Critical Properties of Complex Fitness Landscapes

Bjørn Østman[1], Arend Hintze[1], and Christoph Adami[1]

[1]Keck Graduate Institute, Claremont, CA 91711
bostman@kgi.edu

Abstract

Evolutionary adaptation is the process that increases the fit of a population to the fitness landscape it inhabits. As a consequence, evolutionary dynamics is shaped, constrained, and channeled, by that fitness landscape. Much work has been expended to understand the evolutionary dynamics of adapting populations, but much less is known about the structure of the landscapes. Here, we study the global and local structure of complex fitness landscapes of interacting loci that describe protein folds or sets of interacting genes forming pathways or modules. We find that in these landscapes, high peaks are more likely to be found near other high peaks, corroborating Kauffman's "Massif Central" hypothesis. We study the clusters of peaks as a function of the ruggedness of the landscape and find that this clustering allows peaks to form interconnected networks. These networks undergo a percolation phase transition as a function of minimum peak height, which indicates that evolutionary trajectories that take no more than two mutations to shift from peak to peak can span the entire genetic space. These networks have implications for evolution in rugged landscapes, allowing adaptation to proceed after a local fitness peak has been ascended.

Introduction

The structure of the fitness landscapes that populations find themselves in determines to a large extent how those populations will evolve. In introducing the concept of an adaptive fitness landscape, Sewall Wright (1932) sought to illustrate the idea that some combinations of characters will give rise to very high fitness (peaks) while some others do not (valleys), and to study the processes that allow a population to shift from peak to peak. Evolution in simple smooth landscapes (where each site or locus contributes independently to fitness) is trivial, because the ascent of a single fitness peak is largely deterministic (Tsimring et al., 1996; Kessler et al., 1997). At the other extreme lie "random" landscapes (Derrida and Peliti, 1991; Flyvbjerg and Lautrup, 1992), which are characterized by an absence of any fitness correlations between genotypes, and whose dynamics can likewise be solved using statistical approaches. In between these two extremes lie fitness landscapes that are neither smooth nor random, where mutations at different loci interact in complex patterns, giving rise to variably rugged and highly *epistatic* landscapes (Whitlock et al., 1995; Burch and Chao, 1999; Phillips et al., 2000; Beerenwinkel et al., 2007; Phillips, 2008). Experiments with bacteria and viruses (Elena and Lenski, 2003) have revealed that real fitness landscapes are of this nature: they are neither smooth nor random, and consist of a large number of fitness peaks.

Unfortunately, while experiments with bacteria and viruses have taught us a lot about evolutionary dynamics, they can only probe very limited regions of the fitness landscape, confined to the genotype space surrounding those of living organisms. In artificial landscapes we are not constrained by generation time or the specific genotypic space that organisms happen to occupy, but can place organisms anywhere in the fitness landscape, thus enabling us to examine the statistical properties of fitness landscapes.

If realistic fitness landscapes are neither smooth (a single peak) nor random (very many randomly placed peaks in the landscape), what is the structure of complex landscapes in "peak space"? Are most peaks confined to one region of genotype space, leaving other areas empty? Are peaks clustered or are they evenly distributed? One hypothesis about the structure of fitness landscapes was proposed by Kauffman (1993), who posited that peaks are not evenly distributed, but that high peaks are correlated in space, forming a *Massif Central*, and presented numerical evidence supporting this view. According to this observation, the best place to look for a high fitness peak is near another high fitness peak. A corollary to this hypothesis is that large basins with no peaks surrounds the central massif. If fitness peaks are indeed distributed in this manner, it would have profound implications for the *traversability* of the landscape, and for evolvability in general (Altenberg and Wagner, 1996).

Here we strive to study this question in much more detail, by analyzing all the peaks in a landscape in which the ruggedness can be tuned from smooth to random. In particular, we would like to know whether the highest peaks form clusters of connected walks that can *percolate*, i.e., form connected clusters that span the entire fitness landscape. Such clusters are very different from the neutral net-

works studied elsewhere (van Nimwegen et al., 1999; Wilke, 2001), and we briefly argue that peak networks may be more important for evolvability.

