



Percolation on fitness landscapes: Effects of correlation, phenotype, and incompatibilities

Authors: Janko Gravner, Damien Pitman, & Sergey Gavrilets

NOTICE: this is the author's version of a work that was accepted for publication in Journal of Theoretical Biology. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in [Journal of Theoretical Biology](#), [VOL# 248, ISSUE# 4, (October 2007)] DOI# [10.1016/j.jtbi.2007.07.009](https://doi.org/10.1016/j.jtbi.2007.07.009)

Gravner, Janko, Damien Pitman, and Sergey Gavrilets. "Percolation on Fitness Landscapes: Effects of Correlation, Phenotype, and Incompatibilities." *Journal of Theoretical Biology* 248, no. 4 (October 2007): 627–645. doi:10.1016/j.jtbi.2007.07.009.

Made available through Montana State University's [ScholarWorks](https://scholarworks.montana.edu)
scholarworks.montana.edu

Percolation on fitness landscapes: Effects of correlation, phenotype, and incompatibilities

Janko Gravner^a, Damien Pitman^a, Sergey Gavrilets^{b,c}

^a *Department of Mathematics, University of California, Davis, CA 95616, USA*

^b *Department of Ecology and Evolutionary Biology, University of Tennessee, Knoxville, TN 37996, USA*

^c *Department of Mathematics, University of Tennessee, Knoxville, TN 37996, USA*

Abstract

We study how correlations in the random fitness assignment may affect the structure of fitness landscapes, in three classes of fitness models. The first is a phenotype space in which individuals are characterized by a large number n of continuously varying traits. In a simple model of random fitness assignment, viable phenotypes are likely to form a giant connected cluster percolating throughout the phenotype space provided the viability probability is larger than $1/2^n$. The second model explicitly describes genotype-to-phenotype and phenotype-to-fitness maps, allows for neutrality at both phenotype and fitness levels, and results in a fitness landscape with tunable correlation length. Here, phenotypic neutrality and correlation between fitnesses can reduce the percolation threshold, and correlations at the point of phase transition between local and global are most conducive to the formation of the giant cluster. In the third class of models, particular combinations of alleles or values of phenotypic characters are “incompatible” in the sense that the resulting genotypes or phenotypes have zero fitness. This setting can be viewed as a generalization of the canonical Bateson–Dobzhansky–Muller model of speciation and is related to K -SAT problems, prominent in computer science. We analyze the conditions for the existence of viable genotypes, their number, as well as the structure and the number of connected clusters of viable genotypes. We show that analysis based on expected values can easily lead to wrong conclusions, especially when fitness correlations are strong. We focus on pairwise incompatibilities between diallelic loci, but we also address multiple alleles, complex incompatibilities, and continuous phenotype spaces. In the case of diallelic loci, the number of clusters is stochastically bounded and each cluster contains a very large sub-cube. Finally, we demonstrate that the discrete NK model shares some signature properties of models with high correlations.

Keywords:

Fitness landscapes; Percolation;
Holey landscapes; Nearly neutral
networks; Genetic incompatibilities;
Neutral evolution; Speciation

1. Introduction

The notion of fitness landscapes, introduced by a theoretical evolutionary biologist Wright (1932) (see also Kauffman, 1993; Gavrilets, 2004), has proved extremely useful both in biology and well outside of it. In the standard interpretation, a fitness landscape is a relationship between a set of genes (or a set of quantitative characters) and a measure of fitness (e.g. viability, fertility, or mating success). In Wright’s original formulation the set of genes (or quantitative characters) is the property of an individual. However, the notion of fitness landscapes can be generalized to

the level of a mating pair, or even a population of individuals (Gavrilets, 2004).

To date, most empirical information on fitness landscapes in biological applications has come from studies of RNA (e.g., Schuster, 1995; Huynen et al., 1996; Fontana and Schuster, 1998), proteins (e.g., Lipman and Wilbur, 1991; Martinez et al., 1996; Rost, 1997), viruses (e.g., Burch and Chao, 1999, 2004), bacteria (e.g., Elena and Lenski, 2003; Woods et al., 2006), and artificial life (e.g., Lenski et al., 1999; Wilke et al., 2001). The three paradigmatic landscapes—rugged, single-peak, and flat—emphasizing

particular features of fitness landscapes have been the focus of most of the earlier theoretical work (reviewed in Kauffman, 1993; Gavrilets, 2004). These landscapes have found numerous applications with regards to the dynamics of adaptation (e.g., Kauffman and Levin, 1987; Kauffman, 1993; Orr, 2006a, 2006b) and neutral molecular evolution (e.g., Derrida and Peliti, 1991).

More recently, it was realized that the dimensionality of most biologically interesting fitness landscapes is enormous and that this huge dimensionality brings some new properties which one does not observe in low-dimensional landscapes (e.g. in two- or three-dimensional geographic landscapes). In particular, multidimensional landscapes are generically characterized by the existence of neutral and nearly neutral networks (also referred to as holey fitness landscapes) that extend throughout the landscapes and that can dramatically affect the evolutionary dynamics of the populations (Gavrilets, 1997; Gavrilets and Gravner, 1997; Reidys et al., 1997; Gavrilets, 2004; Reidys et al., 2001; Reidys and Stadler, 2001, 2002).

An important property of fitness landscapes is their correlation pattern. A common measure for the strength of dependence is the *correlation function* ρ measuring the correlation of fitnesses of pairs of individual at a distance (e.g., Hamming) d from each other in the genotype (or phenotype) space

$$\rho(d) = \frac{\text{cov}[w(\cdot), w(\cdot)]_d}{\text{var}(w)} \quad (1)$$

(Eigen et al., 1989). Here, the term in the numerator is the covariance of fitnesses of two individuals at distance d , and $\text{var}(w)$ is the variance in fitness over the whole fitness landscape. For uncorrelated landscapes, $\rho(d) = 0$ for $d > 0$. In contrast, for highly correlated landscapes, $\rho(d)$ decreases with d very slowly.

The aim of this paper is to extend our previous work (Gavrilets and Gravner, 1997) in a number of directions paying special attention to the question of how correlations in the random fitness assignment may affect the structure of fitness landscapes. In particular, we shed some light on issues such as the existence and number of viable genotypes, as well as the number and size of the clusters of viable genotypes. To this end, we introduce a variety of models, which could be divided into two essentially different classes: those with local correlations, and those with global correlations. As we will see, techniques used to analyze these models, and answers we obtain, differ significantly. We use a mixture of analytical and computational techniques; it is perhaps necessary to point out that these models are very far from trivial, and one is quickly led to outstanding open problems in probability theory and computer science.

We start (in Section 2) by briefly reviewing some results from Gavrilets and Gravner (1997). In Section 3 we generalize these results for the case of a continuous phenotype space when individuals are characterized by a large number of continuously varying traits such as size,

weight, color, or the concentrations of some gene products. The latter interpretation of the phenotype space may be particularly relevant given the rise of proteomics and the growing interest in gene regulatory networks.

The main idea behind our local correlations model studies in Section 4 is fitness assignment *conformity*. Namely, one randomly divides the genotype space into components which are forced to have the same phenotype; then, each different phenotype is independently assigned a random fitness. This leads to a simple two-parameter model, in which one parameter determines the density of viable genotypes, and the other the correlations between them. We argue that the probability of existence of a giant cluster (which swallows a positive proportion of all viable genotypes) is a non-monotone function of the correlation parameter and identify the critical surface at which this probability jumps almost from 0 to 1. In Appendix B we also investigate the effects of interaction between conformity structure and fitness assignment.

Section 5 introduces a class of models where genotypes are eliminated due to random *incompatibilities* between alleles. These models are characterized by global correlations. When there are only two alleles at each locus, the models are equivalent to random versions of the SAT problem, which is the canonical constraint satisfaction problem in computer science. In general, a SAT problem involves a set of Boolean variables and their negations that are strung together with OR symbols into *clauses*. The *clauses* are joined by AND symbols into a *formula*. A SAT problem asks one to decide, whether the variables can be assigned values that will make the formula true. Arguably, SAT is the most important class of problems in complexity theory. In fact, the general SAT problem was the first known NP-complete problem and was established as such by Cook (1971). An important special case, *K-SAT*, has the length of each clause fixed at *K*. Even considerable simplifications, such as 3-SAT (see Section 5.4), remain NP-complete, although 2-SAT (see Section 5.1) can be solved efficiently by a simple algorithm. See e.g. Korte and Vygen (2005) for a comprehensive presentation of the theory. Difficulties in analyzing random SAT problems, in which formulas are chosen at random, in many ways mirror their complexity classes, but even random 2-SAT presents significant challenges (de la Vega, 2001; Bollobás et al., 1994). In our present interpretation, the main reason for these difficulties is that correlations are so high that the expected number of viable genotypes may be exponentially large, while at the same time the probability that even one viable genotype exists is very low. In Section 5, we show that the high correlations in the random 2-SAT model essentially force the landscape to either have no viable genotypes or else to have fairly large sub-cubes contained within any cluster of viable genotypes. Moreover, the number of such clusters will be, in a proper interpretation, finite. The theory for higher order incompatibilities is closely related and we briefly review the existing literature for 3-SAT. We also describe pairwise incompatibility

models on multiallelic genotype spaces and continuous phenotype spaces where we address the question of existence of viable individuals.

In Section 6 we demonstrate that the discrete NK model shares some signature properties of models with high correlations. The proofs of our major results are relegated to Appendices A–E.

We summarize our main findings, provide their intuitive explanations, and briefly discuss their biological significance at the end of each section. Our most general results, their theoretical and biological significance, implications, and limitations are discussed in Section 7.

2. The basic case: binary hypercube and independent binary fitness

We begin with a brief review of the basic setup, from Gavrillets and Gravner (1997) and Gavrillets (2004). There are n diallelic loci with alleles 0 and 1. Each genotype is represented by a binary sequence of length n . The *genotype space*, i.e. the set of all possible genotypes, can be represented by an n -dimensional *binary hypercube* \mathcal{G} . In this space, genotypes are linked by edges induced by bit-flips, i.e., *mutations* at a single locus, for example. For $n = 4$, a sequence of mutations might look like

0000 \leftrightarrow 1000 \leftrightarrow 1001 \leftrightarrow 1101 \leftrightarrow 1100.

The (Hamming) *distance* $d(x, y)$ between genotypes x and y is the number of loci in which x and y differ or, equivalently, the least number of mutations which connect x and y .

The *fitness* of each genotype x is denoted by $w(x)$. We will describe several ways to prescribe the fitness w at random, according to some probability measure P on all 2^n possible assignments. Then we say that an event A_n happens *asymptotically almost surely* (a.a.s.) if its probability $P(A_n) \rightarrow 1$ as $n \rightarrow \infty$. Typically, A_n will capture some important property of (random) clusters of genotypes.

We will commonly assume that fitness w takes only two values so that each genotype x is either viable ($w(x) = 1$) or inviable ($w(x) = 0$). As a natural starting point, Gavrillets and Gravner (1997) considered uncorrelated landscapes, in which $w(x)$ is chosen to be 1 with probability p , for each x independently of others. Thus, p can be viewed as the probability that a random combination of genes results in a viable individual. Such binary fitness assignment is not as restrictive as one might expect. If fitness were to be a random number between 0 and 1, then p could be interpreted as the width of an acceptable “fitness band” (Gavrillets and Gravner, 1997; Gavrillets, 2004). We continue to refer to viable genotypes for the rest of this section and note that this is a well-studied problem in mathematical literature, although it presents considerable technical difficulties and some issues are still not completely resolved.

Given a particular fitness assignment, viable genotypes form a subset of \mathcal{G} , which is divided into connected *components* or *clusters*. For example, with $n = 4$, if 0000 is viable, but its four neighbors 1000, 0100, 0010, and 0001 are not, then it is isolated in its own cluster.

Perhaps the most basic result determines the *connectivity threshold* (Toman, 1979): when $p > \frac{1}{2}$, the set of all viable genotypes is connected a.a.s. By contrast, when $p < \frac{1}{2}$, the set of viable genotypes is *not* connected a.a.s. This is easily understood, as the connectedness is closely linked to isolated genotypes, whose expected number is $2^n p(1-p)^n$. This expectation makes a transition from exponentially large in n to exponentially small at $p = \frac{1}{2}$. The events $\{x \text{ is isolated}\}$, $x \in \mathcal{G}$, are only weakly correlated, which implies that when $p < \frac{1}{2}$ there are exponentially many isolated genotypes with high probability, while when $p > \frac{1}{2}$, a separate argument shows that the event that the set of viable genotypes contains no isolated vertex but is not connected becomes very unlikely for large n . This is perhaps the clearest instance of the *local method*: a local property (no isolated genotypes) is a.a.s. equivalent to a global one (connectivity).

