Review

Genotype–phenotype mapping and the end of the ‘genes as blueprint’ metaphor

Massimo Pigliucci*

Department of Philosophy, City University of New York-Lehman, NY, USA

In a now classic paper published in 1991, Alberch introduced the concept of genotype–phenotype ($G \rightarrow P$) mapping to provide a framework for a more sophisticated discussion of the integration between genetics and developmental biology that was then available. The advent of evo-devo first and of the genomic era later would seem to have superseded talk of transitions in phenotypic space and the like, central to Alberch’s approach. On the contrary, this paper shows that recent empirical and theoretical advances have only sharpened the need for a different conceptual treatment of how phenotypes are produced. Old-fashioned metaphors like genetic blueprint and genetic programme are not only woefully inadequate but positively misleading about the nature of $G \rightarrow P$, and are being replaced by an algorithmic approach emerging from the study of a variety of actual $G \rightarrow P$ maps. These include RNA folding, protein function and the study of evolvable software. Some generalities are emerging from these disparate fields of analysis, and I suggest that the concept of ‘developmental encoding’ (as opposed to the classical one of genetic encoding) provides a promising computational–theoretical underpinning to coherently integrate ideas on evolvability, modularity and robustness and foster a fruitful framing of the $G \rightarrow P$ mapping problem.

Keywords: modularity; developmental encoding; evolvability;
genotype–phenotype mapping; robustness

1. INTRODUCTION: GENETIC BLUEPRINTS AND GENOTYPE–PHENOTYPE MAPPING

What is the relationship between genotypes and phenotypes? This question has marked the evolution of evolutionary theory ever since the rediscovery of Mendel’s work at the beginning of the twentieth century, which immediately generated an apparent conflict with the Darwinian view of gradual evolution (Mayr & Provine 1998). Famously, the answer proposed by the architects of the Modern Synthesis is that genes determine phenotypes, as in the oft-cited metaphors of a ‘genetic blueprint’ or a ‘genetic programme’ (for recent examples of usage, see Cuntz *et al.* 2008; Larsen 2008; Shoguchi *et al.* 2008; Papini-Terzi *et al.* 2009). This sort of answer bypasses the process of development, which is treated as an incidental blackbox with no direct causal relevance to the evolutionary process. Given this conceptual framework, it is no wonder that developmental biology was famously left out of the Modern Synthesis, and that it has (partially) re-emerged only recently within the so-called ‘evo-devo’ approach (e.g. Amundson 2005; Minelli & Fusco 2005; Love 2006; Newman *et al.* 2006; Carroll 2008).

In this paper, I will re-examine the question of the relationship between genotype and phenotype by going back to Alberch’s (1991) concept of a

genotype–phenotype ($G \rightarrow P$) ‘map’ and examine what recent research tells us on actual $G \rightarrow P$ maps. As we shall see, computational and empirical studies of three classes of systems (RNA folding, protein function and software development) are yielding important generalizations about the problem, as well as novel insight into the evolutionary process more broadly. One of the consequences of these new lines of research is that the blueprint metaphor is untenable and in fact positively misleading, and should be replaced by the concept of developmental encoding. This re-thinking of the genotype–phenotype relationship and its consequences in terms of the related concepts of robustness, modularity and evolvability (Wagner 2007) are part of an emerging Extended Synthesis in evolutionary biology (Pigliucci & Müller 2010).

Alberch (1991) set up the problem for the Modern-Synthesis-type view of the genotype–phenotype relationship by reminding biologists that genes do not specify development, and much less organismal form, but are instead one of several causal factors that are jointly determinant of the phenotype, with developmental events both being affected by, and in their turn affecting, genetic expression. Alberch (1991) then introduced a different metaphor from the standard blueprint view, the one of a $G \rightarrow P$ ‘mapping function,’ defined by a given parameter space and at least potentially amenable to mathematical description. Parameters defining the function would be developmental in nature, and of course their values would be affected by gene expression. Alberch

*massimo@platofootnote.org

One contribution of 12 to a Theme Issue ‘From polyphenism to complex metazoan life cycles’.

(1991) derived four general conclusions from his conceptualization of the G→P map: (i) the map is (much) more complex than a one-to-one relation between genotype and phenotype, which means that the same phenotype may be obtained from different combinations of genetic informational resources; (ii) the area in parameter space where a particular phenotype exists gives an indication of how stable (in reference to alterations of developmental parameters, and hence to both environmental and genetic perturbation) that phenotype is likely to be; (iii) the parameter space is marked by ‘transformational boundaries’, i.e. areas where a small change in one or more developmental parameters will cause the transition from one phenotypic state to another; and (iv) the phenotypic stability of a given population will depend on which area of the parameter space it occupies, and in particular whether it is close to a transformational boundary or not. In the first case, the population will show polymorphisms or polyphenisms, and in the second case it will be less phenotypically variable.

Alberch’s (1991) famous example of a phenotypic transition that is amenable to be described according to his idea of parameter space and mapping function was the evolution of the number of digits in amphibians. In particular, he showed how salamanders tend to lose their fifth toe every time the digit reduction evolves, while anurans tend to lose their first digit. The difference between the two groups can be recreated experimentally by administration of a mitotic inhibitor, a result that Alberch (1991) interpreted as telling us that anurans and salamanders find themselves in different areas of the parameter space, and in particular that they are located near different transitional boundaries, so that every time the transition happens within one of the two groups it occurs by the same developmental means, but when the two groups are compared the transitions happen by different developmental routes.

