Modularity in Evolution: Some Low-Level Questions *

Lee Altenberg
Information and Computer Sciences
University of Hawai‘i at Manoa
Honolulu, HI U.S.A.
Internet: altenber@hawaii.edu

Abstract

Intuitive notions about the advantages of modularity for evolvability run into the problem of how we parse the organism into traits. In order to resolve the "question of multiplicity", there needs to be a way to get the human observer out of the way, and define modularity in terms of physical processes. I will offer two candidate ideas towards this resolution:

- the dimensionality of phenotypic variation, and
- the causal screening off of phenotypic variables by other phenotypic variables.

With this framework, the evolutionary advantages that have been attributed to modularity do not derive from modularity *per se*. Rather, they require that there be an "alignment" between the spaces of phenotypic variation, and the selection gradients that are available to the organism. Modularity may facilitate such alignment, but it is not sufficient; the appropriate phenotype-fitness map in conjunction with the genotype-phenotype map is also necessary for evolvability.

1 The Question of Multiplicity

A good deal of work in recent years has shown that the structure of the genotype-phenotype map is of fundamental importance to the process of evolution. The variational properties

*Chapter 5 in *Modularity: Understanding the Development and Evolution of Complex Natural Systems*, Werner Callebaut and Diego Rasskin-Gutman, editors. MIT Press, 2005.

of the genotype-phenotype map—how genetic variation maps to phenotypic variation (Altenberg, 1994a, 1995; Wagner & Altenberg, 1996)—largely determine whether mutations and recombination can generate the sequence of phenotypes with increasing fitness that produce adaptation.

A most important property of the genotype-phenotype map is its *modularity*. The concepts of "modularity" and "module" are being employed now in novel contexts in the fields of genetics, behavior, and evolution. Their precise meaning has been fluid. "Modular" will be used for the current discussion to describe a genotype-phenotype map that can be decomposed (or nearly decomposed, Simon (1962, 1969)) into the product of independent genotype-phenotype maps of smaller dimension. The extreme example of modularity would be the idealized model of a genome in which each locus maps to one phenotypic trait. For the converse, the extreme example of non-modularity would be a genotype-phenotype map with uniform "universal pleiotropy" (Wright, 1968), in which every gene has an effect on every phenotypic variable. Real organisms, one could argue, have genotype-phenotype maps that range somewhere in between these extremes.

It may seem intuitively obvious why modularity in the genotype-phenotype map should benefit evolution: if genetic changes tend to map to changes in a small number of phenotypic traits, then the genome can respond to selection on those traits alone, independently of the rest of the phenotype, with a minimum of deleterious pleiotropic side effects. Hence modularity would enhance the ability of the genetic system to generate adaptive variants, which one can refer to as its "evolvability" (Altenberg, 1994a, 1995).

In a genotype-phenotype map with low modularity, where genes have high pleiotropy, a genetic change that produces adaptation in one character may be confounded by maladaptive changes it causes in other characters. To produce adaptive changes, a patchwork of just the right mutations among modifier genes may be necessary to cancel out their overlapping negative pleiotropic effects. Therefore, the problem of pleiotropy points to a solution through polygeny.

Two kinds of constraints may prevent such solutions from being found. First, such patchwork may be be impossible to produce from any combination of genetic changes, so that only an approximation to the optimal phenotype can evolve. In other words, the phenotypes that are possible may span a subspace that does not include the optimum. I refer to this as a "subspace constraint".

Second, the kinetics of mutation, recombination, and selection may make optimal combinations of genetic changes unreachable by evolutionary processes. If coordinated mutations at a number of loci are required in order to produce a fitness advantage, and the single or double mutations along the way are deleterious or neutral, it becomes very improbable that such multiple mutations will ever appear (Riedl, 1975, 1977, 1978; Kauffman & Levin, 1987; Kauffman & Weinberger, 1991; Weinberger, 1991). This is a generic result, notwith-

standing the complications of recombination and neutral networks (van Nimwegen *et al.*, 1999). So in cases where adaptation requires the coordinated change of multiple loci, there may be no selective pathway to reach those changes, and the phenotype can remain stuck at a suboptimal genotype, resulting in a condition called "frustration" in statistical mechanics (McKay *et al.*, 1982), or a "rugged fitness landscape" (Kauffman & Levin, 1987). I refer to this as a "kinetic constraint".

These two mechanisms—subspace and kinetic constraints—may prevent the simultaneous optimization of multiple phenotypic variables. A way to avoid these constraints would appear to be modularity, where genetic variation maps to small numbers of traits.

1.1 A Deconstruction of This Framework

While this explanation for the benefits of modularity may seem straightforward, a number of problems arise when we take a closer look. The advantage of modular genetic variation is seen to come from the small number of traits that are affected. By implication, this advantage is thus premised on the idea that selection tends to act on small numbers of traits alone.

What do we know about the nature of selection as it relates to numbers of traits? Here we find ourselves in a swamp, because the process by which traits are distinguished from one another is a human measurement process, dependent on the instruments and cognitive structures that we possess to parse the organism. How, for example, should we deal with a change in the size of an organism? Is this a change in all of the organism's measurements or, if allometric scaling relationships are maintained, in just one measurement? Is genetic variation modular if it causes just one part of an organism to change size, or if it causes the entire organism to change size? Suppose there were selection for sharper teeth. Would a genetic variant that made all teeth sharper be more or less modular than a genetic variant that made half of the teeth sharper, or just one pair of teeth, or just one tooth?

Consider another situation. Suppose that climate change has caused a simultaneous change in the optimal values of a number of organismal traits. Suppose further—just as a thought experiment—that genetic variation for some physiological variable happens to move many of these traits closer to their new optima. A gene with effects on these many traits would, under the common usage, be called pleiotropic and non-modular. And yet under this circumstance, such a gene, with the ability to move many phenotypic traits closer to their optima, would be an asset to the genome's evolvability.

Does the gene in our example "know" that it has several traits under directional selection rather than just one trait? Does the environment know that it has selected on several traits rather than just one trait? Where does the fact that several traits have been affected appear in the dynamics of this situation? The reality of what has occurred is that

1. there has been a change in climate, and

2. there is an allele that is now at a selective advantage under the new climate.

For "whom", then, are there multiple traits?

We see that when we try to apply our intuitive notion about the advantages of modularity, we run into the problem of how we parse the organism into traits. This is not a new problem—indeed, the problem of how to "carve nature at its joints" has been with us since Plato (c 370 BC, 262b30). Until this problem is resolved, we cannot say whether variation is modular or not.

In order to resolve the "question of multiplicity", there needs to be a way to get the human observer out of the way, and define modularity in terms of physical processes. I will offer two candidate ideas towards this resolution:

1. the dimensionality of phenotypic variation, and

2. the causal screening off of phenotypic variables by other phenotypic variables.

2 Description and Degrees of Freedom

When we say that a gene affects multiple traits, we mean that it changes multiple features of the organism that are measured independently of one another. Each trait constitutes a variable that can take on a variety of values, distinct from the values other traits may take. To represent all the traits simultaneously therefore requires a multidimensional space, which will be the Cartesian product of the space of values each variable can take on. So, references to multiple traits are equivalent to references to multidimensional spaces of descriptive variables.

Thus, if S_1 is the space of possible values for trait x_1, and so forth for x_2, et al., then an organism with trait values $(x_1, x_2, \ldots, x_n)$ corresponds to a point $\boldsymbol{x}$ in the multidimensional space $\mathcal{S} = S_1 \times S_2 \times \cdots \times S_n$.

While our multidimensional representation of the organism allots one degree of freedom for each trait, the critical question is whether these degrees of freedom have any physical reality as dimensions of variation in the organism, or dimensions of variation for selection. We can apply similar reasoning to the environment: a description of the environment can contain many variables, but we must also ask whether these variables correspond to physical dimensions of variability in the environment.

Let us return to our thought experiment about climate change. Suppose the genetic underpinnings are the sort that Waddington (1942) considered for the evolution of canalization, where a physiological adaptive response, involving many phenotypic variables, is

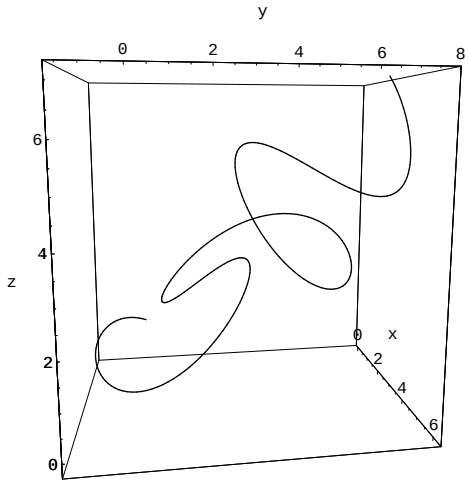

Figure 1: A one-dimensional space of variation embedded in three phenotypic dimensions.

cued both by environmental signals—day length, temperature, etc.—as well as internal signals under genetic control. Genetic changes in how these adaptations are invoked may be capable of moving the whole complex response toward a more optimal match to a changed environment (such as time of flowering, moulting, hybernation, dormancy, budding, quantities of stored metabolites, etc.). While many traits would be observed to change under such genetic variation, there may be in fact only one degree of freedom if there is a single cueing mechanism that is being altered. In contrast to the apparent high dimensionality of the space of traits affected by the gene, the space of variation in this example may be a one-dimensional space merely embedded in the higher dimensions. This is illustrated in Figure 1.

