The transformability of genotype-phenotype landscapes

Malvika Srivastava^{1,2} and Joshua L. Payne^{1,2,*}

¹Institute of Integrative Biology, ETH Zurich, Zurich, Switzerland

²Swiss Institute of Bioinformatics

*joshua.payne@env.ethz.ch

ABSTRACT

The mapping from genotype to phenotype to fitness typically involves multiple nonlinearities that can transform the individual and combined effects of mutations. For example, mutations may contribute additively to a phenotype, but their effects on fitness may combine non-additively because selection favors a low or intermediate value of that phenotype. Here, we study transformability under a non-linear phenotype-to-fitness map. We do so by comparing the topographical properties of genotype-phenotype landscapes, constructed using a diversity of theoretical models and empirical data on transcription factor-DNA interactions, to the topographical properties of fitness landscapes when selection favors a low or intermediate phenotypic value. Using the theoretical models, we prove a number of fundamental results. For example, selection for low or intermediate phenotypic values cannot transform simple sign epistasis into reciprocal sign epistasis, yet it transforms reciprocal sign epistasis into simple sign epistasis and no sign epistasis with equal probability. More broadly, we show that such selection tends to create fitness landscapes that are more rugged than the underlying genotype-phenotype landscape, but surprisingly this increased ruggedness typically does not frustrate adaptive evolution, because the local adaptive peaks in the fitness landscape tend to be nearly as tall as the global peak. Many of these results carry forward to the empirical genotype-phenotype landscapes, which may help to explain why low- and intermediate-affinity transcription factor-DNA interactions are so prevalent in eukaryotic gene regulation.

Introduction

Characterizing the relationship between genotype and phenotype is key to our understanding of evolution^{1,2}. For quantitative phenotypes, such as the expression level of a gene or the enzymatic activity of a protein, this relationship can be formalized as a genotype-phenotype landscape³. In such a landscape, genotypes represent coordinates in an abstract genotype space and their phenotype defines the elevation of each coordinate in this space⁴. The topographical properties of genotype-phenotype

landscapes, such as their ruggedness, are influenced by epistasis⁵ — non-additive interactions between mutations in their contribution to phenotype. These topographical properties have important evolutionary consequences, because they determine how mutation brings forth the phenotypic variation upon which selection acts^{6,7}.

Technological advances are facilitating the construction and analysis of empirical genotype-phenotype landscapes at ever-increasing resolution, scale, and scope⁸. Example phenotypes include the enzymatic activity⁹, binding affinity¹⁰, allosteric profile¹¹, and fluorescence intensity of proteins¹², as well as exon inclusion levels³, the expression levels of genes driven by regulatory elements¹³, the expression patterns of gene regulatory circuits^{14–16}, and flux through metabolic pathways¹⁷. In analyzing the topographical properties of these landscapes and their evolutionary consequences, an assumption is sometimes made that phenotype is a proxy for fitness^{12,18–24}, thus rendering genotype-phenotype landscapes equivalent to fitness landscapes²⁵. While this assumption may be justified under certain conditions, such as in directed protein evolution experiments²⁶, it is often the case that the relationship between phenotype and fitness is not so straightforward. For example, fitness may depend upon more phenotypes than those being assayed²⁷ or the relationship between phenotype and fitness may be inherently nonlinear, for example reflecting a tradeoff between the costs and benefits associated with a phenotype²⁸. In the latter case, selection may favor a low or intermediate phenotypic value^{29,30}; e.g., an intermediate gene expression level,^{31–33} enzyme efficiency³⁴ or protein production rate or activity³⁵. Such non-linearities are a cause of epistasis^{16,36–38}, and they can transform the effects of mutations as they map onto phenotype and fitness³⁷, thus rendering the topographical properties of a fitness landscape qualitatively different from those of its underlying genotype-phenotype landscape (Fig. 1A).