Alberch’s (1991) concept of G→P mapping is well known and yet has seldom been applied in the empirical literature and has suffered from coming across as a bit vague, with talk of ‘developmental parameters’ and admittedly undefined mathematical functions. The genomic era was initially taken to have spelled the definitive retirement of approaches like Alberch’s (1991), thereby marking the ultimate triumph of the genetic blueprint school of thought. As it happens, however, after the initial naively optimistic pronouncements about what biologists were going to be able to do after the first genomes would be sequenced, the complex reality of biological systems began to settle in and confronted investigators with the necessity for an increasing number of ‘omics’ fields (proteomics, metabolomics, even phenomics) to tackle precisely the sort of problem that had led Alberch (1991) to his formulation of the G→P mapping issue. As we shall see, once researchers were actually able to tackle real G→P maps, or at least aspects of them, Alberch’s (1991) intuitive list of general properties of evolution in phenotypic space turned out to be fundamentally correct and can now be expressed on a more firm and precise foundation.

2. G→P MAPPING OF SIMPLE BIOLOGICAL SYSTEMS: RNA FOLDING AND PROTEIN FUNCTION

A good starting point to tackle the G→P mapping problem is to start simple, and the simplest place to start is the growing literature on RNA folding (e.g. Fontana 2002; Cowperthwaite & Meyers 2007; Fernández & Solé 2007; Sumedha *et al.* 2007; Wroe *et al.* 2007; Stich *et al.* 2008; Takeuchi & Hogeweg 2008). RNA folding is relatively well understood at a chemical–physical level, with increasingly sophisticated computer models capable of predicting the three-dimensional folding of a linear sequence of nucleotides based on thermodynamic considerations. Moreover, it is relatively straightforward to verify such predictions experimentally for a subset of simulated folding patterns, and researchers can even carry out competition experiments among RNA molecules for a given catalytic function.

As far as the G→P problem is particularly concerned, the step from genotype to phenotype is in this case as short as it is possible in any biological system, and indeed probably somewhat reflects the ancestral situation in the RNA world hypothesized within the context of the origin of life problem (Ellington *et al.* 2009; Lincoln & Joyce 2009). RNA folding is therefore both an extremely suitable system to begin examining G→P mapping and one that may yield important clues to how historically mapping functions got started and became more complex and indirect. A crucial advantage of RNA folding studies of G→P mapping is that the fitness function is not assumed arbitrarily to follow a particular statistical distribution, but can be studied empirically. In other words, the connections between genotype and phenotype on one hand and between phenotype and fitness on the other hand are explicit, relatively simple and biologically meaningful.

Several important generalizations have emerged from studies of RNA folding, generalizations that are crucial to our understanding of phenotypic evolution beyond the relatively simple framework offered by the Modern Synthesis. Let us consider for instance the study of mutational networks, i.e. of the structure of the genotypic landscape in terms of one-mutation steps surrounding a given focal genotype. The idea goes back to Kauffman & Levin’s (1987) work on genotypic landscapes, itself inspired by Wright’s (1931) classic concept of adaptive landscapes (Pigliucci 2008a). The problem to be tackled is how does evolution explore phenotypic landscapes by moving across a corresponding genotypic landscape in a non-saltatory manner, according to standard Darwinian theory. The solution requires an understanding of the connection between the genotypic and phenotypic landscapes, and in the case of RNA folding one can actually computationally explore the totality of both landscapes for a given short-sequence length, or statistically sample the properties of landscapes defined by longer sequences.

For instance (Cowperthwaite & Meyers 2007), all 30-nucleotide long binary RNA molecules produce about one billion unique sequences, a bewildering genotypic space. This space, however, corresponds to

only 220 000 unique folding shapes in the G/U nucleotide landscape and a mere 1000 shapes in the A/U landscape, the two situations that have been extensively studied. This is a spectacular example of degeneracy (Stich *et al.* 2008), which in turn is a fundamental concept underlying the neutral theory of molecular evolution. Genotypes on these landscapes are connected by mutational networks whose properties can then be explored. An interesting result is that the distribution of phenotypes on RNA mutational networks follows regular patterns (reviewed in Cowperthwaite & Meyers 2007) characterized by a few abundant RNA shapes and a large number of rare ones. The structure of the landscape is such that evolution can explore most or all of the common structures by one-step mutations that preserve structure while moving the population on a neutral path of constant fitness, until it bumps into a novel phenotype with higher fitness (Fernández & Solé 2007; Sumedha *et al.* 2007). Interestingly, most genotypes turn out to be located within a few mutational steps from most of the common phenotypes in the landscape, making it predictable that such phenotypes will in fact be found by natural selection in a relatively short period of time. However, the connectivity on the landscape is always asymmetrical, which means that which particular phenotypes will be reached more easily while starting with a given genotype will be a matter of historical contingency.

Research on the general properties of RNA folding evolution has showed that the G→P function is such that small movements in genotypic space do not necessarily correspond to small movement in phenotypic space (Sumedha *et al.* 2007), a rather flagrant contradiction of one of the standard assumptions of the Modern Synthesis (Futuyma 1998). In particular, if we consider a genotype G with a given phenotype P, it is likely that G is connected to a one-step neighbour associated with a phenotype which is not structurally similar to P. This brings us to a rather surprising general behaviour that emerges from studies of RNA folding (as well as of protein function, micro-organisms and simulated systems), a true ‘punctuated equilibrium’ pattern of evolution that does not result from the usual suspects in terms of underlying causes (Minelli *et al.* 2009).

Punctuated equilibrium, of course, was one of the early challenges to the Modern Synthesis brought about by palaeontologists Eldredge & Gould (1972). The standard explanation for the fossil record pattern of stasis punctuated by occasional rapid shifts in phenotype is that of stabilizing selection (Estes & Arnold 2007). Simulations of RNA folding evolution display the same general pattern that one sees in the fossil record, obviously at a much smaller temporal scale (review and references in Cowperthwaite & Meyers 2007). The mechanism, however, has nothing to do with ‘stabilizing selection’ (a rather vague concept in itself, really simply a way to describe a statistical pattern of constant mean and reduced variance). Rather, the punctuated evolution results from the fact that the population divides itself into smaller chunks, each of which explores a portion of the largely neutral genotypic landscape. From time to time,

a population encounters a new phenotypic optimum and ‘jumps’ on it quickly. Stasis, in this context, is then not the result of selection for a constant phenotype, but rather of the largely neutral structure of the landscape, which allows populations to wander around until they find a new functional phenotype and jump into a nearby neutral network, only to resume their evolutionary wanderings.