2.1 Embeddings and Dimension Reduction

To further illustrate the idea that low-dimensional variation may underlie what appears to be high-dimensional variation, I will draw attention to some recent work on dimension reduction. Dimension reduction has long been a part of morphometrics through the use of principal component analysis (PCA), but this technique assumes a linear form for the lower-dimensional subspaces. When the spaces of variation are nonlinear, other techniques are required to identify these spaces.

Two recent works provide algorithms that can take complex multidimensional data and discover when the variation is restricted to lower-dimensional manifolds, and can charac-

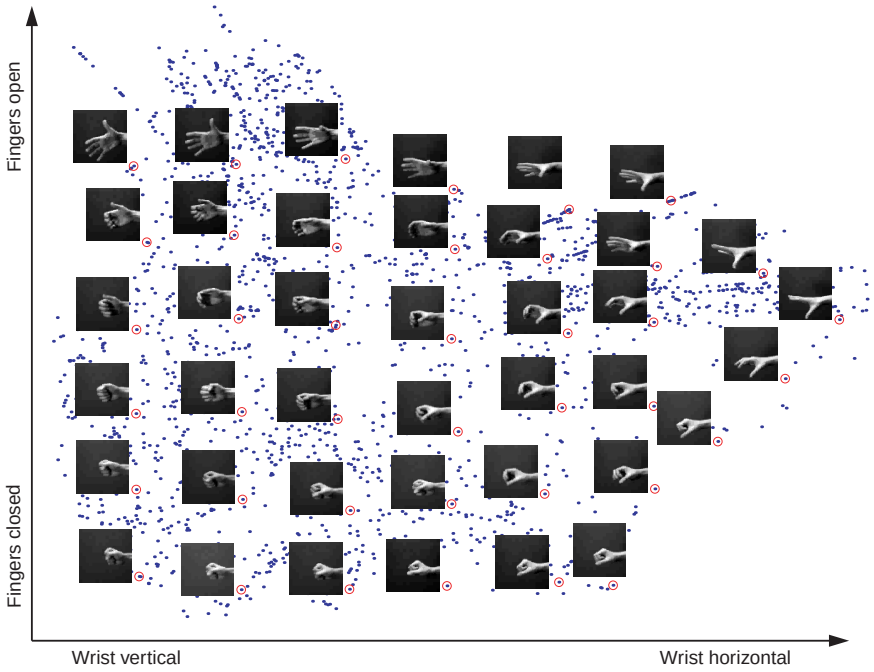

Figure 2: An example of dimension reduction. While the hand can be described by many variables, in this ensemble of states there are really only two dimensions of variability: wrist rotation and finger extension. The two dimensions of variation are recovered from the 4096-dimensional image data by Tenenbaum *et al.* (2000) using their Isomap algorithm. Reprinted with permission from Tenenbaum *et al.* (2000). Copyright 2000 American Association for the Advancement of Science.

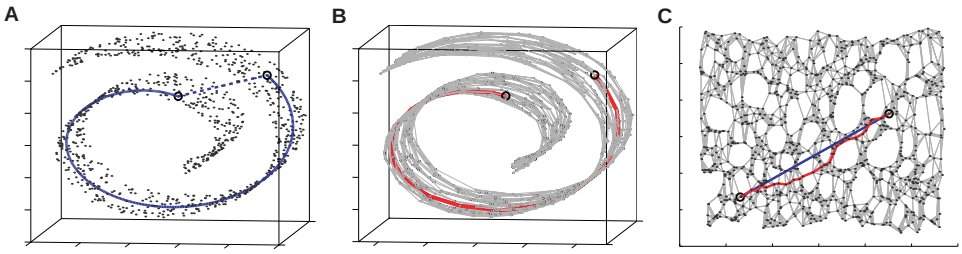

Figure 3: The nonlinear two-dimensional "Swiss Roll" manifold is recovered from its three-dimensional embedding by the Isomap algorithm of Tenenbaum *et al.* (2000). Points that appear close together in the three-dimensional embedding may be far apart in the underlying manifold. Reprinted with permission from Tenenbaum *et al.* (2000). Copyright 2000 American Association for the Advancement of Science.

terize these manifolds (Roweis & Saul, 2000; Tenenbaum *et al.*, 2000).

Two illustrations from Tenenbaum *et al.* (2000) are reproduced here. Figure 2 shows a hand rotating at the wrist or opening its fingers. The hand is a complex object, here described by 64-by-64 pixel photographs, giving 4096 independent variables. All 4096 variables vary as a result of the wrist rotation and finger extension. Yet, if each photograph is mapped to a single point in a 4096-dimensional space, the variation traced out by the set of photographs can be mapped to a two-dimensional manifold embedded in the 4096-dimensional space. This manifold is represented in the plane in Figure 2.

The two-dimensional manifold of variation depicted here is produced by movement of the hand. Instead of a hand moving, we could just as well analyze a set of photographs of hands that represent the morphological variation in extant human populations. The structure and dimensionality of this space of phenotypic variation may very well be revealed by use of the "isomap" (Tenenbaum *et al.*, 2000) or "locally linear embedding" (Roweis & Saul, 2000) methods of nonlinear dimension reduction.

Figure 3 shows data from a two-dimensional manifold that is curled up in a 3-D embedding. Points that appear close in the 3-D embedding may actually be far apart in the manifold, as shown by the geodesic lines. Thus, naive interpretations of the dimensionality and "distances" represented by phenotypic variation may not reflect the real structure of the variation.

The wide application of the isomap, locally linear embedding, and related nonlinear dimension-reduction methods (Verbeek *et al.*, 2002; Agrafiotis & Xu, 2002) to morphometrics might prove fruitful at exposing unknown structures within phenotypic variation. If low-dimensional manifolds are discovered amid the morphological variation found in different organisms, a new window may be opened on the question of "developmental con-

straints". The widespread characterizations of such manifolds across different taxa could provide the basis for a study of "morphomics", as it were (*a la* "genomics" and "proteomics"). To my knowledge, the application of these new methods to morphometrics has not yet been tried.

2.2 Evolvability and Alignment with Selection

I have argued that phenotypic variation which may appear to involve many variables may in fact represent the variation of very few parameters. Geometrical transformations of variables can change what appears to be high pleiotropy into low pleiotropy. We must ask, then, when does the geometry of variation make a difference to evolution? At this point we must consider how selection is involved.

In order for a subspace of phenotypic variation to allow a response to selection, it must pass through a selection gradient. Or, to be more precise, the space of variations must provide, with some reasonable probability, a sequence of genetic operations that produce monotonically increasing fitnesses. If the probability of such a sequence is too low, there is no evolvability. Two causes of such low probabilities are the phenotype being near its constrained optimum, and the situation where "frustration" prevails. Here we find ourselves back at the analysis of Riedl (1977).

Riedl proposes that the solution to the problem of adaptive frustration is the "systemization of the genome". By this he means the creation of new spaces of genetic variation that move the phenotype in directions that are under positive directional selection. Rather than modularity, it is the alignment of the space of variation with selective gradients that is the solution Riedl describes. What I am doing here is describing a geometric interpretation of Riedl's argument.

So, despite the fact that genetic variation may alter a number of phenotypic traits, if there is a selection gradient for that particular dimension or space of variation, then the genotype-phenotype map exhibits high evolvability. In the thought experiment about climate change and a gene that generates change that is adapted to it, I have tried to show that the involvement of multiple traits, per se, in genetic variation does not create the problem that modularity is postulated to solve. Rather, it is the relationship of *selective gradients* to the *space of variation* that is the critical issue.

