Innovation and robustness in complex regulatory gene networks

S. Ciliberti*, O. C. Martin*†, and A. Wagner‡§

*Unité Mixte Recherche 8565, Laboratoire de Physique Théorique et Modèles Statistiques, Université Paris-Sud and Centre National de la Recherche Scientifique, F-91405 Orsay, France; †Unité Mixte de Recherche 820, Laboratoire de Génétique Végétale, L'Institut National de la Recherche Agronomique, Ferme du Moulon, F-91190 Gif-sur-Yvette, France; and ‡Department of Biochemistry, University of Zurich, Y27-J-54, Winterthurerstrasse 190, CH-8057 Zurich, Switzerland

Communicated by Giorgio Parisi, University of Rome, Rome, Italy, June 20, 2007 (received for review November 24, 2006)

The history of life involves countless evolutionary innovations, a steady stream of ingenuity that has been flowing for more than 3 billion years. Very little is known about the principles of biological organization that allow such innovation. Here, we examine these principles for evolutionary innovation in gene expression patterns. To this end, we study a model for the transcriptional regulation networks that are at the heart of embryonic development. A genotype corresponds to a regulatory network of a given topology, and a phenotype corresponds to a steady-state gene expression pattern. Networks with the same phenotype form a connected graph in genotype space, where two networks are immediate neighbors if they differ by one regulatory interaction. We show that an evolutionary search on this graph can reach genotypes that are as different from each other as if they were chosen at random in genotype space, allowing evolutionary access to different kinds of innovation while staying close to a viable phenotype. Thus, although robustness to mutations may hinder innovation in the short term, we conclude that long-term innovation in gene expression patterns can only emerge in the presence of the robustness caused by connected genotype graphs.

evolutionary novelty | evolvability | genotype–phenotype maps

Life's enormous creativity is evident from earth's millions of species with unique life styles, from dazzlingly different modes of development to macromolecules, like proteins and RNA, in which many different molecular functions (catalysis, support, and communication) have evolved. There are many wonderful case studies of individual evolutionary innovations, from the beaks of Darwin's finches (1) to the biochemical innovations represented by the highly refractory eye lens proteins derived from various enzymes (2, 3). These and all other evolutionary innovations are produced by a combination of mutation and natural selection, without apparent foresight and planning. However, mutation and selection do not automatically produce evolutionary innovation. For instance, man-made systems, such as computer hardware and software, seem to be outright incapable of innovation through mutation and selection. Those complex systems exhibit brittleness: Modifying one component often leads to disastrous failure. Diligent research in areas such as "evolvable hardware" (4–6) is needed to understand how complex functionalities can be rendered insensitive to individual component changes, thereby facilitating innovation. It is important to discover what renders living beings so capable of innovation, partly because the lessons learned could be applied to the design of complex systems with specific functions.

Biologists increasingly realize that genetic systems need to be robust to both genetic and nongenetic change (7–14). Robustness means that a system keeps performing its function in the face of perturbations. For example, many proteins can continue to catalyze chemical reactions, regulate transcription, communicate signals, and serve other roles despite mutations changing many amino acids; regulatory gene networks continue to function despite noisy expression of their constituent genes; embryos continue to develop normally even when faced with substantial environmental variation. Mutational robustness means that a system produces little phenotypic variation when subjected to genotypic variation caused by mutations. At first sight, such robustness might pose a problem for evolutionary innovation, because a robust system cannot produce much of the variation that can become the basis for evolutionary innovation.

As we shall see, there is some truth to this appearance, but it is in other respects flawed. Robustness and the ability to innovate cannot only coexist, but the first may be a precondition for the second. Individual case studies of evolutionary innovation are essentially anecdotes based on one or a few observations and would not take us very far in validating this assertion. To examine the relationship between innovation and robustness, we need to examine variation in robustness and in the ability to innovate. That is, we need to examine a great many architectural variants of a system with similar functions, and innovations derived from them. Much innovation (proteins with new catalytic activities or new organismal features like lungs or wings) is surprising, sometimes even in hindsight. To study innovation systematically, one needs to take the element of surprise out of it. To do this, a context is needed where the space of all possible genotypes and phenotypes of a biological system can be characterized, at least in principle. Examples include the sequence (genotype) space of RNA and proteins and the secondary or tertiary structures (phenotypes) they form (15, 16).

We here address the problem of how robustness relates to innovation in a model system completely different from RNA, that of a transcriptional regulation network. In this system, the genotype is a regulatory genotype, a set of interactions among transcriptional regulators. The phenotype is the gene expression pattern produced by these regulatory interactions. We shall be interested in the relationship between robustness and the ability to find new phenotypes, a proxy for the ability to innovate, as a function of the genotype. Despite its level of abstraction, variants of the model we use have proven successful in explaining the regulatory dynamics of early developmental genes in the fruit fly *Drosophila*, as well as in predicting mutant phenotypes (17–20). It has also helped elucidate why mutants often show a release of genetic variation that is cryptic in the wild type and how adaptive evolution of robustness occurs in genetic networks of a given topology (14, 21–27).