RNA-like systems can also be a model for the evolution of ecological communities, thereby beginning to forge a still surprisingly lacking direct link between ecology and evolutionary biology. For instance, Takeuchi & Hogeweg (2008) showed that a population of replicators originally made of just one genotype evolves into a complex system characterized by four functionally distinct groups of genotypes, which the authors call ‘species’. Interestingly, the model also evolved ‘parasites’ which not only were able to coexist with catalytic molecules, but in turn were themselves catalysts for the evolution of further complexity in the system. While Takeuchi & Hogeweg’s (2008) definition of species in this context may appear artificial, the group of genotypes they identified are in fact both ecologically functionally distinct and genealogically related to each other, and a functional–genealogical concept is certainly one of the viable contenders as a definition of biological species (Pigliucci 2003).

Protein functions represent a more complex model of G→P mapping because of the larger number of building blocks than those characteristic of RNAs (amino acids versus nucleotides); yet, they are still amenable to both simulation and experimental studies, which confirm and expand much of what we have learnt from the RNA folding literature (Wroe *et al.* 2007). Again, we find that the genotype space is characterized by large areas of neutrality that facilitate evolvability, and again we find that continuous evolution at the genotype level yields occasional discontinuity at the phenotypic one. In particular, Wroe *et al.* (2007) have shown that new protein functions may arise through what they term the ‘promiscuity’ of existing proteins. This is a phenomenon by which the same protein can perform two functions because of its ability to alternate between different thermodynamically stable forms. Selection can then work to improve the secondary function even while the primary one is being retained, which means that gene duplication can fix the process, rather than initiate it, as in the standard model of molecular evolution (Conant & Wolfe 2008). This is a molecular-level example of West-Eberhard’s (2003) idea that sometimes genes are ‘followers’ in evolution, where the process is initiated by an environmentally induced phenotypic change without the necessity of a genetic change at the onset. Again, this is a significant departure from the accepted scenario within the Modern-Synthesis framework.

Studies of both RNA folding and protein function represent the best we can currently do in terms of direct (as opposed to statistical–quantitative genetic) empirical approaches to G→P mapping, because any other example of G→P map is simply too complex to tackle with current techniques. However, it is

possible to go a step further into the empirical exploration of the properties of $G \rightarrow P$ maps by way of a different approach, making complete use of the new genomic tools to focus not on the phenotypic effects of individual genes but on the properties of genetic networks.

3. NETWORKS, NOT JUST GENES

Several authors have begun to point out that we now have both the empirical and conceptual tools to finally move beyond the type of ‘bean-bag genetics’ that was famously (and controversially) characteristic of the Modern Synthesis, even leading to a famous exchange between two of the architects of the synthesis, Mayr (1963) and Haldane (1964). More constructively, the idea is that the standard population genetic approach by necessity treats genes as equivalent entities, while one of the major insights that is emerging from the study of the properties of genetic networks is that genes have very different evolutionary roles to play, depending on their position inside the networks of which they are a part (Stumpf *et al.* 2007; Chouard 2008; Stern & Orgogozo 2009).

For instance, consider the story of the *Hox*-related *Bicoid* gene in *Drosophila* (Chouard 2008). In the fruitfly, it is a crucial genetic resource to establish body shape, and in particular the segmentation into head, thorax and abdomen. As this segmentation is typical of all insects, and because of the widespread presence of *Hox*-like genes in arthropods, it has been naturally assumed that *Bicoid* would be found playing the same role across a phylogenetically broad range of species. Researchers were therefore stunned when it turned out that most groups of insects other than dipterans do not have the *Bicoid* gene at all (though other *Hox3*-homologues are present throughout the arthropods, often playing different roles)! For instance, in parasitic wasps and flour beetles, the specific role of *Bicoid* in *Drosophila* is compensated for through a series of minor adjustments in the rest of the genetic network of which *Bicoid* would be a part if it were present in these species. Similar insights about how genetic networks adjust themselves to compensate for the lack of an allegedly central gene can be gleaned from research on the *Arabidopsis Frigida* gene (involved in the control of flowering time), or *Drosophila*’s shavenbaby (which affects trichome cell formation; Stern & Orgogozo 2009).

According to a review by Stern & Orgogozo (2009), one major generalization that is emerging from studies of genetic variation in gene networks within and across species is the decidedly non-Modern-Synthesis concept that the genetic bases of interspecific differences tend to be statistically distinct from the genetic bases of within-species variation. In particular, epistatic variants and null alleles are often found within species, but contribute much less to differences among species; conversely, *cis*-regulatory changes are far more typical of interspecific differences than are mutations in the structural, protein-coding regions of those same genes. These findings of course do not mark a sharp boundary between micro- and macro-evolution. We are not talking about Goldschmidt (1940)-type ‘hopeful

monsters’. Nonetheless, criticism of the simplistic view of macro-evolution that has characterized the Modern Synthesis, and that did inspire Goldschmidt (1940) to write *The Material Basis of Evolutionary Change* is finally beginning to be vindicated by new knowledge about those very material bases.

Given what we are learning about gene action, it is becoming increasingly clear that understanding the $G \rightarrow P$ function has to include a focus not so much on what individual genes are doing, but on the emergent properties of gene networks (Chouard 2008). Promising lines of inquiry use naturally occurring genetic variation (Benfey & Mitchell-Olds 2008) and ‘reverse engineering’ (at the network level, Rockman 2008) to understand how exactly evolution transforms the genetic network of a given species into the genetic network of a closely related one, although the computational and methodological challenges are far from having been worked out to everyone’s satisfaction.