Selection for low or intermediate phenotypic values is especially relevant to transcription factor-DNA interactions³⁹. Transcription factors are sequence-specific DNA binding proteins that help regulate gene expression. They do so by binding DNA sequences (transcription factor binding sites) in regulatory regions such as promoters and enhancers to recruit or block the recruitment of RNA polymerase⁴⁰. The regulatory effect of such a binding event depends in part on the affinity with which the DNA sequence is bound by the transcription factor^{41,42}. As such, binding affinity is an important molecular phenotype of transcription factor binding sites, upon which selection acts^{43,44}. While it is commonly assumed that selection increases binding affinity^{21,43–47}, several lines of evidence suggest that selection for low or intermediate binding affinity also influences the evolution of transcription factor binding sites. For example, paralogous transcription factors often bind the same DNA sequences with high affinity, but different DNA sequences with low affinity^{48,49}. If an optimal gene expression pattern requires binding by just one of several transcription factor paralogs, such specificity can be achieved using low-affinity transcription factor binding sites, resulting in selection for low binding affinity⁵⁰. Additional documented cases in which low-affinity binding sites play important regulatory roles include negative auto-regulation by high-copy number transcription factors in *Escherichia coli*⁵¹, where high-affinity binding sites cause suboptimal noise suppression, and developmental patterning in *Ciona intestinalis* embryos^{52,53}, where high-affinity binding sites cause deleterious ectopic gene expression patterns. Moreover, low-affinity binding sites are commonly observed in the regulatory portfolios of a diversity of organisms, including bacteria⁵¹, yeast⁵⁴, fly^{49,55,56}, sea stars and sea urchins⁵⁷, as well as humans⁵⁸.

Figure 1. Transformability. (A) Schematic illustration of how selection for an intermediate phenotypic value w_{opt} transforms a genotype-phenotype landscape into a fitness landscape. (B) An additive three-locus, biallelic genotype-phenotype landscape with a single peak (gray filled circle; $p_{\text{gp}} = 1$). One pairwise interaction is highlighted in blue. It exhibits no magnitude epistasis ($\epsilon_{\text{gp}} = 0$) or sign epistasis. (C) The Gaussian phenotype-fitness map (Eq. 1) is shown for three values of w_{opt} (dashed line, $w_{\text{opt}} = 0$; dotted line, $w_{\text{opt}} = 0.5$; solid line, $w_{\text{opt}} = 1$), with three values of σ shown for $w_{\text{opt}} = 0.5$. (D) Applying the Gaussian phenotype-fitness map with $w_{\text{opt}} = 0$ to the genotype-phenotype landscape results in a single-peaked fitness landscape (gray filled circle; $p_{\text{f}} = 1$). The same pairwise interaction from (B) is highlighted in blue. It exhibits positive epistasis ($\epsilon_{\text{gp}} = 0.266$), but no sign epistasis. (E) Applying the Gaussian phenotype-fitness map with $w_{\text{opt}} = 0.5$ to the genotype-phenotype landscape results in a multi-peaked fitness landscape (gray filled circles; $p_{\text{f}} = 2$). The same pairwise interaction from (B) is highlighted in red. It exhibits negative epistasis ($\epsilon_{\text{gp}} = -1.01$), as well as reciprocal sign epistasis. (F) Applying the Gaussian phenotype-fitness map with $w_{\text{opt}} = 1$ to the genotype-phenotype landscape results in a single-peaked fitness landscape (gray filled circle; $p_{\text{f}} = 1$). The same pairwise interaction from (A) is highlighted in blue. It exhibits positive epistasis ($\epsilon_{\text{gp}} = 0.53$), but no sign epistasis. In panels B,D-F, arrows point from genotypes with lower phenotypic or fitness values to genotypes with higher phenotypic or fitness values. The no sign epistasis motif is highlighted in blue and the reciprocal sign epistasis motif in red. Note the symmetry of landscapes in panels D and F.