The model (Fig. 1*a*) is concerned with a regulatory network of N transcriptional regulators, which are represented by their expression patterns $\mathbf{S}(t) = (S_1(t), S_2(t), \ldots, S_N(t))$ at some time

Author contributions: S.C. and O.C.M. contributed equally to this work; A.W. designed research; S.C., O.C.M., and A.W. performed research; S.C. and O.C.M. contributed new reagents/analytic tools; S.C., O.C.M., and A.W. analyzed data; and S.C., O.C.M., and A.W. wrote the paper.

The authors declare no conflict of interest.

§To whom correspondence should be addressed. E-mail: aw@bioc.uzh.ch.

This article contains supporting information online at www.pnas.org/cgi/content/full/0705396104/DC1.

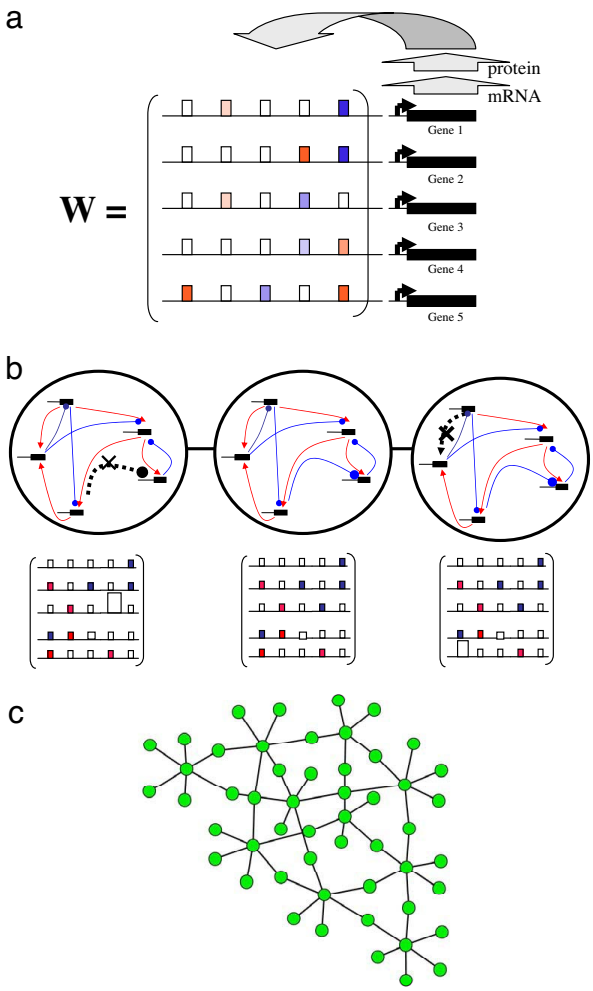

Fig. 1. Neutral networks in transcriptional regulation. (*a*) A transcriptional regulation network. Solid black bars indicate genes that encode transcriptional regulators in a hypothetical network of five genes. Each gene is expressed at a level that is influenced by the transcriptional regulators in the network. This influence is usually exerted through the binding of a transcriptional regulator to a gene's regulatory region (horizontal line). The model represents the regulatory interactions between transcription factors j and genes i through a matrix $w = (w_{ij})$. A regulator's effect can be activating ($w_{ij} > 0$, red rectangles) or repressing ($w_{ij} < 0$, blue rectangles). Any given gene's expression may be unaffected by most regulators in the network ($w_{ij} = 0$, white rectangles). The different hues of red and blue correspond to different magnitudes of w_{ij}. The highly regular correspondence of matrix entries to binding sites serves the purpose of illustration and is not normally found, because transcription factor binding sites usually function, regardless of their position in a regulatory region. (*b*) The topology on the space of genotypes induced by single mutations. The center network shows a hypothetical network of five genes (*Upper*) and its matrix of regulatory interactions w (*Lower*), if genes are numbered clockwise from the uppermost gene. Red arrows indicate activating interactions, and blue lines terminating in a circle indicate repressive interactions. The leftmost network and the center network differ in one repressive interaction from gene 4 to gene 3 (dashed gray line, black cross, and large open rectangle). The rightmost network and the middle network differ in one activating interaction from gene 1 to gene 5 (dashed line, black cross, and large white rectangle). Each of the three networks corresponds to one node in a graph as indicated by the large circle around the networks. These circles are connected because the respective networks are neighbors; i.e., they differ by one regulatory interaction. [*a* and *b* were reproduced with permission from Ciliberti *et al.* (31) (Copyright 2007, Public Library of Science)]. (*c*) The neutral network for a given phenotype. Each node corresponds to a network of a given topology, and two nodes are connected by an edge if they differ at one regulatory interaction ($n = 3$ genes, $4 \leq M \leq 5$ regulatory interactions, and Hamming distance of $\mathbf{S}(0)$ and $\mathbf{S}_\infty$ of $d = 2/3$). This neutral network is connected and the number of edges incident on a node is highly variable. Note that neutral networks for greater numbers of genes typically have a huge number of nodes. The number of nodes in a neutral network can be counted, because different nodes differ only in the signs of their regulatory interactions.

t during a developmental or cell biological process and in one cell or domain of an embryo. These transcriptional regulators can influence each other's expression through cross-regulatory and autoregulatory interactions, which are encapsulated in a matrix $w = (w_{ij})$. The elements w_{ij} of this matrix indicate the strength of the regulatory influence that gene j has on gene i (Fig. 1*a*). This influence can be either activating ($w_{ij} > 0$), repressing ($w_{ij} < 0$), or absent. These regulatory interactions can change the expression state of the network $\mathbf{S}(t)$ as time t progresses, according to the difference equation $S_i(t + \tau) = \sigma[\Sigma_{j=1}^{N} w_{ij} S_j(t)]$, where τ is a constant, and $\sigma(.)$ is a sigmoidal function whose values lie in the interval $(-1, +1)$. This function reflects cooperative regulation of gene i's expression by other genes. We focus on the strong cooperation limit where σ becomes the sign function, and thus S_i assumes values ± 1.