A spectacular case study comes from the work on evolution of gene networks regulating mating in ascomycete fungi, and particularly in yeast. Tsong *et al.* (2006) began by identifying a set of genes that are regulated by a repressor in *Saccharomyces cerevisiae* but by an activator in the related species *Candida albicans*, an instance in which evolution has somehow reversed the type of regulation of an entire gene network. This change, however, surprisingly did not alter the ‘logical output’ of the gene network, meaning that the phenotypic outcome is unchanged. Tsong *et al.* (2006) then examined a group of 16 closely phylogenetically related species to find out how evolution might have accomplished this sort of switch while maintaining fitness in the affected organisms. The authors identified several specific changes in regulatory elements (both *cis* and *trans*) that bridged the evolutionary gap between *S. cerevisiae* and *C. albicans*. They used comparative phylogenetic methods to map the changes in gene regulation on the available phylogeny for the group, providing the first example of a detailed reconstruction of the evolutionary rewiring of a genetic network. This in turn gives us an important glimpse into how $G \rightarrow P$ maps (or significant portions of them at any rate) change over long evolutionary times.

Despite spectacular advances on direct $G \rightarrow P$ maps like those represented by RNA folding and protein functionality, and insights gained by studying the emergent properties of gene/protein networks that represent a crucial component of organismal $G \rightarrow P$ mapping, we also need a more theoretically firm grasp of how the genotype–phenotype relationship is built and under what constraints it evolves. Such a theoretical understanding is emerging from the parallel field of computational science, both in cases in which it is directly inspired by biological problems and in those in which it is the independent result of research into more practical problems posed by software engineering.

4. COMPUTATIONAL APPROACHES TO $G \rightarrow P$ MAPPING: MODULARITY, ROBUSTNESS AND EVOLVABILITY

Discussions of the evolution of the $G \rightarrow P$ map cannot prescind from discussions of three related concepts:

modularity, robustness and evolvability. There is a significant degree of scepticism among some population geneticists about the meaningfulness of these concepts (Lynch 2007). Such scepticism, however, is misplaced (Pigliucci 2009), as these three ideas are foundational to our understanding of the evolution of genotype–phenotype relationships. Much of what follows is contributed by research that makes use of computational approaches, sometimes not explicitly meant to address biological questions, but highly germane to our problem nonetheless.

There are several partially overlapping concepts of evolvability (Pigliucci 2008b), but Crombach & Hogeweg (2008) helpfully defined it as the efficiency of an evolving system at finding beneficial mutations (or, more generally, as new phenotypes can be found by recombination (Szöllosi & Derényi 2008), beneficial solutions to a phenotypic search problem). Crombach & Hogeweg (2008) suggested that $G \rightarrow P$ functions in fact evolve to increase evolvability. This is made possible by the ‘hubs-and-spoke’ structure of most genetic networks discussed above, which leads to evolution spending most of the time in large neutral zones of genotypic space, and occasionally producing a rather rapid switch from one ‘basin of attraction’ to another. This is the computational equivalent of (and thereby in a sense a theoretical foundation for) the biological evolution of regulatory networks studied by Tsong *et al.* (2006) in yeast and as mentioned above.

The question is what the relationship is among evolvability and the other two properties of the $G \rightarrow P$ map (figure 1): modularity (which, at the genetic level, is the degree of interconnectedness among components of a gene network, related to Kauffman & Levin’s (1987) NK-systems, i.e. systems made of N parts with K connections) and robustness (the ability of a system to maintain functionality in the face of internal, i.e. genetic, and external, i.e. environmental perturbations—a concept akin to the older ideas of homeostasis and canalization). A helpful approach to this question is illustrated by the computational work of Gjuvsland *et al.* (2007), who simulated and explored the properties of a simple network constituted by three regulatory genes with one downstream element. They found that both robustness and the ability to accumulate hidden genetic variation (Bergman & Siegal 2003) are properties of the evolved $G \rightarrow P$ map. Robustness in turn arises as a result of the regulatory feedbacks that are typical of gene networks, which means that it can be at least indirectly favoured by natural selection.

Another important part of the modularity–robustness–evolvability puzzle is provided by Ciliberti *et al.* (2007), where they showed that robustness is in fact a precondition for evolvability (which they conceptualize as the ability to innovate). Their simulations show that there is a tradeoff between robustness and evolvability: the more robust a system is, the more it is by definition resistant to change, and therefore less likely to hit upon a phenotypic novelty. Then again, too little robustness (which would increase evolvability) runs into the problem of making the system too labile to either internal or external perturbation, which means that too little

robustness would be selected against. The solution is to be found in the fact that genetic networks with intermediate degrees of connectedness strike a balance between too much robustness and too little evolvability, and vice versa. As natural selection can act on the degree of connectedness of a gene network, this means that it can alter both robustness and evolvability.

For the solutions of Gjuvsland *et al.* (2007) and Ciliberti *et al.* (2007) to the evolution of modularity–robustness–evolvability to work, however, the evolutionary process has to take place within genotypic spaces that are largely neutral or quasi-neutral, because these permit the type of exploratory dynamics that are most conducive to both maintaining robustness and increasing evolvability. Interestingly, independent research on the intrinsic properties of highly dimensional genotypic landscapes has shown convincingly that these are (unlike the low-dimensional ones that have been treated by most population genetic models throughout the twentieth century) in fact characterized by large areas of neutral or quasi-neutral space (Gravner *et al.* 2007).

All the above casts a significant shadow on the whole idea that genes are in any meaningful sense the equivalent of standard computer programs, and even less blueprints for organisms. As we shall see subsequently, one more piece of the puzzle remains to be addressed, however, and the answer again seems to come from computational approaches aimed at exploiting the characteristics of living organisms to solve general classes of problems in software engineering.

