2nd edn.), Chapman and Hall, London.
- Hall, B.K.: 2000, ‘Eco-Devo or Devo-Evo—Does It Matter?’, *Evolution and Development* **2**, 177–178.
- Kauffman, S. and Levin, S.: 1987, ‘Towards a General Theory of Adaptive Walks on Rugged Landscapes’, *Journal of Theoretical Biology* **128**, 11–45.
- Kimura, M.: 1983, *The Neutral Theory of Molecular Evolution*, Cambridge University Press, Cambridge.
- Macken, C.A. and Perelson, A.S.: 1989, ‘Protein evolution on rugged landscapes’, *Proceedings of the National Academy of Sciences (USA)* **86**, 6191–6195.
- Sarkar, S.: 1990, ‘On Adaptation: A Reduction of the Kauffman-Levin Model to a Problem in Graph Theory and Its Consequences’, *Biology and Philosophy* **5**, 127–148.
- Sarkar, S.: 2003, ‘Evolutionary Theory in the 1920s: The Nature of the “Synthesis”’, *Philosophy of Science*.
- Sarkar, S. and Fuller, T.: 2003, ‘Generalized Norms of Reaction for Ecological Developmental Biology’, *Evolution and Development* **5**, 106–115.
- Sarkar, S., Garson, J. and Wang, L.: in preparation, ‘Developmental Maps and the Directionality of Evolution’.
- Sober, S.: 1984, *The Nature of Selection*, MIT Press, Cambridge, MA.
- Stadler, B.M.R., Stadler, P.F., Wagner, G.P. and Fontana, W.: 2001, ‘The Topology of the Possible: Formal Spaces Underlying Patterns of Evolutionary Change’, *Journal of Theoretical Biology* **213**, 241–274.
- Wagner, G.P.: 2000, ‘What Is the Promise of Developmental Evolution? Part I: Why Is Developmental Biology Necessary to Explain Evolutionary Innovations?’, *Journal of Experimental Zoology (Molecular Development and Evolution)* **288**, 95–98.
- Weinberger, E.: 1988, ‘A More Rigorous Derivation of Some Properties of Uncorrelated Fitness Landscapes’, *Journal of Theoretical Biology* **134**, 125–129.