We are concerned here with networks whose expression state starts from a prespecified initial state $\mathbf{S}(0)$ at some time $t = 0$ during development, and arrives at a prespecified stable equilibrium state $\mathbf{S}_\infty$. We will call such a network a viable network. The initial state is determined by regulatory factors upstream of the network, which may represent signals from the cell's environment or from nearby domains of an embryo. Transcriptional regulators that are expressed in the stable equilibrium state $\mathbf{S}_\infty$ affect the expression of genes downstream of the network, and thus the course of development. The matrix w represents the (regulatory) genotype of this system, and the expression state $\mathbf{S}_\infty$ its phenotype. We here examine variation in the network genotype w through variation in the topology of a network, the "who interacts with whom," represented by values of w_{ij} that are different from zero (Fig. 1*b*). Part of the motivation to focus on topologies is biological: Because biochemical parameters determining the behavior of cellular circuitry change incessantly and are difficult to measure, circuit topologies (instead of different parameters within one topology) are becoming an increasingly important subject of study (8, 28, 29). Changes in topology correspond to the loss of a regulatory interaction ($w_{ij} \to 0$), or to the appearance of a new regulatory interaction that was previously absent. Such topological changes can occur on very short evolutionary time scales, in particular in higher eukaryotes with large regulatory regions (30).

For this model, we will show that genotype space can be traversed in small steps without changing the phenotype, a property that is crucial for evolutionary innovation in gene expression patterns. Furthermore, different novel gene expression patterns become accessible in different parts of genotype space. Our use of an abstract model of a biological system permits a clearer understanding of the relationship between robustness and evolutionary innovation and the properties of the genotype to phenotype mapping.

Results

Long-Distance Travel in a Vast Network Space. Networks with different topologies can be thought of as existing in a space that has as many dimensions (N^2) as there are regulatory interactions. We showed previously (31) that all or most networks with the same gene expression pattern $\mathbf{S}_\infty$ form a connected graph (Fig. 1 *b* and *c*) in this space. We had previously called this graph a metagraph (a graph of graphs) because each network can be viewed itself as a graph (31). However, for consistency with established terminology, we here refer to this graph as a neutral network (15). To avoid confusion between a neutral network and

its nodes, which are themselves (regulatory) networks, we will often refer to the regulatory networks as genotypes. Two genotypes are neighbors in the neutral network if they differ by only one regulatory interaction. The neutral network can be traversed in small evolutionary steps that change regulatory regions on DNA. Each such step affects just one regulatory interaction at a time, and leaves the network's gene expression pattern $\mathbf{S}_\infty$ unchanged. To encapsulate a more general notion of distance among regulatory networks within this space, we introduce a measure of the distance of genotypes (network topologies). Specifically, we define the distance D of two network topologies w and w' as

$$D(w, w') = \frac{1}{2M_+} \sum_{i,j} |\mathrm{sign}(w_{ij}) - \mathrm{sign}(w'_{ij})|.$$

Here, M_+ is the maximum number of regulatory interactions, which we restrict to explore how our results depend on this number, and sign(x) is the sign function [sign(x) = ± 1 depending on the sign of x and sign(x) = 0 for $x = 0$]. This distance function ranges from $D = 0$ for identical networks to $D = 1$ for networks with completely different topologies or organization.

Equipped with this distance measure, we first ask how far we can travel in genotype space without affecting the phenotype, that is, the equilibrium gene expression pattern $\mathbf{S}_\infty$ attained in response to an initial gene expression pattern $\mathbf{S}(0)$. In a previous contribution, we have shown that the number of viable networks, networks with the same phenotype $\mathbf{S}_\infty$, is astronomically large even for moderate N (it grows exponentially with N^2) (31). Do these networks occur in a small, localized area of genotype space? To find out, we first randomly sampled [see supporting information (SI) *Text*] networks with the same phenotype $\mathbf{S}_\infty$, and determined the distribution of genotype distances D between them. SI Fig. 4*a* shows that the mean distance of random networks is greater than $D = 0.8$. This means that randomly chosen pairs of networks with the same phenotype have vastly different organization, suggesting that networks with the same phenotype can be found in very distant "corners" of network space. This observation is not a peculiarity of the particular distance measure we use: It also holds, for a second, different distance measure D', as SI Fig. 5 shows.

The suggestion that very distant networks can have the same phenotype is confirmed by a complementary analysis, which asks about the maximal genotype distance in a large random sample of network pairs with the same phenotype. This maximum distance of randomly chosen network pairs is a lower bound for the maximum possible genotype distance for networks of the same genotype. SI Table 1 shows that this maximum distance is generally large and often equal to the maximum possible distance $D = 1$.

In summary, most or all of the vast space of network topologies can be traversed in small steps without changing a network's phenotype (gene expression pattern $\mathbf{S}_\infty$).