Connectivity is clearly too much to ask for, as p above $\frac{1}{2}$ is not biologically realistic. Instead, one should look for a weaker property which has a chance of occurring at small p . Such a property is *percolation*, a.k.a. existence of the *giant component*. For this, we scale $p = \lambda/n$, for a constant λ . When $\lambda > 1$, the set of viable genotypes percolates, that is, it a.a.s. contains a component of at least $c \cdot n^{-1} 2^n$ genotypes, with all other components of at most polynomial (in n) size. When $\lambda < 1$, the largest component is a.a.s. of size Cn . Here and below, c and C are some constants. These are results from Bollobás et al. (1994).

The local method that correctly identifies the percolation threshold is a little more sophisticated than the one for the connectivity threshold, and uses branching processes with Poisson offspring distribution—hence we introduce notation $\text{Poisson}(\lambda)$ for a Poisson distribution with mean λ . Viewed from, say, genotype $0 \dots 0$, the binary hypercube locally approximates a tree with uniform degree n . Thus viable genotypes approximate a branching process in which every node has the number of successors distributed binomially with parameters $n - 1$ and p , hence this random number has mean about λ and is approximately $\text{Poisson}(\lambda)$. When $\lambda > 1$, such a branching process survives forever with probability $1 - \delta > 0$, where δ depends on λ and is given by the implicit equation

$$\delta = e^{\lambda(\delta-1)}. \quad (2)$$

(e.g., Athreya and Ney, 1971). Large trees of viable genotypes created by the branching processes which emanate from viable genotypes merge into a very large (“giant”) connected set. On the other hand, when $\lambda < 1$ the branching process dies out with probability 1.

The condition $\lambda > 1$ for the existence of the giant component can be loosely rewritten as

$$p > \frac{1}{n}. \quad (3)$$

This shows that the larger the dimensionality n of the genotype space, the smaller values of the probability of being viable p will result in the existence of the giant component. Biological populations can evolve along this giant cluster by mutation and random drift and can diverge dramatically without the need to cross any fitness valleys or any help from selection. See Gavrillets and Gravner (1997); Gavrillets (1997, 2004); Skipper (2004); Pigliucci and Kaplan (2006); Wilkins (2007) for discussions of biological and philosophical significance and implications of this important result.

3. Percolation in a continuous phenotype space

In this section we assume that individuals are characterized by n continuous traits, labeled $1, 2, \dots, n$ (such as size, weight, color, or concentrations of particular gene products), each of which can have values between 0 and 1. A phenotype is therefore given by n trait-values, that is, by numbers $z_1, \dots, z_n \in [0, 1]$. Let $\mathcal{P} = [0, 1]^n$ be the *phenotype space*, i.e., the set of all possible phenotypes.

We begin with a modification of the model of independent binary fitness assignment used in the previous section. First, we choose N points $x_1, \dots, x_N \in \mathcal{P}$ uniformly at random where N is a $\text{Poisson}(\lambda)$ random variable. (This construction is known as *Poisson point location*.) Points x_1, \dots, x_N will be interpreted as “peaks” of equal height in the fitness landscape. Note that in this section, parameter λ gives the expected number of peaks in the phenotype space. Second, we declare any phenotype within r of one of the peaks, where r is a small positive number, viable and any phenotype not within r of one of the peaks inviable. Parameter r can be interpreted as measuring how harsh the environment is. For simplicity, we will assume “within r ” to mean that “every coordinate differs by at most r .” Note that this makes the set of viable phenotypes correlated, albeit the range of correlations is limited to $2r$.

Our most basic question is whether a positive proportion of viable phenotypes is connected together into a giant cluster. Note that the probability p that a random point in \mathcal{P} is viable is equal to the probability that there is a “peak” within r from this point. Therefore,

$$p = 1 - \exp[-\lambda(2r)^n] \approx \lambda(2r)^n.$$

This is also the expected combined volume of viable phenotypes.

We will consider peaks x_i and x_j to be *neighbors* if they share a viable phenotype, that is, if their r -neighborhoods overlap; and we consider them to be *connected* if they are linked by a chain of consecutively neighboring peaks. By the standard branching process comparison, the necessary condition for the existence of a giant cluster is that a peak x

neighbors more than one other peak on the average. All peaks within $2r$ of the focal peak are its neighbors. Therefore, the expected number of peaks neighboring x is

$$v = \lambda \cdot (4r)^n,$$

and $v > 1$ is necessary for percolation. As demonstrated by Penrose (1996) (for a different choice of the norm, but the proof is the same), this condition becomes sufficient when n is large.

If $v > 1$ and fixed, then a.a.s. a positive proportion of all peaks (that is, cN peaks, where $c = c(v) > 0$) are connected in one “giant” component, while the remaining connected components all have size of order $\log N$. On the other hand, if $v < 1$, all components a.a.s. have size of order $\log N$.

Note that the expected number λ of peaks in \mathcal{P} can be written as $v \cdot (4r)^{-n}$. Thus, the condition $v > 1$ for the existence of the giant component of viable phenotypes can be loosely rewritten as

$$p > \frac{1}{2^n}. \quad (4)$$

This shows that viable phenotypes are likely to form a large connected cluster even when one is *very* unlikely to hit one of them at random, if n is even moderately large. The same conclusion and the same threshold are valid if instead of n -cubes we use n -spheres of a constant radius.

Thus, dramatic divergence in phenotype without any loss in fitness is possible if the dimensionality of the phenotype space is sufficiently large. The percolation threshold in the continuous phenotype space given by inequality (4) is much smaller than that in the discrete genotype space which is given by inequality (3). For example, the percolation threshold in the 10-dimensional continuous space \mathcal{P} is similar to that in the genotype space \mathcal{G} corresponding to 1024 diallelic loci. An intuitive reason for this is that continuous space offers a viable point a much greater opportunity to be connected to a large cluster. Indeed, in the discrete genotype space \mathcal{G} there are n neighbors per each genotype. In contrast, in the continuous phenotype space \mathcal{P} , the ratio of the volume of the space where neighboring peaks can be located (which has radius $2r$) to the volume of the focal n -cube (which has radius r) is 2^n .

4. Percolation in a correlated landscape with phenotypic neutrality

The standard paradigm in biology is that the relationship between genotype and fitness is mediated by phenotype (i.e., observable characteristics of individuals). Both the genotype-to-phenotype and phenotype-to-fitness maps are typically not one-to-one. Here, we formulate a simple model capturing these properties which also results in a correlated fitness landscape. Below we will call mutations that do not change phenotype *conformist*. These mutations represent a subset of *neutral* mutations that do not change fitness.

We propose the following two-step model. To begin the *first step*, we make each *pair* of genotypes x and y in a binary hypercube \mathcal{G} independently *conformist* with probability q_d where $d = d(x, y)$ is the Hamming distance between x and y . We then declare any pair x and y to belong to the same *conformist cluster* if they are linked by a chain of conformist pairs. This version of long-range percolation model (cf., Berger, 2004; Biskup, 2004) divides the set of genotypes \mathcal{G} into conformist clusters. We postulate that all genotypes in the same conformist cluster have the same phenotype. Therefore, genetic changes represented by a change from one member of a conformist cluster to another (i.e., single or multiple mutations) are phenotypically neutral. In the *second step*, we make each conformist cluster independently *viable* with probability $p = \lambda/n$. This generates a random set of viable genotypes, and we aim to investigate when this set has a giant connected component.

To illustrate our model, we can interpret the “genotype” as a linear RNA sequence. This sequence folds into a three-dimensional molecule which has a particular structure, and corresponds to our “phenotype.” Finally, the molecule itself has a particular function, e.g., to bind to a specific part of the cell or to another molecule. A measure of how well this can be accomplished is represented by our “fitness.”

The distribution of conformist clusters depends on the probabilities q_1, q_2, q_3, \dots which determine how the conformity probability varies with genetic distance. We will study the case when $q_1 = q > 0, q_2 = q_3 = \dots = 0$ (Häggström, 2001). That is, q is the probability that a pair of nearest neighbors is a conformist pair. Thus, we can talk about *conformist edges* or equivalently *conformist mutations*, and q is the probability that a mutation is conformist. (Note however that it is possible that nearest neighbors x and y are in the same conformist cluster even if the edge between them is non-conformist.) Fig. 1 illustrates our 2-step procedure in a four-dimensional example.

We expect that a more general model with q_k declining fast enough with k is just a smeared version of this basic one, and its properties are not likely to differ from those of the simpler model. We conjecture that for our purposes, “fast enough” decrease should be exponential with a rate logarithmically increasing in the dimension n , e.g. for large k , $q_k \leq \exp(-\alpha(\log n)k)$, for some $\alpha > 1$. (This is expected to be so because in this case the expected number of neighbors of the focal genotype is finite.)

We observe that the first step of our procedure is an edge version of the percolation model discussed in the second section, with a similar giant component transition (Bollobás et al., 1992). Namely, let $q = \mu/n$. Then, if $\mu > 1$, there is a.a.s. one giant conformist cluster of size $c \cdot 2^n$, with all others of size at most Cn . In contrast, if $\mu < 1$ all conformist clusters are of size at most Cn . Note that the number of conformist clusters is always on the order 2^n . In fact, even the number of “non-conformist” (i.e., isolated)

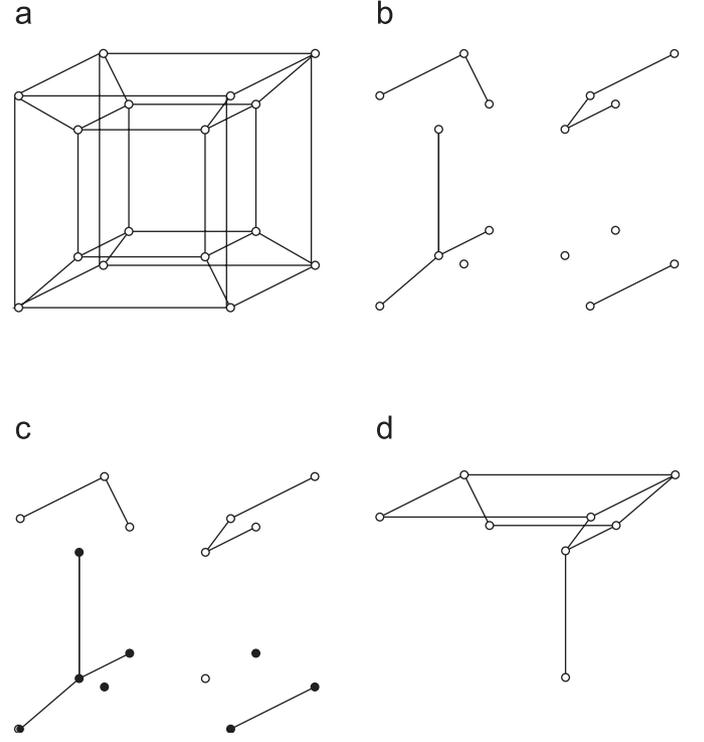


Fig. 1. A four-dimensional example. (a) Start with the four-dimensional hypercube. (b) On the first step, create conformist clusters by randomly eliminating each edge with probability $1 - q$. In the example, there are seven conformist clusters left. (c) On the second step, remove each conformist cluster with probability $1 - p$; the vertices to be removed are black. In the example, there are three viable conformist clusters left. (d) The remaining viable genotypes have nearest neighbors connected by edges. In the example, all remaining viable genotypes are connected in a single component.

clusters is a.a.s. asymptotic to $e^{-\mu 2^n}$, as the probability $P(x \text{ is isolated}) = (1 - \mu/n)^n$.

Denote by $x \leftrightarrow y$ (resp. $x \not\leftrightarrow y$) the event that x and y are (resp. are not) in the same conformist cluster. First, we note that the probability $P(x \leftrightarrow y)$ that two genotypes belong to the same conformist cluster depends on the Hamming distance $d(x, y)$ between them, and on $q = \mu/n$. In particular, we show in Appendix A that, if $\mu < 1$ and $d(x, y) = k$ is fixed, then

$$k!q^k(1 - O(n^{-2})) \leq P(x \leftrightarrow y) \leq k!q^k(1 + O(n^{-1} \log n)). \quad (5)$$

The dominant contribution $k!q^k$ is simply the expected number of conformist pathways between x and y that are of shortest possible length.

It is also important to note that, for every $x \in \mathcal{G}$, the probability $P(x \text{ is viable}) = p$, therefore it does not depend on q . Moreover, for $x, y \in \mathcal{G}$,

$$\begin{aligned} P(x \text{ and } y \text{ viable}) &= p^2 \\ &= P(x \text{ and } y \text{ viable, } x \leftrightarrow y) \\ &\quad + P(x \text{ and } y \text{ viable, } x \not\leftrightarrow y) - p^2 \\ &= pP(x \leftrightarrow y) + p^2 \cdot P(x \not\leftrightarrow y) - p^2 \\ &= p(1 - p)P(x \leftrightarrow y) \geq 0. \end{aligned}$$

Therefore, the correlation function (1) of fitness is

$$\rho(x, y) \approx P(x \leftrightarrow y), \quad (6)$$

which clearly increases with q and, thus, with μ . The correlation function $\rho(x, y)$ decreases exponentially with distance $d(x, y)$ when $\mu < 1$, and is bounded below when $\mu > 1$. Nevertheless, as we will see below, we can effectively use local methods for all values of μ . Therefore, in this model the probability of being viable is controlled by parameter λ while the correlation structure of the landscape is controlled by parameter μ .