NK Landscape

Kauffman's NK model (Kauffman and Levin, 1987, see also Altenberg, 1997) has been used extensively to study evolution because it is a computationally tractable model of N binary interacting loci where the ruggedness of the landscape can be tuned by varying K, the number of loci that each locus interacts with. Typically N is of the order of 10-30, but larger sets can be studied if a complete enumeration of genotypes is not necessary. If $K = 0$, the smooth landscape limit is reached, because if loci do not interact, then there is a single peak in the landscape that can be reached by optimizing each locus independently. If $K = N - 1$, on the other hand, the model reproduces the random energy model of Derrida (Derrida and Peliti, 1991). The N loci are usually thought of as occupying sites on a circular genome, while the interactions occur between *adjacent* sites (see Fig. 1), but the identity of the interactors are immaterial and the results do not depend on their physical location on the genome. The example genome in Fig. 1 shows the interactions between loci in an $N = 20$ and $K = 2$ model, where the width and darkness of the lines reflects the strength of the epistatic interactions between sites for the global peak of that landscape.

While clearly the NK model should not be thought of as describing the genome of whole organisms, the model has been used extensively to study the evolution of a smaller set of sites, such as the residues in a protein (Macken and Perelson, 1989; Perelson and Macken, 1995; Hayashi et al., 2006; Carneiro and Hartl, 2010) or the set of interacting genes coding for a pathway or a module (Kauffman and Weinberger, 1989; Sole et al., 2003; Yukilevich et al., 2008; Østman et al., 2010).

In the original NK model, the fitness contribution of each locus is calculated as the arithmetic mean of the fitness contributions of each locus $w(x_i)$, which itself is a function of the value of the bit at that locus ('1' if the gene is expressed, '0' if it is silent) and the allele of the K genes it interacts with. This fitness landscape is constructed by obtaining uniformly distributed independent random numbers for all the possible combinations of the $K + 1$ sites (2^{K+1} numbers for each locus), so that the fitness contribution for any combinations of alleles can simply be found by looking up that value in the table. Here, we modify this model slightly, by replacing the customary arithmetic mean by the geometric one, so that the fitness of genotype $\vec{x} = (x_1, ..., x_N)$ is given by

$$W(\vec{x}) = \left(\prod_{i=1}^{N} w(x_i) \right)^{1/N} . \qquad (1)$$

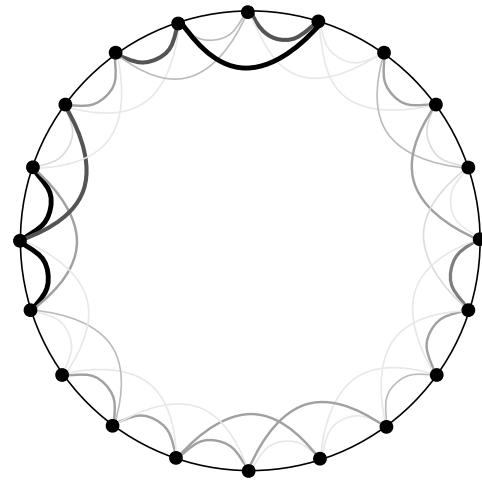This modification better captures the nature of real genetic

Figure 1: Genome and epistatic interactions between sites for the peak genotype of an $N = 20$ and $K = 2$ model. While all sites within a "radius" of two interact (light grey), the strength of interaction can be very different depending on the actual landscape that was formed. Here, the strength of epistatic interactions was calculated by performing all single-site and pairwise knockouts on the global peak genotype, and calculating the deviation of independence using a standard method (Bonhoeffer et al., 2004; Elena and Lenski, 1997; Østman et al., 2010).

interactions (see, e.g., St Onge et al., 2007), and it makes it possible to introduce lethal mutations by setting one or more numbers in the fitness lookup-table to zero. Taking the geometric mean skews the distribution of genotype fitness to the left, resulting in a mean of about 0.4, rather than the value of 0.5 when using the arithmetic mean (see Fig. 2). Of course the logarithm of $W(\vec{x})$ reduces to the usual arithmetic mean of the log-transformed fitnesses.