Distant Travel Is Necessary to Find All New Phenotypes. In the context of our model, evolutionary innovation is innovation in gene expression patterns $\mathbf{S}_\infty$. For a network with N genes, the total number of such gene expression patterns is large (2^N). Here, we ask the following: Given one viable network, how far do we have to travel in the space of genotypes to encounter new phenotypes, especially phenotypes that are very different from the network's original phenotype $\mathbf{S}_\infty$? To address this question, we define as a distance measure of two phenotypes $\mathbf{S}_\infty$ and $\mathbf{S}'_\infty$ the Hamming distance $d = 1 - \Sigma_j\, \delta[\mathbf{S}_\infty(j), \mathbf{S}'_\infty(j)]/N$, which ranges from zero to one. (δ is the Kronecker δ-function, which is equal to one if its two arguments are equal and zero otherwise.)

A first step toward answering this question is to examine the fraction of new phenotypes $\mathbf{S}_\infty$ (relative to the total number of phenotypes, 2^N) that occur in networks at a given distance D to a reference network w. This requires us to examine networks that differ from w in k regulatory interactions (a k-neighborhood of w). For small k, one can easily visualize graphically the appearance of new phenotypes as a function of k (SI Fig. 6). For larger k, a quantitative analysis is needed (SI Fig. 4*b*). To obtain the data shown in SI Fig. 4*b* we chose, at random, an initial state $\mathbf{S}(0)$ and, also at random, a reference network that starting from $\mathbf{S}(0)$ arrives at some equilibrium state $\mathbf{S}^r_\infty$. For each given genotype distance D, we enumerated all networks at this distance from w and their equilibrium states $\mathbf{S}_\infty$, if any; then we divided the number of distinct equilibrium states found in these networks by 2^N (the total number of possible equilibrium states) to obtain the fraction of distinct equilibrium states in the neighborhood considered. To obtain reliable estimates, we repeated this procedure for at least 25 further reference networks. SI Fig. 4*b* shows the fractions of distinct equilibrium states, averaged over our reference networks. The figure demonstrates that one needs to travel far from any one network, that is, through about half the genotype space ($D = 0.5$), to find the great majority of the 2^N possible phenotypes. We note that the set of networks at a fixed distance D from any network contains only a small fraction of all networks (e.g., a fraction 10^{-6} for $D = 0.5$ and $n = 10$).

In a complementary analysis, we asked about the phenotype distance of networks whose genotype distance is D. Specifically, we pursued a sampling strategy that estimates the probability $P_D(d)$ that two viable networks with genotype distance D have phenotypes with distance d. The data (SI Fig. 7 *a* and *b*) clearly show that evolving networks have substantial "memory" of past phenotypes. Networks at genotypic distances up to $D = 0.5$ from any one network w are much more likely to have phenotypes similar to that of w than one would expect by chance alone; this is clear in particular from the increase with D of the mean of d (see SI Fig. 7*b Inset*). Only for genotype distances greater than that are network phenotypes no longer significantly correlated with that of w. Taken together, these observations mean that one has to travel far through genotype space to find all or most novel gene expression patterns that networks can produce and to erase any memory of past gene expression states.

Robustness and the Ability to Innovate Are Negatively Correlated. Robustness to environmental change and internal noise on the one hand and to mutations on the other hand are two different aspects of robustness in the circuits we study. In both cases, the robust feature is the network's equilibrium gene expression pattern $\mathbf{S}_\infty$. Robustness to noise corresponds to robustness of $\mathbf{S}_\infty$ to changes in the initial expression pattern $\mathbf{S}(0)$ or to perturbations in the time evolution of the dynamics. Robustness to mutations corresponds to robustness of $\mathbf{S}_\infty$ to changes in regulatory interactions. A network's mutational robustness R_μ is simply the fraction of its immediate neighbors in genotype space that have the same phenotype. (In the graph representation of Fig. 1*c*, it is the degree of the node in the neutral network divided by the total number of neighbors, viable or not.) We previously showed (31) that the different measures of robustness are strongly correlated, so from here on we shall focus mainly on mutational robustness, because its geometric interpretation is simplest. We also found that the value of robustness varies widely among networks with the same phenotype (31). This is already the case for very small networks, as exemplified by Fig. 1*c*.

We next asked how network robustness relates to the ability to find new phenotypes through regulatory mutations. To address this question, we studied a sample of 3,000 randomly chosen networks. Each of these networks has an initial gene expression state $\mathbf{S}(0)$ and arrives at a given equilibrium state $\mathbf{S}_\infty$. For each network w, we then determined its mutational robustness $R_\mu(w)$. In addition, we examined all networks at a distance