What is the relationship between the topography of a genotype-phenotype landscape and the topography of a fitness landscape when selection favors a low or intermediate phenotypic value? How does this relationship depend on the ruggedness of the genotype-phenotype landscape? These are important questions, because the topography of a fitness landscape has implications for several evolutionary phenomena, including the evolution of genetic diversity⁵⁹, reproductive isolation⁶⁰, and sex⁶¹, as well as the predictability of the evolutionary process itself⁶². Despite decades of research on fitness landscapes, these questions have not been addressed even in the context of classical theoretical models, such as Mt. Fuji⁶³, House-of-Cards⁶⁴, or NK landscapes⁶⁵. They have also not been addressed in the context of biophysical models of genotype-phenotype landscapes or empirical genotype-phenotype landscapes. Here, we fill this knowledge gap by defining local and global measures of landscape transformability that describe the topographical differences between a genotype-phenotype landscape and the

corresponding fitness landscape when selection favors a low or intermediate phenotypic value. We use these measures to study landscape transformability in the context of the aforementioned theoretical models^{63–65}, a biophysical model of transcription factor-DNA interactions⁶⁶, and 1,137 empirical genotype-phenotype landscapes of transcription factor-DNA interactions, wherein genotypes are transcription factor binding sites and the phenotype is a measure of relative binding affinity²¹. We study transcription factor-DNA interactions because there is strong biological motivation for studying selection for low or intermediate binding affinity, as discussed above, and because a large number of empirical genotype-phenotype landscapes of transcription factor-DNA interactions are publicly available²¹, thus facilitating the statistical analysis of their transformability.

Results

We first present our measures of landscape transformability and the phenotype-to-fitness map used to study transformability under selection for a low or intermediate phenotypic value. We use these measures and this map to study the transformability of randomly generated genotype-phenotype landscapes, specifically two-locus biallelic landscapes and multi-locus biallelic landscapes. For the latter, we use NK landscapes⁶⁵, which include the corner cases of Mt. Fuji⁶³ ($K = 0$) and House-Of-Cards⁶⁴ ($K = N - 1$) landscapes. We then apply the principles learned from these model genotype-phenotype landscapes to genotype-phenotype landscapes of transcription factor-DNA interactions, first in the context of the mismatch model⁶⁶ and then in the context of empirical measurements of transcription factor-DNA interactions²¹. Finally, we study the evolutionary consequences of landscape transformability.

Landscape transformability

Our goal is to study landscape transformability, which we define in terms of the topographical differences between a fitness landscape and its underlying genotype-phenotype landscape. Whereas we study landscape transformability under selection for low or intermediate phenotypic values, the measures of transformability we use can be applied to any pairing of landscapes, so long as they are defined over the same set of genotypes. For example, they can be used to compare genotype-phenotype landscapes across environmental conditions^{67,68} or across homologous proteins^{69,70}, as well to compare fitness landscapes in different environmental conditions⁷¹.

We quantify landscape transformability at both a local and a global scale. At a local scale, we quantify differences in pairwise epistasis amongst loci in the genotype-phenotype landscape relative to the same loci in the fitness landscape (Methods). We do so by classifying the type of magnitude epistasis (ϵ) as additive (i.e., no epistasis; $\epsilon = 0$), positive ($\epsilon > 0$), or negative ($\epsilon < 0$). For a pair of loci, we then compare the type of epistasis in the genotype-phenotype landscape (ϵ_{gp}) to the type of epistasis in the fitness landscape (ϵ_f), and report whether this is the same in the two landscapes. We do the same for an additional classification of epistasis based on the absence or presence of sign epistasis. This results in three categories – no sign epistasis, simple sign epistasis, or reciprocal sign epistasis⁷². At a global scale, we quantify transformability as the absolute difference in the number of peaks in the fitness landscape (p_f) relative to the genotype-phenotype landscape (p_{gp}). Taken together, these local

and global measures allow us to determine the extent to which selection for a low or intermediate phenotypic value increases or decreases the ruggedness of the fitness landscape, relative to the genotype-phenotype landscape.

Phenotype-to-fitness map

We use a Gaussian phenotype-to-fitness map $F(w)$ centred around an optimal phenotype w_{opt} ,

$$F(w) = \exp \left[- \left(\frac{w - w_{\text{opt}}}{\sigma} \right)^2 \right]. \quad (1)$$

The parameter σ determines the strength of selection, because it controls how rapidly fitness decreases as the phenotype w deviates from the optimal phenotype w_{opt} . Increasing σ decreases the strength of selection, because it broadens the fitness map around w_{opt} and thus decreases the fitness differences between similar phenotypes. However, the rank ordering of fitness values only depends on w_{opt} and is independent of σ . Similar maps are commonly used in evolutionary modeling frameworks, such as Fisher's Geometric model^{73,74} and models of speciation⁷⁵, as well as in biophysical models of intermolecular interactions²⁹, including transcription factor-DNA interactions⁴⁶.