How does the earlier idea, that a non-modular genotype-phenotype map produces frustration, hold up after this deconstruction? Frustration occurs because none of the spaces of phenotypic variability are able to provide, with sufficient probability, a sequence of genotypes that traverse the selective gradients that may happen to be present—i.e. they are not aligned with selection gradients. Hence, the genome is unable to access regions of the phenotype that may be adaptive. Modularity, if it is to be a means to attain evolvability, must

somehow imply an alignment between the spaces of phenotypic variation and the selection gradients. One can conclude either that:

1. Modularity is one means to such an alignment; or that
2. Modularity should be *defined* in terms of such alignments.

3 The Underlying Degrees of Freedom

I have talked about geometrical aspects of variation, and its relation to selection, without delving into the possible causes behind such properties. Here I will explore this issue a little further. I will propose that the notion of causal "screening off" (Salmon, 1971, 1984; Brandon, 1984, 1990) can be used to describe the sort of modularity in the genotype-phenotype map that matters to evolvability.

The fundamental dimensions of variation in the genotype are determined by the spectrum of genetic changes that can occur. These include:

- point mutation, in which one nucleotide is replaced by another nucleotide;
- deletions and insertions;
- gene duplication, in which a sequence of nucleotides copied from an existing sequence is inserted in a chromosome;
- gene conversion;
- polyploidy, in which an entire genome is duplicated one or more times;
- translocation;
- transposition;
- recombination;
- segregation and syngamy;
- methylation change;
- horizontal genetic transmission (e.g. plasmid exchange); and
- a variety of taxon-specific genetic mechanisms.

Each of these variation processes produces its own space of genetic variation (Stadler *et al.*, 2002), distinguished not so much by the nature of the phenotypic changes that they produce, as by the evolutionary paths they make possible between genotypes.

Let me be more concrete in describing the spaces of genetic variation. A genome of L nucleotides can be represented as a point in the genotype space $\mathcal{S} = \{A, T, C, G\}^L$ (ignoring for the sake of discussion the meta-sequence properties such as methylation, chromosome structure, etc.). Under the action of point mutation, there exist L degrees of freedom for the genotype. The magnitude of L varies from being on the order of 10^6 for prokaryotes to 10^{11} for lungfish and trumpet lilies. A million to a hundred billion is clearly a vast number of degrees of freedom for point mutations, but each degree of freedom constitutes only a miniscule space—comprising only four points in fact, the four nucleotide bases A, T, C, and G. It really makes no sense to even speak of directional selection or a selection "gradient" on a space of four discrete points; direction is undefined.

However, the process of gene expression groups these individual degrees of freedom into new spaces of variation with fewer degrees of freedom, but many more elements. To begin with, DNA triplets in transcribed sequences map to the space of amino acids, $\{A, T, C, G\}^3 \mapsto \mathcal{P}$ (where $\mathcal{P}$ includes the twenty amino acids and the stop codons).

The dynamics of protein folding and molecular interactions in turn group the amino acids in a protein into a new set of variables that characterize the protein and its interactions. This is the first point in this chain of "decoding" where real-valued variables enter, such as the geometry of the protein fold, the kinetic rates for interaction with other molecules, binding energies, catalytic rates, thermal stability, hydrophobicity, etc. Non-transcribed DNA has different mechanisms of expression, and real-valued variables can be seen to emerge immediately in characterizing its phenotypic effects, such as its affinities for binding with regulatory molecules, methylation enzymes, replication and transcription complexes, etc.

3.1 Screening Off

While a great many real-valued variables are needed to describe, for example, a protein, it is typically the case that only a small subset of variables is needed to describe the causal effects of a gene on the organism—a catalytic rate, binding constants, levels of expression, the timing of expression, half-life, and so forth. Variation in a gene will not cause phenotypic variation except in how it varies these variables. In other words, there is some set of variables that *screen off* (Salmon, 1971, 1984; Brandon, 1982, 1984, 1990, 2002) the causal impact of genetic variation: if one knows the values of these variables, there is nothing more one needs to know about the gene in order to determine its effect on the organism. By "small number" I mean small relative to the typical number of nucleotides in a gene,

which ranges from 10^3 to 10^5 in eukaryotes.

The processes of gene expression, ontogeny, and physiology convert the large number of essentially "digital" degrees of freedom in the genome into degrees of freedom of a smaller set of real-valued variables (similar, really, to what happens in electronic digital-to-analog conversion). The variables that screen off the properties of a gene may themselves be screened off by other variables that summarize their effects on other functions in the organism. For example, many factors contribute to levels of cortisol in vertebrates. But to the extent that they affect the organism through the action of cortisol, the cortisol level contains all the information about their effect.

In common usage, when people refer to the "function" of part of an organism, they may mean one of two things: "What does the part *do*?", or "What is the part *for*?" In both cases, nevertheless, the positing of a "what" implicitly uses the concept of screening off. The "what" refers to a function that screens off the detailed characteristics of this organismal part, a variable that summarizes one of its causal consequences for the organism, or one of its purposes, respectively.

Regarding the latter notion—purpose—I will invoke the popular rejoinder: "Let's not go there." Much has been written—indeed books—on the notion of function as "what something is for". In particular I note the line of thought about "proper functions" developed by Ruth G. Millikan (1984). Throughout this chapter, it will be "what something does" that I mean when I refer to "function".

The idea of phenotypic variables that screen off other variables can be expressed mathematically by saying that there is a set of functions, $\{F_i(\boldsymbol{g}_k)\}$, that forms a complete description of the causal consequences for the organism due to variation in the gene $\boldsymbol{g}_k$. To be complete, and account for gene-environment interactions and epistasis, these functions will need to have other arguments that include genes, and environmental variables, ϕ_j, and therefore be of the form $F_i(\{\boldsymbol{g}_k\}, \{\phi_j\})$. These functions in turn may be screened off for their organismal effects by other sets of functions, $\{G_j(F_{i_1}, F_{i_2}, \ldots)\}$. Organisms have many chains of such dependence, which could be said to form a "function network". A simplistic illustration of the "function network" can be seen in Figure 4.

Ultimately, one arrives at the variables that describe the rates of mortality and fertility of an organism as functions of its interactions with the physical and biotic environment. These variables screen off all other phenotypic properties of the organism in determining natural selection on the organism. Examples of such variables would be efficiency of nutrient absorption, mating success, offspring number, death rates due to predation, infection, injury, etc. When we know the value of such variables, there is no additional information that can tell us anything further about an organism's fitness.

If we consider what is meant by "directional selection", then it is clear that this last tier of variables defines the "directions" under selection. Directional selection, as conceived,

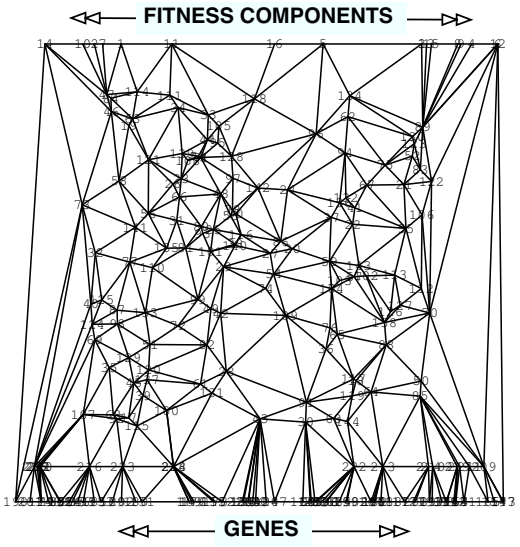

Figure 4: The "function network", showing the relationship between genes, variables that screen off the gene's causal effects, and variables that screen off these variables, etc., all the way to fitness components.

means that there is some phenotypic property which alone can confer a fitness advantage if it changes in the right direction. Stabilizing selection, as conceived, means that there is some phenotypic property which alone can impose a fitness disadvantage if it changes in any direction away from its current value. This is precisely how I have defined the highest level of variables that screen off all other phenotypic variables with respect to selection.

Attempts to describe these function networks can be found in the literature. Dullemeijer (1974), for example, presents a graph of the top layer of the function network in a study of the cranial feeding system of a crotalid snake (cited by Schwenk (2001) in the collection Wagner (2001)) .

The top level of screening off functions defines what matters to each component of an organism's fitness, and thus defines what needs to be optimized by selection. The degree to which each of these top-level screening-off functions is optimized, in any particular organism in a particular environment, can be expected to fall along a spectrum: the functions that are nearly optimal will be sources of stabilizing selection, while those that are suboptimal will be the sources of directional selection.

Let us examine whether it is useful to define modularity in terms of the relation between spaces of phenotypic variation and these top-level screening-off functions. The genotype-phenotype map is defined as modular if very few of these functions are affected throughout a space of phenotypic variation. When would such modularity enhance evolvability? First, it is of no use to evolvability for there to be a modular genotype-phenotype map when all the modules are under stabilizing selection, since there is no adaptive opportunity no matter how it is sliced (Wagner, 1996). Modularity among functions under stabilizing selection may nevertheless have other kinetic population genetic consequences, as found in the study by Waxman & Peck (1998).