Proceeding by local branching process heuristics analogous to those used in Section 2, we reason that a surviving node on the branching tree can have two types of descendants: those that are connected by conformist mutations and those that are in different conformist clusters and viable independently. Therefore the number of descendants in the branching process is approximately $\text{Poisson}(\mu + \lambda)$.

The above argument can only work when $\mu < 1$, as otherwise the correlations are global. If $\mu > 1$, we need to eliminate the entire conformist giant component, which is a.a.s. inviable (because p is small). Locally, we condition on the (supercritical) branching process of the supposed descendant to die out. Such conditioned process is a subcritical branching process, with $\text{Poisson}(\mu\delta)$ distribution of successors (Athreya and Ney, 1971) where $\delta = \delta(\mu) < 1$ is given by the equation $\delta = \exp(\mu(\delta - 1))$, analogous to Eq. (2). This gives the conformist contribution, to which we add the independent $\text{Poisson}(\lambda\delta)$ contribution.

To have a convenient summary of the conclusions above, assume that μ is fixed and let $\zeta(\mu)$ be the smallest λ which a.a.s. ensures the giant component, i.e.,

$$\zeta(\mu) = \inf\{\lambda : \text{a cluster of at least } cn^{-1}2^n \text{ viable genotypes exists a.a.s. for some } c > 0\}.$$

One would expect that for $\lambda < \zeta(\mu)$ all components are a.a.s. of size at most Cn . The asymptotic critical curve is given by $\lambda = \zeta(\mu)$, where

$$\zeta(\mu) = \begin{cases} 1 - \mu & \text{if } \mu \in [0, 1], \\ \frac{1}{\delta} - \mu & \text{if } \mu \in [1, \infty). \end{cases} \quad (7)$$

The asymptotic critical curve $\zeta(\mu)$ reaches the smallest value of 0 at $\mu = 1$. To intuitively understand why percolation occurs the easiest with $\mu \approx 1$ (that is, when the probability of a conformist mutation is $\approx 1/n$), it helps to think of the model as a branching process on clusters rather than on genotypes. Suppose we fix a particular set of k genotypes as a viable conformist cluster. The number of genotypes neighboring this cluster is then fixed, but the number of neighboring conformist clusters is random. For any $\mu < 1$, it follows from Eq. (5) that the expected number of neighboring clusters is asymptotically the same as the number of neighboring genotypes. Consequently, we expect the overall number of descendants in the branching

process to be greater if the size of the neighboring clusters is greater on average; which is exactly what happens as μ increases towards 1. If $\mu > 1$, then there is a positive proportion of the neighboring genotypes that are in the giant cluster. This giant cluster is likely to be inviable, so the parameter λ must be greater to compensate for its loss.

In fact, one can prove rigorously (Pitman, unpub.) that $1 - \mu$ is a lower bound on the critical surface by comparison with an edge percolation model in which each edge is retained independently with probability $p + q - pq$. We also remark that the number of viable clusters is on the same order as the number of all viable genotypes, that is, $2^n/n$. Since $p = \lambda/n$, it follows that the number of viable genotypes is on the order of $2^n/n$.

We have not been able to find an analytic proof that $\zeta(\mu)$ (7) is also an upper bound, i.e., that the actual critical surface is given by ζ . Thus, we resort to computer simulations for confirmation. For this, we indicate global connectivity with the event A that a genotype within distance $d = 2$ of $0 \dots 0$ is connected (through viable genotypes) to a genotype within distance $d = 2$ of $1 \dots 1$. We make this choice because the distance 2 is the smallest that works with asymptotic certainty. Indeed, the genotypes $0 \dots 0$ and $1 \dots 1$ are likely to be inviable. Even the number of viable genotypes within distance one of each of these is only of constant order. It follows that, for any $\mu > 0$, the probability of connectivity between a viable genotype within distance one of $0 \dots 0$ and a viable one within distance one of $1 \dots 1$ does not converge to 1 but is of a nontrivial constant order. By contrast, there are about n^2 vertices within distance 2 of $0 \dots 0$ among which of order n are viable.

When $\lambda > \zeta(\mu)$ the probability of the event A should therefore be (exponentially) close to 1. On the other hand, when $\lambda < \zeta(\mu)$ the probability that a connected component within distance 2 of either $0 \dots 0$ or $1 \dots 1$ extends for distance of the order n is exponentially small. We further define the critical curves

$$\lambda_m = \text{the smallest } \lambda \text{ for which } P(A) > 0.1,$$

$$\lambda_M = \text{the largest } \lambda \text{ for which } P(A) < 0.9.$$

We approximated λ_m and λ_M for $n = 10, \dots, 20$ and $\mu = 0(0.1)2$, with 1000 independent realizations of each choice of n , μ , and λ . We used the linear cluster algorithm described in Sedgewick (1997). Unfortunately, simulations above $n \approx 20$ are not feasible. The results are depicted in Fig. 2.

We make the following observations relating to Fig. 2:

- Even for low n , both critical curves approximate well the overall shape of the theoretical limit curve ζ .
- The two curves λ_m and λ_M appear to be approaching a limiting curve away from the theoretical limiting curve ζ . However, this need not be interpreted as evidence against ζ being the actual limiting curve. Indeed, this behavior is an instance of a *sharp threshold* (Friedgut,

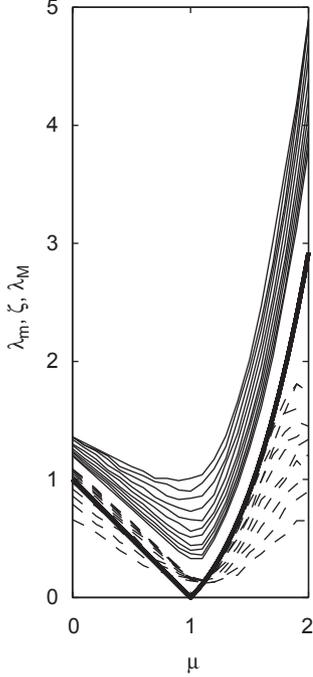


Fig. 2. Simulated lower bounds λ_m (dashed curves) and upper bounds λ_M (solid curves), and ζ (thick curve) plotted against μ , for $n = 10, \dots, 20$. Lower bounds increase with n , and upper bounds decrease, for this range of n .

1999): if one fixes a μ , and increases λ from 0 to ∞ , then the global connectivity probability $P(A)$ rises from 0.1 to 0.9 within a very short interval (of length $\lambda_M - \lambda_m$) even for a relatively low n . The *location* of this sharp increase gets close to ζ much later; in symbols, $\lambda_M - \lambda_m \ll |\zeta - \lambda_m|$ for large n .

- For $\mu < 1$, λ_m tends to be above the limit curve. This is not surprising because the local argument always gives an upper bound on the probability $P(A)$ of event A .
- The approximation of λ_m appears to deteriorate near $\mu = 2$, which stems from the increased probability of survival of the giant component for larger μ .

Summarizing, we have shown that this more realistic model with phenotypic neutrality and fitness correlation shares the most important property of simpler models considered above—the existence of a percolating giant cluster of viable genotypes extending throughout the genotype space. More specifically, the number of clusters of viable genotypes is on the order of $2^n/n$ for any λ and μ , and percolation occurs for λ larger than $\zeta(\mu)$ given by Eqs. (7). In terms of the viability probability p , the percolation threshold is thus inversely proportional to the dimensionality n of the genotype space (because $p = \lambda/n$). Conformist mutations are most advantageous at about $\mu = 1$, which is the point of phase transition between local and global conformist connectivity. In particular, for $\mu \sim 1$, percolation could potentially occur for p on a scale smaller than $1/n$. What is also clear from the heuristics and simulations is that conformist clustering, and thus correla-

tions, can help or hinder connectivity in fitness landscapes. Indeed, the critical probability for global connectivity in the uncorrelated landscape is about $\lambda = 1$ (Section 2). In the correlated landscape model considered here, the critical value $\zeta < 1$ approximately for $\mu < 1.5$ (Fig. 2). However, for large values of μ , correlations hinder connectivity.

In the model studied above, the viability probability p was independent of the size of the conformist clusters. In Appendix B, we consider a simple generalization that can be tuned to create positive or negative correlations between conformist cluster size and viability.

5. Percolation in incompatibility models

In the model considered in the previous section correlations rapidly decreased with distance. This property made local analysis possible. The models we introduce now are fundamentally different in the sense that correlations are so high that the local method gives a wrong answer.

In the previous sections, in constructing fitness landscapes we were assigning fitness to individual genotypes or phenotypes. Here, we make certain assumptions about “fitness” of particular combinations of alleles or the values of continuously varying phenotypic characters. Specifically, we will assume that some of these combinations are “incompatible” in the sense that the resulting genotypes or phenotypes under given environmental conditions have reduced (or zero) fitness (Orr (1995); Orr and Orr (1996); Gavrilets (2004)). The resulting models can be viewed as a generalization of the Bateson–Dobzhansky–Muller model (Cabot et al., 1994; Orr, 1995; Orr and Orr, 1996; Orr, 1997; Orr and Turelli, 2001; Gavrilets and Hastings, 1996; Gavrilets, 1997; Gavrilets and Gravner, 1997; Gavrilets, 2003, 2004; Coyne and Orr, 2004) which represents a canonical model of speciation.

5.1. Diallelic loci

We begin by considering n diallelic loci and assuming that each pair of alleles is independently incompatible with probability

$$p = \frac{c}{2n}.$$

We let a single pair of incompatible alleles make an individual inviable. Thus, the variable p characterizes the probability of being inviable rather than viable as in the previous sections.

There are $4 \cdot \binom{n}{2}$ allele pairs in all, and we form a random list F of those that are incompatible. We denote the allele pair u at locus i and v at locus j by (u_i, v_j) . In this nonstandard notation, $(0_1, 0_2) \in F$, for example, means that allele 0 at locus 1 and allele 0 at locus 2 are incompatible. In general, if $(u_i, v_j) \in F$, all genotypes with u in position i and v in position j are inviable. A genotype x is then inviable if and only if there exist i and j , with $i < j$, so that u and v are, respectively, the alleles of x at loci i and j ,

and $(u_i, v_j) \in F$. For example, if $F_1 = \{(0_1, 0_2), (1_2, 0_3), (1_1, 1_2)\}$, viable genotypes may have 011, 100, and 101 as their first three alleles. For $F_2 = F_1 \cup \{(0_1, 1_3), (1_1, 0_2)\}$, no viable genotype remains.

Incompatibility $(0_1, 0_2)$ is equivalent to two implications: $0_1 \Rightarrow 1_2$ and $0_2 \Rightarrow 1_1$ or to the single OR statement $1_1 \text{ OR } 1_2$. In this interpretation, the problem of whether, for a given list of incompatibilities F , there is a viable genotype is known as the 2-SAT problem (Korte and Vygen, 2005). The associated *digraph* D_F is a graph on $2n$ vertices x_i , $i = 1, \dots, n$, $x = 0, 1$, with oriented edges determined by the implications. A well-known theorem (Korte and Vygen, 2005) states that a viable genotype exists iff D_F contains no oriented cycle from 0_i to 1_i and back to 0_i for any $i = 1, \dots, n$ in D_F . For example, for the incompatibilities F_2 as above, one such cycle is $0_1 \rightarrow 1_2 \rightarrow 1_3 \rightarrow 1_1 \rightarrow 1_2 \rightarrow 0_1$.

5.1.1. Existence of viable genotypes

In this model, viable genotypes exist only if the probability of pairwise incompatibility p is smaller than $1/(2n)$. More precisely, let N be the number of viable genotypes. Then

- if $c > 1$, then a.a.s. $N = 0$.
- if $c < 1$, then a.a.s. $N > 0$.

This result first appeared in the computer science literature in the 1990s (see de la Vega (2001) for a review), and it is an extension of the celebrated Erdős-Rényi random graph results (Bollobás, 2001; Janson et al., 2000) to the oriented case.

Note that the expectation $E(N) = 2^n(1-p)^{\binom{2n}{2}} \approx 2^n e^{-cn/4}$, which grows exponentially whenever $c < 4 \log 2 \approx 2.77$. Neglecting correlations would therefore suggest a wrong threshold for the existence of viable genotypes. The local method based on analysis of the averages (e.g., used in Gavrillets, 2004, Chapter 6) is even farther off, as it suggests an a.a.s. giant component when $p < (1-\varepsilon) \log n/n$ for any $\varepsilon > 0$.

5.1.2. The number of viable genotypes

Assume that $c < 1$. Sophisticated, but not mathematically rigorous methods based on *replica symmetry* (Monasson and Zecchina, 1997; Biroli et al., 2000) from statistical physics suggest that, as $n \rightarrow \infty$, $\lim n^{-1} \log N$ varies almost linearly between $\log 2 \approx 0.69$ (for small c , when, as we prove below, this limit is $\log 2 + \mathcal{O}(c)$) and about 0.38 (for c close to 1). One can however prove that $n^{-1} \log N$ is for large n sharply concentrated around its mean (de la Vega, 2001).