Fig. 1B-F illustrates the application of this phenotype-to-fitness map to a simple three-locus, biallelic genotype-phenotype landscape. This landscape is purely additive, and as a consequence, it exhibits no epistasis and has only one peak (Fig. 1B). Applying the phenotype-to-fitness map (Fig. 1C) to this landscape can change the amount and type of epistasis, as well as the location and number of peaks in the resulting fitness landscape, relative to the genotype-phenotype landscape (Fig. 1D-F). It can therefore transform the genotype-phenotype landscape.

Two-locus biallelic genotype-phenotype landscapes

We first study transformability using the simplest form of genotype-phenotype landscape that is capable of exhibiting epistasis: a two-locus biallelic landscape. We represent genotypes as binary strings of length $L = 2$ and randomly assign a phenotype w_i to each genotype i , which we draw from a uniform distribution between (0,1). We then apply the Gaussian phenotype-fitness map (Eq. 1) to generate the corresponding fitness landscape. We repeat this process 10,000 times for $w_{\text{opt}} \in [0, 1]$ (in increments of 0.01), and report the probability that the type of epistasis in the genotype-phenotype landscape is the same as in the fitness landscape. We first differentiate between no magnitude epistasis, positive epistasis, and negative epistasis, and then between no sign epistasis, simple sign epistasis, and reciprocal sign epistasis.

To determine whether the type of magnitude epistasis is the same in the genotype-phenotype landscape and fitness landscape, we calculate the product $\epsilon_{\text{gp}} \cdot \epsilon_{\text{f}}$, which will be positive when the two landscapes have the same type of epistasis. We use this product, rather than a discrete categorization, to ensure analytical tractability. We obtain an analytical expression for $\epsilon_{\text{gp}} \cdot \epsilon_{\text{f}}$ by assuming σ to be large (Supplementary Material, Derivation 1). It is:

$$\epsilon_{\text{f}} \cdot \epsilon_{\text{gp}} = \frac{1}{\sigma^2} \left[2\epsilon_m \cdot \epsilon_{\text{gp}} + (2w_{\text{opt}} - \sum_i w_i) \cdot \epsilon_{\text{gp}}^2 \right], \quad (2)$$

Figure 2. Local transformability: magnitude epistasis. (A) The probability of retaining the type of epistasis, shown in relation to the optimal phenotypic value w_{opt} . The black line shows the theoretical prediction and the dots show the results from randomly generated genotype-phenotype landscapes for different values of σ . The theoretical approximation agrees well with the results from randomly generated genotype-phenotype landscapes for large σ (i.e. $\sigma \geq 1$). (B) The probability of observing zero magnitude epistasis in the fitness landscape, shown in relation to w_{opt} , for different values of σ , which were selected to show the range of variation in the probability of observing zero epistasis.

where w_i is the phenotypic value of genotype i with $i \in \{0, 1\}^2$ and $\epsilon_m = w_{00}w_{11} - w_{01}w_{10}$, which is also known as multiplicative epistasis⁷⁶.