Suppose, on the other hand, that the space of variation maps only to functions under *directional selection*. The interactions of these functions in determining fitness would define the selection gradient on the space of variation.

In order for modularity to enhance evolvability, it must parcellate the functions under stabilizing selection from those under directional selection. The importance of this cleavage between stabilizing and directional selection has been recognized for some time as an important element of evolvability, and motivated the development of the "corridor model" (Wagner, 1984, 1988; Bürger, 1986). Mechanisms that that produce the evolution of such a cleavage are discussed by Altenberg (1995). Evolvability is enhanced when there are spaces of phenotypic variation that fall narrowly within the functions under directional selection, and remain orthogonal to the functions under stabilizing selection. This is the kind of modularity that is implicit in the naive framework that I described at the beginning of this chapter. Genes with that sort of modularity would look like the ones in Figure 4 with direct connections to variables on the top level that are under directional selection, and few

connections—direct or indirect—to top level of variables under stabilizing selection.

Hence, these top-level variables provide a way of describing the sort of modularity that is important to evolvability. There is nothing to prevent one from defining pleiotropy and modularity in terms of the map between a gene and the variables at any level in the function network—or for that matter, between a gene and any observer-defined phenotypic characters. But pleiotropy or modularity so defined will not say anything about whether the spaces of variation are aligned with selection gradients, and so will not be relevant to evolvability. The pleiotropy that is relevant to evolvability is that which applies to the map between the gene and the top-level screening-off variables.

Let us now return to my earlier hypothetical question about selection for sharper teeth. Would a genetic variant that made all teeth sharper be more or less modular than a genetic variant that made half of the teeth sharper, or just one pair of teeth, or just one tooth? We now have some machinery to answer this question.

What are the organismal functions that screen off the causal effects of tooth morphology with respect to selection? Such functions would include the rate of catching and killing of prey, the size of food particles sent to the stomach, the amount of flesh removed from a carcass, the success rate for defenses against attack, morbidity and mortality due to tooth and gum infections, mating success, and so on. We must ask which of these functions would be altered by the different spaces of tooth variation.

Suppose that there was directional selection for a stronger mouth grip on prey. A genetic variant that sharpened just the front half of the teeth would improve that quantity, and would leave alone the grinding function of the back teeth. A genetic variant that sharpened all the teeth might also improve the prey-grabbing function, but make it harder to grind food, and thus increase the particle size of food in the stomach, decrease nutrient absorption, and adversely affect fitness. It would affect two top-level functions instead of just one, involving both directional and stabilizing selection instead of just directional selection, and would thus be more pleiotropic and less modular than the mutation of just the front teeth, even though fewer characters (teeth) were altered.

Any number of variants of this example can be posed and analyzed in the same way. Whether genetic variation has a modular effect depends on how it affects the top-level screening-off functions.

3.2 Spaces of Environmental Variation

Can this conceptual framework for modularity give us any guidelines as to what we should expect from nature regarding modularity in the genotype-phenotype map? We have three principal features to consider: the top-level screening-off functions, the partitioning of these functions into those under stabilizing selection and those under directional selection,

and the modularity of the genotype-phenotype map with respect to this partition.

The amount of modularity with respect to directional selection actually exhibited by an organism will depend on the particular set of functions that are under directional selection. If it happens to be a set for which the organism has a modular genotype-phenotype map, then it will show a modular relation to selection.

What determines which functions are under directional selection? Clearly, environmental change—biotic and abiotic—would be the principal cause of directional selection by dislodging the phenotypic optima. So the possession of a modular genotype-phenotype map—in the way that matters to evolvability—would appear to depend on the vagaries of environmental change. The study of modularity in the genotype-phenotype map of organisms as it pertains to evolvability would thus be somewhat of a haphazard subject.

However, there may be processes that give modularity a more systematic existence than the vagaries of environmental change would lead one to expect. I later describe population genetic mechanisms that can lead to the evolution of modularity that enhances evolvability. From the foregoing discussion, we would expect that such modularity would evolve for functions that were under recurrent directional selection. This brings us to the spaces of variation in the environment.

The environment is analogous to the phenotype in that it takes vast numbers of variables to describe it, yet its degrees of freedom for variation are few in comparison. If we go forward with the idea that low-dimensional manifolds characterize the variation of the environment, then each of these spaces will be characterized by different fluctuation statistics. The ones that are highly variable will induce recurring directional selection on those screening off functions of the organism that are sensitive to these environmental variables. Modularity for these variables will enhance the organism's ability to respond evolutionarily to this recurring directional selection. Therefore, if evolvability-enhancing modularity can evolve as a response to directional selection, it will be most well developed for those functions that are under recurring directional selection (Altenberg, 1995; Wagner, 1996).

The upshot is that the spaces of variation in the organism may come to mirror the spaces of variation in the environment. This idea is really only a technical revision of the idea originally proposed by Riedl, that "the epigenetic system copies the functional interdependencies of the phene system." (Riedl 1978, p. 93).

How great a degree of mirroring we can expect depends on the quantitative details of the processes that would produce the evolution of modularity. Such details are left for another day, and here are merely proposed as a possibility that merits investigation.

4 Mutational Kinetics, Modularity, and Evolvability

To pursue the foregoing discussion with a specific example, I utilize the "B-matrix" model of Wagner (1989). This model contains all the ingredients discussed thus far:

- Genes control multiple phenotypic variables, creating the dimensions of variation for the phenotype.
- Phenotypic variables are controlled by multiple genes.
- A fitness function is defined on the phenotypic variables.

In this model, the "function-network" has only this one, top, level of phenotypic variables; there are no other phenotypic variables that are screened off by these fitness-defining variables.

Because the B-matrix model is simple and well-defined, we can answer the question of how the alignment between the dimensions of variation and selection gradients affect evolvability, and derive a means to define modularity as a property intrinsic to the model, not imposed by subjective parsing of the phenotype. We shall discover that a critical feature for defining modularity turns out to be the magnitude of mutation effects.

4.1 Wagner's B-Matrix Model

In the B-matrix model of Wagner (1989), selection is optimizing, acting on multiple traits controlled additively by multiple loci. There are three spaces in this model: genotype, phenotype, and fitness. Genetic variables are mapped to phenotypic variables, and these in turn are mapped to fitness. Each phenotypic variable has an optimal value, and fitness is defined as a Gaussian function of the departure of the traits from the optimum. I describe each of these mappings.

The phenotype-fitness Map. The optimal value of each phenotypic variable is set to 0 for simplicity. Letting $\boldsymbol{x}$ represent the vector of phenotypic variables, the fitness w is defined to be:

$$w = \exp(-\frac{1}{2}\boldsymbol{x}^\top \boldsymbol{M}\boldsymbol{x}),$$

where $\boldsymbol{M}$ is a positive definite matrix (positive definiteness assures that fitness decreases as one departs from the optimum). Here, the function $\delta(\boldsymbol{x}, \boldsymbol{M}) = \boldsymbol{x}^\top \boldsymbol{M}\boldsymbol{x}$ is the sole top-level screening-off function in this system, since if we know δ, the fitness is $w = \exp(-\delta/2)$, and there is no additional information that $\boldsymbol{x}$ gives about fitness.

The Genotype-Phenotype Map. The vector of phenotypic variables $\boldsymbol{x}$ is itself a linear function of the underlying genetic variables $\boldsymbol{y}$:

$$\boldsymbol{x} = \boldsymbol{B}\boldsymbol{y}$$

The scalar value y_i can be interpreted as the lowest-level screening-off function for gene i, which summarizes the entire causal effect that gene i has on the organism.

The fitness function, expressed in terms of $\boldsymbol{y}$, is:

$$w(\boldsymbol{y}) = \exp(-\frac{1}{2}\boldsymbol{y}^\top \boldsymbol{B}^\top \boldsymbol{M}\boldsymbol{B}\boldsymbol{y}), \tag{1}$$

In this model, the spaces of phenotypic variation are simple straight lines defined by the columns of $\boldsymbol{B}$,

$$\boldsymbol{b}_k = \begin{bmatrix} B_{1k} \\ B_{2k} \\ \vdots \\ B_{Lk} \end{bmatrix},$$

where L is the number of genetic variables (loci). So the space of phenotypic variation produced by variation at locus k is the line, $S_k = \{y_k \boldsymbol{b}_k : y_k \in \Re\}$.

4.2 Finite and Infinitesimal Models for Mutational Kinetics

In order for a genotype with these dimensions of variation to respond to selection, the earlier discussion claims that there must be selection gradients along the spaces of variation. When we wish to analyze the evolutionary dynamics, we see immediately that we must know something more about the magnitude of variation produced by mutation of the genotypic variables. In the quantitative genetic literature, we find two main kinetic models (Bürger, 2000) for the production of variation:

- the "random-walk" mutation model (Crow & Kimura, 1964) and
- the "house-of-cards" mutation model (Kingman, 1977, 1978).