Upper and lower bounds on N can also be obtained rigorously. For example, if X is a number of incompatibilities which involve *disjoint* pairs of loci (i.e., those for which every locus is represented at most once among the incompatibilities), then $N \leq \exp(n \log 2 + X \log(\frac{3}{4}))$, as each of the X incompatibilities reduces the number of viable genotypes by the factor $\frac{3}{4}$. If we imagine adding

incompatibilities one by one at random until there are about cn of them, then after we have k incompatibilities on disjoint pairs of loci the waiting time (measured by the number of incompatibilities added) for a new disjoint one is geometric with expectation $\binom{n}{2} / \binom{n-2k}{2}$. Therefore, X is a.a.s. at least Kn , where K solves the approximate equation

$$\binom{n}{2} \left(\sum_{k=0}^{Kn} \frac{1}{\binom{n-2k}{2}} \right) \sim cn,$$

or

$$\int_0^{Kn} \frac{1}{(n-2k)^2} dk \sim \frac{c}{n},$$

which reduces to $K = c/(1+2c)$. This implies that in the limit of large n the upper bound on N can be written as

$$\frac{1}{n} \log N \leq \frac{1}{1+2c} \log 2 + \frac{c}{1+2c} \log 3. \quad (8)$$

A lower bound is even easier to obtain. Namely, the probability that a fixed location (i.e., locus) i does not appear in F is $(1-p)^{4(n-1)} \rightarrow e^{-2c}$, and then it is easy to see that the number of loci represented in F is asymptotically $(1-e^{-2c})n$. As the other loci are neutral (in the sense that changing their alleles does not affect fitness), $n^{-1} \log N$ is asymptotically at least $e^{-2c} \log 2$. Clearly, this gives a lower bound on the exponential size of any cluster of viable genotypes.

If this was an accurate bound, it would imply that the space of genotypes is rather simple, in that almost all its entropy would come from neutral loci. (By analogy with physical systems, entropy is a measure of the amount of choice (“number of degrees of freedom”) one has in selecting a viable genotype.) Appendix C presents two arguments which will demonstrate that this is not the case.

5.1.3. The structure of clusters

The derivations in Appendix C show that every viable genotype is connected through mutation to a fairly substantial viable sub-cube. In this sub-cube, alleles on at most a proportion $r_u(c) < 1$ of loci are fixed (to 0 or 1) while the remaining proportion $1 - r_u(c)$ could be varied without effect on fitness. Note from Fig. 4 in Appendix C that $1 - r_u(c) \geq 0.3$ for all c , and that such a phenomenon is extremely unlikely on uncorrelated landscapes. Namely, in the setting of Section 2, the probability that a viable subcube of size k exists anywhere is at most

$$\binom{n}{k} \cdot p^{2k} \leq \exp(k \log n + 2^k \log p),$$

so even at the connectivity threshold $p = \frac{1}{2}$ the size of largest viable subcube is a.a.s. of order $\log \log n$. Note also that, for $c < 1$, $N \geq 2^{(1-r_u(c))n}$ a.a.s. and so in the limit of

large n the lower bound on N can be written as

$$\frac{1}{n} \log N \geq (1 - r_u(c)) \log 2. \quad (9)$$

5.1.4. The number of clusters

The natural next question concerns the number of clusters R when $c < 1$. This again has quite a surprising answer, unparalleled in landscapes with rapidly decaying correlations. Namely, R is *stochastically bounded*, that is, for every $\varepsilon > 0$ there exists an $z = z(\varepsilon)$ such that $P(R \leq z) > 1 - \varepsilon$ for all n . This follows from a result in Pitman (unpub.), which shows that the distribution of R converges in distribution to 2^X , where X is distributed Poisson(λ) for

$$\lambda = -\frac{1}{2}(\log(1 - c) + c).$$

The probability that there is a unique cluster thus converges to $e^{-\lambda} = \sqrt{(1 - c)e^c}$.

Asymptotically, a unique cluster has a better than even chance of occurring for c below about 0.9, and is *very* likely to occur for small c , though of course not a.a.s. so. To confirm, we have done simulations for $n = 20$ and $c = 0.01(0.01)1$ (again 1000 trials in each case) and got distribution of clusters depicted in Fig. 3. The results suggest that the convergence to limiting distribution is rather slow for c close to 1, and that the likelihood of a unique cluster increases for low n .

To summarize, in the presence of random pairwise incompatibilities, the set of viable genotypes is, when nonempty, divided into a stochastically bounded number of connected clusters, where a unique cluster is usually the most likely possibility. These clusters are all of exponentially large size (with bounds given by Eqs. (8) and (9)), in fact they all contain sub-cubes of dimension at least $(1 - r_u(c))n$. However, the proportion of viable genotypes among all 2^n genotypes is exponentially small, by equation (8).

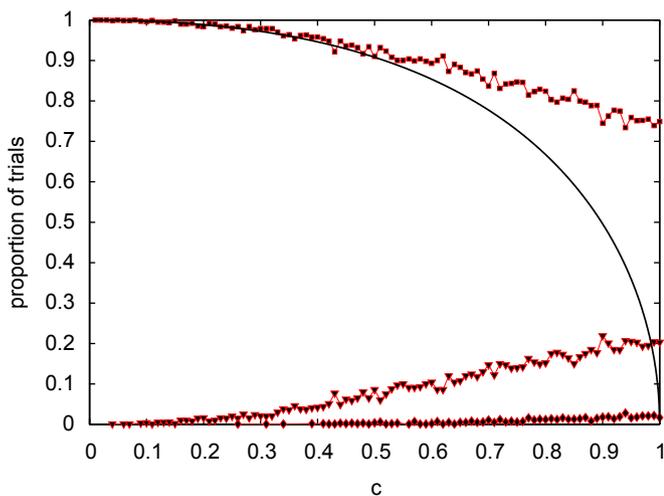


Fig. 3. Simulated number of clusters, vs. c for $n = 20$, with the proportion (out of 1000) of trials with exactly one (squares), exactly two (triangles), and at least three (diamonds) clusters. The solid curve is $\sqrt{(1 - c)e^c}$.

In fitness landscapes based on incompatibilities, high correlations make a significant proportion of genetic information irrelevant in the sense that genotypes can be changed by mutations at many loci without affecting fitness. This might be expected from the perspective that the model of lethal allele pairs (i.e., pairwise incompatibilities) is a natural generalization of the much simpler model of lethal alleles in which each of $2n$ alleles is independently lethal with probability p . In this model, the proper scale is $p = c/\sqrt{n}$. A viable genotype exists when there is no locus at which both alleles are lethal, which happens with probability $(1 - p^2)^n \sim e^{-c^2}$. (Note that there is not a sharp threshold in this mode, i.e., taking the existence probability from 1 to 0 requires varying c from 0 to ∞ .) Moreover, all the loci for which neither allele is lethal can be changed by mutation without affecting fitness. The number of such loci is about $(1 - p)^2 n \geq n - 2c\sqrt{n}$, so nearly all loci are neutral. It is also interesting to point out that the uncorrelated model, where viability of genotypes is independently determined is at the other extreme—viability of genotypes is based on lethal combinations of all n alleles.

In the following sections, we study the conditions for the existence of viable genotypes in several related models starting with complex incompatibilities.

5.2. Complex incompatibilities

Here we assume that incompatibilities involve K (≥ 2) diallelic loci (Cabot et al., 1994; Orr and Orr, 1996; Gavrillets, 2004). Determining whether a viable combination of genes exists is then equivalent to the K -SAT problem (Korte and Vygen, 2005). Even for $K = 3$, this is an NP-complete problem (Korte and Vygen, 2005), so there is no known polynomial algorithm to answer this question. The random case, which we now describe, is also much harder to analyze than the 2-SAT one. Let F be a random set to which any of the $2^K \binom{n}{K}$ incompatibilities belong independently with probability

$$p = \frac{K!}{2^K} \cdot \frac{c}{n^{K-1}}.$$

Here $c = c(K)$ is a constant, and the above form has been chosen to make the number of incompatibilities in F asymptotically cn (Note also the agreement with the definition of p in Section 5.1 when $K = 2$). For a fixed K , it has been proved (Friedgut, 1999) that the probability that viable genotype exists jumps sharply from 0 to 1 as c varies. However, the location of the jump has not been proved to converge as $n \rightarrow \infty$. Instead, a lot of effort has been invested in obtaining good bounds. For example (Achlioptas and Peres, 2004), for $K = 3$, $c < 3.42$ implies a.a.s. existence of viable genotype, while $c > 4.51$ implies a.a.s. nonexistence (while the sharp constant is estimated to be about 4.48, see e.g. Biroli et al., 2000). For $K = 4$ the best current bounds are 7.91 and 10.23. For large K , the transition occurs at $c = 2^K \log 2 - \mathcal{O}(K)$ (Achlioptas and Peres, 2004).

Techniques from statistical physics (Biroli et al., 2000) strongly suggest that, for $K \geq 3$, there is another phase transition, which for $K = 3$ occurs at about $c = 3.96$. For smaller c , the viable genotypes are conjectured to be contained in a *single* cluster. For larger c , the space of viable genotypes (if nonempty) is divided into exponentially many connected clusters. It appears that 3-SAT is already noticeably closer than 2-SAT to the uncorrelated model in Section 2 where the number of clusters is on the order of the number of viable genotypes and to the weakly correlated model of Section 4.

Perhaps more relevant to genetic incompatibilities is the following *mixed* model (commonly known as $(2+p)$ -SAT, Monasson and Zecchina, 1997). Assume that every 2-incompatibility is present with probability $c_2/(2n)$, while every 3-incompatibility is present with probability $3c_3/(4n^2)$. The normalizations are chosen so that the numbers of 2-incompatibilities and 3-incompatibilities are asymptotically c_2n and c_3n , respectively.

If c_2 (resp. c_3) is very small, then the respective incompatibility set affects a very small proportion of loci, therefore c_3 (resp. c_2) determines whether a viable genotype is likely to exist. Intuitively, one also expects that 2-incompatibilities should be more important than 3-incompatibilities as one of the former type excludes more genotypes than one of the latter type. A careful analysis confirms this. First observe that $c_2 > 1$ implies a.s. non-existence of a viable genotype. The surprise (Monasson and Zecchina, 1997; Achlioptas et al., 2001) is that if c_3 is small enough, $c_2 < 1$ implies a.s. existence of viable genotypes, so the 3-incompatibilities do not change the threshold. This is established in Monasson and Zecchina (1997) by a physics argument for $c_3 < 0.703$, while Achlioptas et al. (2001) gives a rigorous argument for $c_3 < \frac{2}{3}$. Therefore, even if their numbers are on the same scale, if the more complex incompatibilities are rare enough compared to the pairwise ones, their contribution to the structure of the space of viable genotypes is not essential.

5.3. Multiallelic loci

Here we assume that at each locus there can be a (≥ 2) alleles (cf., Reidys, 2006). In this case, the genotype space is the generalized hypercube $\mathcal{G}_a = \{0, \dots, a-1\}^n$. For $a = 3$ this could be interpreted as the genotype space of diploid organisms without *cis-trans* effects (Gavrilets and Gravner, 1997), $a = 4$ corresponds to DNA sequences, and $a = 20$ corresponds to proteins. Much larger values of a can correspond to a number of alleles at a protein coding locus and we will see later that there is not much difference between this model and a natural continuous space model.

We will assume that each pair of alleles, out of total number of $a^2 \binom{n}{2}$ is independently incompatible with probability

$$p = \frac{c}{2n}.$$

The main question we are interested in here is for which values of c viable genotypes exist a.s.

Clearly, if N is the number of viable genotypes, then the expectation

$$E(N) = a^n (1-p)^{\binom{n}{2}} \approx \exp(n \log a - \frac{1}{4}cn)$$

and so there are a.s. no viable genotypes when $c > 4 \log a$. On the other hand, clearly there are viable genotypes (with all positions filled by 0's and 1's) when $c < 1$. It turns out that the first of these trivial bounds is much closer to the critical value when a is large. Before we proceed, however, we state a sharp threshold result from Molloy (2003): there exists a function $\gamma = \gamma(n, a)$ so that for every $\varepsilon > 0$,

- if $c > \gamma + \varepsilon$, then a.s. $N = 0$.
- if $c < \gamma - \varepsilon$, then a.s. $N > 0$.

In words, for a fixed a , the probability of the event that $N \geq 1$ transitions sharply from large to small as c increases through γ . As it is not proved that $\lim_{n \rightarrow \infty} \gamma(n, a)$ exists, it is in principle possible that the place of this sharp transition fluctuates as n increases (although it must of course remain within $[1, 4 \log a]$).

Our main result in this section is

$$\gamma = 4 \log a - o(1) \quad \text{as } a \rightarrow \infty. \quad (10)$$

This somewhat surprising result is proven in Appendix D by the second moment method, as developed in Achlioptas and Moore (2004) and Achlioptas and Peres (2004).