In the NK model we can easily compute the fitness of all genotypes as long as N and K are not too large, and we can also identify *fitness peaks* as those genotypes whose N one-mutation neighbors all have lower fitness. Increasing K creates landscapes that are increasingly rugged, containing more and higher peaks with deeper valleys in between. The waiting time to new mutations becomes a determining factor in how much the population can evolve before it risks becoming stuck on a peak of suboptimal fitness. Visualizing natural fitness landscapes is difficult since it requires probing genotype-space by measuring the fitness of organisms whose genomes are fully sequenced. Even worse, natural fitness landscapes are rarely static, making such an endeavor even more futile. In computational models all genotypes can sometimes be enumerated, and we can thus learn about the global properties of the fitness landscape. This exciting possibility is muted by the fact that we cannot easily visualize high-dimensional spaces, and we are forced to resorting to

statistical methods to probe the landscape.

How Peaks Cluster

In Fig. 2 we show the fitness distribution of all genotypes of an $N = 20, K = 4$ landscape (this distribution is virtually identical for different realizations of landscapes with the same N and K). Of those 2^{20} genotypes, less than 0.07% are peaks (this fraction depends on the particular realization of the landscape), and are also roughly normally distributed in fitness. Note that while the highest-fitness genotypes are very likely peaks, there are peaks whose fitness is significantly smaller, down to the mean fitness of genotypes in the landscape. The number of peaks scales approximately exponentially with N (when K is fixed), but only about linearly with K for K sufficiently large, and at fixed N (data not shown).
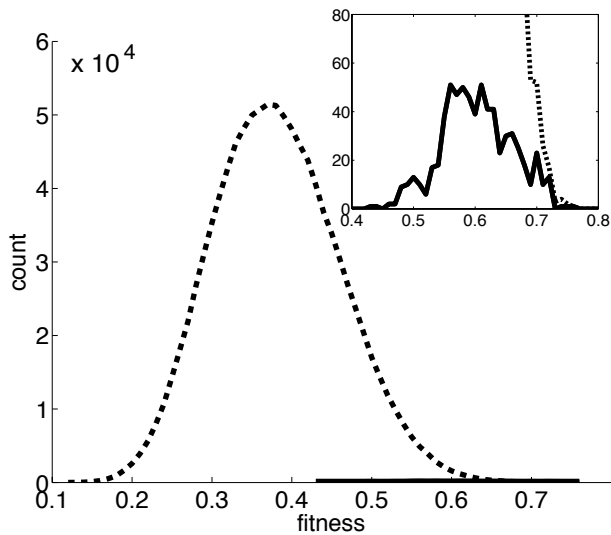

Figure 2: Fitness distribution of all $1,048,576$ genotypes (dashed line) in a typical landscape of $N = 20$ and $K = 4$. This landscape contains 679 peaks whose fitness distribution is shown as a solid black line. In the inset we have zoomed in on the peaks.

Pairwise distances

Because the "Massif Central" hypothesis says that the neighborhoods of high peaks are the best places to look for other high peaks, it is natural to also look at the pairwise distance of all peaks in a landscape. As we now know the genotypes of all the peaks in the landscape, we can ask whether peaks have a tendency to be located close to each other by studying the distribution of *Hamming distances* between peaks, which counts the number of differences in the binary representation of the sequences. In fact, this is how Kauffman validated his hypothesis: by plotting the fitness of peaks as

a function of the Hamming distance of all peaks to the highest peak he found (Kauffman (1993), page 61), for a landscape with $N = 96$ and $K = 2, 4,$ and 8. As it is not possible to enumerate $2^{96} \approx 8 \cdot 10^{28}$ genotypes, Kauffman found high peaks using random uphill walks. Here, we instead use $N = 20$, for which we can compute the fitness of all genotypes and thus locate all peaks. After computing the Hamming distance between all pairs of peaks, we can compare the distribution of these distances to a control distribution constructed with the same number of *random* genotypes, which are not expected to show any bias in the distribution of distances. (It is easy to see that the distribution of pairwise distances of random binary sequences of length $N = 20$ peaks at $d = 10$.)
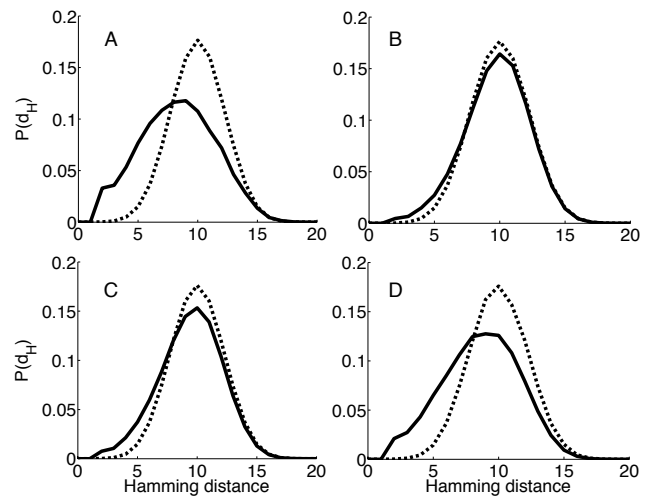

