



Population and Quantitative Genetics of Regulatory Networks

STEVEN A. FRANK*

*Department of Ecology and Evolutionary Biology, University of California, Irvine,
CA 92697-2525, U.S.A.*

(Received on 20 July 1998, Accepted in revised form on 12 November 1998)

I evolved boolean regulatory networks in a computer simulation. I varied mutation, recombination, the size of the network, and the number of connections per node. I measured the performance of networks and the heritability and epistasis of genetic effects. Networks of intermediate connectivity performed best. The distinction between metabolic and quantitative genetic additivity explained some of the variation in performance. Metabolic additivity describes the interaction between changes in a single network, whereas quantitative genetic additivity measures the consistency of phenotypic effect caused by gene substitution in randomly chosen members of the population. I analysed metabolic additivity by the distribution of epistatic effects of pairs of mutations in individual networks. I measured quantitative genetic additivity by heritability. Highly connected networks had greater metabolic additivity for perturbations to individual networks, but had lower additivity when measured by the average effect of a gene substitution (heritability). The lower heritability of highly connected nets appeared to reduce the effectiveness of recombination in searching evolutionary space.

© 1999 Academic Press

Introduction

Gene networks control the expression of characters. Such networks can be studied from a mechanistic point of view. What are the particular gene products involved? How do control factors turn on and off the expression of individual genes? How is the timing of expression controlled? What is the relation between regulatory dynamics and character expression?

These mechanistic questions can be complemented by population-level questions. How variable are characters? How variable are particular regulatory loci? What is the relationship between variability in the regulatory cascade and

variable expression of the character? In the statistical language of quantitative genetics, what is the nature of heritability, dominance, and epistasis?

Mechanistic and population-level questions concern description of regulatory networks. What about the evolutionary processes that have shaped gene networks? Are there design properties of regulatory cascades that are the hallmarks of natural selection or of evolutionary history? Are there predictable mechanistic and population-level features?

I analyse a simple model of a regulatory network. The network controls a quantitative character that is favored to be highly expressed through the first half of life and turned off during the second half of life. This can be thought of as

* E-mail: safrank@uci.edu.

a fitness function that favors high expression during early development and juvenile growth followed by shutting down of expression in adult life. I begin with a population of individuals, each of which has a randomly connected network. I evolve this population through a typical cycle that includes selection, mating, recombination, and mutation.

I use this model to address the final question from the above framework: are there predictable population-level features of regulatory networks designed by natural selection? My model cannot, of course, address even this very limited question in a general way. But the model does provide some insight into how regulatory networks accumulate mutations and epistatic interactions. Perhaps more importantly, such models are helpful in clarifying how to proceed with the interesting questions that define the subject.

Background

BOOLEAN MODELS

Kauffman (1969, 1974) initiated a series of studies on boolean networks as models of gene expression and control of developmental pattern (Fig. 1 defines boolean networks). These models assume that all components of a regulatory network can be approximated by a boolean function of digital states. Although Kauffman emphasized that this is clearly inexact, he outlined three benefits of his approach (for synthesis, see Kauffman, 1993). First, many kinetic control components exert their influence according to thresholds which are well approximated by a boolean function. Second, boolean networks have great computational power—sufficient to model the essential features of most control structures. Third, arbitrary control structures are too complex to draw any general theoretical conclusions; abstraction to purely boolean control allows some theoretical analysis.

Kauffman (1993) characterized networks by the number of nodes, n , and the number of inputs to each node, k . He analysed the dynamics of networks with randomly chosen boolean functions. The dynamics can be studied in classical fashion by the extent to which a small perturbation changes the subsequent trajectory.

pro	pro	Xis	Int
1	1	0	0
1	1	1	0
1	0	0	1
1	1	0	1
0	0	0	0
0	0	1	0
0	0	1	1
0	1	1	1

FIG. 1. Example boolean function. The left column is the output of the function (0 or 1) given the input states (0 or 1) in the other three columns. The table shows one step in the regulatory control of the lysis–lysogeny switch of the bacteriophage λ (McAdams & Shapiro, 1995). The virus is in the integrated prophage (pro) state of the lysogeny pathway in the next time step (a) if the virus is currently a prophage and the concentration of the Int regulatory protein is below a threshold (top two rows) or (b) if Int concentration is above its threshold and the concentration of the Xis regulatory protein is below its threshold (middle two rows).

Kauffman showed that weakly connected networks with relatively small k tended to be frozen, that is, small perturbations had little or no effect on subsequent trajectories. Such limited effects occurred because perturbations were propagated slowly or not at all through the network. By contrast, strongly connected nets with relatively high k tended to be chaotic, that is, small perturbations cause large, rapid deviations in subsequent trajectories. Networks with intermediate connectivity respond to small changes in network control with small deviations in subsequent trajectories.

Kauffman argued that natural networks must evolve to an intermediate level of connectivity. If k is too high, then small changes in control structure lead to chaotic, unpredictable consequences, clearly a poor property of a control system. If k is too low, then adjustments in output are difficult because control information is not propagated effectively. Kauffman used the colorful phrase “edge of chaos” to describe intermediate connectivity. I prefer “zone of linearity” because Kauffman is in fact emphasizing that control structures are most useful when changes in input have approximately linear

effects on output. Thus the key is to find the zone of linearity within control structures such as boolean nets that, for most topologies, tend to propagate in a highly nonlinear fashion. The models I present below clarify the meaning of linearity in this context and connect these ideas to traditional quantitative genetic concepts such as heritability.