Eq. (10) implies that viable genotypes exist only if the incompatibility probability

$$p < \frac{2 \log a}{n}, \quad (11)$$

which decreases linearly with the number of the loci n but increases only logarithmically with the number of alleles a . That is, multiple alleles widen the conditions for the existence of viable genotypes but the effect is not too strong.

5.4. Continuous phenotype spaces

Here we extend the model of pair incompatibilities for the case of continuous phenotypic space $\mathcal{P} = [0, 1]^n$. Again, each phenotype is characterized by values z_1, \dots, z_n , each between 0 and 1, which specify n continuous traits. This time we also have a small $r > 0$ as a parameter. For each pair of traits (i, j) , $i < j$, we independently choose a Poisson(λ) number of points Π_{ij} uniformly at random in the unit square $[0, 1] \times [0, 1]$. These points will be interpreted as “gorges” of equal depth in the fitness landscapes. As in the discrete case, we assume that $\lambda = c/(2n)$. Then we declare each phenotype $z = (z_1, \dots, z_n) \in \mathcal{P}$ inviable if it has two traits $i < j$ so that (z_i, z_j) is within r of the “gorge” Π_{ij} . Our procedure can be visualized as

throwing a random number of $(n - 2)$ -dimensional square tubes of inviable phenotypes into the phenotype space.

Our main result here is that the existence threshold is on the order $c \approx -\log r/r^2$. Namely, we prove in Appendix E that there exists a constant $C > 0$ so that for small enough r ,

- if $c > 4 - \log r/r^2$, then a.a.s. $N = 0$.
- if $c < -\log r - C/r^2$, then a.a.s. $N > 0$.

In the basic and continuous models in Sections 2 and 3, the difference in percolation thresholds (given by Eqs. (3) and (4)) was dramatic. In contrast, the continuous and discrete cases here are a lot closer. The scaling of the critical probability is $1/n$ for both cases, and as explained in Appendix E, the continuous case described here is related to the discrete one (from Section 5.3) with $a \approx 1/r$ and the incompatibility probability $p \approx r^2/n$.

Overall, the results presented in Section 5 show that the structure of fitness landscapes in incompatibility models are very different from that in models considered in earlier sections. Although both types of models result in large clusters of viable genotypes, the former allow for more neutral changes on local scales (because incompatibilities are rare) but for less extensive genetic divergence on larger scales (because viable clusters are represented by subcubes rather than trees) than the latter.

6. Notes on neutral clusters in the discrete NK model

The model considered here is a special case of the discretized NK model (Kauffman, 1993), introduced in Newman and Engelhardt (1998a). This model features n diallelic loci each of which interacts with K other loci. To have a concrete example, assume that the loci are arranged on a circle, so that $n + 1 \equiv 1$, $n + 1 \equiv 2$, etc., and let the interaction *neighborhood* of the i th locus consist of itself and K loci to its right $i + 1, \dots, i + K$. For a given genotype $x \in \mathcal{G} = \{0, 1\}^n$, the neighborhood configuration of the i 'th locus is then given by $\mathcal{N}_i(x) = (x_i, x_{i+1}, \dots, x_{i+K}) \in \{0, 1\}^{K+1}$. To each locus and to each possible configuration in its neighborhood we independently assign a binary fitness contribution. To be more precise, we choose the $2^{K+1}n$ numbers $v_i(y)$, $i = 1, \dots, n$ and $y \in \{0, 1\}^{K+1}$, to be independently 0 or 1 with equal probability, and interpret $v_i(y)$ as the fitness contribution of locus i when its neighborhood configuration is y . The fitness of a genotype x is then the sum of contributions from each locus

$$w(x) = \sum_{i=1}^n v_i(\mathcal{N}_i(x)).$$

In Kauffman (1993), the values v_i were taken from a continuous distribution. In Newman and Engelhardt (1998a), these values were integers in the range $[0, F - 1]$ so that our model is a special case $F = 2$. The NK model results in a correlated fitness landscape; the corresponding

correlation function (1) can be found explicitly (Fontana et al., 1993; Campos et al., 2002; Gavrillets, 2004).

We define *neutral clusters* as connected components of same fitness. The $K = 0$ case, which corresponds to an additive fitness model, is easy but nevertheless illustrative. Namely, a mutation at locus i will not change fitness iff $v_i(0) = v_i(1)$; let D be the number of such loci. Since the fitnesses contributions are assigned randomly and with equal probabilities, each locus has a 50% chance to be in this category. Then $D \sim n/2$ a.a.s., the number of different fitnesses is $n - D$, each neutral cluster is a sub-cube of dimension D , and there are exactly 2^{n-D} neutral clusters.

The next simplest situation is when $K = 1$, so that each locus interacts with one other locus. Let D_1 be the number of loci i for which the fitness contribution v_i is constant, i.e., does not depend on the configuration in the i th neighborhood \mathcal{N}_i . Then $D_1 \sim n/8$ a.a.s. (because out of the $2^4 = 16$ possible functions v_i , two are constant, the 0 and the 1), and each neutral cluster contains a sub-cube of dimension D_1 . Moreover, let D_2 be the number of loci i for which $v_i(00) = v_i(01) \neq v_i(10) = v_i(11)$. Note that there are again two functions with this property, and that any genotypes that differ at such a locus i must belong to a different neutral cluster. Consequently, the number of different neutral clusters is at least 2^{D_2} and there are exponentially many of them, as again $D_2 \sim n/8$ a.a.s. This division of genotype space into exponentially many clusters of exponential size persists for every K , although the distribution of numbers and sizes of these clusters is not well understood (see Newman and Engelhardt, 1998a for simulations for $n = 20$).

Overall, neutral clusters are a prominent feature of the NK model. We also mention that the question of whether a genotype with the maximal possible fitness n exists for a given K is in many way related to issues in incompatibilities models (Choi et al., 2005).

7. Discussion

In this section we summarize our major findings and provide their biological interpretation.

The previous work on neutral and nearly neutral networks in multidimensional fitness landscapes has concentrated exclusively on genotype spaces in which each individual (or a group of individuals) is characterized by a discrete set of genes. However many features of biological organisms that are actually observable and/or measurable are described by continuously varying variables such as size, weight, color, or concentration. A question of particular biological interest is whether (nearly) neutral networks are as prominent in a continuous phenotype space as they are in the discrete genotype space. Our results provide an affirmative answer to this question. Specifically, we have shown that in a simple model of random fitness assignment, viable phenotypes are likely to form a large connected cluster even if their overall frequency is very low provided the dimensionality of the phenotype space, n , is

sufficiently large. In fact, the percolation threshold for the probability of being viable scales with n as $1/2^n$ and, thus, decreases much faster than $1/n$ which is characteristic of the analogous discrete genotype space model.

Earlier work on nearly neutral networks has been limited to consideration of the relationship between genotype and fitness. Any phenotypic properties that usually mediate this relationship in real biological organisms have been neglected. In Section 4, we proposed a novel model in which phenotype is introduced explicitly. In our model, the relationships both between genotype and phenotype and between phenotype and fitness are of many-to-one type, so that neutrality is present at both the phenotype and fitness levels. Moreover, this model results in a correlated fitness landscape in which the correlation function can be found explicitly. We showed that phenotypic neutrality and correlation between fitnesses can reduce the percolation threshold. Our results suggest that the most conducive conditions for the formation of the giant component is when the correlations are at the point of phase transition between local and global. To explore the robustness of our conclusions, we then look at a simplistic but mathematically illuminating model in which there is a correlation between conformity (i.e., phenotypic neutrality) and fitness. The model has supported our conclusions.

In Section 5, we studied a number of models that have been recently proposed and explored within the context of studying speciation. In these models, fitness is assigned to particular gene/trait combinations and the fitness of the whole organisms depends on the presence or absence of incompatible combinations of genes or traits. In these models, the correlations of fitnesses are so high that local methods lead to wrong conclusions. First, we established the connection between these models and SAT problems, prominent in computer science. Then we analyzed the conditions for the existence of viable genotypes, their number, as well as the structure and the number of clusters of viable genotypes. These questions have not been studied previously. The majority of our results are for the case of pairwise incompatibilities between diallelic loci, but we also looked at multiple alleles and complex incompatibilities. In the case of diallelic loci we showed that the number of clusters is stochastically bounded and each cluster contains a very large sub-cube. Our results suggest (in the context of the mixed model) that more complex incompatibilities are less important than less complex incompatibilities in controlling the structure of the space of viable genotypes. However, complex incompatibilities appear to allow for exponentially many clusters of viable genotypes. Finally, we generalized some of our findings to continuous phenotype spaces.

In Section 6, we provided some additional results on the size, number and structure of neutral clusters in the discrete NK model which represent a popular tool for studying adaptation (Kauffman and Levin, 1987; Kauffman, 1993; Newman and Engelhardt, 1998b; Welch and Waxman, 2005).

In the majority of models studied here, we assumed for simplicity that fitness can take only two values: zero and one. In Section 2, we mentioned the relationship between viable genotypes and fitness bands. This concept is explored in some detail in our previous work (Gavrilets, 1997; Gavrilets and Gravner, 1997; Gavrilets, 2004) where we showed that many results on viable genotypes in uncorrelated landscapes easily generalize to the case of continuously changing fitness. The same is true of the models in Sections 3 and 4. For simplicity, our analysis has been restricted to very symmetric models. However, we expect that many conclusions of this paper can be generalized to nonsymmetric fitness assignments provided the dimensionality of fitness landscapes is very large.

Overall, our results reinforce the previous conclusion (Gavrilets, 1997; Gavrilets and Gravner, 1997; Reidys et al., 1997; Gavrilets, 2004; Reidys et al., 2001; Reidys and Stadler, 2001, 2002) that percolating networks of genotypes with approximately similar fitnesses (holey landscapes) is a general feature of *multidimensional* fitness landscapes (both uncorrelated and correlated and both in genotype and phenotype spaces). The evolution along such networks can readily proceed by mutation and random genetic drift (Kimura, 1983; Ohta, 1992, 1998; Lynch, 2007) and result in the accumulation of reproductive isolation between diverging lineages and in allopatric, parapatric or sympatric speciation (Cabot et al., 1994; Orr, 1995; Orr and Orr, 1996; Orr, 1997; Orr and Turelli, 2001; Gavrilets and Hastings, 1996; Gavrilets, 1997; Gavrilets and Gravner, 1997; Gavrilets, 2003, 2004; Coyne and Orr, 2004). Selection (e.g., for local adaptation or sexual) can either help or hinder these processes depending on whether its net effect is diversifying or stabilizing.

Some more general theoretical lessons of our work are that

- Correlations may help or hinder connectivity in fitness landscapes. Even when correlations are positive and tunable by a single parameter, it may be advantageous (for higher connectivity) to increase them only to a limited extent. One reason (see Section 4.2) is that high correlations may result in a loss of very large phenotypically similar clusters.
- Averages (i.e., expected values) can easily lead to wrong conclusions, especially when correlations are strong (see Section 5.1). Nevertheless, they may still be useful with a crafty choice of relevant statistics (see Achlioptas and Peres, 2004 and Appendix D).
- Very high correlations may fundamentally change the structure of connected clusters. For example, clusters may look locally more like cubes (Section 5) than trees (Sections 2–4) and their number may be reduced dramatically (Section 5).
- Necessary analytic techniques may be unexpected and quite sophisticated; for example, they may require detailed understanding of random graphs, spin-glass machinery, or decision algorithms.

Acknowledgments

This work was supported by the Defense Advanced Research Projects Agency (DARPA), by National Institutes of Health (Grant GM56693), by the National Science Foundation (Grants DMS-0204376 and DMS-0135345), and by Republic of Slovenia's Ministry of Science (Program P1-285).

Appendix A. Proof of Eq. (5)

To prove Eq. (5), we assume that $\mu < 1$ and show that for a fixed k (which does not grow with n), the event that x and y at distance k are in the same conformist cluster is most likely to occur because x and y are connected via the shortest possible path. Indeed, the dominant term $k!q^k$ is the expected number of conformist pathways between x and y that are of shortest possible length k . This easily follows from the observation that on a shortest path there is no opportunity to backtrack; each mutation must be toward the other genotype. We can assume that x is the all 0's genotype and y is the genotype with 1's in the first k positions and 0's elsewhere. There are $k!$ orders in which the 1's can be added.

To obtain the lower bound we use inclusion-exclusion on the probability that $x \leftrightarrow y$ through a shortest path. Let $\mathcal{J}_k = \mathcal{J}_k(x, y)$ be the set of all paths of length k between x and y . Then

$$P(x \leftrightarrow y) \geq \sum_{\alpha \in \mathcal{J}_k} P(A_\alpha) - \sum_{\alpha \neq \beta \in \mathcal{J}_k} P(A_\alpha \cap A_\beta)$$

where A_α is the event that a particular path α consists entirely of conformist edges. Notice that two distinct paths of the same length differ by at least two edges. Thus, we get the following upper bound

$$\sum_{\alpha, \beta} P(A_\alpha \cap A_\beta) < (k!)^2 q^{k+2},$$

and the lower bound in Eq. (5) follows.