The probability that the type of epistasis is the same in the genotype-phenotype and fitness landscapes, $P(\epsilon_{gp} \cdot \epsilon_f > 0)$ can be computed using Monte-Carlo methods and yields results in good agreement with the randomly generated genotype-phenotype landscapes when σ is large. This is shown in Fig. 2, where three trends are immediately apparent. First, for large σ , the probability that the type of epistasis is the same in the genotype-phenotype landscape and the fitness landscape increases as the optimal phenotype w_{opt} increases, in agreement with the intuition that selection for large phenotypic values leaves the genotype-phenotype landscape mostly unchanged, except for the nonlinear rescaling introduced by the phenotype-fitness map. Second, even when $w_{opt} = 1$, the probability that the type of epistasis is the same in the genotype-phenotype landscape and the fitness landscape is less than one. The reason is the nonlinear rescaling introduced by the phenotype-to-fitness map does not guarantee conservation of the type of magnitude epistasis, even though it does preserve the rank ordering of fitness values. Results obtained with the randomly generated genotype-phenotype landscapes show this effect becomes even more pronounced as σ decreases. This means that as the strength of selection for w_{opt} increases, so does the likelihood of landscape transformation. Finally, as σ decreases, the probability of retaining the type of epistasis from the genotype-phenotype landscape in the fitness landscape is not maximized at $w_{opt} = 1$, but rather at a smaller w_{opt} (e.g., at $w_{opt} \approx 0.8$ when $\sigma = 0.01$, Fig. 2A). The reason is as σ decreases, the fitness function becomes extremely narrow and more phenotypic values are mapped to zero fitness (within computer precision), resulting in cases where the genotype-phenotype landscape exhibits epistasis (i.e., $\epsilon_{gp} \neq 0$), but the fitness landscape does not (i.e., $\epsilon_f = 0$), because all fitness values are zero. Fig. 2B shows this is more likely to occur when σ is small and as w_{opt} approaches its extreme values of 0 or 1. In these randomly generated landscapes, the probability of obtaining negative epistasis is the same as the probability of obtaining positive epistasis, and the probability of conserving the

type of epistasis is independent of the type of epistasis in the genotype-phenotype landscape. In sum, these results show that selection for low or intermediate phenotypic values can transform the genotype-phenotype landscape, such that the resulting fitness landscape exhibits a different type of epistasis, and this effect is most pronounced when selection is strong and the optimal phenotypic value is low.

Next, we categorized landscapes as exhibiting no sign epistasis, simple sign epistasis, or reciprocal sign epistasis. These three motifs are shown in the center panel of Fig. 3, where arrows point from genotypes with a lower phenotypic or fitness value to genotypes with a higher phenotypic or fitness value. Because the presence of sign epistasis in the genotype-phenotype landscape only depends upon the partial ordering of the phenotypic values, we expect to retain the motif from the genotype-phenotype landscape in the fitness landscape as $w_{\text{opt}} \rightarrow 1$. We also expect to retain the motif as $w_{\text{opt}} \rightarrow 0$, because selecting for $w_{\text{opt}} = 0$ simply flips all the arrows from the genotype-phenotype landscape in the fitness landscape, which does not change the categorization of the motif. Thus, we expect the probability of retaining the motif from the genotype-phenotype landscape in the fitness landscape to be "U" shaped and symmetric about $w_{\text{opt}} = 0.5$.

Figure 3. Local transformability: sign epistasis. The probability of retaining or changing the type of epistasis in the genotype-phenotype landscape, relative to the fitness landscape, shown in relation to the optimal phenotypic value w_{opt} . Data are grouped based on whether the genotype-phenotype landscapes exhibits (A) no sign epistasis (blue), (B) simple sign epistasis (brown), or (C) reciprocal sign epistasis (red). The colours of the lines represent the type of epistasis in the resulting fitness landscape. These results are independent of σ , because they only depend on the rank ordering of fitness values. Notice the "U" shape of the probability of retaining the type of epistasis in each panel.

Fig. 3 shows the probability of retaining or changing the motif from the genotype-phenotype landscape in the fitness landscape for 10,000 randomly generated two-locus biallelic landscapes, grouped according to the motif in the genotype-

phenotype landscape. These results confirm the expected “U” shape of the probability of retaining the type of epistasis from the genotype-phenotype landscape in the fitness landscape as w_{opt} is varied from 0 to 1. When the genotype-phenotype landscape has the no sign epistasis motif, selection for an intermediate phenotypic value can transform the landscape into the simple sign epistasis motif or the reciprocal sign epistasis motif (Fig. 3A). Similar results have recently been uncovered for Fisher’s Geometric model⁷⁷. Next, when the genotype-phenotype landscape has the simple sign epistasis motif, selection for an intermediate phenotypic value can transform the landscape into a no sign epistasis motif, but not into a reciprocal sign epistasis motif (Fig. 3B; Supplementary Material, Proof 1). Finally, when the genotype-phenotype landscape has the reciprocal sign epistasis motif, selection for an intermediate phenotypic value transforms the landscape into the no sign epistasis motif or the simple sign epistasis motif with equal probability (Fig. 3C; Supplementary Material, Proof 2). Moreover, the probability of retaining the motif from the genotype-phenotype landscape in the fitness landscape is always higher than the probability of changing it when the genotype-phenotype landscape has the simple sign epistasis motif (Fig. 3B) or the reciprocal sign epistasis motif (Fig. 3C). This is not always true when the genotype-phenotype landscape has the no sign epistasis motif (Fig. 3A), because at intermediate w_{opt} the landscape is most likely to transform into the reciprocal sign epistasis motif. Taken together, these results show that selection for intermediate phenotypic values can transform genotype-phenotype landscapes with no sign epistasis into fitness landscapes with sign epistasis and vice versa.