5. WHY DEVELOPMENT? GENE NETWORKS, DEVELOPMENTAL ENCODING AND THE END OF THE BLUEPRINT METAPHOR

Notoriously, developmental biology was essentially left out of the Modern Synthesis of the 1940s that gave us the current structure of evolutionary theory (Mayr & Provine 1998). Part of the reason for this is that it has never been conceptually clear what exactly the role of development in evolution is. Mayr (1963) famously made a distinction—arching back to Aristotle—between proximate and ultimate causes in biology, with the genetic bases of phenotypes counting as proximate causes and the evolutionary processes that brought those phenotypes about considered as ultimate causes (Ariew 2003). Even if one accepts Mayr’s (1963) framework, it is not clear whether development should be considered a proximate or an ultimate cause.

The onset of evo-devo and calls for an Extended Synthesis in biology (Love 2006; Müller 2007; Pigliucci & Müller 2010) have reopened that question. The answer is emerging from research on the structure of $G \rightarrow P$ maps, and in particular from a parallel literature in computational science that attempts to exploit the characteristics of biological development to produce a new generation of ‘evolvable hardware’. The picture that is forming out of these efforts, as we shall see, is that development is a necessary link between proximate and ultimate causality, and that

Figure 1. A concept map summarizing the relationships among modularity, robustness and evolvability, as well as how natural selection, neutral genotypic spaces and the structure of genetic networks affect the evolution of these three fundamental characteristics of the G→P map (see text for details).

in a sense the G→P map is whatever specific type of 'developmental encoding' (as opposed to the classic genetic encoding) a given species of organism uses to produce environmentally apt phenotypes.

Several authors have pointed out the limitations of both direct genetic encoding of 'information' and of the blueprint metaphor that results from it. Ciliberti *et al.* (2007), for instance, have referred to human-engineered systems as being characterized by 'brittleness', i.e. the unfortunate property that if one component ceases functioning properly, there is a high probability that the whole system will unravel. This is most clearly not what happens with biological organisms, as we have seen above, which means that the oft-made analogy (ironically, by both some biologists and proposers of intelligent design creationism) between living organisms and 'machines' or 'programmes' is profoundly misleading. Along similar lines, Stanley (2007) reiterated that the amount of direct genetic information present in, say, the human genome (now estimated to be around 30 000 protein-coding genes) is orders of magnitude below what would be necessary to actually specify the spatial location, functionality and connectivity among the trillions of cells that make up a human brain. The answer must be in the local deployment of information that is possible through developmental processes, where the 'instructions' can be used in a way that is sensitive (and therefore capable of adjusting) to both the internal and external environments.

According to Hartmann *et al.* (2007), artificial development is increasingly being used to solve computational problems outside of biology by direct analogy with biological systems. The results indicate that replacing direct genetic encoding with indirect developmental encoding dramatically reduces the search space for evolutionary algorithms. Moreover, the resulting systems are less complex and yet more robust ('fault-tolerant' in engineering jargon) than those obtained by evolving standard genetic algorithms. Another way to put the point is that direct genetic encoding is limited by the fact that the length of the genetic string grows proportionally to the complexity of the phenotype, thereby quickly encountering severe limitations in search space (Roggen *et al.* 2007). With developmental encoding, instead, the evolving system can take advantage of a small number of genetic instructions mapping to a large number of phenotypic outcomes, because those outcomes are determined by the (local) interactions among parts of the system and by interactions of the system with the environment (which means, of course, that information is distributed among genes, developmental mechanisms and the external environment; Oyama *et al.* 2003).

Roggen *et al.* (2007) explained that developmental encoding, again by analogy with biological systems, is based on the deployment of two processes: a signalling phase where information is communicated locally within a given circuit and an expression phase where local cells/components of the circuit adopt a particular

functional stage depending on the signal that they have received. Simulations comparing the evolution of standard genetic systems of information encoding with systems based on developmental encoding clearly show that genetic systems reach a maximum level of fitness for low levels of complexity; at higher levels of complexity developmental encoding 'scales' much better, with developmental systems being capable of achieving high fitness more quickly and efficiently. Moreover, developmental encoding leads to the artificial evolution of systems that are both significantly more robust to internal disruptions and significantly more flexible in response to external environmental conditions than standard genetic systems. This is an interesting situation whereby a research area parallel to evolutionary biology, computational science, draws inspiration from the actual structure of biological systems and ends up providing a theoretical underpinning for why, in fact, those biological systems are structured the way they are.

It is also worth noting that developmental encoding leaves wide open the possibility of constructing a theoretical–computational model of the role of phenotypic plasticity in evolution (Pigliucci 2001; West-Eberhard 2003). While for reasons of space I have to leave any treatment of plasticity in the context of $G \rightarrow P$ mapping aside for the moment, the phenomenon represents the obvious link between developmental biology and a serious consideration of the role of the external environment in shaping organic evolution. In this view, developmental encoding can be thought of as the mechanistic link between genetics and ecology, thereby providing us with a first glimpse of a truly comprehensive formal theory of evolution.

6. CONCLUSION: FROM BEAN-BAG GENETICS TO DEVELOPMENT AS AN EVOLUTIONARY MECHANISM

The conceptual and mathematical foundations of evolutionary theory are evolving from a simple beginning as bean-bag genetics, Mayr's (1963) derogatory term for population genetics theory which Haldane (1964) felt compelled to defend, to a sophisticated 'patchwork' that draws from population genetics, quantitative genetics, bioinformatics and computational science. Medawar & Medawar (1983) famously said that 'genetics proposes, epigenetics disposes', where epigenetics here means the whole of developmental processes, a way to highlight that evolutionary theory finally needs a good conceptual understanding of development, and not just of genetics. As I have argued here, such a broadened theoretical framework cannot come from population genetics, but benefits from the input of computational research both on simple biological examples of $G \rightarrow P$ maps, such as those underlying RNA folding and protein function, and on broader issues such as the properties of large neutral networks in genotypic space and of developmental versus genetic-encoding systems.