Figure 3: Distributions of pairwise Hamming distances between all peaks (solid) and between random "control" genotypes (dashed). The distributions shown are the averages of 50 different landscapes with genomes of length $N = 20$. (A) $K = 2$ landscapes containing an average of 98 peaks. (B) $K = 4$ landscapes containing an average of 720 peaks. (C) $K = 4$ landscapes including only an average of 363 peaks with a fitness above a threshold: $W \geq \Theta = 0.60$. (D) $K = 4$ landscapes including only an average of 95 peaks with a fitness above a threshold of $\Theta = 0.66$. As the samples include fewer and higher peaks, the pairwise distributions of $K = 4$ landscapes begin to resemble that of the $K = 2$ landscapes, suggesting that the highest peaks do cluster in genotype space, whereas the distribution of lower peaks is less biased.

We find that for $K = 2$, peaks are generally closer to each other than expected, indicating that peaks cluster in genotype space (see Fig. 3A). This alone does not tell us whether high peaks are more frequently associated with other high peaks (as opposed to peaks of lower fitness). Moreover, when examining $K = 4$ landscapes (that contain over seven

times as many peaks on average as for $K = 2$) we notice that the tendency for peaks to cluster close to each other is nearly gone, that is, the distribution closely resembles the random control (Fig. 3B). However, the bias reappears when we filter the peaks so that we only include those of high fitness (Figs. 3C and D), reaffirming the hypothesis that in complex epistatic landscapes, there is something special about being a high peak, genotypically speaking.
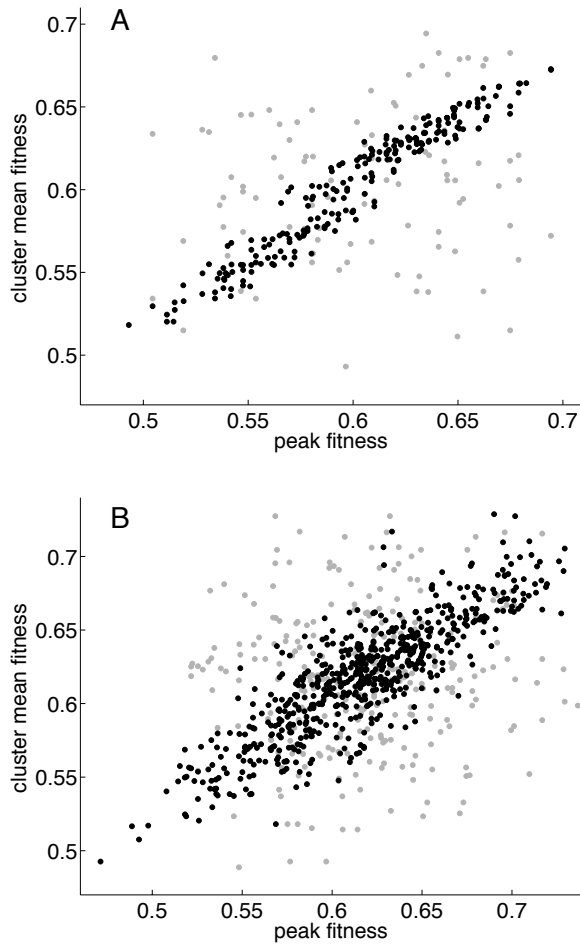

Figure 4: Mean fitness of peaks in circular clusters of radius $d = 2$ as a function of the fitness of the peak in the center of the cluster. (A) One landscape of $K = 2$ with 166 peaks (black dots). All landscapes show a strong correlation between cluster mean fitness and peak fitness, while the same analysis of assigning random genotypes to the peaks (but keeping the fitness) shows no such correlation (gray dots). The random data are from ten samplings. (B) One landscape of $K = 4$ with 679 peaks (black dots), and random genotypes (gray dots) obtained by sampling four times.