The inferences about pairwise interactions can be carried forward to multi-locus biallelic landscapes because their genotype spaces, which are L -dimensional hypercubes, are composed of two-dimensional squares. Due to the adjacency of squares, in the three-locus case, the motifs of four out of the six squares are sufficient to determine the motifs of the rest of the squares, and for any L , the motifs of only $2^{L-2} \cdot (L-1)$ squares are necessary to determine the motifs of all of the remaining squares. This is only a fraction $2/L$ of all the squares in the hypercube (because the total number of faces in an L -dimensional hypercube is $2^{L-2} \cdot \binom{L}{2}$), which is clearly minuscule for large L . However, pairwise interactions are not sufficient to predict peak patterns and other higher dimensional features of multi-locus biallelic landscapes, which may result from higher-order interactions⁷⁸. We study these in the next section.

Multi-locus biallelic genotype-phenotype landscapes

We use the NK model⁶⁵ to study the transformability of multi-locus biallelic genotype-phenotype landscapes (Methods). In this model, each locus in a genotype of length L epistatically interacts with K other loci (whereas N is typically used to denote the number of loci in this model, we use L for consistency with the rest of our text). As corner cases, this model includes Mt. Fuji landscapes⁶³ when $K = 0$ and House-of-Cards landscapes⁶⁴ when $K = L - 1$. For each combination of L and K , we use this model to randomly generate a genotype-phenotype landscape. We then apply the Gaussian phenotype-fitness map (Eq. 1) to generate a fitness landscape. We repeat this process 10,000 times for $w_{\text{opt}} \in [0, 1]$ (in increments of 0.01), and report the average of the absolute change in the number of peaks, i.e., $\langle |p_f - p_{\text{gp}}| \rangle$, where p_f is the number of peaks in the fitness landscape and p_{gp} is the number of peaks in the genotype-phenotype landscape.

We begin with Mt. Fuji genotype-phenotype landscapes. Because these are single-peaked, selection for a low or intermediate phenotypic value can only maintain or increase the number of peaks from the genotype-phenotype landscape in the fitness landscape. Fig. 4A shows this change in the number of peaks, and Fig. 4B shows the probability that the number of peaks changes, in relation to w_{opt} for landscapes with $L = 2$ to $L = 8$ loci. These trends are symmetric about $w_{\text{opt}} = 0.5$, because Mt. Fuji landscapes are additive, so selecting for $w_{\text{opt}} = 0$ is equivalent to selecting for $w_{\text{opt}} = 1$ with regard to the change in the number of peaks. The reason is that selecting for $w_{\text{opt}} = 0$ flips all of the arrows in the fitness landscape, relative to the genotype-phenotype landscape, which changes the location of the peak, but does not change the number of peaks. This is illustrated in Fig. 1B,D. In Fig. 4A, the plot of $\langle |p_f - p_{\text{gp}}| \rangle$ with respect to w_{opt} exhibits $L - 1$ maxima, each separated from one another by a distance $1/L$ in w_{opt} (Supplementary Material, Note 1). For $L \geq 4$, the maxima are difficult to visually discern. More obvious is the increase in the number of peaks, and the probability that the number of peaks increases, as L increases, the latter converging to one for all values of w_{opt} except the extreme cases of $w_{\text{opt}} = 0$ and $w_{\text{opt}} = 1$. Note, however, that a high probability of increase in the number of peaks does not necessarily correspond to a high increase in the number of peaks, as can be seen from the different positions of the maxima in Fig. 4A,B (Supplementary Material, Note 1). For large L (> 10), the expected number of peaks in the fitness landscape increases exponentially with L ⁷⁷. In sum, these results show that selection for intermediate phenotypic values readily transforms Mt. Fuji genotype-phenotype landscapes, which are smooth and single-peaked, into rugged fitness landscapes, and that this effect is most pronounced for large L and intermediate w_{opt} .