In less than the two decades that have elapsed since Alberch's (1991) milestone paper that introduced the very concept of $G \rightarrow P$ mapping, we have achieved a spectacular number of conceptual insights into the

problem of how genotype and phenotype are related to each other, the most crucial of which are summarized in table 1; of course, much remains to be done. To begin with, the theoretical insights I have been discussing come from what I referred to as a patchwork of disciplines and approaches, some even outside of biology. What is needed is a more organic mathematical–theoretical framework that maintains the valuable contributions of classical population genetics while at the same time expanding the horizons of theoretical biology to include the new computational approaches. One promising development in this direction is provided by the concept of 'holey' adaptive landscapes (Gavrilets 1999), which takes seriously that biologically relevant landscapes are highly multidimensional and uses a combination of analytical and computational approaches to characterize their properties. The results of this line of inquiry are rather encouragingly converging with the so-far parallel literature on robustness and evolvability (Ciliberti *et al.* 2007; Draghi & Wagner 2008; figure 1).

One factor that may come to complicate things further than they already are is the possibility that $G \rightarrow P$ maps of different types of biological systems may turn out to have different general (i.e. not just local) properties. For instance, Ciliberti *et al.* (2007) discussed the property of 'long memory' that is characteristic of complex gene networks. This is the fact that often changing a genotype at random will produce very similar phenotypes, with the system having a sort of 'memory' of past phenotypes. This, as Ciliberti *et al.* (2007) pointed out, is not the case for RNA folding $G \rightarrow P$ maps, where few changes can essentially randomize the structure of the molecule (the phenotype). Similarly, Sumedha *et al.* (2007) discussed a property of RNA neutral networks they call 'shape-space covering'. This consists of the fact that most structures in phenotype space can be found within a relatively small 'ball' of genotypic space. Shape-space covering, as the authors argue, applies to RNA networks but not to proteins or transcriptional regulatory networks, again pointing to the very real possibility that while we may achieve some generalizations that hold for a given class of $G \rightarrow P$ maps, not all those generalizations can automatically be assumed to hold for all classes of maps.

It should be obvious even from a limited discussion such as this one that a deficiency of the field of $G \rightarrow P$ mapping studies is that the ratio of theoretical to empirical results is fairly high. While this is indeed typical of other sciences (think physics), it is rather unusual in biology, and it is a situation that biologists may feel needs remediation. The problem, of course, is that we are already running out of experimentally approachable systems. The reason that research on RNA folding maps is so relatively advanced is because it is the simplest actual example of $G \rightarrow P$ map; next comes protein function, which is accordingly the second area of publication in this respect. Beyond that, we have been making good progress studying the properties of partial $G \rightarrow P$ maps, such as those of gene regulatory networks. But a truly satisfactory empirical understanding of $G \rightarrow P$ relations in complex organisms may be forever beyond our grasp because of practical epistemic limitations.

Table 1. Summary of important conceptual advances in the study of $G \rightarrow P$ mapping. References are meant to be useful for further readings, and they are by far not exhaustive.

research area	conceptual advance	references
computational science	modularity, robustness and evolvability are tightly related, with an intermediate degree of modularity allowing for a compromise of high robustness and high evolvability (the two are otherwise negatively related)	Ciliberti <i>et al.</i> (2007), Gjuvsland <i>et al.</i> (2007) and Crombach & Hogeweg (2008)
computational science	$G \rightarrow P$ functions evolve to increase evolvability, which is made possible by the 'hub-and-spoke' structure of genetic networks	Crombach & Hogeweg (2008)
computational science	robustness arises because of the general properties of feedback systems that are built into gene networks; this means that natural selection can indirectly affect robustness	Gjuvsland <i>et al.</i> (2007)
computational science	the large neutral or quasi-neutral areas of genotypic space that are necessary for the evolution of robustness/evolvability are intrinsic properties of highly dimensional genotypic spaces	Gravner <i>et al.</i> (2007)
computational science	artificial development, used to generate 'evolvable hardware,' provides a good model and theoretical foundation for why developmental encoding evolved on top of genetic encoding	Hartmann <i>et al.</i> (2007) and Roggen <i>et al.</i> (2007)
gene networks	a gene's position in a network determines its evolutionary role; genes are not equivalent entities as in standard population genetic models	Stumpf <i>et al.</i> (2007), Chouard (2008) and Stern & Orgogozo (2009)
gene networks	evolution can turn a gene network into another by altering the regulation of individual components while leaving the 'logical output' unchanged	Tsong <i>et al.</i> (2006)
general theory	introduction of $G \rightarrow P$ mapping concept	Alberch (1991)
general theory	properties of mutational one-step networks determine evolutionary paths	Kauffman & Levin (1987)
general theory	$G \rightarrow P$ mapping results in punctuated equilibria patterns	Cowperthwaite & Meyers (2007)
general theory	large neutral areas of genotypic landscape allow for evolutionary search of novel phenotypes (evolvability)	Cowperthwaite & Meyers (2007), Fernández & Solé (2007), Sumedha <i>et al.</i> (2007) and Wroe <i>et al.</i> (2007)
general theory	developmental encoding allows for the evolution of complex phenotypes by decoupling the proportionality between length of genetic instructions and complexity of phenotypes; the evolutionary search space becomes smaller; in this sense, development is the $G \rightarrow P$ map	Ciliberti <i>et al.</i> (2007) and Stanley (2007)
RNA folding	RNA folding is an empirically tractable model of $G \rightarrow P$ mapping	Fontana (2002)
RNA folding	$G \rightarrow P$ mapping is characterized by high degrees of 'degeneracy' or genetic redundancy, where the same phenotype is generated by a variety of not necessarily similar genotypes	Stich <i>et al.</i> (2008)
RNA folding	connectivity in genotypic space is asymmetrical, so evolution depends to some extent on contingent events	Cowperthwaite & Meyers (2007)
RNA folding	genotypic one-step neighbours are often connected to phenotypes that are not structurally similar	Sumedha <i>et al.</i> (2007)
RNA folding	'species' and specialized ecological functions (including parasitism) evolve spontaneously as emergent properties of populations of replicators	Takeuchi & Hogeweg (2008)

Despite these problems, the undeniable progress we have made in understanding $G \rightarrow P$ maps, both empirically and theoretically, is such that one should hope that evolutionary biology has reached the point of forever being past simplistic ideas like genetic programmes and blueprints, embracing instead a more nuanced understanding of the complexity and variety of life.