Kauffman & Smith (1986) evolved boolean networks to analyse the evolutionary approach to fitness optima. In particular, they used a fitness function that measures the number of matching bits (states of nodes as either 0 or 1) to some arbitrary pattern. They compared the ability of networks with $k = 2$ and $k = 10$ inputs per node to approach a local fitness optimum. Interestingly, there was no difference for low and high k with respect to the best network of a population—in both cases the best networks were able to match about 70% of the target bits. When the best network of a population was mutated at a single component, an interesting difference was observed between weakly ($k = 2$) and strongly ($k = 10$) connected networks. The weakly connected networks suffered small, continuous fitness degradation as additional mutations were added to the best network, whereas very few mutations reduced the best of the strongly connected networks to the fitness of an average member of the population. These observations support Kauffman's (1993) theoretical analysis of random boolean nets, in which $k = 2$ is typically associated with the zone of linearity.

POPULATION AND QUANTITATIVE GENETICS

Three distinct points of view are commonly used to discuss genetic effects on regulatory networks—metabolic, population genetic, and quantitative genetic (Phillips, 1998). It is important to distinguish clearly the aims, methods, and language used by each approach.

Metabolic studies analyse the effects of each gene product in a biochemical network of interactions (Kacser & Burns, 1979; Hartl *et al.*, 1985). Consider, for example, the statement that two gene products act additively to influence the total flux through the pathway. This may mean that, if each mutated gene increases flux by $x\%$, then both mutants together increase flux by

$2x\%$. This statement does not necessarily have any relationship to fitness or to the pattern of genetic variability found at various steps in the biochemical network. The measurements assume a particular wild-type set of gene products against which the performance of variants is measured.

Population genetic analyses assume, or measure, the relationship between variants and fitness. A full analysis must be able to assign a fitness value to each genotype. If genetic variation is rare, then each genotype is likely to have only one or two allelic variants in a particular biochemical network when compared with the common (wild) type. Then statements about fitness can be closely related to the metabolic descriptions in the previous paragraph. For example, if variants have particular effects on flux, then fitnesses must be assigned to different levels of flux for each variant alone and in combination. If genetic variation is common, then there may be almost as many genotypes as individuals in the population. Assignment of fitnesses for each genotype requires almost complete knowledge of the biochemical outcome of simultaneous variation at many steps in the network and the relationship of the biochemical outcome to fitness. This kind of population genetic analysis is therefore limited to small networks of two or three components or to rare variants in an otherwise homogeneous population.

Many regulatory networks appear to have significant genetic variation. Metabolic and population genetic analyses are therefore of limited use at the population level. Instead, a statistical or quantitative genetic approach is used to characterize populations (Keightley, 1989; Clark, 1991). Most importantly, a statistical rather than metabolic notion of additivity plays a central role. I emphasize two measures.

The first measure of additivity is based on Fisher's (1958) average effect of a gene substitution. The average effect is the average phenotypic consequence of replacing an allele, in each individual in the population, by a particular genetic variant. This measure is important because, in a sexually reproducing population, parents transmit alleles rather than whole genotypes. A parent's contribution to the future

of the population depends on the average effects of its alleles over the distribution of genetic backgrounds in the population. The amount of phenotypic variation in the population explained by average effects of alleles is called the additive genetic variance. The ratio of additive variance to total phenotypic variance is the fraction of variability that is transmissible to future generations, thus this ratio is also called the heritability (Falconer & Mackay, 1996; Lynch & Walsh, 1998).

The second key measure is the amount of non-additivity, or epistasis, that occurs between pairs of alleles. Consider a network in which two control coefficients are varied, for example, an individual with two mutations affecting enzymes of a metabolic pathway. For many theories of population genetics, it is most useful to define such interaction on a logarithmic scale (Kondrashov, 1988, 1993; Charlesworth, 1990). First, define $W(x)$ as the performance of a mutant genotype divided by the performance of a standard, non-mutated genotype. Given this measure of relative performance, the expected relative performance of a double mutant with both mutations i and j is defined as

$$\begin{aligned} \text{Expected}(\log[W(i, j)]) &= \log[W(i)W(j)] \\ &= \log[W(i)] + \log[W(j)] \end{aligned}$$

under the assumption that the effects of mutations are additive on the logarithmic scale. Here $W(i, j)$ is the performance of the double mutant and $W(i)$ and $W(j)$ are the performances of the single mutants.

With these definitions we can express the interaction, or deviation from additivity, for two mutations, i and j , as

$$I(i, j) = \log[W(i, j)] - \log[W(i)] - \log[W(j)]. \quad (1)$$

When $I = 0$, the mutations act additively on a logarithmic scale; $I > 0$ implies that the mutations have a positive interaction, performing better when together than expected from the sum of their individual effects; and $I < 0$ implies negative interaction and poorer performance for the pair than expected from their individual effects. Population genetic theories suggest that whether the typical value of I tends to be positive

or negative has a strong influence on the evolution of genetic systems (Kondrashov, 1988, 1993; Charlesworth, 1990). Although this measure can be applied to the metabolic effects in a particular individual, the key measure in evolutionary theories of genetic systems is the distribution of I for all pairs of mutations over all individuals in the population.