The random-walk model embodies the assumption that mutation perturbs the genetic variable away from its current value by a random variable ϵ, giving

$$x_i \rightarrow x_i + \epsilon.$$

Typically, ϵ is distributed symmetrically around 0, having a Gaussian, exponential, or Γ distribution. The transition probability (or density) is

$$T(x_i \leftarrow x_j) = u(x_i - x_j).$$

The house-of-cards model assumes that mutation "topples the house of cards" that adaptation has built up, producing a new phenotype that is independent of the old, with a value that is sampled from the same distribution regardless of the original value, giving

$$x_i \to \epsilon.$$

The transition probability (or density) is

$$T(x_i \leftarrow x_j) = u(x_i).$$

A key difference between the models becomes apparent when they are adapted to the multivariate context. In the random-walk model, the perturbation caused by each individual mutation is taken to be small, and thus nearly neutral. Finite perturbations are taken to be the result of multiple small mutations. Under this process, multiple infinitesimal mutations can accumulate before selection can differentiate them, giving for the random-walk model:

$$\boldsymbol{x} \to \boldsymbol{x} + \sum_i \epsilon_{\kappa_i} \mathbf{1}_{\kappa_i},$$

where $\kappa_i \in \{1, 2, 3, \ldots, L\}$ are independent random variables designating the index of the locus to be mutated, and $\mathbf{1}_{\kappa_i}$ is a vector for the ith mutation that has a single non-zero entry:

$$\mathbf{1}_k = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{bmatrix} \leftarrow k\text{th entry},$$

so $\mathbf{1}_k$ has all 0 entries except for the kth entry, which is 1.

By the law of large numbers, $\sum_i \epsilon_{\kappa_i} \mathbf{1}_{\kappa_i}$ approaches a multivariate Gaussian random variable $\boldsymbol{\epsilon}$, so the mutation process gives:

$$\boldsymbol{x} \to \boldsymbol{x} + \sum_i \epsilon_{\kappa_i} \mathbf{1}_{\kappa_i} \approx \boldsymbol{x} + \boldsymbol{\epsilon}.$$

The mutation process will diffuse away from any "wild-type" genotypes and produce a cloud of genotypes surrounding it.

In the house-of-cards model, on the other hand, mutations are not infinitesimal in size, since the mutant genotype value is sampled from a fixed distribution independent of its current value. In the mutant genotype, a random locus κ is mutated, which replaces element

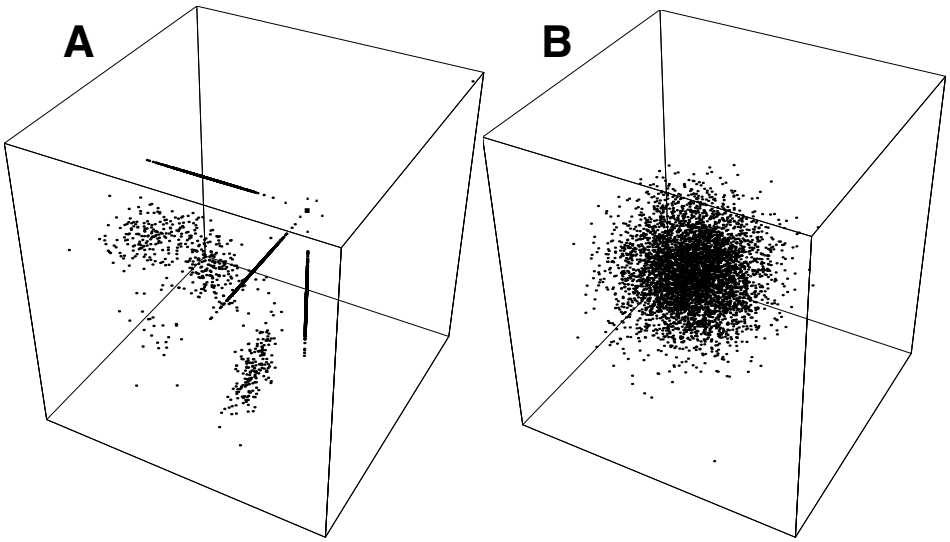

Figure 5: A. Distribution of mutational effects under the "house-of-cards" assumptions. B. Distribution of mutational effects under the "random-walk" assumptions. 5000 points are sampled.

y_κ by ϵ_κ in the vector $\boldsymbol{y}'$, leaving the other positions alone. A way to express the mutant genotype, $\boldsymbol{y}'$, is:

$$\boldsymbol{y}' = \boldsymbol{y} \circ (\mathbf{1} - \mathbf{1}_\kappa) + \epsilon_\kappa \mathbf{1}_\kappa = \boldsymbol{y} + (\epsilon_\kappa - y_\kappa)\mathbf{1}_\kappa, \tag{2}$$

where ϵ_κ is the random variable for the new genotype value, sampled from the distribution $u(x_\kappa)$.

As selection moves the population toward fitter genetic values in the house-of-cards model, a smaller and smaller fraction of the fixed distributions $u(x_i)$ is closer to the optimum, hence the probability of generating fitter mutants trails off with increasing adaptation. Contrary to the random-walk model, since mutations are not infinitesimal, they will not be nearly-neutral, so selection will start to fix or purge mutations as soon as they occur. Thus it will not be possible to build up small mutations at multiple loci before selection acts. This has a significant impact, because the mutation process will no longer produce a multivariate Gaussian perturbation. Instead, the frequencies of single-, double-, triple-locus mutants, etc., will decrease exponentially (at a rate which is a function of the mutation rate and selection magnitude).

The house-of-cards and random-walk models are illustrated in Figure 5. The points in Figure 5A show the spectrum of phenotypes that are accessible under the house-of-cards

assumptions. This graph is produced with the assumption that loci mutate independently, so single, double, triple, mutations occur with frequencies proportional to powers of the mutation rate. What matters, however, is how the frequency spectrum of multiple mutations affects the accessibility of the space. Most variants are single mutants, which fall along the axes of variation produced by each gene. Rarer double mutants fall along the planes defined by each pair of single-mutant axes. Even rarer triple mutants fall in the interior. The "wild type" is at the intersection of the single-mutant axes. The points in Figure 5B show the spectrum of phenotypes that are accessible under the random-walk model, where multiple infinitesimal mutations allow access to the entire space around the wild-type phenotype.

I have presented the house-of-cards model as a paradigm for a mutational kinetics that generates variation along low-dimensional spaces as depicted in Figure 5A. However, what is critical to this result is not the "house-of-cards" assumption itself. Rather, it is that mutation is finite rather than infinitesimal in effect. Any mutational distribution that is dominated by finite effects will result in a population distribution similar to Figure 5A in which variation falls along the single-mutant axes. Therefore, I will refer to the two distinct paradigms for mutational kinetics as "finite" versus "infinitesimal" models.

These two different models for mutational distributions have very different implications for the issue of modularity. In the finite-effects model, there must be selection gradients along the single-mutant axes in order for adaptation to occur. In the infinitesimal-effects model, on the other hand, multiple single-mutant axes combine to span a higher-dimensional linear subspace, and evolution can follow any selection gradient within this subspace.

It should be noted that in the infinitesimal-effects model, the multiple-mutant subspaces may impose their own constraints upon adaptation if they do not span the entire space of phenotypes. Translated, this means that no genotype exists that can produce an optimal phenotype. In an infinite-dimensional trait such as a growth curve, this is a generic situation (Kirkpatrick & Lofsvold, 1992). It is the expectation whenever there are fewer dimensions to the genotype space than there are to the phenotype space, or if the columns of the B-matrix are linearly dependent. In either case, the result is that the B-matrix will not be not full-rank, and the genetic variance-covariance matrix, $\boldsymbol{B}\,\boldsymbol{B}^{\top}$, will be singular.

When the B-matrix is not full rank, the generic outcome of evolution under an infinitesimal-effects model is that the phenotype will reach a constrained optimum within the space spanned by the B-matrix, at some distance from the global optimum (Kirkpatrick & Lofsvold, 1992; Altenberg, 1995). At this constrained optimum, additive genetic variation may exist for each phenotypic variable individually, but the reduced dimensionality of their joint variation will prevent any response to selection. There will remain a "latent" directional selection orthogonal to the space of variation (Altenberg, 1995).

The finite-effects model makes possible a form of constraint — frustration — above and beyond the constraint caused by a non-full-rank B-matrix. Frustration may prevent even the constrained optimum from being reached.