Peak neighborhood

If we want to know whether peaks with high fitness are likely to be found near other such peaks, we should study the mean fitness of peaks within a specified radius of that peak. These "circular" clusters contain all peaks within a Hamming distance d of a chosen peak (not counting the peak at the center). For the smallest possible distance between peaks $d = 2$, the size of a cluster is limited to 210 genotypes, but since peaks must be at least two mutations away from each other, there can be at most 190 peaks within a Hamming distance of two.

Fig. 4A depicts the mean fitness of adjacent peaks in circular clusters of radius $d = 2$ (black dots, for $K = 2$), showing a tight correlation between peak fitness and average adjacent peak fitness that indicates that the immediate neighborhood of high peaks is populated by other peaks of high fitness. On the contrary, when we randomize the location of the 166 peaks in genotype space without changing their height, this relationship vanishes (light gray dots in Fig. 4A). For $K = 2$ random peaks are far apart, resulting in only very few peaks within a distance $d = 2$ of each other. The $K = 4$ landscape has four times as many peaks as the $K = 2$ landscape, and the effect persists (Fig. 4B). The observed relation between mean fitness of these circular clusters and peak fitness persists even when the radius in increased to $d = 6$ (data not shown). We observe a similar correlation between mean cluster fitness and maximum peak height in network clusters (data not shown).
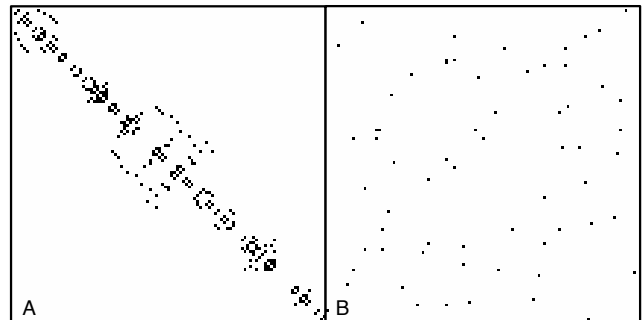
Adjacency matrices

Figure 5: Adjacency matrices showing clusters of peaks. (A) Single $K = 4$ landscape with peaks of Hamming distance $d = 2$ connected. The peaks are ordered according to which network cluster they belong to. This landscape consists of 109 peaks with fitness above $\Theta = 0.66$ that are grouped into nine clusters (not counting singletons). (B) Random $K = 4$ landscape with $d = 4$ and $\Theta = 0$, showing only the first 109 genotypes.

While circular clusters can tell us whether high peaks are surrounded by peaks that are higher than expected, they do

not allow us to examine certain critical properties of the landscape. To do this, we should think of peaks in the genetic landscape as nodes in a random graph, and study the size of clusters of peaks that are formed by connecting all those peaks that are within a distance d of each other. Connecting such *networks clusters* of peaks creates a *percolation problem* (see, e.g., Bollobas and Riordan (2006)). In statistical physics, systems where nodes are connected by edges that are placed with a fixed probability undergo a geometric phase transition as a function of the edge placement probability. One of the quantities studied in percolation theory is the size of the largest cluster, because this variable rises dramatically at the critical point so that it takes up most of the system once past the critical point. If the largest cluster takes up most of the nodes, the system is said to "percolate", which implies that the cluster spans the entire system (allowing you to walk across connected nodes from any part to any other in the system). We will study the percolation properties of the fitness landscape by using the peak height as the critical parameter. Clearly, if only the highest few peaks are considered the system is far from percolation, as these peaks are unlikely to be connected. But if the highest peaks are closer to each other than expected in a random control, then the peaks could percolate far earlier.