Model

Each individual in the population is a boolean network with n nodes. Each node contains: the current state, 0 or 1; a list of k inputs from other nodes, where any particular input can be unconnected and register a constant value of 0; and a boolean function that maps each of the 2^k possible input states to an output of 0 or 1. A bit string of length 2^k encodes the boolean function. Because each the 2^k possible input states corresponds to a variable output bit, there are 2^{2^k} possible boolean functions for k inputs (Kauffman, 1993).

I encoded the information for each of the n nodes of an individual along a single linear chromosome (haploid genetics). Each individual expresses two characters. I designated the first eight nodes as character one and the second eight nodes as character two. For each character, the position of the eight output nodes is given by $i = 0, \dots, 7$, and the phenotype is $z = \sum_i 2^i S_i$, where S_i is the state of the i -th node. Each character is an unsigned byte with range $0, \dots, 255$. This encoding yields quantitative characters with control that varies in effect by seven steps of magnitude two, providing an opportunity for major effects as well as smaller modifying effects.

The genotype of each individual encodes the list of inputs and a boolean function for each of the n nodes. The phenotype is the trajectory of bivariate character states over D developmental steps, in which the network is updated in each step by assigning the state for each node from its input state and boolean function. The initial state for all nodes is zero. The developmental period is divided into two phases. In the first phase, selection favors maximal expression of character one and minimal expression of character two. Selection favors reversal of character




FIG. 2. Optimal trajectory for two characters. The development of each individual occurs over $D = 50$ time steps. Character 1 is favored to be maximally expressed in the first half of life and off in the second half, whereas character 2 is favored to be expressed only in the second half of life. The deviation of fitness from the optimum is measured by the average distance of the bivariate trajectory from the optimal trajectory.

expression in the second phase, that is, minimal expression of character two and maximal expression of character one. The bivariate trajectory with maximal fitness is shown in Fig. 2. In general, if z_{ij} is the average expression of the j -th character during the i -th phase of development, then fitness is proportional to $F = (z_{11} - z_{12}) +$

$(z_{22} - z_{21})$. Each character ranges over $0, \dots, 255$, so the maximum of F is 510. Figure 3 shows a sample developmental sequence.

This simple fitness function has two advantages. First, regulatory problems in development often concern timing of expression—many morphogens and hormones must be expressed for

Time	Character 1	Character 2	Character 1	Character 2	
0:	1 0 1 1 0 1 1 1	0 0 0 0 0 0 0 0	0 0 1 0 0 1 0 0 1 0 0 1 1 0 0 0 0 0 0 0 0 0 1 0 0 0 1 1 1 0 0 0 0 0 0 0	237	0
1:	0 0 0 0 1 1 0 1	0 0 0 0 0 0 0 0	0 0 0 0 0 1 0 1 0 0 0 0 0 0 0 1 0 0 0 0 1 0 1 1 1 0 1 0 0 0 0 0 0 0 0	176	0
2:	0 0 1 0 0 1 1 1	0 0 0 0 0 0 0 0	0 1 1 0 0 1 0 0 1 0 0 1 1 0 0 0 0 0 1 0 0 0 1 1 0 0 1 0 0 0 0 0 0 0 0 0	228	0
3:	0 0 0 1 1 1 1 1	0 0 0 0 0 0 0 0	0 0 0 0 0 1 0 1 1 0 0 0 0 1 0 0 0 0 0 0 0 0 1 1 1 1 1 1 0 0 0 0 0 0 0 0	248	0
4:	0 0 0 0 0 1 0 1	0 0 0 0 0 0 0 0	0 0 0 0 0 1 0 1 1 0 0 0 0 0 0 1 0 0 0 0 1 0 1 1 1 1 1 1 0 1 0 0 0 0 0 0	160	0
20:	1 0 0 0 0 0 1 1	0 0 0 0 0 0 0 0	0 0 0 0 0 1 0 1 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 1 1 1 0 1 0 0 0 0 1	193	0
21:	1 0 0 0 0 0 1 1	0 0 0 0 0 0 0 0	0 0 0 0 0 1 0 1 1 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 1 1 0 0 0 0 0 0 0	193	0
22:	1 0 0 0 0 0 1 1	0 0 0 0 0 0 1 0	0 0 0 0 0 1 0 1 1 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 1 1 0 0 0 0 0 0 0	193	64
23:	1 0 0 0 0 0 1 1	0 0 0 1 0 0 1 0	0 0 0 0 0 1 0 1 0 0 0 1 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 1 1 0 0 0 0 0 0 0	193	72
24:	1 0 0 0 0 0 1 1	0 1 0 1 0 1 1 0	0 0 1 0 0 1 0 0 1 0 0 1 0 0 1 0 1 0 0 0 1 0 0 1 0 1 1 0 0 0 0 0 0 0 0 0 0 0 0	193	106
25:	1 0 0 0 0 0 0 0	0 1 0 1 0 1 1 0	0 0 1 0 0 1 0 0 1 0 0 1 0 1 1 0 1 0 1 0 1 0 0 1 0 1 1 0 0 0 0 0 0 1 0 0 0 0 0 0	1	106
26:	1 0 0 0 0 0 0 0	0 1 1 1 0 1 1 1	0 0 1 0 0 1 0 0 1 0 0 0 1 1 1 0 1 0 1 0 1 0 0 1 0 0 1 0 0 0 0 0 0 1 0 0 0 0 0 0	1	238
27:	1 0 0 0 0 0 0 0	0 1 1 1 1 1 1 1	0 0 1 0 0 1 0 0 1 0 0 0 1 1 1 0 1 0 1 0 1 0 0 1 1 0 1 0 0 0 0 0 0 0 0 0 0 0 0	1	254
28:	1 0 0 0 0 0 0 0	0 1 1 1 1 1 1 1	0 0 1 0 0 1 1 0 1 0 0 0 1 1 1 0 1 0 1 0 1 0 0 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0	1	254