Figure 4. Global transformability: Mt. Fuji and House-of-Cards genotype-phenotype landscapes. The absolute change in the number of peaks and the probability that the number of peaks changes in the fitness landscape, relative to (A,B) Mt. Fuji and (C,D) House-of-Cards genotype-phenotype landscapes, shown in relation to w_{opt} for $L \in \{2, 3, \dots, 8\}$. These results are independent of σ , because they only depend on the rank ordering of fitness values.

We next study the transformability of House-of-Cards genotype-phenotype landscapes. These landscapes are highly rugged, with an average of $\frac{2^L}{L+1}$ peaks⁷⁹, whereas the maximum possible number of peaks is 2^{L-1} . As such, selection for a low or intermediate phenotypic value can either increase or decrease the number of peaks in the fitness landscape, relative to the genotype-phenotype landscape. Fig. 4C shows this change in the number of peaks, and Fig. 4D shows the probability that the number of peaks changes, in relation to w_{opt} for landscapes with $L = 2$ to $L = 8$ loci. The change in the number of peaks is symmetric about $w_{\text{opt}} = 0.5$ for the two-locus case, where the number of peaks does not change as $w_{\text{opt}} \rightarrow 0$ or $w_{\text{opt}} \rightarrow 1$. However, this symmetry is lost for $L > 2$. The reason is that although the phenotype-fitness map flips all of the arrows in the fitness landscape when $w_{\text{opt}} = 0$, relative to the genotype-phenotype landscape, this does not guarantee conservation of the number of peaks. The number of peaks is jointly determined by the adjacent faces of the hypercube and thus, only very specific reorderings of arrows in the genotype-phenotype landscape guarantees conservation of the number of peaks in the fitness landscape (Supplementary Fig. 3). However, in contrast to Mt. Fuji genotype-phenotype landscapes, the magnitude of change in the number of peaks increases very little with L , despite an exponential increase in the maximum number of possible peaks. Moreover, the probability that the number of peaks changes is still less than one for large L . Finally, for large L , both the change in the number of peaks and the probability that the number of peaks changes are independent of w_{opt} , so long as w_{opt} is sufficiently less than one. This is a distribution dependent observation. Because the 2^L phenotypes in the NK model are drawn from a uniform distribution, nearly the same number of these phenotypes will be close to the optimal phenotype w_{opt} , so long as L is sufficiently large. Thus, averaging over all possible configurations of the genotype-phenotype landscape yields the same value of $\langle |p_f - p_{\text{gp}}| \rangle$ for every w_{opt} . For sufficiently large L , this value is given by (Supplementary Material, Derivation 2):

$$\lim_{L \rightarrow \infty} \langle |p_f - p_{\text{gp}}| \rangle = \sqrt{\frac{2^L \cdot (L-1)}{2\pi(L+1)^2}}. \quad (3)$$

So far we have focused on the absolute change $\langle |p_f - p_{\text{gp}}| \rangle$ in the number of peaks in the fitness landscape, relative to the genotype-phenotype landscape. For House-of-Cards genotype-phenotype landscapes, selection for a low or intermediate phenotypic value can either increase or decrease the number of peaks. We were therefore interested in finding out which outcome is more likely. While one might expect a decrease to be more likely, due to the extreme ruggedness of House-of-Cards genotype-phenotype landscapes, we find that the number of peaks is equally likely to increase or decrease (Supplementary Material, Proof 3). In sum, these results show that in House-of-Cards genotype-phenotype landscapes, selection for a low or intermediate phenotypic value increases or decreases the number of peaks in the fitness landscape with equal probability, and the severity as well as the probability of this change increases with L and is largely independent of w_{opt} .