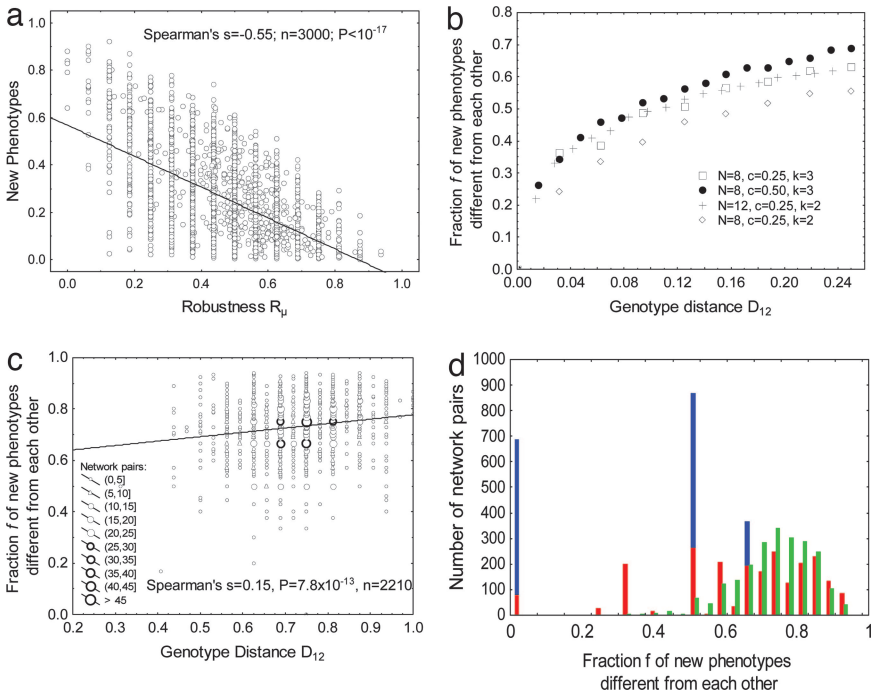

Fig. 2. The ability to innovate depends on a genotype's position in the neutral network. (*a*) A tradeoff between robustness and innovation. The horizontal axis shows mutational robustness R_μ, the fraction of a network's w topological neighbors that share the same equilibrium expression state, S_∞, with w. For each network w whose robustness is displayed on the horizontal axis, the vertical axis shows the fraction of networks of genotype distance $D < 0.1$ around w, whose equilibrium state is different from S_∞. This fraction declines with increasing robustness. $n = 8$, $c = 0.25$, and $d = 0.5$. (*b*) The horizontal axis shows genotype distance D_{12} of two networks (w_1 and w_2) with the same phenotype. The vertical axis shows the mean fraction f of unique new phenotypes, as defined in the main text, found in a k-neighborhood (see legend for k) around these networks. If f is close to zero, then all or most of the phenotypes of networks in the two neighborhoods are identical. If f is close to one, then almost all phenotypes in the two neighborhoods are different. Standard deviations around each data point are no greater than 8×10^{-3}. (*c*) Like *b*, except for a sample of 2,210 network pairs (w_1 and w_2) chosen at random from the neutral network and with mutational robustness R_μ in the interval (0.45, 0.60). $n = 8$, $c = 0.25$, $d = 0.5$, and $k = 3$. As opposed to the strong and positive statistical association between genotype distance and f for networks at small D_{12}, this association is considerably weaker at larger distances. Notice the large fraction of unique new phenotypes for almost all network pairs shown (mean $f = 0.73$). (*d*) Histogram of f for 1-neighbors (blue), 2-neighbors (red), and 3-neighbors (green) of 2,210 randomly chosen network pairs with R_μ in the interval (0.45, 0.60). Data are shown for $n = 8$, $c = 0.25$, and $d = 0.5$. For one-mutant neighbors, the robustness R_μ of a network w is the fraction of a network's neighbors that has the same gene expression pattern S_∞. For k-neighbors with $k > 1$, we define R_μ as the fraction of all networks that differ from w by no more than k regulatory interactions and that have the same gene expression pattern S_∞.

no greater than D around w in genotype space, and identified all networks w' in this neighborhood that have some equilibrium gene expression state $\mathbf{S}_\infty$. We then determined the fraction $f(w)$ of the networks w' in this neighborhood for which $\mathbf{S}_\infty$ was different from the $\mathbf{S}_\infty$ of w. This fraction indicates how readily changes in w can lead to new phenotypes. By repeating this procedure for all networks w in our sample, we arrive at 3,000 values for $[R_\mu(w), f(w)]$. Fig. 2*a* illustrates the results for the neighborhood of radius $D = 0.1$; we see a strong negative association between robustness R_μ and the indicator $f(w)$ of the ability to innovate (Spearman's $s = -0.55$, $P < 10^{-17}$; $n = 8$, $c = 0.25$, and $d = 0.5$). A highly significant negative association exists also when $R_\mu(w)$ is replaced by either of two measures of robustness to noise: (*i*) the likelihood that a change in a single gene's expression pattern in $\mathbf{S}(0)$ changes $\mathbf{S}_\infty$ (Spearman's $s = -0.29$, $P < 10^{-17}$) or (*ii*) the fraction of genes whose expression pattern needs to change in $\mathbf{S}(0)$ such that the probability of arriving at $\mathbf{S}_\infty$ falls to one-half (Spearman's $s = -0.38$, $P < 10^{-17}$). Negative associations between robustness and innovative potential $f(w)$ are also observed for different values of N, c, D, and d (results not shown).