FIG. 3. Sample developmental sequence. The left column is the time step in the sequence. The bit strings show the state of the $n = 50$ nodes in each time step. The first eight bits define character 1; the second eight bits define character two. The right columns show the decimal values for characters one and two. Initially, all states are set to zero. The first updated time step is labeled zero in the left column. The first half of the sequence is steps 0–24; the second half is steps 25–49. The first few steps and the middle steps near the transition are shown. Character one is maintained at a high level and character two at a low level during the first half as favored by selection. The optimal time for switching is between steps 24 and 25, after which high expression of character two and low expression of character one are favored. This particular genotype produced a fitness of 437, or 86% of the maximum value of 510.

a particular period during early life and then turned off while some other regulatory molecule is turned on for the remainder of life. Second, this formulation expresses the problem of regulatory control very clearly in terms of the trajectory of network dynamics. Optimal fitness concerns tuning a transient leading to an attractor; the fitness function neatly summarizes the distance of variant trajectories perturbed from this optimal path.

A simulation run begins with a population of individuals. Each individual is initialized with a random genotype, which includes for each node the k input connections to other nodes and a boolean function. A generation consists of evaluating the fitness of each individual, choosing two parents to mate by sampling in proportion to fitness, mutating and recombining parental genotypes to form an offspring, and continuing this process to create a full population of offspring. The parents are then discarded and the offspring generation is used to start the cycle again. In each generation parents are chosen with replacement for mating.

The base mutation probability is μ . Mutation occurs in two steps. First, for each of the n nodes, mutation of one input connection occurs with probability μk , where k is the number of inputs. Mutation of an input occurs by first selecting with uniform probability one of the k connections. Then that input is set to a null connection with probability one-half, otherwise the connection goes to one of the n nodes with uniform probability. Null connections provide constant zero input. The second mutational phase affects the boolean function, which is a string of 2^k bits. A mutation occurs by flipping a particular bit, that is, by changing the output state for a given input state. The expected number of bits flipped for each node is $\mu 2^k$. This mutational scheme makes the probability that a particular input or output changes independent of n and k .

The haploid offspring genotype is created by recombining the two haploid parental chromosomes. The first of n nodes is taken from a randomly chosen parent. The next node is taken from that parent with probability $1 - r$ and from the other parent with probability r , where r is the recombination fraction. Each successive node is taken from the same parent as the last node with

probability $1 - r$. Recombination does not occur within nodes.

In summary, the key parameters are: n , the nodes per net; k , the inputs per node; p , the population size; D , the number of developmental steps to determine phenotype; g , the number of generations per run; μ , the mutation rate; and r , the recombination fraction between nodes.

Analysis

The parameter space is too large to vary all parameters in an orthogonal way. For all runs, I set the number of generations as $g = 10000$, the number of developmental steps as $D = 50$, and the population size as $p = 500$. My analysis focuses on the remaining four parameters each varied over three levels in a 3^4 factorial design. The levels are: the number of nodes, $n = 50, 100, 200$; the maximum number of connections per node, $k = 2, 4, 6$; the mutation rate, $\mu = 10^{-4}, 10^{-3}, 10^{-2}$; and the recombination rate, $r = 0.005, 0.05, 0.5$. To test the role of population size, I repeated this design with $p = 1000$ and $n = 50$. The larger population size had no significant effect, so I confine my analysis to the main 3^4 design.

The mutation rates may seem high relative to the typically quoted values on the order of 10^{-5} or 10^{-6} per locus. But in this simulation the character under study directly determines fitness, thus selection is approximately two orders of magnitude more intense than for a typical quantitative trait. The ratio of mutation to selection in the simulations is roughly consistent with the norm. This ratio determines many aspects of population-level characteristics in theoretical analyses (Barton & Turelli, 1989). In addition, approximately 70–90% of single mutations had no effect on phenotype. Thus the effective mutation rate was about one order of magnitude less than μ .

PERFORMANCE

Most evolutionary simulations of boolean networks, neural networks, and other control systems focus on the best parameters for quickly and effectively discovering optimal performance. By contrast, my goal is to evaluate the attributes of a population shaped by natural selection.




FIG. 4. Distribution of phenotypes in three runs. Plots show the cumulative distribution function for each case. Parameters are $n = 50$; $\mu = 10^{-3}$; $r = 0.05$; and k shown on each curve. The parameter k is the maximum number of inputs per node. The actual number of inputs per node for the best phenotype in the population was 1.52 for $k = 2$ (—), 2.75 for $k = 4$ (---), and 3.92 for $k = 6$ (..).

Nonetheless, the performance, measured by the distribution of fitness in the population, is a good place to begin my analysis. This allows connection to the optimization of evolutionary performance in other studies and a point of departure for describing the state of the population created by natural selection.

I measure performance by the character value for each individual as defined above. Analysis of variance shows that the mutation rate determines most of the variation in performance. The recombination rate and the maximum number of inputs, k , also have some effect, as does the interaction between k and the mutation rate. The number of nodes, n , has negligible effect.