Frustration is a *kinetic constraint*, in that genotypes with the optimal phenotype may be possible, but the probability of generating them is minute because it requires multiple simultaneous mutations away from the wild type.

Riedl (1977) delves into the issue of finite versus infinitesimal effects in his discussion of alternative theories for the evolution of complex phenotypes. One which he calls the "storage theory" proposes that in cases where multiple mutations are needed to produce a particular adaptation, these mutations can be stored in the gene pool until they are brought together by recombination or hybridization. But this requires that the mutations, not valuable individually, be nearly neutral so as not to be expunged by selection. In order to be nearly neutral, they must be of extremely small effect. The storage theory, then, is an infinitesimal-effects model for mutational kinetics.

4.3 Alignment with Selection Gradients

With this distinction between these two models for mutational kinetics now spelled out, let us return to the thesis described at the beginning of this chapter about the advantage of modularity for evolvability. As should now be obvious, these conventional ideas about modularity have as a core assumption that mutation follows a finite-effects kinetics.

Recalling that the fitness function in the B-matrix model is:

$$w(\boldsymbol{y}) = \exp(-\frac{1}{2}\boldsymbol{y}^\top \boldsymbol{B}^\top \boldsymbol{M}\boldsymbol{B}\boldsymbol{y}),$$

then in the house-of-cards model, with locus κ mutated, the fitness of the mutant genotype $\boldsymbol{y}' = \boldsymbol{y} + (\epsilon_\kappa - y_\kappa)\mathbf{1}_\kappa$ (from equation 2) is:

$$\begin{aligned} w(\boldsymbol{y}') &= \exp(-\frac{1}{2}\boldsymbol{y}'^\top \boldsymbol{B}^\top \boldsymbol{M}\boldsymbol{B}\boldsymbol{y}') \\ &= \exp\left(-\frac{1}{2}\left[\boldsymbol{y}^\top \boldsymbol{B}^\top \boldsymbol{M}\boldsymbol{B}\boldsymbol{y} + 2(\epsilon_\kappa - y_\kappa)\,\boldsymbol{y}^\top \boldsymbol{B}^\top \boldsymbol{M}\boldsymbol{b}_\kappa + (\epsilon_\kappa - y_\kappa)^2 \boldsymbol{b}_\kappa{}^\top \boldsymbol{M}\boldsymbol{b}_\kappa\right]\right) \\ &= w(\boldsymbol{y})\,\exp\left[(y_\kappa - \epsilon_\kappa)\,\boldsymbol{y}^\top \boldsymbol{B}^\top \boldsymbol{M}\boldsymbol{b}_\kappa - \frac{1}{2}(y_\kappa - \epsilon_\kappa)^2 \boldsymbol{b}_\kappa{}^\top \boldsymbol{M}\boldsymbol{b}_\kappa\right]. \end{aligned}$$

So we see that whether the mutation is adaptive or not depends on the relationship of the column vectors $\boldsymbol{b}_\kappa$ with the matrix $\boldsymbol{M}$ and the current genotype $\boldsymbol{y}$.

One could exactly quantify the magnitude of evolvability for this model by specifying the sampling distribution of ϵ, and deriving the probability that $w(\boldsymbol{y}') > w(\boldsymbol{y})$. This, however, would go beyond the purpose of this chapter, which is merely to delineate the

relationship between the different factors described here: modularity, spaces of variation, selection gradients, and evolvability.

I have claimed that the natural notion of modularity—a genotype-phenotype map that is decomposable into the product of lower-dimensional genotype-phenotype maps—is no more than a means (nor the only means) to enhance evolvability by making it easier to align the spaces of variation with selection gradients. In the B-matrix model, a modular genotype-phenotype map corresponds, in the extreme degree, to $\boldsymbol{B}$ being a diagonal matrix. Under this condition, we have:

$$\boldsymbol{y}^\top \boldsymbol{B}^\top \boldsymbol{M} \boldsymbol{b}_\kappa = \sum_j y_j B_{jj} M_{j\kappa} B_{\kappa\kappa},$$

and

$$\boldsymbol{b}_\kappa{}^\top \boldsymbol{M} \boldsymbol{b}_\kappa = B_{\kappa\kappa}^2 M_{\kappa\kappa},$$

hence

$$w(\boldsymbol{y}') = w(\boldsymbol{y}) \exp\left[(y_\kappa - \epsilon_\kappa) \sum_j y_j B_{jj} M_{j\kappa} B_{\kappa\kappa} - \frac{1}{2}(y_\kappa - \epsilon_\kappa)^2 B_{\kappa\kappa}^2 M_{\kappa\kappa} \right].$$

We notice that, despite the modularity of the genotype-phenotype map, there are interaction terms $y_j B_{jj} M_{j\kappa} B_{\kappa\kappa}$ that signify epistasis between loci—i.e. whether mutation at locus κ can generate adaptation depends on the state of the other loci, y_j. In fact, the situation with a modular genotype-phenotype map is really no different from the situation with a non-modular genotype-phenotype map, because we can write:

$$w(\boldsymbol{y}) = \exp(-\frac{1}{2}\boldsymbol{y}^\top \boldsymbol{M}' \boldsymbol{y}),$$

where $\boldsymbol{M}' = \boldsymbol{B}^\top \boldsymbol{M} \boldsymbol{B}$ is a positive definite matrix, which is the same form as if $\boldsymbol{B}$ were the identity matrix.

Therefore, a modular genotype-phenotype map is not sufficient to ensure any special evolutionary capabilities of the variation generating system. What is further required is that the M matrix itself be a diagonal. In that case, we obtain:

$$\begin{aligned} w(\boldsymbol{y}') &= w(\boldsymbol{y}) \exp\left[(y_\kappa - \epsilon_\kappa)\, B_{\kappa\kappa}^2 M_{\kappa\kappa} [y_\kappa - \frac{1}{2}(y_\kappa - \epsilon_\kappa)] \right] \\ &= w(\boldsymbol{y}) \exp\left[\frac{1}{2}(y_\kappa^2 - \epsilon_\kappa^2)\, B_{\kappa\kappa}^2 M_{\kappa\kappa} \right]. \end{aligned}$$

Here, we see that the ability of a new mutation to produce a fitness increase depends solely on whether the new genotypic value, ϵ_κ, is closer to the optimum than the old genotypic

value, y_κ. No other loci are involved. But we see that "modularity" here cannot be defined solely in terms of the genotype-phenotype map; it must involve also the matrix $\boldsymbol{M}$, which describes how phenotypes map to selection. So again, what is more fundamental to evolvability than modularity in the genotype-phenotype map is the relationship between the spaces of genetic variation and the selection gradient.

5 Discussion

In this chapter I have tried to focus on some of the low-level issues that arise when trying to approach the issue of modularity in evolution. I have not delved at all into the question of how evolutionary dynamics may affect modularity and the alignment of spaces of variation with selective gradients. I will offer some comments on the evolutionary dynamics affecting modularity.

5.1 Constructional Selection

The role of gene origin in sculpting the modularity of the genotype-phenotype map is explored in Riedl's work (Riedl, 1975, 1977, 1978), and in several of my own papers (Altenberg & Brutlag, 1986; Altenberg, 1994b, 1995).

The central idea of this work is that we expect the dimensions of variation in the genome to be enriched with spaces that are in alignment with selective gradients. This enrichment process is a systematic outcome of the dynamics of genome growth. New genes that happen to change the phenotype along a positive selection gradient are much more likely to be preserved by selection than genes which produce variation that randomly perturbs the phenotype and is thus likely to be detrimental. Thus, the degrees of freedom in the genome should grow in the direction of greater evolvability. My shorthand term for this process is "constructional selection" because it pertains to the construction of the genome.

Modularity is one means, though not the only means, to achieve the correct alignment of the space of variation with respect to selection. So modularity is one feature that we expect to be enriched by the process of genome growth. Clear examples of this sort of modularity are the separation of regulatory function from coding function in eukaryotic genomes. Such separation is not a functional necessity, as seen in non-modular genes where sequences carry both coding and regulatory function. But separation of these functions permits one of the dimensions of genetic variation—sequence duplication—to explore combinatorial spaces which preserve the regulatory and coding functions of the gene fragments. By maintaining these functions, but bringing them together in a new combination, such modular genetic elements have a greater likelihood to produce a selective advantage, and thus be kept by the genome. Therefore, the genome should become more enriched for

such elements as it grows. This same process would also apply to elements within regulatory regions, or within coding regions. And we find that many proteins are mosaics of function recombined from other genes (Hegyi & Bork, 1997).