The upper bound is a little more difficult to obtain (it is only here that we use $\mu < 1$) and we need some notation. Each genotype can be identified with the set of 1's that it contains, so for any two genotypes u and v we let $u \Delta v$ denote the set of loci on which they differ. Notice that if $u \Delta v$ is even (resp. odd) then every path between u and v is of even (resp. odd) length because each mutation which alters the allele at a locus not in $u \Delta v$ must later be compensated for.

To estimate the expected number of conformist pathways, we will need to bound the number of paths of length l between x and y . This is given by

$$k! \binom{l}{m} m! n^m \quad \text{where } m = \frac{l-k}{2}.$$

We show this via the methods of Bollobás et al. (1992). They obtain an estimate for the number of cycles of a given length through a fixed vertex of the cube.

Given a path, say $x = v_0, v_1, \dots, v_l = y$, between x and y , let us associate the sequence $(\varepsilon_1 i_1, \dots, \varepsilon_l i_l)$ where

$$v_j \Delta v_{j-1} = \{i_j\} \quad \text{and} \quad \varepsilon_j = \begin{cases} +1 & \text{if } v_j = v_{j-1} \cup i_j, \\ -1 & \text{if } v_j = v_{j-1} \setminus \{i_j\}, \end{cases}$$

where $j = 1, \dots, l$. Since distinct paths will have distinct sequences we can bound the number of paths by finding an upper bound for the number of sequences.

Note that there must be $m+k$ positive entries, which occur at $\binom{l}{m+k} = \binom{l}{m}$ possible locations. The absolute values of m of these entries are chosen freely from $\{1, \dots, n\}$, while the remaining k must be the integers $1, \dots, k$. There are $n^m k!$ ways to do this. We are free to order the m negative entries and the bound follows.

We now assume that $d(x, y)$ is even and relabel $d(x, y) = 2k$. We omit the similar calculation for odd distances. Define $b = -3k/(2 \log \mu)$ and $t = \lfloor b \log n \rfloor$. Then the expected number of conformist paths between x and y can be expressed as

$$\begin{aligned} & \sum_{l \geq k+1} \sum_{\mathcal{J}_{2l}} q^{2l} \\ &= \sum_{k+1 \leq l < t} \sum_{\mathcal{J}_{2l}} q^{2l} + \sum_{l \geq t} \sum_{\mathcal{J}_{2l}} q^{2l} \\ &< \sum_{k+1 \leq l < t} \binom{2l}{l-k} n^{l-k} (l-k)! (2k)! q^{2l} + \sum_{l \geq t} n^{2l} q^{2l} \\ &= \sum_{k+1 \leq l < t} (2l)^{l-k} n^{l-k} q^{2(l-k)} (2k)! q^{2k} + \sum_{l \geq t} \mu^{2l} \\ &< (2k)! q^{2k} \sum_{l \geq k+1} (2b\mu q \log n)^{l-k} + O(\mu^{2b \log n}) \\ &= k(2k)! q^{2k} O(q \log n) + O(n^{2b \log \mu}) \\ &= k(2k)! q^{2k} O(n^{-1} \log n). \end{aligned}$$

Appendix B. Correlations between conformity and viability

Assume now that conformist clusters are formed as in Section 4 (i.e., with edges being conformist with probability $q = \mu/n$), are still independently viable, but now the probability of their viability depends on their size. We will consider the simple case when an isolated genotype (one might call it *non-conformist*) is viable with probability $p_0 = \lambda_0/n$, while a conformist cluster of size larger than 1 is viable with probability $p_1 = \lambda_1/n$.

In this case, the probability that a random genotype is viable,

$$\begin{aligned} P(x \text{ is viable}) &= (1-q)^n p_0 + (1-(1-q)^n) p_1 \\ &\sim \frac{1}{n} (e^{-\mu} \lambda_0 + (1-e^{-\mu}) \lambda_1). \end{aligned}$$

Moreover, by a similar calculation as before,

$$\begin{aligned} & P(x \text{ and } y \text{ viable}) - P(x \text{ viable})^2 \\ &= p_1(1-p_1)P(x \leftrightarrow y) \\ &\quad + P(x \text{ non-conformist})^2 q(p_0 - p_1)^2 \cdot 1_{\{d(x,y)=1\}}. \end{aligned}$$

Here, the last factor is the indicator of the set $\{(x, y), d(x, y) = 1\}$, which equals 1 if $d(x, y) = 1$ and 0 otherwise. Therefore, for $d(x, y) \geq 2$, the correlation function (1) is

$$\rho(x, y) \sim \frac{\lambda_1}{e^{-\mu}\lambda_0 + (1 - e^{-\mu})\lambda_1} P(x \leftrightarrow y),$$

which is smaller than before iff $\lambda_1 < \lambda_0$. However, it has the same asymptotic properties unless $\lambda_1 = 0$.

Assume first that $\mu < 1$. The local analysis now leads to a *multi-type* branching process (Athreya and Ney, 1971) with three types: NC (non-conformist node), CI (non-isolated node independently viable, so no conformist edge is accounted for), and CC (non-isolated node viable by conformity, so a conformist edge is accounted for). Note that while the type NC is a new feature, the difference between CI and CC was the key to our earlier local analysis in Section 4.1—a neighbor of a viable genotype may be viable because of a conformist edge between them or because it belongs to an independently viable conformist cluster.

Note first that a genotype is non-conformist with probability about $e^{-\mu}$. Hence a node of any of the three types creates a Poisson ($e^{-\mu}\lambda_1$) number of type NC descendants, and a Poisson ($(1 - e^{-\mu})\lambda_1$) number of type CI descendants. In addition, the type CI creates a Poisson(μ), conditioned on being nonzero, number of descendants of type CC and type CC creates a Poisson(μ) number of descendants of type CC. Thus the matrix of expectations, in which the ij th entry is the expectation of the number of type j descendants from type i , is

$$M = \begin{bmatrix} e^{-\mu}\lambda_0 & (1 - e^{-\mu})\lambda_1 & 0 \\ e^{-\mu}\lambda_0 & (1 - e^{-\mu})\lambda_1 & \mu/(1 - e^{-\mu}) \\ e^{-\mu}\lambda_0 & (1 - e^{-\mu})\lambda_1 & \mu \end{bmatrix}.$$

When $\mu > 1$, μ needs to be replaced by $\mu\delta$, and λ_1 by $\lambda_1\delta$, where $\delta = \delta(\mu)$ is given by Eq. (2).

It follows from the theory of multi-type branching processes (Athreya and Ney, 1971) that the critical surface for survival of a multi-type branching process is given by $\det(M - 1) = 0$.

The simplest case is when only non-conformist genotypes may be viable, i.e., $\lambda_1 = 0$. In this case the critical surface is given by $\lambda_0 e^{-\mu} = 1$ (Pitman, unpub.). Not surprisingly, the critical λ_0 necessary to achieve global connectivity strictly increases with μ , which is the result of negative correlations between conformity and viability.

The other extreme is when non-conformist genotypes are inviable, i.e., $\lambda_0 = 0$. As an easy computation demonstrates, the critical curve is now given by $\lambda_1 = \zeta(\mu)$, where

$$\zeta(\mu) = \begin{cases} \frac{1 - \mu}{\mu e^{-\mu} + 1 - e^{-\mu}} & \text{if } \mu \in [0, 1], \\ \frac{\rho^{-1} - \mu}{\mu e^{-\mu} + 1 - e^{-\mu\rho}} & \text{if } \mu \in [1, \infty). \end{cases} \quad (12)$$

Note that $\zeta(\mu) \rightarrow \infty$ as $\mu \rightarrow 0$. We carried out exactly the same simulations as before. These are also featured in

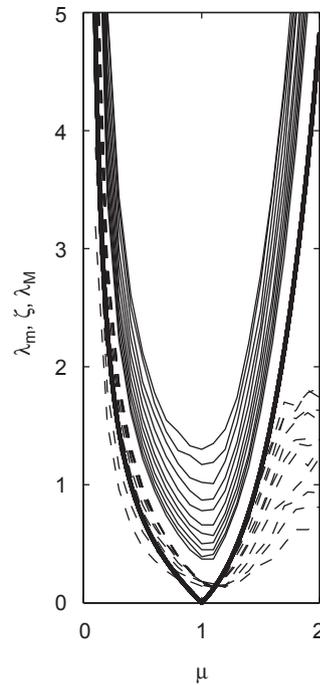


Fig. 4. Simulated lower bounds λ_m (dashed curves) and upper bounds λ_M (solid curves), and ζ (thick curve) plotted against μ , for $n = 10, \dots, 20$ in the model from Appendix B. Lower bounds increase with n , and upper bounds decrease, for this range of n . Compare to Fig. 2.

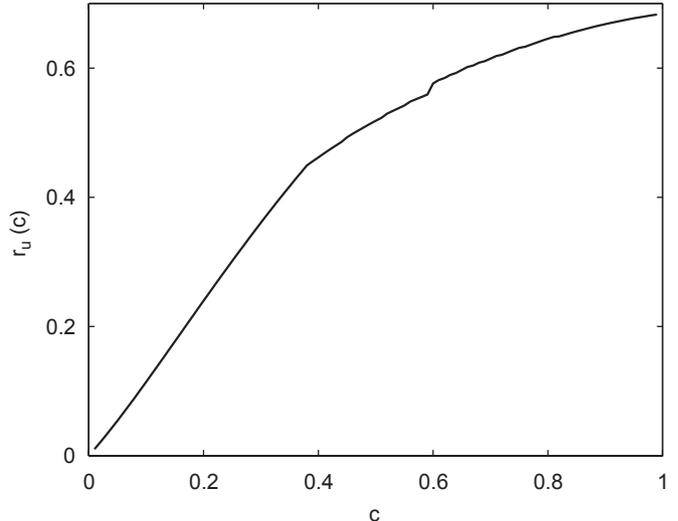


Fig. 5. The upper bound $r_u(c)$ for the number of fixed loci (non-*'s) in the implicant of shortest length included in every cluster of viable genotypes, plotted against c . As explained in the text, $1 - r_u(c)$ is a lower bound on the dimension of the subcube specified by the implicant.

Fig. 4, and again confirm our local heuristics. We conclude that positive correlations between viability and conformity tend to lead to a V-shaped critical curve, whose sharpness at critical conformity $\mu = 1$ increases with the size of correlations. In short, then, correlations help more if viability probability increases with size of conformist clusters (Fig. 5).

Appendix C. Cluster structure under random pair incompatibilities

Here we show that, under random pairwise incompatibilities model introduced in Section 5.1, connected clusters include large subcubes. The basic idea comes from Boufkhad and Dubois (1999). A configuration $a \in \{0, 1, *\}^n$ is a way to specify a sub-cube of \mathcal{G} , if $*$'s are thought of as places which could be filled by either a 0 or a 1. The number of non- $*$'s is the *length* of a . Call a an *implicant* if the entire sub-cube specified by a is viable.

We present two arguments, beginning with the one which works better for small c . Let the auxiliary random variable X be the number of pairs of loci (i, j) , $i < j$, for which:

- (E1) There is exactly one incompatibility involving alleles on i and j .
- (E2) There is no incompatibility involving an allele on either i or j , and an allele on $k \notin \{i, j\}$.

Assume, without loss of generality, that the incompatibility which satisfies (E1) is $(1_i, 1_j)$. Then fitness of all genotypes which have any of the allele assignments $0_i 0_j$, $0_i 1_j$ and $1_i 0_j$, and agree on other loci, is the same. Note also that all pairs of loci which satisfy (E1) and (E2) must be disjoint. Therefore, if x is any viable genotype, its cluster contains an implicant with the number of $*$'s at least X plus the number of free loci. To determine the size of X , note that the expectation

$$E(X) = \binom{n}{2} 4p(1-p)^3(1-p)^{8(n-2)} \sim ce^{-4c}n$$

and furthermore, by an equally easy computation,

$$E(X^2) - E(X)^2 = \mathcal{O}(n),$$

so that $X \sim ce^{-4c}n$ a.a.s. It follows that every cluster contains a.a.s. at least $\exp((e^{-2c} + ce^{-4c}) \log 2 - \varepsilon)n$, viable genotypes, for any $\varepsilon > 0$.

The second argument is a refinement of the one in Boufkhad and Dubois (1999) and only works better for larger c . Call an implicant a a *prime implicant (PI)* if at any locus i , replacement of either 0_i or 1_i by $*$ results in a non-implicant. Moreover, we call a the *least prime implicant (LPI)* if it is a PI, and the following two conditions are satisfied. First, if all the $*$'s are changed to 0's, then no change from 1_i to 0_i results in a viable genotype. Second, no change $*_i 1_j$ to $1_i *$, where $i < j$, results in an indicator.

Now, every viable genotype must have an LPI in its cluster. To see this, assume we have a PI for which the first condition is not satisfied. Make the indicated change, then replace some 0's and 1's by $*$'s until you get a prime indicator. If the second condition is violated, make the resulting switch, then again make some replacement by $*$'s until you arrive at a PI. Either of these two operations moves within the same cluster, and keeps the number of 1's nonincreasing and their positions more to the left. Therefore, the procedure must at some point end, resulting in an LPI in the same cluster.