The distribution of character values at the end of a run describes the performance of the population. Figure 4 shows distributions for three runs. It is difficult to draw plots for the distributions over many parameter values, thus I summarize performance by examining percentiles of the distribution. There was little difference between the trends when analysing the median or the upper percentiles. I show in Fig. 5 the 96th percentile to represent the performance of the best few percent in each population.

Figure 5 is a matrix of nine plots showing the performance over three levels of three parameters. Each point is the 96th percentile of character value at the end of a particular run.

The x -axis shows the three levels of mutation, which determine the greatest portion of the variation. Intermediate mutation of 10^{-3} produces the best performance. The rows show varying recombination rate. High recombination (top row) yielded lower performance than the lower treatment values. The columns show varying levels of k . An intermediate value of $k = 4$ yielded the highest performance.

The best performance was observed for intermediate values of the three key parameters: $\mu = 10^{-3}$, $r = 0.05$, and $k = 4$. The full distribution for this combination, with $n = 50$, was shown in Fig. 4 along with the distributions for $k = 2$ and $k = 6$.

I described above Kauffman's theory concerning the connectivity per node and the expected evolutionary performance of boolean networks. According to Kauffman's theory, randomly connected networks have the greatest evolutionary potential when the connections per node is near 2. Lower connectivity causes little response to change and higher connectivity yields chaotic response. In my simulations, k is the maximum number of inputs per node. My nets can be shaped by selection to have fewer connections, so k represents the potential rather than actual connectivity. Over parameter values $k = 2, 4, 6$, my simulated networks performed best with $k = 4$. Figure 6 shows the actual connectivity of the nets when compared with the maximum connectivity, k .

I measured the connectivity values in Fig. 6 as follows. I found the best performing network in the population at the end of a run. From each of the 16 output nodes, I traced the network back through each input. This yielded the set of connected nodes in the network, which was less than n , the number of nodes available to form the network. For each node in the network, including the output nodes, I calculated the average number of inputs per node.

Interestingly, when measured by the best individual in the population, the best performing networks were the $k = 4$ class. These had about three active inputs per node. However, the fitnesses of the $k = 4$ and $k = 6$ classes dropped off more quickly when moving from the best individuals in a population to the median individual (Fig. 4). In particular, the sharp jump




FIG. 5. Influence of mutation, recombination, and number of inputs on the evolutionary performance of a population. Performance is measured here by the 96th percentile of the phenotypic distribution. Each panel shows variation in response to changes in mutation, each row shows variation in response to recombination, and each column shows variation in response to the maximum number of inputs.

in the $k = 2$ class in Fig. 4 near the maximum shows that most individuals in that population had fitnesses close to the best individual. By contrast, fitnesses dropped off more quickly in the $k = 4$ and $k = 6$ populations. I return to this subject in the section on the distribution of mutational effects.

HERITABILITY

There are several ways to measure and interpret heritability (Falconer & Mackay, 1996; Lynch & Walsh, 1998). I used the regression coefficient of offspring phenotype on the average phenotype of the pair of parents. This offspring–midparent regression measures the extent to which offspring phenotype is predicted by parental phenotype. The interpretation of interest here concerns how well one can predict the effect of individual alleles when the alleles are placed in a randomized genetic background. This

measures the statistical additivity of allelic effects in the context of the existing genetic variation in the population. This statistical additivity does not necessarily imply metabolic additivity, as discussed above in the section on Population and Quantitative Genetics.

I measured heritability as follows. At the end of a run, I calculated the phenotypic value of each individual in the population. I truncated the lower percentage τ of the population to simulate the fact that the weakest individuals in the population rarely survive to reproduce. I defined as T the phenotypic value of the population that corresponds to the τ -th percentile—the truncation point in phenotypic units. I then made 5000 babies by choosing pairs of parents randomly with replacement.

I created an offspring in the standard way of this model, by combining parental chromosomes to form a diploid stage, and then following a

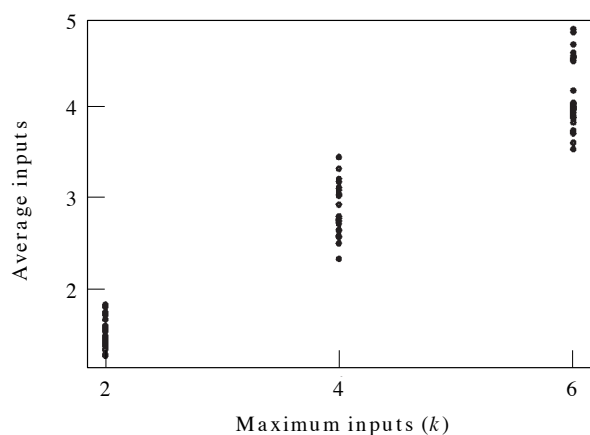


FIG. 6. Number of inputs per node. The x -axis shows the maximum number of inputs, k . The y -axis shows the average number of inputs per node for the best performing network in the population. The points represent all of the parameter combinations for the 3^4 design.

meiotic step with recombination to make a haploid offspring. I assumed the recombination probability was 0.5 for measuring heritability independently of the recombination rate for that run. Higher recombination provides better mixture of alleles and a better estimate of the effect of each allele in a random genetic background. I then calculated the offspring's phenotype; if less than T the offspring was discarded under the assumption that it would be unlikely to survive. If the offspring survived, then the midparent and offspring phenotypes were added to the data vector for the regression calculation. I used various values of τ . I report here for $\tau = 5\%$.