REFERENCES

- Alberch, P. 1991 From genes to phenotype: dynamical systems and evolvability. *Genetica* **84**, 5–11. (doi:10.1007/BF00123979)
- Amundson, R. 2005 *The changing role of the embryo in evolutionary thought: roots of evo-devo*. Cambridge, UK: Cambridge University Press.
- Ariew, A. 2003 Ernst Mayr's 'ultimate/proximate' distinction reconsidered and reconstructed. *Biol. Phil.* **18**, 553–565. (doi:10.1023/A:1025565119032)
- Benfey, P. N. & Mitchell-Olds, T. 2008 From genotype to phenotype: systems biology meets natural variation. *Science* **320**, 495–497. (doi:10.1126/science.1153716)
- Bergman, A. & Siegal, M. L. 2003 Evolutionary capacitance as a general feature of complex gene networks. *Nature* **424**, 549–552. (doi:10.1038/nature01765)
- Carroll, S. B. 2008 Evo-devo and an expanding evolutionary synthesis: a genetic theory of morphological evolution. *Cell* **134**, 25–36. (doi:10.1016/j.cell.2008.06.030)

- Chouard, T. 2008 Darwin 200: beneath the surface. *Nature* **456**, 300–303. (doi:10.1038/456300a)
- Ciliberti, S., Martin, O. C. & Wagner, A. 2007 Innovation and robustness in complex regulatory gene networks. *Proc. Natl Acad. Sci. USA* **104**, 13 591–13 596. (doi:10.1073/pnas.0705396104)
- Conant, G. C. & Wolfe, K. H. 2008 Turning a hobby into a job: how duplicated genes find new functions. *Nat. Rev. Genet.* **9**, 938–950. (doi:10.1038/nrg2482)
- Cowperthwaite, M. C. & Meyers, L. A. 2007 How mutational networks shape evolution: lessons from RNA models. *Ann. Rev. Ecol. Evol. Syst.* **38**, 203–230. (doi:10.1146/annurev.ecolsys.38.091206.095507)
- Crombach, A. & Hogeweg, P. 2008 Evolution of evolvability in gene regulatory networks. *PLoS Comput. Biol.* **4**, e1000112. (doi:10.1371/journal.pcbi.1000112)
- Cuntz, H., Forstner, F., Haag, J. & Borst, A. 2008 The morphological identity of insect dendrites. *PLoS Comput. Biol.* **4**, e1000251. (doi:10.1371/journal.pcbi.1000251)
- Draghi, J. & Wagner, G. P. 2008 Evolution of evolvability in a developmental model. *Evolution* **62**, 301–315. (doi:10.1111/j.1558-5646.2007.00303.x)
- Eldredge, N. & Gould, S. J. 1972 Punctuated equilibria: an alternative to phyletic gradualism. In *Models in paleobiology* (ed. T. J. M. Schopf), pp. 82–115. San Francisco, CA: Freeman, Cooper & Co.
- Ellington, A. D., Chen, X., Robertson, M. & Syrett, A. 2009 Evolutionary origins and directed evolution of RNA. *Int. J. Biochem. Cell Biol.* **41**, 254–265. (doi:10.1016/j.biocel.2008.08.015)
- Estes, S. & Arnold, S. J. 2007 Resolving the paradox of stasis: models with stabilizing selection explain evolutionary divergence on all timescales. *Am. Nat.* **169**, 227–244. (doi:10.1086/510633)
- Fernández, P. & Solé, R. V. 2007 Neutral fitness landscapes in signalling networks. *J. R. Soc. Interface* **4**, 41–47. (doi:10.1098/rsif.2006.0152)
- Fontana, W. 2002 Modeling evo-devo with RNA. *BioEssays* **24**, 1164–1177. (doi:10.1002/bies.10190)
- Futuyma, D. J. 1998 In *Evolutionary biology*, 3rd edn. Sunderland, MA: Sinauer Associates.
- Gavrilets, S. 1999 A dynamical theory of speciation on holey adaptive landscapes. *Am. Nat.* **154**, 1–22. (doi:10.1086/303217)
- Gjrvsland, A. B., Plahte, E. & Omholt, S. W. 2007 Threshold-dominated regulation hides genetic variation in gene expression networks. *BMC Syst. Biol.* **1**, 57. (doi:10.1186/1752-0509-1-57)
- Goldschmidt, R. B. 1940 In *The material basis of evolution*. New Haven, CT: Yale University Press.
- Gravner, J., Pitman, D. & Gavrillets, S. 2007 Percolation on fitness landscapes: effects of correlation, phenotype, and incompatibilities. *J. Theor. Biol.* **248**, 627–645. (doi:10.1016/j.jtbi.2007.07.009)
- Haldane, J. B. 1964 A defense of beanbag genetics. *Perspect. Biol. Med.* **7**, 343–359.
- Hartmann, M., Haddow, P. C. & Lehre, P. K. 2007 The genotypic complexity of evolved fault-tolerant and noise-robust circuits. *Biosystems* **87**, 224–232. (doi:10.1016/j.biosystems.2006.09.017)
- Kauffman, S. & Levin, S. 1987 Towards a general theory of adaptive walks on rugged landscapes. *J. Theor. Biol.* **128**, 11–45. (doi:10.1016/S0022-5193(87)80029-2)
- Larsen, P. 2008 Genetic blueprint and environment—which has the bigger impact? *Gerontologist* **48**, 324.
- Lincoln, T. A. & Joyce, G. F. 2009 Self-sustained replication of an RNA enzyme. *Science* **323**, 1229–1232. (doi:10.1126/science.1167856)
- Love, A. C. 2006 Evolutionary morphology and evo-devo: hierarchy and novelty. *Theory Biosci.* **124**, 317–333. (doi:10.1016/j.thbio.2005.11.006)
- Lynch, M. 2007 The frailty of adaptive hypotheses for the origins of organismal complexity. *Proc. Natl Acad. Sci. USA* **104**(Suppl. 1), 8597–8604.
- Mayr, E. 1963 *Animal species and evolution*. Cambridge, MA: Harvard University Press.
- Mayr, E. & Provine, W. B. 1998 *The evolutionary synthesis*. Cambridge, MA: Harvard University Press.
- Medawar, P. B. & Medawar, J. S. 1983 *Aristotle to zoos: a philosophical dictionary of biology*. Cambridge, MA: Harvard University Press.
- Minelli, A. & Fusco, G. 2005 Conserved versus innovative features in animal body organization. *J. Exp. Zool. B (Mol. Dev. Evol.)* **304B**, 520–525. (doi:10.1002/jez.b.21044)
- Minelli, A., Chagas-Júnior, A. & Edgecombe, G. D. 2009 Saltational evolution of trunk segment number in centipedes. *Evol. Dev.* **11**, 318–322. (doi:10.1111/j.1525-142X.2009.00334.x)
- Müller, G. B. 2007 Evo-devo: extending the evolutionary synthesis. *Nat. Rev. Genet.* **8**, 943–949. (doi:10.1038/nrg2219)
- Newman, S. A., Forgacs, G. & Müller, G. B. 2006 Before programs: the physical origination of multicellular forms. *Int. J. Dev. Biol.* **50**, 289–299. (doi:10.1387/ijdb.052049sn)
- Oyama, S., Griffiths, P. & Gray, R. (eds) 2003 *Cycles of contingency: developmental systems and evolution*. Cambridge, MA: MIT Press.
- Papini-Terzi, F. S. et al. 2009 Sugarcane genes associated with sucrose content. *BMC Genomics* **10**, 120. (doi:10.1186/1471-2164-10-120)
- Pigliucci, M. 2001 *Phenotypic plasticity: beyond nature and nurture*. Baltimore, MD: Johns Hopkins University Press.
- Pigliucci, M. 2003 Species as family resemblance concepts: the (dis-)solution of the species problem? *BioEssays* **25**, 596–602. (doi:10.1002/bies.10284)
- Pigliucci, M. 2008a Sewall Wright's adaptive landscapes: 1932 vs. 1988. *Biol. Phil.* **23**, 591–603. (doi:10.1007/s10539-008-9124-z)
- Pigliucci, M. 2008b Is evolvability evolvable? *Nat. Rev. Genet.* **9**, 75–82. (doi:10.1038/nrg2278)
- Pigliucci, M. 2009 The proper role of population genetics in modern evolutionary theory. *Biol. Theory* **3**, 316–324. (doi:10.1162/biot.2008.3.4.316)
- Pigliucci, M. & Müller, G. B. (eds) 2010 *Evolution: the extended synthesis*. Cambridge, MA: MIT Press.
- Rockman, M. V. 2008 Reverse engineering the genotype–phenotype map with natural genetic variation. *Nature* **456**, 738–744. (doi:10.1038/nature07633)
- Roggen, D., Federici, D. & Floreano, D. 2007 Evolutionary morphogenesis for multi-cellular systems. *Genet. Program. Evolvable Mach.* **8**, 61–96. (doi:10.1007/s10710-006-9019-1)
- Shoguchi, E., Hamaguchi, M. & Satoh, N. 2008 Genome-wide network of regulatory genes for construction of a chordate embryo. *Dev. Biol.* **316**, 498–509. (doi:10.1016/j.ydbio.2008.01.009)
- Stanley, K. O. 2007 Compositional pattern producing networks: a novel abstraction of development. *Genet. Program. Evolvable Mach.* **8**, 131–162. (doi:10.1007/s10710-007-9028-8)
- Stern, D. L. & Orgogozo, V. 2009 Is genetic evolution predictable? *Science* **323**, 746–751. (doi:10.1126/science.1158997)
- Stich, M., Briones, C. & Manrubia, S. C. 2008 On the structural repertoire of pools of short, random RNA sequences. *J. Theor. Biol.* **252**, 750–763. (doi:10.1016/j.jtbi.2008.02.018)