Let us begin by computing the Hamming distance between all pairs of peaks with fitness greater than Θ, and connect those peaks that are a distance of no more than d away from each other. In Fig. 5A, we show the *adjacency matrix* of clusters, which we obtained by placing a dot for every two peaks that are with a distance d (that is, immediately adjacent). Peaks are ordered in such a way that peaks that fall into the same cluster are placed next to each other. This procedure allows us to the visualize the structure of clustered peaks in the landscape. In contrast, if the same peaks are assigned random locations in the landscape, there is no apparent structure, and clusters of peaks are on average very small (Fig. 5B). For $K = 4$ and $d = 2$ very few peaks are connected in a random landscape, and because of this the adjacency matrix shown in Fig. 5B is for $d = 4$, and includes peaks of any height. Only the first 109 peaks are shown.

Percolation phase-transition

In Fig. 6 we show the average relative size of the largest network cluster as a function of the peak threshold Θ, defined as the ratio of the largest number of connected peaks with fitness above Θ to the total number of peaks in the landscape. The relative size of the largest connected component (also called the "giant cluster" in percolation theory) increases dramatically as the critical threshold is reached, much like the size of the giant component increases when the critical probability of edges is reached in percolation theory. But what is remarkable about this transition is that it only occurs because the high peaks in the landscape occur near other high peaks: if the peaks were
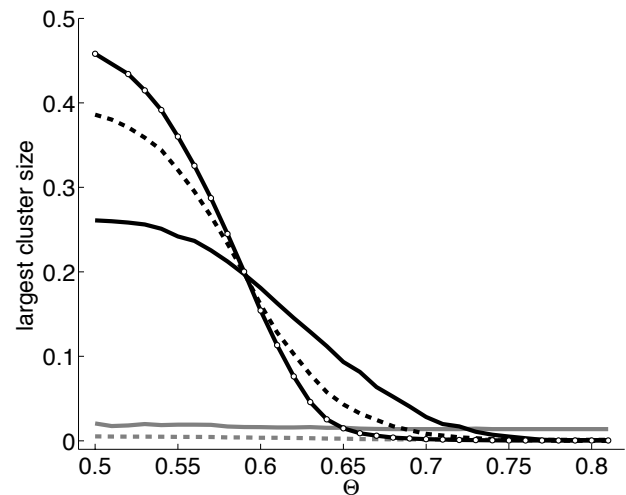

Figure 6: Size of the largest network cluster in the landscape averaged over 50 landscapes for each K as a function of fitness threshold, Θ. $K = 2$ (solid black line), $K = 4$ (dashed black line), and $K = 6$ (solid black line with white circles). The more rugged the landscapes are, the more abrupt the transition is from small network clusters to one cluster dominating the landscape. Random genotypes for $K = 2$ (solid gray line) and $K = 4$ (dashed gray line) show no increase in cluster size.

not clustered, the largest network cluster size would not increase when we lower Θ, as is the case when we reassign peaks to random genotypes (gray lines in Fig. 6).

When we include enough peaks, either by setting Θ low for $K = 4$ (or else for $K = 6$ or higher) we find that for $d = 2$ there are always *two* largest network clusters, while the third largest cluster contains significantly fewer peaks. Both large clusters percolate genotype space and the diameter of both graphs is 18, not 20 (in general, $N - 2$), while the shortest distance between the two clusters is always 3. This is peculiar to the way clusters are formed in this particular percolation problem. It is a rewarding exercise to determine the root cause of this peculiarity, which we leave to the interested reader. The transition seen in Fig. 6 suggests that in more rugged landscapes there are several clusters containing high peaks (high Θ), and that these high-peak clusters are connected by the peaks of lower fitness (lower Θ).

The percolation of genetic space by peaks with a sufficiently low height is reminiscent of the percolation of genetic space by arbitrary shapes in the RNA folding problem (Grüner et al., 1996), except that in that case structures with different genotypes form a neutral network that can be traversed by single point mutations. The giant cluster of peaks in the NK landscapes cannot be traversed like that: rather, it requires a minimum of two mutations to jump

from peak to peak, and because some of the peaks have inferior fitness, such mutations can only be tolerated for a finite amount of time–long enough to jump to the next highest peak. Thus, deleterious mutations are likely to be important to reach distant areas in genotype space, and the importance of these is slowly being realized (Lenski et al., 2003, 2006; Cowperthwaite et al., 2006; Østman et al., 2010).