Mutation explains most of the variation in heritability, with an interaction between mutation and the maximum number of inputs, k (Fig. 7). Note that heritabilities are highest with an intermediate mutation rate of 10^{-3} , which corresponds to the mutation rate that yielded the best performance. Put another way, best performance corresponds to the highest level of statistical additivity for allelic effects. It may be that with less mutation the nets are not able to evolve to the zone of additivity, and with more mutation the nets are not able to maintain this additive zone.

DISTRIBUTION OF MUTATIONAL EFFECTS

Mutational analysis is a classical method in genetics for the study of regulatory pathways.

Perturbation reveals the individual steps in the cascade and the interaction between steps. In quantitative genetics, the statistical distribution of mutational effects influences the patterns of genetic variation and the level of heritability.

I analysed the distribution of mutational effects by performing the following steps on each




FIG. 7. Heritability as a function of the mutation rate and the number of inputs, k .

individual in the population: (1) copy the original genotype and save the copy; (2) calculate phenotype of the original genotype; (3) mutate the original genotype once, choosing the location and mutational change at random; (4) calculate the phenotype of the mutated copy; (5) copy the unmutated original again, mutate the new copy, and calculate the new mutant's fitness; (6) start with the original genotype and create an individual that has both of the new mutations that were studied in prior steps; (7) measure the fitness of the double mutant. I repeated this cycle five times for each individual in the population. This section reports data on the single mutants. The following section analyses the double mutants.

Figure 8 shows the distribution of mutational effects for three runs. The x -axis is the fitness of an individual with one mutation divided by the fitness of the original genotype, presented on a logarithmic scale. The right panel shows the same distribution as the left panel but restricts the range of fitness effects to those of 5% or less. Note that a sizable fraction of the mutations have small fitness effects and that nets with higher connectivity (k) have more mutations of small effect.

Networks that respond to single mutations with small changes in output are in the zone of linearity, in which small changes in input cause correspondingly small deviations in dynamical trajectory. Interestingly, networks with the

highest connectivity ($k = 6$) generated the most mutations of small, positive effect, yet did not have the highest heritability or performance. This positive relationship between connectivity and the fraction of mutations with small effect counters Kauffman's theory that high connectivity leads to a large, often chaotic response to perturbation. This may point to one difference between the randomly connected nets studied by Kauffman and the evolved nets analysed here.

I abbreviate the percentage of mutations that cause no phenotypic effects as the percent zeros. Figure 9 shows the percent zeros in response to the mutation rate and the maximum number of inputs, k . The other parameters had little effect on variation in the percent zeros.

The percent zeros provide further clues about the distinction between evolved nets and random nets. When the maximum number of inputs, k , is high, an increase in the mutation rate causes a strong decline in the percent zeros. One possibility is that increased mutation causes a rise in the actual number of inputs per node and thus a greater probability that a changed node has a phenotypic effect. But the actual number of inputs per node is not strongly affected by the mutation rate. For $k = 6$ and log 10 mutation rates of -4 , -3 , and -2 , the median numbers of actual inputs are 4.5, 3.7, and 3.9, respectively. Thus higher k leads to more mutations with non-zero effect (Fig. 9) but each mutation has, on average, a smaller effect (Fig. 8). This further




FIG. 8. Cumulative distribution of the effects of single mutations. The effect of a mutation is measured on a logarithmic scale as the fitness of the mutant relative to the original genotype. Each curve is for a single run with parameters $p = 500$, $n = 50$, $\mu = 10^{-3}$, and $r = 0.05$. The values of k for each run are shown. Mutations with zero effect were not included in the distributions. The heritability (β) and fraction of mutations with zero effect (γ) were: for $k = 2$, $\beta = 1.02$ and $\gamma = 0.77$; for $k = 4$, $\beta = 0.79$ and $\gamma = 0.79$; and for $k = 6$, $\beta = 0.87$ and $\gamma = 0.65$.




FIG. 9. Fraction of mutations with zero effect on phenotype.

supports the conclusion that the evolved nets have different properties from random nets with respect to the role of connectivity and dynamics.

The fraction of mutations of small effect measures the extent to which small perturbations cause small dynamical changes. Figure 10 shows the fraction of mutations with effects in the range $(-0.05, 0.05)$, given that the mutations have non-zero effect. For $k = 2$, mutations have relatively large effects independently of the mutation rate. Low connectivity apparently yields fragile nets.

The higher connectivity levels are more interesting. Intermediate mutation rate (10^{-3}) yields the most mutations of small effect, corresponding to nets with the best performance. It may be that lower mutation rates lead to nets that are poorly adapted and fragile to perturbation, whereas higher mutation rates disrupt the nets and maintain a population of poorly adapted individuals. This corresponds to Kauffman's suggestion that small response to perturbation (mutation), or the zone of linearity, is indeed a property favored by selection. Interestingly, the higher connectivity ($k = 6$) yielded the most mutations of small effect at

intermediate mutation rate but had lower average performance than the $k = 4$ nets. This contradicts Kauffman's (1993) suggestion that increasing connectivity leads to small perturbations creating large effects. It is consistent, however, with Kauffman & Smith's (1986) suggestion that intermediate connectivity provides a smoother landscape for evolutionary progress.

The fraction of mutations that has small, positive effect is another interesting measure. Figure 11 shows the fraction of mutations with effects in the range $(0, 0.05)$ relative to the fraction with effects $(-0.05, 0.05)$, where all mutations have non-zero effect. Higher mutation rate leads to a more symmetrical distribution of small mutational effects. This again suggests that, under low mutation, the nets are stuck near a relatively low adaptive peak whereas, under high mutation, the nets are pushed away from a peak by mutation pressure.