For a sub-cube a to be an LPI, the following conditions need to be satisfied:

- (I1) Every non- $*$ has to be compatible with every other non- $*$, and with both 0 and 1 on each of the $*$'s.
- (I2) Any of the four 0,1 combinations on any pair of $*$'s must be compatible.
- (LPI1) Pick an i with allele 1, that is, a 1_i . Then 0_i must be incompatible with at least one non- $*$, or at least one 0 on a $*$. Furthermore, if 0_i has an incompatibility with a 0 on a $*$ to its left, it has to have another incompatibility, either with a non- $*$, or with a 0 or a 1 on a $*$.
- (LPI2) Pick a 0_i . Then 1_i must be incompatible with a non- $*$, or a 0 or a 1 on a $*$.

The first two conditions make a an implicant, and the last two an LPI. Note also that these conditions are independent.

Let now X be the number of LPI of length m . We will identify a function $L_4 = L_4(r, c)$ such that

$$\frac{1}{n} \log E(X) \leq L_4.$$

Let

$$L_1 = L_1(\beta, p, z) = z(\beta \log p + (1 - \beta) \log(1 - p) - \beta \log \beta - (1 - \beta) \log(1 - \beta)).$$

This is the exponential rate for the probability that in zn Bernoulli trials with success probability p there are exactly βn successes, i.e., this probability is $\approx \exp(L_1 n)$. Further, if $\kappa, \varepsilon, \delta \in (0, 1)$ are fixed, then among sub-cubes with rn non- $*$'s and αn 1's ($\alpha \leq r$), the proportion which have εn 1's in $[\kappa n, n]$ and δn $*$'s in $[1, \kappa n]$ has exponential rate

$$L_2 = L_2(r, c, \kappa, \alpha, \varepsilon, \delta) \\ = L_1((\alpha - \varepsilon)/\kappa, \alpha, \kappa) + L_1(\varepsilon/(1 - \kappa), \alpha, 1 - \kappa) \\ + L_1(\delta/(\kappa - \alpha + \varepsilon), 1 - r, \kappa - \alpha + \varepsilon) \\ + L_1((1 - r - \delta)/(1 - \kappa - \varepsilon), 1 - r, 1 - \kappa - \varepsilon).$$

(Here all four first arguments in L_1 are in $[0, 1]$, or else the rate is $-\infty$.)

The expected number of LPI, with $r, \kappa, \varepsilon, \delta$ given as above, has exponential rate at most (and this is only an upper bound)

$$L_3 = L_3(r, c, \kappa, \alpha, \varepsilon, \delta) \\ = -(1 - r) \log(1 - r) - \alpha \log \alpha - (r - \alpha) \log(r - \alpha) \\ - c(1 - r/2)^2 + (r - \alpha) \log(1 - \exp(-c(1 - r/2))) \\ + (\alpha - \varepsilon) \log(1 - \exp(-c/2)) + \varepsilon \log(1 - \exp(-c/2)) \\ - \frac{1}{2} \delta c \exp(-c(1 - r/2)) + L_2(r, c, \kappa, \alpha, \varepsilon, \delta).$$

The next to last line is obtained from (LPI1), as εn 1's must have δn $*$'s on their left.

It follows that L_4 can be obtained by

$$L_4(r, c) = \inf_{\kappa} \sup_{\alpha, \varepsilon, \delta} L_3(r, c, \kappa, \alpha, \varepsilon, \delta).$$

If $L_4(r, c) < 0$, all LPI (for this c) a.a.s. have length at most r . Numerical computations show that this gives a better bound than $1 - e^{-2c} - ce^{-4c}$ for $c \geq 0.38$. Let us denote the best upper bound from the two estimates by $r_u(c)$. This function is computed numerically and plotted in Fig. 3.

Appendix D. Proof of Eq. (10)

In this section we assume that genotypes have multi-allelic loci, which are subject to random pair incompatibilities. The model introduced in Section 5.2 is the most natural, but is not best suited for our second moment approach. Instead, we will work with the equivalent modified model with m pair incompatibilities, each chosen independently at random, and the first and the second member of each pair chosen independently from the an available alleles. We will assume that $m = \frac{1}{4}ca^2n$, label $c' = \frac{1}{4}c$, and denote, as usual, the resulting set of incompatibilities by F .

To see that these two models are equivalent for our purposes, first note that the number of incompatibilities which are *not legitimate*, in the sense that the two alleles are chosen from the same locus, is stochastically bounded in n . (In fact, it converges in distribution to a Poisson($c'a^2$) random variable.) Moreover, by the Poisson approximation to the birthday problem (Barbour et al., 1992), the number of pairs of choices which result in the same incompatibility in this model is asymptotically Poisson($c'a^2/2$). In short, then, the procedure results in the number $m - \mathcal{O}(1)$ of different legitimate incompatibilities. If m in the modified model is increased to, say, $m' = m + n^{2/3}$, then the two models could be coupled so that the incompatibilities in the original model are included in those in the modified model. As the existence of a viable phenotype becomes less likely when m is increased, this demonstrates that Eq. (10) will follow once we show the following for the modified model: for every $\varepsilon > 0$ there exists a large enough a so that $c' < \log a - \varepsilon$ implies that $N \geq 1$ a.a.s.

To show this, we introduce the auxiliary random variable

$$A = \frac{(a-1)^\alpha}{a \cdot \alpha^\alpha (1-\alpha)^{1-\alpha}} \times \frac{(w_0^2(1 - \frac{1+\alpha}{a})^2 + 4w_0w_1\frac{\alpha}{a}(1 - \frac{1+\alpha}{a}) + 2w_1^2(\frac{1-\alpha}{a}(1 - \frac{1+\alpha}{a}) + (\frac{\alpha}{a})^2))^{c'a^2}}{(w_0(\frac{a-1}{a})^2 + w_1\frac{2(a-1)}{a^2})^{2c'a^2}}.$$

$$X = \sum_{\sigma \in \mathcal{G}_a} \prod_{I \in F} (w_0 1_{\{|I \cap \sigma| = 0\}} + w_1 1_{\{|I \cap \sigma| = 1\}}),$$

where 1_A is the indicator of the set A . The size of the intersection $I \cap \sigma$ is computed by transforming both the incompatibility I and the genotype σ to sets of (indexed)

alleles, and the weights w_0 and w_1 will be chosen later. To intuitively understand the statistic X , note that when $w_0 = w_1 = 1$, the product is exactly the indicator of the event that σ is viable and X is then the number of viable genotypes N . In general, X gives different scores to different viable genotypes—however, the crucial fact to note is that $X > 0$ iff $N > 0$. Therefore

$$P(N > 0) = P(X > 0) \geq (E(X))^2 / E(X^2),$$

which is how the second moment method is used (Achlioptas and Moore, 2004).

As

$$P(|\sigma \cap I| = 0) = \left(\frac{a-1}{a}\right)^2,$$

$$P(|\sigma \cap I| = 1) = \frac{2(a-1)}{a^2},$$

we have

$$E(X) = a^n \left(w_0 \left(\frac{a-1}{a}\right)^2 + w_1 \frac{2(a-1)}{a^2} \right)^m.$$

Moreover

$$E(X^2) = \sum_{k=0}^n a^n \binom{n}{k} (a-1)^k \times (w_0^2 P(00) + 2w_0w_1 P(01) + w_1^2 P(11)),$$

where $P(01)$ is the probability that I has intersection of size 0 with $\sigma = 0_1 \dots 0_k 0_{k+1} \dots 0_n$ and of size 1 with $\tau = 1_1 \dots 1_k 0_{k+1} \dots 0_n$, and $P(00)$ and $P(11)$ are defined analogously. Thus, if $k = \alpha n$,

$$P(00) = \left(1 - \frac{1+\alpha}{a}\right)^2,$$

$$P(01) = \frac{2\alpha}{a} \left(1 - \frac{1+\alpha}{a}\right),$$

$$P(11) = \frac{2(1-\alpha)}{a} \left(1 - \frac{1+\alpha}{a}\right) + 2\left(\frac{\alpha}{a}\right)^2.$$

Let $A = A_{a, w_0, w_1}(\alpha)$ be the n th root of the $k = (\alpha n)$ th term in the sum for $E(X^2)$, divided by $E(X)^2$. Hence

Let $\alpha^* = (a-1)/a$. A short computation shows that $A = 1$ when $\alpha = \alpha^*$.

If $A > 1$ for some α , then $E(X^2)/(E(X))^2$ increases exponentially and the method fails (as we will see below, this always happens when $w_0 = w_1 = 1$, i.e., when $X = N$). On the other hand, if $A < 1$ for $\alpha \neq \alpha^*$, and $\frac{d^2 A}{d\alpha^2}(\alpha^*) < 0$, then

Lemma 3 from Achlioptas and Moore (2004) implies that $E(X^2)/(E(X))^2 \leq C$ for some constant C , which in turn implies that $P(N > 0) \geq 1/C$. The sharp threshold result then finishes off the proof of Eq. (10).

Our aim then is to show that w_0 and w_1 can be chosen so that, for $c' = \log a - \varepsilon$, A has the properties described in the above paragraph. We have thus reduced the proof of Eq. (10) to a calculus problem.

Certainly the necessary condition is that $\frac{dA}{d\alpha}(\alpha^*) = 0$, and

$$\frac{dA}{d\alpha}(\alpha^*) = -\frac{2}{a^3}(w_0(a-1) - w_1(a-2))^2,$$

so we choose $w_0 = a - 2$ and $w_1 = a - 1$. (Only the quotient between w_0 and w_1 matters, so a single equation is enough.) This simplifies A to

$$A = A_a(\alpha) = \frac{(a-1)^2}{a\alpha^2(1-\alpha)^{1-\alpha}} \cdot \frac{\left(\alpha - \frac{a-1}{a}\right)^2 - \frac{(a-1)^4}{a^2}}{\left(\frac{(a-1)^2}{a}\right)^{2c'a^2}}.$$

Let $\varphi = \log A$. We need to demonstrate that $\varphi < 0$ for $\alpha \in [0, \alpha^*) \cup (\alpha^*, 1]$ and that $\varphi''(\alpha^*) < 0$. A further simplification can be obtained by using $x - Cx^2 \leq \log(1+x) \leq x$ (valid for all nonnegative x), which enables us to transform φ (without changing the notation) to

$$\varphi(\alpha) = c' \frac{a^4}{(a-1)^4} \left(\alpha - \frac{a-1}{a}\right)^2 - \alpha \log \alpha - (1-\alpha) \log(1-\alpha) + \alpha \log(a-1) - \log a.$$

Now

$$\varphi''(\alpha) = 2c' \frac{a^4}{(a-1)^4} - \frac{1}{\alpha(1-\alpha)}.$$

Table 1

The lower bounds on γ obtained by the method described in text, compared to the easy upper bounds $4 \log a$

a	l.b.on γ	$4 \log a$
3	1.679	4.395
4	2.841	5.546
5	3.848	6.438
6	4.714	7.168
7	5.467	7.784
8	6.128	8.318
9	6.715	8.789
10	7.242	9.211
20	10.672	11.983
30	12.608	13.605
40	13.944	14.756
50	14.960	15.649
100	18.017	18.421
200	20.982	21.194
300	22.663	22.816
400	23.846	23.966
500	24.759	24.859

So automatically, for c' large but $c' = o(a)$, $\varphi''(\alpha^*) < 0$ for large a . Moreover, φ cannot have another local maximum when $\varphi'' > 0$. If $\varphi(\alpha) \geq 0$ for some $\alpha \neq \alpha^*$, then this must happen for an α in one of the two intervals $[0, 1/(2c') + \mathcal{O}((c')^{-2})]$ or $[1 - 1/(2c') - \mathcal{O}((c')^{-2}), 1]$. Now, φ has a unique maximum at α^* in the second interval. In the first interval, a short computation shows that

$$\varphi(\alpha) = -\varepsilon - \alpha \log a + \mathcal{O}\left(\frac{\log \log a}{\log a}\right),$$

which is negative for large a . This ends the proof. \square

This method yields nontrivial lower bounds for γ for all $a \geq 3$, cf. Table 1.

Appendix E. Existence of viable phenotypes

In this section we describe a comparison between models from Sections 5.2 and 5.3 that will yield the result in Section 5.3. We discretize the continuous phenotype space and in the process obtain slightly dependent incompatibilities which cause a loss of precision by the factor 4.

We begin by assuming that $a = 1/r$ is an integer, which we can do without loss of generality. Divide the i th coordinate interval $[0, 1]$ into a disjoint intervals $I_{i0}, \dots, I_{i,a-1}$ of length r . Note that in our notation, the first subscript of I specifies a particular trait while the second subscript specifies a particular interval of the trait's values. To a phenotype $x \in \mathcal{P}$ we assign an element $\Delta(x)$ of the generalized hypercube \mathcal{G}_a by $\Delta(x)_i = j$ iff $x_i \in I_{ij}$.