- Stumpf, M. P., Kelly, W. P., Thorne, T. & Wiuf, C. 2007 Evolution at the system level: the natural history of protein interaction networks. *Trends Ecol. Evol.* **22**, 366–373. (doi:10.1016/j.tree.2007.04.004)
- Sumedha, Martin, O. C. & Wagner, A. 2007 New structural variation in evolutionary searches of RNA neutral networks. *BioSystems* **90**, 475–485. (doi:10.1016/j.biosystems.2006.11.007)
- Szöllosi, G. J. & Derényi, I. 2008 The effect of recombination on the neutral evolution of genetic robustness. *Math. Biosci.* **214**, 58–62. (doi:10.1016/j.mbs.2008.03.010)
- Takeuchi, N. & Hogeweg, P. 2008 Evolution of complexity in RNA-like replicator systems. *Biol. Direct.* **3**, 11. (doi:10.1186/1745-6150-3-11)
- Tsong, A. E., Tuch, B. B., Li, H. & Johnson, A. D. 2006 Evolution of alternative transcriptional circuits with identical logic. *Nature* **443**, 415–420. (doi:10.1038/nature05099)
- Wagner, A. 2007 *Robustness and evolvability in living systems*. Princeton, NJ: Princeton University Press.
- West-Eberhard, M. J. 2003 *Developmental plasticity and evolution*. New York, NY: Oxford University Press.
- Wright, S. 1931 Evolution in Mendelian populations. *Genetics* **16**, 97–159.
- Wroe, R., Chan, H. S. & Bornberg-Bauer, E. 2007 A structural model of latent evolutionary potentials underlying neutral networks in proteins. *HFSP J.* **1**, 79–87. (doi:10.2976/1.2739116)