DISTRIBUTION OF EPISTATIC EFFECTS

The previous section outlined the methods for mutational analysis and the results for single mutants. This section presents results for the




FIG. 10. Fraction of mutations with small effect on phenotype, in the range $(-0.05, 0.05)$.

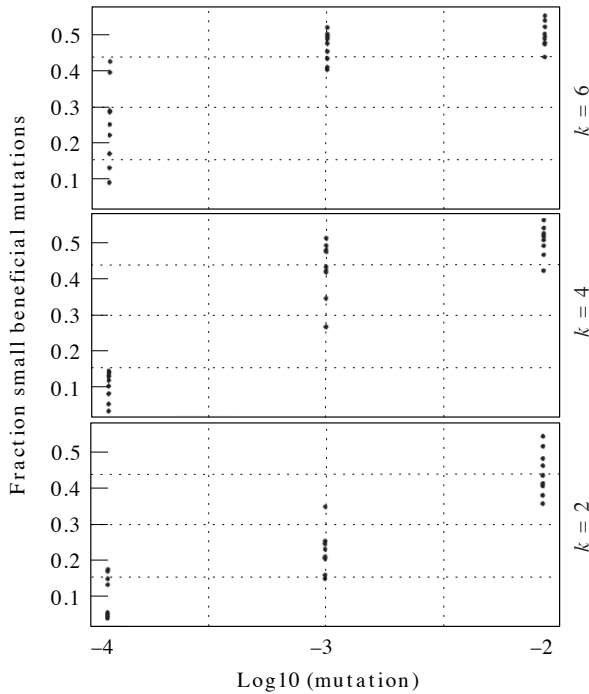


FIG. 11. Fraction of mutations with small positive effect on phenotype.

phenotypic effects of double mutants. These results provide information about the degree of non-additivity (epistasis) between pairs of alleles.

There are many ways to measure non-additivity, as discussed in the earlier section on Population and Quantitative Genetics. I use the




FIG. 12. Measurement of epistasis. Expected fitness of a double mutant is calculated under the assumption of additivity of single mutant effects when measured on a logarithmic scale (see text). When the observed fitness of double mutants is less than expected, the negative value for the interaction is measured by the deviation from the line as shown in the figure. Negative interactions are sometimes called synergistic epistasis because two deleterious mutations have a worse effect when together than expected from their effects when alone.

method shown in Fig. 12, which corresponds to eqn (1). Following that equation, the expected fitness of the double mutant is $\log_{10}(M_1/P) + \log_{10}(M_2/P)$, where P is the fitness of the unmutated parental genotype, M_1 is the fitness of the parental genotype with a single additional mutation, and M_2 is the fitness of the parental genotype with a single mutation that differs from M_1 . The double mutant is the parental genotype with the mutations of both M_1 and M_2 . The observed fitness of the double mutant is the actual fitness of this doubly mutated parental genotype. The difference between the observed and expected fitnesses of double mutants is $I(i, j)$ of eqn (1) and the deviation from the line in Fig. 12.

Figure 13 shows the distribution of epistatic values for three runs. The left panel shows the full distribution which spans several orders of magnitude. The right panel shows epistatic values of 5% or less. Two points deserve mention. First, more strongly connected nets (higher k) have smaller epistatic effects. This matches the greater additivity of strongly connected nets described in an earlier section. The second point is that all distributions are nearly symmetric with medians close to zero. (See the right panel, where the median, or 50th percentile, is emphasized by the intersection of the horizontal line and the distributions.)

The distribution of epistatic values was centered very close to zero for all runs. None of the parameters explained a significant fraction of the variance in the median epistatic value. The median epistatic value was less than 1% in 75% of all runs.

The epistatic distributions were skewed toward negative (synergistic) epistasis. Comparison of the first and third quartiles of the epistatic distribution provides a simple measure of skew. The median over all runs of the first quartile was -0.22 . The median of the third quartile was 0.04 . On the whole, the epistatic distributions were centered near zero and skewed toward negative (synergistic) epistasis.

Conclusion

Are there predictable population-level features of regulatory networks designed by natural