The Large Diameter of Neutral Networks is Critical for Phenotypic Diversity. In our final analysis, we examined pairs of networks with the same phenotype and similar robustness R_μ but at different positions in genotype space. Consider such a pair of networks, w_1 and w_2, which have some genotype distance D_{12}. In a small neighborhood with radius D around each of these two networks in genotype space, we will find two sets of new phenotypes called $\{P_1\}$ and $\{P_2\}$. We are interested in the phenotypes that are "unique," i.e., that appear in $\{P_1\}$ or $\{P_2\}$ but not in both. Let f be the fraction of these phenotypes, calculated as $f = 1 - |P_1 \cap P_2|/(|P_1| + |P_2| - |P_1 \cap P_2|)$, where $|P|$ indicates the number of elements in the set $\{P\}$. If this fraction is small, even for networks of large genotype distance, then we would conclude that most networks of the same phenotype produce the same or similar new phenotypes when mutated. In contrast, if we found that this fraction is large, we would conclude that a network's position in genotype space critically determines what kind of innovations (new phenotypes) can be produced from it. In this case, the large typical genotype distances of nodes on the neutral network (SI Fig. 4*a*) could lead to very diverse new phenotypes, some fraction of which could lead to innovations.

Fig. 2*b* shows that the fraction f of unique new phenotypes increases very rapidly with increasing genotype distance of two networks w_1 and w_2. For instance, for networks of $n = 8$ genes, at a genotype distance of merely $D_{12} = 0.06$, 34% of new phenotypes found in the 2-neighborhood of two networks are different. As the genotype distance increases, this positive association between genotype distance and the diversity of innovation decreases, as Fig. 2*c* shows. In Fig. 2*c*, each genotype pair is chosen at random from a neutral network, the only constraint being that the mutational robustness R_μ of both

genotypes lies in the narrow interval (0.45, 0.60). However, at these large (and typical) distances for two genotypes on the neutral network, the mean fraction of unique phenotypes for the two genotypes is high: $f = 0.73$. In other words, most phenotypes found in a neighborhood of these genotypes are unique. Finally, Fig. 2*d* shows histograms of f for the same sample of genotypes used in Fig. 2*c*. Blue, red, and green bars represent f in 1-, 2-, and 3-neighborhoods, respectively. An appreciable fraction of network pairs has the same phenotypes in the 1-neighborhood ($f = 0$). Nonetheless, even in 1-neighborhoods the average fraction of unique phenotypes is >0.4. In other words, in two randomly chosen genotypes on a neutral network, a random mutation that produces new phenotypes has a $>40\%$ chance to produce new phenotypes that differ from each other. In the 2- and 3-neighborhoods, the distribution of f is shifted to the right, indicating even greater diversity of new phenotypes. Qualitatively identical patterns are observed for different network sizes N, fraction c of interacting genes, and mutational robustness R_μ.

In summary, two networks with exactly the same phenotype may produce very different innovations, depending on their organization, i.e., their position in genotype space. On a final note, we emphasize that, although we reported numerical work only for networks in which regulatory interactions can take one of three values ($w_{ij} = \pm 1, 0$), our key results also hold for continuously valued regulatory interactions. Specifically, vastly different genotypes can have the same phenotype, genotypes close together can have uncorrelated phenotypes, and genotypes in different positions on the neutral network can produce very different new phenotypes.

Discussion

In summary, we have shown that networks with vastly different organizations can have the same phenotype. In contrast, two networks with completely unrelated phenotypes can be found very close to each other in genotype space, even though changing a genotype at random will often lead to highly similar phenotypes. This latter property, a genotype's long "memory" of past phenotypes, is not self-evident. For instance, it is not observed in another kind of biological system, RNA, where the relationship between genotype (nucleotide sequence) and one aspect of phenotype (secondary structure) has been thoroughly explored (15, 32–36). Indeed, in that system, even very few changes in a molecule's nucleotide sequence can completely randomize the molecule's structure (15). Within our models of gene networks, the long genotypic memory, together with the existence of regions of high robustness in a neutral network (31), relate to the phenomenon of developmental constraints (37), where genotypic variation leads to little or no variation in one aspect of an organism's phenotype, that is, its development and its body plan. This relation between genotype memory and constraint is intriguing, because our model abstractly represents the kinds of transcriptional regulation networks that pattern some of the most constrained body plan features of organisms, including the early segmentation genes in flies, and the *Hox* genes involved in axial patterning of most animal phyla (19, 38).

Under point mutations, regulatory networks of identical phenotype can produce very different new phenotypes, depending on their location in genotype space. Taken together with our previous results (19), this means that both the connectedness of a neutral network (with the robustness that this implies) and the fact that a neutral network spans genotype space are crucially important for both robustness and evolutionary innovation. Fig. 3 illustrates that neither feature separately would achieve both robustness and evolutionary innovation. If there were multiple genotypes that produced the same phenotype, but if these genotypes were isolated from each other (Fig. 3 *Left*), a network would be neither robust nor capable of producing many different evolutionary innovations. If the genotypes were connected but