What is important to remember is that the modularity that can result from selective genome growth is defined *in terms of the genetic operators* producing the genetic variation, in this case the processes of sequence duplication. So, for example, if sequence duplication happened to be restricted to a certain range of sequence lengths, it would be on that length scale that genome growth would select for modularity. And modularity is selected only with respect to its ability to increase the likelihood that the sequence duplication event is beneficial. All structural features that we would call "modular" are defined in terms of this probability rather than any *a priori* structural definitions that we might impose.

Failure to appreciate this essential point is a source of confusion when discussing the issue of modularity of exons (Logsdon, 1998; De Souza *et al.*, 1998). Modularity with respect to exon shuffling can be achieved when protein domain boundaries correspond to exon boundaries. But a lack of correspondence is not in itself evidence against modularity. If functional properties of an exon are maintained after exon shuffling, then this exon exhibits modularity. It may not be necessary for domain and exon boundaries to correspond in order for the functional properties to withstand exon shuffling—other properties of the sequence can stabilize the functional elements. This distinction is subject to empirical testing because modifications of splice sites in exons with evolved modularity would be expected to decrease their modularity, whether or not the splice sites fall between protein domains.

5.2 Other Sources of Modularity in the Genotype-Phenotype Map

In addition to genome growth processes, there may be other sources that produce modularity in the genotype-phenotype map. These deserve some mention here.

5.2.1 Modularity "For Free"

There may be generic features of biology, chemistry, or physics that provide modularity in the genotype-phenotype map "for free"—to borrow the phrase from Kauffman (1995, Chapter 4). Kauffman speaks of "order for free", that is, order in living organisms that arises not from Darwinian selection (order at a cost), but as a generic outcome of physical self-organizing processes. Similarly, there may be examples of "modularity for free" in the genotype-phenotype map that have a similar origin. In other words, there may be circumstances when we expect modularity to be a generic property of organisms that does not require natural selection to establish or maintain.

One obvious candidate source for modularity without natural selection is the branching structure of the cellular genealogy in multicellular organisms. Multicellular organisms

arise from the repeated division of cells. The ancestral state for the cells of multicellular organisms is the single-celled organism, which is the epitome of a module. The tendency of unicellular organisms to separate, disperse, and become independent after replication is a generic property that makes them modular. Many single-celled organisms can nevertheless have aggregate properties (e.g., production of biofilms), while multicellular organisms have adaptations that counteract independence after replication, and maintain proximity and interaction to varying degrees. However, a certain amount of separation and independence is inescapable among the cells in multicellular organisms. This would be a fundamental source of "modularity for free" in multicellular organisms.

Vascular plants maintain a close parallel between their physical structure and their genealogical structure, because their cells have less mobility than cells in animals. So cells which are genealogically distant also tend to be physically distant. This physical distance makes modularity in the genotype-phenotype map more easily realized, because phenotypic alterations in the structure of, say, a flower, may have fewer physical interactions with, for example, a root.

In complex animals, there is less isolation between genealogically distant cells because of cellular mobility and physiological integration. Multiple tissue lineages participate in the construction of integrated organs. Hormonal and neuronal communication integrates genealogically distant cells in their function. Therefore, in animals one would expect to find significantly less "modularity for free" from the cellular genealogy.

5.2.2 Modularity "Included"

It is possible that modularity in the genotype-phenotype map can "hitchhike" (Maynard Smith, 1974) along with traits under natural selection. This is what I mean by "modularity included"—it doesn't come free, but is included as a side effect of natural selection for traits under selection. A paradigmatic example of "modularity included" is the work on selection for robustness in RNA structures by Ancel & Fontana (2000), and Chapter 6 of this volume (Callebaut & Rasskin-Gutman, 2004). They find that as greater stability evolves in their molecular structures, most of the molecular sites become structurally neutral, while structural sensitivity to mutation concentrates in a tightly integrated core of sites.

It is possible that there is a physical explanation for this phenomenon, which may make it a generic property of molecular interactions. Structural stability depends on strong molecular bonding, and strong bonding requires physical proximity of bonding sites. Such physical proximity, however, can be shared by only a limited number of sites. Therefore, the strongest bonding interactions are expected to be limited to a selected set of sites, screening off other sites from these high-energy bonds. Thus selection for strong bonding can have the side effect that these high bond energies become concentrated among a small number of sites.

This correlation between the strength of interaction and the specificity of interaction may be a generic feature of a wide class of molecules, especially ones where the interaction is specified by shape, such as proteins, nucleotides, receptors, and enzymatic reactions. There are obvious exceptions, such as peroxides, that achieve strong interaction with little specificity. But many biological molecules, especially proteins and nucleotides, may receive "modularity included" in selection for structural stability because of this correlation.

This mechanism for "modularity included" would also apply to spatial compartmentalization (Weng *et al.*, 1999). Compartmentalization of reacting molecules increases the strength of interaction simply by increasing concentrations, but because of the conservation of matter, it decreases concentrations elsewhere, and thereby increases the specificity of interactions. Selection for high concentration of molecules may thereby bring along modularity as a side effect.

5.2.3 Direct Selection for Modularity

Specificity of interaction may be a side effect of selection for strong interaction, but it may also be a target of selection in its own right. Coordination of activities from the scale of the chromosome to the entire organism, or even an entire population, requires precise specificity between signals and receptors. Specificity is needed so that the control of different processes in the organism has the degrees of freedom needed to optimize their coordination. This specificity of interaction can translate directly into specificity for the phenotypic effects of genetic variation, a.k.a. modularity.

5.3 Subfunctionalization

Where are we to place the phenomenon of subfunctionalization (Force *et al.*, 1999; Lynch & Force, 2000) within this categorization scheme? Subfunctionalization is a process in which duplicate genes make themselves necessary to the organism by losing, rather than gaining, function. In the classical thinking about the fate of gene duplications, the duplicates had to gain new functions in order to avoid being redundant and eventually silenced by mutation (Ohno, 1970). However, if genes carry out multiple functions, and these functions can be silenced independently of one another, then a different set of functions can be silenced in each gene, and the remaining functions of each gene can be preserved by selection. In essence, after subfunctionalization, there is still only one gene functioning, but it is split up into two different loci, and involves two different transcripts with complementary function. The complementation must therefore be *trans*-acting.

It should be immediately clear that subfunctionalization is not a means to produce modularity, but rather the reverse: it requires that functions of the gene already be modular, in that the gene has independent degrees of freedom for the loss of each function. The process

of gene duplication and subfunctionalization will exhaust itself when the modules inherent in the original gene have been completely parceled out among the duplicate genes. A further gene duplication will not be able to simultaneously lose part of its function and complement the losses in other genes. It will either be redundant or necessary as a whole.

Subfunctionalization thus faces a finite limit on the process, which distinguishes it from constructional selection. In constructional selection, the amplification of modular elements in the genome is limited only by the selective opportunity for new combinations of modules. Subfunctionalization, on the other hand, is effectively conservative for module number—spreading out modules among multiple loci but not creating them. Therefore it cannot explain module origin, and thus is a consequence, rather than a source, of module-creating processes, such as constructional selection, genetic modification, and selection for properties that have "modularity included".

6 Conclusion

I have endeavored in this chapter to delve into some of the low-level conceptual issues associated with the idea of modularity in the genotype-phenotype map. My main proposal is that the evolutionary advantages that have been attributed to modularity do not derive from modularity *per se*. Rather, they require that there be an "alignment" between the spaces of phenotypic variation, and the selection gradients that are available to the organism. Modularity in the genotype-phenotype map may make such an alignment more readily attained, but it is not sufficient; the appropriate phenotype-fitness map in conjunction with the genotype-phenotype map is also necessary for evolvability.

Acknowledgements

My greatest thanks to the Konrad Lorenz Institute for inviting me to participate in the workshop and creating a very stimulating scientific environment; to the editors Diego Rasskin-Gutman and Werner Callebaut for their forbearance; to the Department of Information and Computer Sciences at the University of Hawaii at Manoa; and to my family and friends on general principle.

Bibliography

Agrafiotis, D. K. and Xu, H. (2002). A self-organizing principle for learning nonlinear manifolds. *Proceedings Of The National Academy Of Sciences Of The United States Of America 99*, 15869–15872.

Altenberg, L. and Brutlag, D. L. (1986). Selection for modularity in the genome. Unpublished. Cited in Doolittle (1987), and Tomita *et al.* (1996).

Altenberg, L. (1994a). The evolution of evolvability in genetic programming. In Kinnear, K. E. (Ed.) *Advances in Genetic Programming* (pp. 47–74). Cambridge, MA: MIT Press.

Altenberg, L. Evolving better representations through selective genome growth. In *Proceedings of the 1st IEEE Conference on Evolutionary Computation. Part 1 (of 2)* (pp. 182–187) Piscataway N.J. 1994. IEEE.