Note that $\Delta(x)_i = \Delta(y)_i$ implies that $|x_i - y_i| \leq r$. Therefore, as soon as $I_{i_1 j_1} \times I_{i_2 j_2}$ contains a point $z \in \Pi_{i_1 i_2}$, every x with $\Delta(x)_{i_1} = j_1$ and $\Delta(x)_{i_2} = j_2$ has $|x_{i_1} - z_{i_1}| \leq r$ and $|x_{i_2} - z_{i_2}| \leq r$ and is therefore inviable by definition. Moreover, by the fundamental property of a Poisson point location (e.g., Penrose, 1996), each product $I_{i_1 j_1} \times I_{i_2 j_2}$ intersects $\Pi_{i_1 i_2}$ independently with probability $1 - \exp(-\lambda r^2) \sim cr^2/(2n)$, for large n . The result from Section 5.2 implies that, when $cr^2 > 4 \log a = -4 \log r$, a.a.s. no $\Delta(x)$ is viable and thus there is a.a.s. no viable genotype.

On the other hand, let I^ε be the closed ε -neighborhood of the interval I in $[0, 1]$ (the set of points within ε of I), and consider the events that the square $I_{i_1 j_1}^{r/2} \times I_{i_2 j_2}^{r/2}$ contains a point in $\Pi_{i_1 i_2}$. These events are independent if we restrict j_1, j_2 to even integers, as these squares are then disjoint. Moreover, each event has probability $1 - \exp(-4\lambda r^2) \sim 4cr^2/(2n)$, for large n . We thus look for a viable genotype with $\Delta(x)_i$ even for all i . This reduces the number of traits a by a factor of 2, and it follows from Section 5.2 that a viable genotype x with all $\Delta(x)_i$ even a.a.s. exists as soon as $4cr^2 < 4(\log(a/2) - o(1)) = (-4 \log r - \log 2 - o(1))$.

References

Achlioptas, D., Moore, C., 2004. Random k-SAT: two moments suffice to cross a sharp threshold. SIAM J. Comput. 17, 947–973.

- Achlioptas, D., Peres, Y., 2004. The threshold for random k -SAT is $2^k \log 2 - o(k)$. *J. Am. Math. Soc.* 17, 947–973.
- Achlioptas, D., Kirousis, L.M., Kranakis, E., Krizanc, D., 2001. Rigorous results for $(2 + p)$ -SAT. *Theoret. Comput. Sci.* 265, 109–129.
- Athreya, K., Ney, P., 1971. *Branching Processes*. Springer, Berlin (reprinted by Dover 2004).
- Barbour, A.D., Holst, L., Janson, S., 1992. *Poisson Approximation*. Oxford University Press.
- Berger, N., 2004. A lower bound for the chemical distance in sparse long-range percolation models. (<http://arxiv.org/abs/math/0409021>).
- Biroli, G., Monasson, R., Weigt, M., 2000. A variational description of the ground state structure in random satisfiability problems. *Eur. Phys. J. B-Condensed Matter* 14, 551–568.
- Biskup, M., 2004. On the scaling of the chemical distance in long-range percolation models. *Ann. Probab.* 32, 2938–2977.
- Bollobás, B., 2001. *Random Graphs*. Cambridge University Press, Cambridge.
- Bollobás, B., Kohayakawa, Y., Łuczak, T., 1992. The evolution of random subgraphs of the cube. *Random Struct. Algorithms* 3, 55–90.
- Bollobás, B., Kohayakawa, Y., Łuczak, T., 1994. On the evolution of random Boolean functions. In: *Extremal Problems for Finite Sets (Visegrád, 1991)*, Bolyai Society Mathematical Studies, vol. 3. János Bolyai Mathematical Society, Budapest, pp. 137–156.
- Boufkhad, Y., Dubois, O., 1999. Length of prime implicants and number of solutions of random CNF formulae. *Theoret. Comput. Sci.* 215, 1–30.
- Burch, C.L., Chao, L., 1999. Evolution by small steps and rugged landscapes in the RNA virus phi 6. *Genetics* 151, 921–927.
- Burch, C.L., Chao, L., 2004. Epistasis and its relationship to canalization in the RNA virus phi 6. *Genetics* 167, 559–567.
- Cabot, E.L., Davis, A.W., Johnson, N.A., Wu, C.-I., 1994. Genetics of reproductive isolation in the *Drosophila simulans* clade: complex epistasis underlying hybrid male sterility. *Genetics* 137, 175–189.
- Campos, P.R.A., Adami, C., Wilke, C.O., 2002. Optimal adaptive performance and delocalization in NK fitness landscapes. *Physica A* 304, 495–506.
- Choi, S.-S., Jung, K., Kim, J.H., 2005. Phase transition in a random NK landscape model. In: *Proceedings of the 2005 Conference on Genetic and Evolutionary Computation*. ACM Press, New York, pp. 1241–1248.
- Cook, S.A., 1971. The complexity of theorem proving procedures. In: *Proceedings of the Third Annual ACM Symposium on the Theory of Computing*. ACM, New York, pp. 151–158.
- Coyne, J., Orr, H.A., 2004. *Speciation*. Sinauer Associates, Inc., Sunderland, MA.
- de la Vega, W.F., 2001. Random 2-SAT: results and problems. *Theoret. Comput. Sci.* 265, 131–146.
- Derrida, B., Peliti, L., 1991. Evolution in flat landscapes. *Bull. Math. Biol.* 53, 255–282.
- Eigen, M., McCaskill, J., Schuster, P., 1989. The molecular quasispecies. *Adv. Chem. Phys.* 75, 149–263.
- Elena, S.F., Lenski, R.E., 2003. Evolution experiments with microorganisms: The dynamics and genetic bases of adaptation. *Nature Rev. Genetics* 4, 457–469.
- Fontana, W., Schuster, P., 1998. Continuity in evolution: on the nature of transitions. *Science* 280, 1451–1455.
- Fontana, W., Stadler, P.F., Tarazona, P., Weinberger, E.O., Schuster, P., 1993. RNA folding and combinatorial landscapes. *Phys. Rev. E* 47, 2083–2099.
- Friedgut, E., 1999. Necessary and sufficient conditions for sharp thresholds of graph properties, and the k -SAT problem. *J. Am. Math. Soc.* 12, 1017–1054.
- Gavrilets, S., 1997. Evolution and speciation on holey adaptive landscapes. *Trends in Ecology and Evolution* 12, 307–312.
- Gavrilets, S., 2003. Models of speciation: what have we learned in 40 years? *Evolution* 57, 2197–2215.
- Gavrilets, S., 2004. *Fitness Landscapes and the Origin of Species*. Princeton University Press, Princeton, NJ.
- Gavrilets, S., Gravner, J., 1997. Percolation on the fitness hypercube and the evolution of reproductive isolation. *J. Theoret. Biol.* 184, 51–64.
- Gavrilets, S., Hastings, A., 1996. Founder effect speciation: a theoretical reassessment. *Am. Nat.* 147, 466–491.
- Häggström, O., 2001. Coloring percolation clusters at random. *Stochastic Process. Appl.* 96, 213–242.
- Huynen, M.A., Stadler, P.F., Fontana, W., 1996. Smoothness within ruggedness: the role of neutrality in adaptation. *Proc. Nat. Acad. Sci. USA* 93, 397–401.
- Janson, S., Łuczak, T., Rucinski, A., 2000. *Random Graphs*. Wiley, New York.
- Kauffman, S.A., 1993. *The Origins of Order*. Oxford University Press, Oxford.
- Kauffman, S.A., Levin, S., 1987. Towards a general theory of adaptive walks on rugged landscapes. *J. Theoret. Biol.* 128, 11–45.
- Kimura, M., 1983. *The Neutral Theory of Molecular Evolution*. Cambridge University Press, New York.
- Korte, B., Vygen, J., 2005. *Combinatorial Optimization, Theory and Algorithms*, third ed. Springer, Berlin.
- Lenski, R.E., Ofria, C., Collier, T.C., Adami, C., 1999. Genome complexity, robustness and genetic interactions in digital organisms. *Nature* 400, 661–664.
- Lipman, D.J., Wilbur, W.J., 1991. Modeling neutral and selective evolution of protein folding. *Proc. R. Soc. London B* 245, 7–11.
- Lynch, M., 2007. *The Origins of Genome Architectures*. Sinauer Associates, Sunderland, MA.
- Martinez, M.A., Pezo, V., Marlière, P., Wain-Hobson, S., 1996. Exploring the functional robustness of an enzyme by *in vitro* evolution. *EMBO J.* 15, 1203–1210.
- Molloy, M., 2003. Models for random constraint satisfaction problems. *SIAM J. Comput.* 32, 935–949.
- Monasson, R., Zecchina, R., 1997. Statistical mechanics of the random K-satisfiability model. *Phys. Rev. E* 56, 1357–1370.
- Newman, M.E.J., Engelhardt, R., 1998a. Effects of selective neutrality on the evolution of molecular species. *Proc. R. Soc. London B* 265, 1333–1338.
- Newman, M.E.J., Engelhardt, R., 1998b. Effects of selective neutrality on the evolution of molecular species. *Proc. R. Soc. London B* 265, 1333–1338.
- Ohta, T., 1992. The nearly neutral theory of molecular evolution. *Annu. Rev. Ecol. Systematics* 23, 263–286.
- Ohta, T., 1998. Evolution by nearly-neutral mutations. *Genetica* 102/103, 83–90.
- Orr, H.A., 1995. The population genetics of speciation: the evolution of hybrid incompatibilities. *Genetics* 139, 1803–1813.
- Orr, H.A., 1997. Dobzhansky Bateson and the genetics of speciation. *Genetics* 144, 1331–1335.
- Orr, H.A., 2006a. The distribution of fitness effects among beneficial mutations in Fisher’s geometric model of adaptation. *J. Theoret. Biol.* 238, 279–285.
- Orr, H.A., 2006b. The population genetics of adaptation on correlated fitness landscapes: The block model. *Evolution* 60, 1113–1124.
- Orr, H.A., Orr, L.H., 1996. Waiting for speciation: the effect of population subdivision on the waiting time to speciation. *Evolution* 50, 1742–1749.
- Orr, H.A., Turelli, M., 2001. The evolution of postzygotic isolation: accumulating Dobzhansky-Muller incompatibilities. *Evolution* 55, 1085–1094.
- Penrose, M.D., 1996. Continuum percolation and Euclidean minimal spanning trees in high dimensions. *Ann. Appl. Probab.* 6, 528–544.
- Pigliucci, M., Kaplan, J., 2006. *Making Sense of Evolution: The Conceptual Foundations of Evolutionary Biology*. University of Chicago Press, Chicago.
- Reidys, C.M., 2006. Combinatorics of genotype-phenotype maps: an RNA case study. In: Percus, A., Istrate, G., Moore, C. (Eds.), *Computational Complexity and Statistical Physics*. Oxford University Press, Oxford, pp. 271–284.

- Reidys, C.M., Stadler, P.F., 2001. Neutrality in fitness landscapes. *Appl. Math. Comput.* 117, 321–350.
- Reidys, C.M., Stadler, P.F., 2002. Combinatorial landscapes. *SIAM Rev.* 44, 3–54.
- Reidys, C.M., Stadler, P.F., Schuster, P., 1997. Generic properties of combinatory maps: neutral networks of RNA secondary structures. *Bull. Math. Biol.* 59, 339–397.
- Reidys, C.M., Forst, C.V., Schuster, P., 2001. Replication and mutation on neutral networks. *Bull. Math. Biol.* 63, 57–94.
- Rost, B., 1997. Protein structures sustain evolutionary drift. *Folding & Design* 2, S19–S24.
- Schuster, P., 1995. How to search for RNA structures. theoretical concepts in evolutionary biotechnology. *J. Biotechnol.* 41, 239–257.
- Sedgewick, R., 1997. Algorithms in C, Parts 1–4: Fundamentals, Data Structures, Sorting, Searching. Addison-Wesley, Reading, MA.
- Skipper, R.A., 2004. The heuristic role of Sewall Wright's 1932 adaptive landscape diagram. *Phil. Sci.* 71, 1176–1188.
- Toman, E., 1979. The geometric structure of random boolean functions. *Problemy Kibernet.* (in Russian) 35, 111–132.
- Welch, J., Waxman, D., 2005. The *NK* model and population genetics. *J. Theoret. Biol.* 234, 329–340.
- Wilke, C.O., Wang, J.L., Ofria, C., Lenski, R.E., Adami, C., 2001. Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature* 412, 331–333.
- Wilkins, J., 2007. The dimensions, modes and definitions of species and speciation. *Biol. Phil.* 22, 247–266.
- Woods, R., Schneider, D., Winkworth, C.L., Riley, M.A., Lenski, R.E., 2006. Tests of parallel molecular evolution in a long-term experiment with *Escherichia coli*. *Proc. Natl. Acad. Sci. USA* 103, 9107–9112.
- Wright, S., 1932. The roles of mutation, inbreeding, crossbreeding and selection in evolution. in: Jones, D.F. (Ed.), *Proceedings of the Sixth International Congress on Genetics*, vol. 1, Austin, Texas, pp. 356–366.