Discussion

Using several methods we have shown that the rugged fitness landscapes that epistatic interactions create in the NK model consist of fitness peaks that are distributed in a manner that strongly affects evolution. High peaks are more likely to be found near other high peaks, rather than near lower peaks or far from peaks altogether. Similarly, lower peaks are predominantly located near each other in genotype space. Cluster analysis reveals that peaks tend to cluster (as compared to the same peaks placed randomly in genetic space) giving rise to large basins of attraction that are effectively devoid of peaks. This feature is especially prominent for moderately rugged landscapes ($K = 2$), while the addition of many more smaller peaks in more rugged landscapes ($K = 4$ or higher) makes this trend less significant. To the extent that we think that the NK landscape is an accurate model for real fitness landscapes of proteins and genetic pathways or modules, the discovery that these landscapes possess a remarkable structure that appears to be conducive to adaptation is highly informative about the process of evolution. Clustering of peaks makes a difference when the environment changes in a way that is unfavorable to the population, and forcing the population to adapt anew. If the landscape consists of evenly distributed peaks, then the risk of becoming stuck on a low fitness peak is high, and the population risks extinction. On the other hand, if peaks are unevenly distributed, then the ascent of one peak may not be where adaptation ends, making it possible to locate the global peak or another high fitness peak.

The more rugged a landscape is, the more peaks it contains, and the larger the space of genotypes that the largest network cluster spans. In smooth landscapes with only one or a few peaks, populations can evolve from genotypes of low fitness and move across genotype space toward high fitness. In rugged landscapes, the population always risks becoming stuck on a suboptimal peak. However, networks of closely connected peaks that percolate genotype space may still make it possible to traverse the fitness landscape jumping from peak to peak (given a sufficiently high mutation rate). If peaks are evenly distributed in genotype space, the chance to jump from peak to peak and thereby eventually locate the global peak is virtually nil. It is important, however, to remember that there are limits to the realism of the NK landscape as a model of realistic genetic or protein landscapes. For example, it is known that a significant percentage of substitutions in proteins or mutations in genetic pathways are neutral, while the NK landscape has virtually no neutrality (even though most mutations do not change the fitness significantly). Neutrality plays an important role to enhance traversability, and will facilitate the transition between peaks so that deleterious mutations are not essential for the shift from peak to peak. However, one could maintain that deleterious mutations are more promising for adaptation than neutral mutations are, because they may be what separate important phenotypes (Lenski et al., 2006).

The observation that peaks form clustered networks, and that these networks percolate, implies that the risk of becoming stuck on a suboptimal peak is significantly mitigated, because all it takes is the two right mutations to locate a new peak. Thus, it appears that evolvability comes for free in complex rugged landscapes of interacting loci. We should note, however, that the reason *why* peaks cluster in landscapes with epistatic interactions is not immediately apparent, and is a subject of ongoing investigations.

Acknowledgements

References

Altenberg, L. (1997). Nk fitness landscapes. In Back, T., Fogel, D., and Michalewitcz, Z., editors, *The Handbook of Evolutionary Computation*, New York. Oxford University Press.

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

Beerenwinkel, N., Pachter, L., and Sturmfels, B. (2007). Epistasis and shapes of fitness landscapes. *Statistica Sinica*, 17:1317–1342.

Bollobas, B. and Riordan, O. (2006). *Percolation*. Cambridge University Press, Cambridge, UK.

Bonhoeffer, S., Chappey, C., Parkin, N. T., Whitcomb, J. M., and Petropoulos, C. J. (2004). Evidence for positive epistasis in HIV-1. *Science*, 306:1547–1550.

Burch, C. L. and Chao, L. (1999). Evolution by small steps and rugged landscapes in the RNA virus $\phi 6$. *Genetics*, 151(3):921–927.

Carneiro, M. and Hartl, D. L. (2010). Adaptive landscapes and protein evolution. *Proc Natl Acad Sci U S A*, 107 Suppl 1:1747–1751.

Cowperthwaite, M. C., Bull, J. J., and Meyers, L. A. (2006). From bad to good: Fitness reversals and the ascent of deleterious mutations. *PLoS Computational Biology*, 2:1292–1300.

Derrida, B. and Peliti, L. (1991). Evolution in a flat landscape. *Bulletin of Mathematical Biology*, 53:355–382.