Altenberg, L. (1995). Genome growth and the evolution of the genotype-phenotype map. In Banzhaf, W. and Eeckman, F. H. (Eds.) *Evolution and Biocomputation: Computational Models of Evolution* volume Lecture Notes in Computer Science 899 (pp. 205–259). Berlin: Springer-Verlag.

Ancel, L. W. and Fontana, W. (2000). Plasticity, evolvability and modularity in RNA. *Journal of Experimental Zoology (Molecular and Developmental Evolution) 288*, 242–283.

Ancel, L. W. and Fontana, W. (2004). Evolutionary lock-in and the origin of modularity in RNA structure. In Callebaut, W. and Rasskin-Gutman, D. (Eds.) *Modularity: Understanding the Development and Evolution of Natural Complex Systems* chapter 6. MIT Press.

Brandon, R. N. (1982). The levels of selection. In Asquith, P. and Nickles, T. (Eds.) *PSA 1982* volume 1 (pp. 315–323). East Lansing, MI: Philosophy of Science Association.

Brandon, R. N. (1984). The levels of selection. In Brandon, R. and Burian, R. (Eds.) *Genes, Organisms, Populations: Controversies Over the Units of Selection* (pp. 133–141). Cambridge, Massachusetts: M.I.T. Press.

Brandon, R. N. (1990). *Adaptation and Environment*. Princeton: Princeton University Press.

Brandon, R. N. (2002). Philosophy of selection: Units and levels. *Human Nature Review 2*, 373. http://www.spee.ch/archive/evolution_philoso–phy_of_selection_units_and_levels.htm.

Bürger, R. (1986). Constraints for the evolution of functionally coupled characters: A nonlinear analysis of a phenotpyic model. *Evolution 40*, 182–193.

Bürger, R. (2000). *The Mathematical Theory of Selection, Recombination, and Mutation*. Chichester: John Wiley & Sons, LTD.

Callebaut, W. and Rasskin-Gutman, D. (Eds.) (2004). *Modularity: Understanding the Development and Evolution of Complex Natural Systems.* Cambridge, MA: MIT Press.

Crow, J. F. and Kimura, M. The theory of genetic loads. In *Proceedings of the XIth International Congress of Genetics* volume 2 (pp. 495–505) Oxford 1964. Pergamon Press.

De Souza, S. J., Long, M., Klein, R. J., Roy, S., Lin, S., and Gilbert, W. (1998). Toward a resolution of the introns early/late debate: Only phase zero introns are correlated with the structure of ancient proteins. *Proceedings of the National Academy of Sciences U.S.A. 95*, 5094–5099.

Doolittle, W. F. (1987). The origin and function of intervening sequences in DNA: A review. *American Naturalist 130*, 915–928.

Dullemeijer, P. (1974). *Concepts and Approaches in Animal Morphology*. Assen, the Netherlands: Van Gorcum.

Force, A., Lynch, M., Pickett, F. B., Amores, A., Yan, Y. L., and Postlethwait, J. (1999). Preservation of duplicate genes by complementary, degenerative mutations. *Genetics 151*, 1531–1545.

Hegyi, H. and Bork, P. (1997). On the classification of evolution of protein modules. *Journal of Protein Chemistry 16*, 545–551.

Kauffman, S. A. and Levin, S. (1987). Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology 128*, 11–45.

Kauffman, S. A. and Weinberger, E. D. (1991). The NK model of rugged fitness landscapes and its application to maturation of the immune response. In Perelson, A. S. and Kauffman, S. A. (Eds.) *Molecular Evolution on Rugged Landscapes: Proteins, RNA and the Immune System* (pp. 135–175). Redwood City: Addison-Wesley.

Kauffman, S. A. (1995). *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*. Oxford: Oxford University Press.

Kingman, J. F. C. (1977). On the properties of bilinear models for the balance between genetic mutation and selection. *Mathematical Proceedings of the Cambridge Philosphical Society 81*, 443–453.

Kingman, J. F. C. (1978). A simple model for the balance between selection and mutation. *Journal of Applied Probability 15*, 1–12.

Kirkpatrick, M. and Lofsvold, D. (1992). Measuring selection and constraint in the evolution of growth. *Evolution 46*(4), 954–971.

Logsdon, J. M. (1998). The recent origins of spliceosomal introns revisited. *Current Opinion in Genetics and Development 8*, 637–648.

Lynch, M. and Force, A. (2000). The probability of duplicate gene preservation by subfunctionalization. *Genetics 154*, 459–473.

Maynard Smith, J. (1974). The hitch-hiking effect. *Heredity 33*, 130.

McKay, S. R., Berker, A. N., and Kirkpatrick, S. (1982). Spin-glass behavior in frustrated ising models with chaotic renormalization-group trajectories. *Phys. Rev. Lett. 48*, 767.

Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories :New Foundation for Realism*. Cambridge, Mass: M.I.T. Press.

Ohno, S. (1970). *Evolution by Gene Duplication*. Berlin: Springer Verlag.

Plato (c. 370 B.C.). Statesman. In Skemp, J. (Ed.) *Plato's Statesman: A Translation of the Politicus of Plato with Introductory Essays and Footnotes*. London: Routledge and Kegan Paul (1952).

Riedl, R. J. (1975). *Die Ordnung des Lebendigen: Systembedingungen der Evolution*. Hamburg and Berlin: Parey.

Riedl, R. J. (1977). A systems-analytical approach to macroevolutionary phenomena. *Quarterly Review of Biology 52*, 351–370.

Riedl, R. J. (1978). *Order in Living Organisms: A Systems Analysis of Evolution*. Chichester: John Wiley and Sons. A translation of Riedl (1975) by R.P.S. Jefferies.

Roweis, S. T. and Saul, L. K. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science 290*, 2323–2367.

Salmon, W. C. (1971). *Statistical Explanation and Statistical Relevance*. Pittsburgh: University of Pittsburgh Press.

Salmon, W. C. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton, N. J.: Princeton University Press.

Schwenk, K. (2001). Functional units and their evolution. In Wagner, G. P. (Ed.) *The Character Concept in Evolutionary Biology* (pp. 165–198). Boston: Academic Press.

Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society 106*, 467–482. Reprinted in Simon (1969).

Simon, H. A. (1969). *The Sciences of the Artificial*. Cambridge: MIT Press.

Stadler, B. M. R., Stadler, P. F., Shpak, M., and Wagner, G. P. (2002). Recombination spaces, metrics, and pretopologies. *Zeitschrift für Physikalische Chemie 216*, 217–234.

Tenenbaum, J. B., de Silva, V., and Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science 290*, 2319–2323.

Tomita, M., Shimizu, N., and Brutlag, D. L. (1996). Introns and reading frames: Correlations between splice sites and their codon positions. *Molecular Biology and Evolution 13*(9), 1219–1223.

van Nimwegen, E. J., Crutchfield, J. P., and Huynen, M. (1999). Metastable evolutionary dynamics: Crossing fitness barriers or escaping via neutral paths? *Bulletin of Mathematical Biology 62*, 799–848.

Verbeek, J. J., Vlassis, N., and Kröse, B. Fast nonlinear dimensionality reduction with topology preserving networks. In Verleysen, M. (Ed.) *Proceedings of the 10th European Symposium on Artificial Neural Networks* (pp. 193–198) Bruges, Belgium 2002. D-Side Publications.

Waddington, C. H. (1942). Canalization of development and the inheritance of acquired characters. *Nature 150*, 563–565.

Wagner, G. P. and Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution 50*(3), 967–976.

Wagner, G. P. (1984). Coevolution of functionally constrained characters: Prerequisites of adaptive versatility. *BioSystems 17*, 51–55.

Wagner, G. P. (1988). The influence of variation and of developmental constraints on the rate of multivariate phenotypic evolution. *Journal of Evolutionary Biology 1*, 45–66.

Wagner, G. P. (1989). Multivariate mutation-selection balance with constrained pleiotropic effects. *Genetics 122*, 223–234.

Wagner, G. P. (1996). Homologues, natural kinds and the evolution of modularity. *American Zoologist 36*, 36–43.

Wagner, G. P. (Ed.) (2001). *The Character Concept in Evolutionary Biology*. Boston: Academic Press.

Waxman, D. and Peck, J. (1998). Pleiotropy and the preservation of perfection. *Science 279*, 1210–1213.

Weinberger, E. D. (1991). Local properties of Kauffman's N-k model, a tuneably rugged energy landscape. *Physical Review A 44*(10), 6399–6413.

Weng, G., Bhalla, U. S., and Iyengar, R. (1999). Complexity in biological signaling systems. *Science 284*, 92–96.

Wright, S. (1968). *Evolution and the genetics of populations*. Chicago: University of Chicago Press.