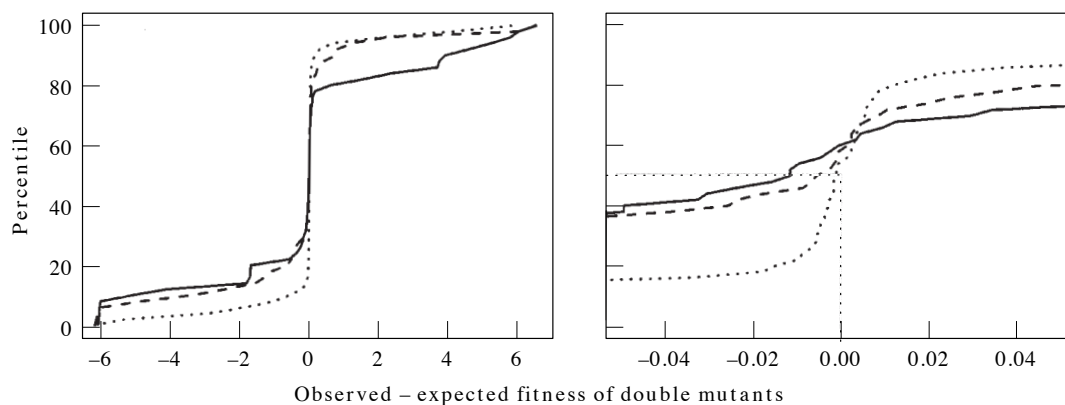


FIG. 13. Cumulative distribution of epistatic effects. The values plotted are $I(i, j)$ from eqn (1). Each curve is for a single run with parameters $p = 500$, $n = 50$, $\mu = 10^{-3}$, and $r = 0.05$. Values are not included for pure additivity, which implies zero epistasis with $I = 0$. The percentages of zeros for $k = 2$ (—), 4 (---), 6 (···) are, respectively, 89, 87, and 72. These percentages exaggerate the degree of additivity because many single mutations have zero effect. Some of these mutations may be unconnected to the outputs and therefore regarded as null mutants.

selection? Kauffman (1993) argued that the number of inputs per node is the crucial factor in network architecture. His theory of random nets predicts that, in response to perturbation, low connectivity causes limited output response and high connectivity causes large, unpredictable response. Intermediate connectivity should therefore be favored by selection.

My results supported the prediction that intermediate connectivity leads to higher evolutionary performance (Fig. 5). But the reasons for this do not seem to match Kauffman's theory of random nets. In my simulations of evolved networks, the more highly connected the network, the smaller the effect of mutational perturbation on phenotype (Figs 8 and 10 for low and intermediate mutation rates).

What causes the greater performance of intermediate connectivity? The distinction between metabolic and quantitative genetic additivity provides a clue. Metabolic additivity describes the interaction between changes to a single network, whereas quantitative genetic additivity measures the consistency of phenotypic effect caused by gene substitution in randomly chosen members of the population. I analysed metabolic additivity by the distribution of epistatic effects of pairs of mutations to individual networks. I measured quantitative genetic additivity by heritability.

Heritability declined with increasing connectivity of networks (Fig. 7). Interestingly, more

highly connected nets produced more mutations of small effect (Fig. 10 for low and intermediate mutation rates) and had lower levels of metabolic epistasis (Fig. 13). Thus highly connected networks had greater metabolic additivity for perturbations to individual networks, but had lower additivity when measured by the average effect of a gene substitution (heritability). The lower heritability of highly connected nets appears to reduce the effectiveness of recombination in searching evolutionary space because new gene combinations formed by recombination interact non-additively. Thus particular alleles will often be advantageous in one genetic background and disadvantageous in another.

I used a simple boolean control structure and a simple environmental challenge. It would be interesting to expand the theoretical analysis to a variety of control structures and challenges in order to develop a comparative theory of networks. Szathmary (1992) has made an important contribution to a comparative theory, but little else has been done along these lines.

My research is supported by National Science Foundation grant DEB-9627259.

REFERENCES

- BARTON, N. H. & TURELLI, M. (1989). Evolutionary quantitative genetics: how little do we know? *Ann. Rev. Genet.* **23**, 337–370.

- CHARLESWORTH, B. (1990). Mutation-selection balance and the evolutionary advantage of sex and recombination. *Gen. Res.* **55**, 199–221.
- CLARK, A. G. (1991). Mutation-selection balance and metabolic control theory. *Genetics* **129**, 909–923.
- FALCONER, D. S. & MACKAY, T. F. C. (1996). *Introduction to Quantitative Genetics*, 4th Edn. Essex, England: Longman.
- FISHER, R. A. (1958). *The Genetical Theory of Natural Selection*, 2nd Edn. New York: Dover.
- HARTL, D. L., DYKHUIZEN, D. E. & DEAN, A. M. (1985). Limits of adaptation: the evolution of selective neutrality. *Genetics* **111**, 655–674.
- KACSER, H. & BURNS, J. A. (1979). Molecular democracy: who shares the controls. *Biochem. Soc. Trans.* **7**, 1149–1160.
- KAUFFMAN, S. A. (1969). Metabolic stability and epigenesis in randomly connected nets. *J. theor. Biol.* **22**, 437–467.
- KAUFFMAN, S. A. (1974). The large-scale structure and dynamics of gene control circuits: an ensemble approach. *J. theor. Biol.* **44**, 167–190.
- KAUFFMAN, S. A. (1993). *The Origins of Order*. Oxford: Oxford University Press.
- KAUFFMAN, S. A. & SMITH, R. G. (1986). Adaptive automata based on darwinian selection. *Physica* **22D**, 68–82.
- KEIGHTLEY, P. D. (1989). Models of quantitative variation of flux in metabolic pathways. *Genetics* **121**, 869–876.
- KONDRASHOV, A. S. (1988). Deleterious mutations and the evolution of sexual reproduction. *Nature* **336**, 435–440.
- KONDRASHOV, A. S. (1993). Classification of hypotheses on the advantage of amphimixis. *J. Heredity* **84**, 372–387.
- LYNCH, M. & WALSH, B. (1998). *Genetics and Analysis of Quantitative Traits*. Sunderland, MA: Sinauer Associates.
- MCADAMS, H. H. & SHAPIRO, L. (1995). Circuit simulation of genetic networks. *Science* **269**, 650–656.
- PHILLIPS, P. C. (1998). The language of gene interaction. *Genetics* **149**, 1167–1171.
- SZATHMÁRY, E. (1992). Do deleterious mutations act synergistically: metabolic control theory provides a partial answer. *Genetics* **133**, 127–132.