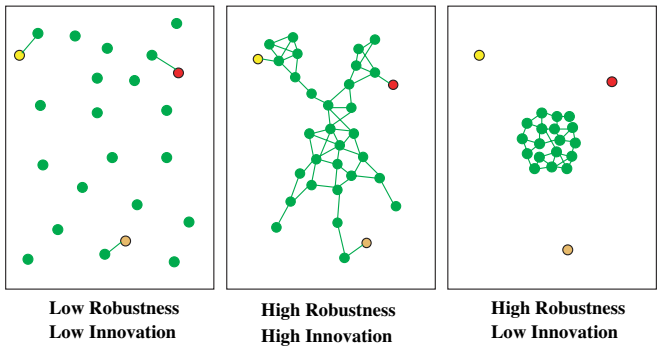

Fig. 3. Conditions for high robustness and the ability to innovate. Each rectangle shows a hypothetical genotype space. Individual genotypes with identical phenotypes (regulatory networks that produce identical gene expression patterns) are shown as circles in this space. Nodes of the neutral network are green. Other colors indicate novel phenotypes. Lines connect genotypes that are nearest neighbors in this space, corresponding in our case to networks that differ in one regulatory interaction. (*Left*) Genotypes are widely scattered and isolated in this space. (*Center*) The genotypes are widely scattered but also connected in this space. (*Right*) The genotypes occur in a small region of the space, and they are connected. We note that this visualization is for expository purposes only. Actual genotype spaces may have hundreds of dimensions, and there may be an astronomical number of genotypes with the same phenotype.

highly localized in genotype space (Fig. 3 *Right*), robustness would be high, but the ability to innovate would be limited because the neighbors of these genotypes would produce very few novel phenotypes. Only when paths through genotype space connect many different networks with identical phenotypes (Fig. 3 *Center*) are both robustness and evolutionary innovation achieved. Then evolving networks can reach different locations in genotype space, which makes the generation of diverse new phenotypes possible. The intermediate to high robustness implied by the connectedness of the neutral network is thus a prerequisite for the ability to innovate. It is important to appreciate that this conclusion would not emerge from studying individual networks and their close neighborhoods: There, robustness and innovation are necessarily antagonistic because having more neighbors of the same phenotype leaves fewer possibilities for neighbors with new phenotypes. To appreciate the innovative potential of our model regulatory networks, one needs to consider the system at a higher level, namely that of the neutral network. If the patterns we observe hold in general for biological systems, then the ability of living organisms to innovate is an emergent property, a feature typical of many complex systems.

Among the two factors that influence the outcome of biological evolution, selection and the production of variation through mutation, we here focused entirely on variation. In doing so, we did not intend to diminish the role of natural selection, because we are acutely aware that the vast majority of mutations in any genetic system are deleterious and that only a small fraction may lead to evolutionary innovations. However, we emphasize that only a genetic architecture like the one of Fig. 3 *Center* can explore the great diversity of new phenotypes needed to sift potential innovations through natural selection.

We note that the phenomenon we describe would not be the only determinant of a biological system's ability to innovate. Other candidates include the mutation rate and the modular organization of biological systems (3, 39–42). With few exceptions, however, our understanding of evolutionary innovation comes from a large number of individual case studies. Albeit beautiful examples of natural history, they may not add up to fundamental evolutionary principles that allow innovations to

emerge. Abstraction, with all of the healthy skepticism that it arouses, may be the only access road to such principles.

Methods

The numerical procedures we used fall into three classes. First, iteration of discrete time equations for the gene expression states, S_i, which is straightforward and has no conceptual subtleties. Second, exact enumeration of elements in a finite space by "brute force." In *SI Text*, we show how combinatorial techniques can make such enumerations computationally efficient. Third, unbiased sampling of elements in a finite space. This is the object of Monte Carlo algorithms that are standard tools of computational physics (43, 44).

To characterize classes of gene regulatory networks (genotypes) and their gene expression states (phenotypes), one must be able (*i*) to sample each class uniformly and (*ii*) to search a given neighborhood of a network. Here, we give an overview of the approaches we took (see *SI Text* for details).

To sample viable networks with equal probability, we simply generate interaction matrices w at random and determine whether they attain the desired $\mathbf{S}_\infty$. This approach becomes ineffective for networks with substantially more than $n = 10$ genes. For such larger networks, we thus use a permutation transformation to reduce the complexity of the task (31). Because the total number of networks (regardless of their expression states) is easily calculated analytically, a numerical estimate of the probability that a random network is viable suffices to estimate the total number of viable networks.

Second, we need to characterize the networks in some k-neighborhood of a given viable network w. (A k-neighborhood comprises all networks that differ from w by at most k-regulatory interactions.) The networks in this neighborhood may also be viable; they may attain some other equilibrium gene expression pattern; or no equilibrium at all. For small k, we can enumerate all k-neighbors exhaustively using an extension of the Monte-Carlo technique. For larger networks, we use a sampling approach that is guided by a combinatorial representation of the total number of networks at a given distance k from w. This last approach is rigorous but the cost is greater computational complexity (see *SI Text* for details).

To characterize the neutral network(s) associated with one or more phenotypes, we determine (*i*) the distribution of distances of randomly chosen genotype pairs having the same phenotype; (*ii*) the maximum distance of two such genotypes; and (*iii*) the minimum distance between two neutral networks. We do so by Monte Carlo sampling of the genotypes that can be reached from a given reference genotype through a series of point changes (modifying one interaction at a time). Specifically, we perform a random walk which follows the Metropolis rule (45), thereby enforcing that a neutral network is sampled uniformly.

The Laboratoire de Physique Théorique et Modèles Statistiques is a Unite de Recherche de l'Universite Paris-Sud and Centre National de la Recherche Scientifique. This work was supported by National Institutes of Health Grant GM063882, Swiss National Foundation Grant 315200-116814 (to A.W.), the European Economic Community FP6 program under Contract IST-001935 (EVERGROW), the European Economic Community's Health Protection and Promotion Project under Contracts HPRN-CT-2002-00307 (to the DYGLAGEMEM European Network) and HPRN-CT-2002-00319 (to the STIPCO European Network), and the Santa Fe Institute.