Elena, S. F. and Lenski, R. (1997). Test of synergistic interactions among deleterious mutations in bacteria. *Nature*, 390:395–397.

Elena, S. F. and Lenski, R. E. (2003). Evolution experiments with microorganisms: the dynamics and genetic bases of adaptation. *Nat Rev Genet*, 4:457–469.

Flyvbjerg, H. and Lautrup, B. (1992). Evolution in a rugged fitness landscape. *Physical Review*, A 46:6714–6723.

Grüner, W., Giegerich, R., Strothmann, D., Reidys, C., Weber, J., Stadler, I. L. H. P. F., and Schuster, P. (1996). Analysis of RNA sequence structure maps by exhaustive enumeration II. structures of neutral networks and shape space covering. *Monatshefte für Chemie*, 127:375–389.

Hayashi, Y., Aita, T., Toyota, H., Husimi, Y., Urabe, I., and Yomo, T. (2006). Experimental rugged fitness landscape in protein sequence space. *PLoS One*, 1:e96.

Kauffman, S. and Levin, S. (1987). Towards a general theory of adaptive walks on rugged landscapes. *J Theor Biol*, 128(1):11–45.

Kauffman, S. A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press US.

Kauffman, S. A. and Weinberger, E. D. (1989). The NK model of rugged fitness landscapes and its application to maturation of the immune response. *J Theor Biol*, 141(2):211–245.

Kessler, D., Levine, H., Ridgway, D., and Tsimring, L. (1997). Evolution on a smooth landscape. *Journal of Statistical Physics*, 87:519–544.

Lenski, R. E., Barrick, J. E., and Ofria, C. (2006). Balancing robustness and evolvability. *PLoS Biology*, 4:e428.

Lenski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423(6936):139–144.

Macken, C. A. and Perelson, A. S. (1989). Protein evolution on rugged landscapes. *Proc Natl Acad Sci U S A*, 86(16):6191–6195.

Østman, B., Hintze, A., and Adami, C. (2010). Impact of epistasis and pleiotropy on evolutionary adaptation. arxiv.org: arXiv:0909.3506v2.

Perelson, A. S. and Macken, C. A. (1995). Protein evolution on partially correlated landscapes. *Proc Natl Acad Sci U S A*, 92(21):9657–9661.

Phillips, P., Otto, S., and Whitlock, M. (2000). Beyond the average, the evolutionary importance of gene interactions and variability of epistatic effects. In Wolf, J., Brodie III, E., and Wade, M., editors, *Epistasis and the Evolutionary Process*, pages 20–38. Oxford University Press.

Phillips, P. C. (2008). Epistasis - the essential role of gene interactions in the structure and evolution of genetic systems. *Nature Reviews Genetics*, 9:855–867.

Sole, R. V., Fernandez, P., and Kauffman, S. A. (2003). Adaptive walks in a gene network model of morphogenesis: insights into the Cambrian explosion. *Int J Dev Biol*, 47(7-8):685–693.

St Onge, R. P., Mani, R., Oh, J., Proctor, M., Fung, E., Davis, R. W., Nislow, C., Roth, F. P., and Giaever, G. (2007). Systematic pathway analysis using high-resolution fitness profiling of combinatorial gene deletions. *Nat Genet*, 39(2):199–206.

Tsimring, L., Levine, H., and Kessler, D. (1996). RNA virus evolution via a fitness-space model. *Phys Rev Lett*, 76(23):4440–4443.

van Nimwegen, E., Crutchfield, J. P., and Huynen, M. (1999). Neutral evolution of mutational robustness. *Proc. Natl. Acad. Sci. USA*, 96:9716–9720.

Whitlock, M. C., Phillips, P. C., Moore, F. B.-G., and Tonsor, S. J. (1995). Multiple fitness peaks and epistasis. *Annu. Rev. Ecol. Syst.*, 26:601–29.

Wilke, C. O. (2001). Adaptive evolution on neutral networks. *Bulletin of Mathematical Biology*, 63:715–730.

Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In *Proceedings of the Sixth International Congress on Genetics*, volume 1, pages 355–366.

Yukilevich, R., Lachance, J., Aoki, F., and True, J. R. (2008). Long-term adaptation of epistatic genetic networks. *Evolution*, 62(9):2215–2235.