1. Grant P (1986) *Ecology and Evolution of Darwin's Finches* (Princeton Univ Press, Princeton).
2. Piatigorsky J, Wistow GJ (1989) *Cell* 57:197–199.
3. Gerhart J, Kirschner M (1998) *Cells, Embryos, and Evolution* (Blackwell, Boston).
4. Rechenberg I (1973) *Evolutionsstrategie* (Frommann–Holzboog, Stuttgart, Germany).
5. Keymeulen D, Zebulum R, Jin Y, Stoica A (2000) *IEEE Trans Rel* 49:305–316.
6. Sanchez E, Tomassini, M. eds. (1996) *Towards Evolvable Hardware: The Evolutionary Engineering Approach* (Springer, New York).
7. Wagner A (2005) *Robustness and Evolvability in Living Systems* (Princeton Univ Press, Princeton).
8. Meir E, von Dassow G, Munro E, Odell G (2002) *Curr Biol* 12:778–786.
9. Gonze D, Halloy J, Goldbeter A (2002) *Proc Natl Acad Sci USA* 99:673–678.
10. Edwards JS, Palsson BO (2000) *Biotechnol Prog* 16:927–939.
11. van Nimwegen E, Crutchfield J, Huynen M (1999) *Proc Natl Acad Sci USA* 96:9716–9720.
12. Alon U, Surette MG, Barkai N, Leibler S (1999) *Nature* 397:168–171.
13. Kirschner M, Gerhart J (1998) *Proc Natl Acad Sci USA* 95:8420–8427.
14. Bornholdt S, Sneppen K (2000) *Proc R Soc London Ser B* 267:2281–2286.
15. Schuster P, Fontana W, Stadler P, Hofacker I (1994) *Proc R Soc London Ser B* 255:279–284.
16. Lipman D, Wilbur W (1991) *Proc R Soc London Ser B* 245:7–11.
17. Jaeger J, Reinitz J (2006) *BioEssays* 28:1102–1111.
18. Sharp DH, Reinitz J (1998) *Biosystems* 47:79–90.
19. Jaeger J, Surkova S, Blagov M, Janssens H, Kosman D, Kozlov K, Manu, Myasnikova E, Vanario-Alonso C, Samsonova M, *et al.* (2004) *Nature* 430:368–371.
20. Mjolsness E, Sharp DH, Reinitz J (1991) *J Theor Biol* 152:429–453.
21. Wagner A (1996) *Evolution (Lawrence, Kans)* 50:1008–1023.
22. Bergman A, Siegal M (2003) *Nature* 424:549–552.
23. Siegal M, Bergman A (2002) *Proc Natl Acad Sci USA* 99:10528–10532.
24. Masel J (2004) *J Evol Biol* 17:1106–1110.
25. Azevedo RBR, Lohaus R, Srinivasan S, Dang KK, Burch CL (2006) *Nature* 440:87–90.
26. Aldana M, Cluzel P (2003) *Proc Natl Acad Sci USA* 100:8710–8714.
27. Nochomovitz YD, Li H (2006) *Proc Natl Acad Sci USA* 103:4180–4185.
28. von Dassow G, Meir E, Munro E, Odell G (2000) *Nature* 406:188–192.
29. Ingolia NT (2004) *PLoS Biol* 2:805–815.
30. Stone J, Wray G (2001) *Mol Biol Evol* 18:1764–1770.
31. Ciliberti S, Martin O, Wagner A (2007) *PLoS Comp Biol* 3(2):e15.
32. Ancel LW, Fontana W (2000) *J Exp Zool/Mol Dev Evol* 288:242–283.
33. Reidys C, Stadler P, Schuster P (1997) *Bull Math Biol* 59:339–397.
34. Fontana W, Konings D, Stadler P, Schuster P (1993) *BioPolymers* 33:1389–1404.
35. Huynen M (1996) *J Mol Evol* 43:165–169.
36. Fontana W, Schuster P (1998) *Science* 280:1451–1455.
37. Maynard-Smith J, Burian R, Kauffman S, Alberch P, Campbell J, Goodwin B, Lande R, Raup D, Wolpert L (1985) *Q Rev Biol* 60:265–287.
38. Carroll SB, Grenier JK, Weatherbee SD (2001) *From DNA to Diversity: Molecular Genetics and the Evolution of Animal Design* (Blackwell, Malden, MA).
39. Schlosser G, Wagner G, eds (2004) *Modularity in Development and Evolution* (Univ of Chicago Press, Chicago).
40. Wagner GP, Altenberg L (1996) *Evolution (Lawrence, Kans)* 50:967–976.
41. Thieffry D, Romero D (1999) *Biosystems* 50:49–59.
42. Csete M, Doyle J (2002) *Science* 295:1664–1669.
43. Hammersley J, Handscomb D (1964) *Monte Carlo Methods* (Chapman & Hall, London).
44. Krauth W (1998) in *Advances in Computer Simulation*, eds Kertesz J, Kondor I (Springer, Berlin).
45. Metropolis N, Rosenbluth A, Rosenbluth M, Teller A, Teller E (1953) *J Chem Phys* 21:1087–1092.