Finally, we study the transformability of NK genotype-phenotype landscapes, which bridge the gap between Mt. Fuji and House-of-Cards landscapes in terms of ruggedness, as K increases from 0 to $L-1$. Fig. 5 shows the absolute change in the number of peaks in the fitness landscape relative to the genotype-phenotype landscape for genotypes of length $L = 5$ and $L = 8$, as K is increased from 0 to $L-1$. Note the gradual transition from the trends observed for Mt. Fuji genotype-phenotype

landscapes (Fig. 4A) to those observed for House-of-Cards genotype-phenotype landscapes (Fig. 4C) as K increases. From these trends, we conclude four principles of how the ruggedness of an NK genotype-phenotype landscape influences its transformability. As an NK genotype-phenotype landscape becomes more rugged, transformability (1) loses symmetry about $w_{\text{opt}} = 0.5$, (2) becomes less sensitive to w_{opt} , and (3) decreases in severity, at least in terms of the absolute change in the number of peaks. Finally, the probability of increasing the number of peaks is always greater than or equal to the probability of decreasing the number of peaks, with the equality holding for House-of-Cards genotype-phenotype landscapes. This last principle is both intuitive and informative — it tells us that on average, selection for low or intermediate values is more likely to increase the ruggedness of a fitness landscape, relative to the genotype-phenotype landscape. Thus, selection for low or intermediate values is more likely to break than to create phenotypic correlations between mutationally similar genotypes as they map onto fitness, rendering fitness landscapes more rugged than their underlying genotype-phenotype landscapes. In the subsequent sections we address whether and how these results apply to genotype-phenotype landscapes with more than two alleles per locus, specifically in the context of a biophysical model and experimental measurements of transcription factor-DNA interactions.

Figure 5. Global transformability: NK genotype-phenotype landscapes. The absolute change in the number of peaks in the fitness landscape, relative to the genotype-phenotype landscape, is shown in relation to w_{opt} for genotypes of length (A) $L = 5$ and (C) $L = 8$, as K increases from zero to $L - 1$. The corresponding probability of change in the number of peaks is shown in relation to w_{opt} for genotypes of length (B) $L = 5$ and (D) $L = 8$, as K increases from zero to $L - 1$. These results are independent of σ , because they only depend on the rank ordering of fitness values.

Genotype-phenotype landscapes of transcription factor-DNA interactions

Motivated by the common usage of low- and intermediate-affinity transcription factor binding sites in the regulatory portfolios of a diversity of organisms^{49,51,54–57}, we now study the transformability of genotype-phenotype landscapes of transcription factor-DNA interactions. In these landscapes, genotypes represent DNA sequences — transcription factor binding sites — and the phenotype of a DNA sequence is the affinity with which it binds a transcription factor²¹. Because the regulatory effects of transcription factor-DNA interactions are partly determined by binding affinity^{41,42} and mutations to transcription factor binding sites can alter binding affinity^{48,80}, the topographies of genotype-phenotype landscapes of transcription factor-DNA interactions have important implications for the evolution of gene regulation²¹. We study these landscapes using both a biophysical model and experimental measurements of transcription factor-DNA interactions. We focus on transcription factor binding sites of length $L = 8$, because this is the length of the binding sites assayed by protein binding microarrays^{80,81} — the data used to construct the empirical genotype-phenotype landscapes of transcription factor-DNA interactions²¹.

The mismatch model

We first study the transformability of genotype-phenotype landscapes of transcription factor-DNA interactions generated using the so-called mismatch model^{46,47,66}. The key assumption of this model is that the binding affinity of a DNA sequence is a linear function of the number of mismatches between the sequence (genotype) and a transcription factor’s consensus sequence — the sequence it binds with the highest affinity. Further, each mismatch is assumed to have the same energetic cost and these costs combine additively to determine binding affinity. This model results in a Mt. Fuji-like, permutation-invariant genotype-phenotype landscape, wherein the phenotype — binding affinity — only depends on the number of differences between the genotype and the consensus sequence, but not on which loci in the genotype differ from the consensus sequence.

We assume a binding affinity of 1 for the genotype that is identical to the consensus sequence. For each mismatch between a genotype and the consensus sequence, we deduct a small positive value e , such that the binding affinity of a genotype with m mismatches is $A_m = 1 - m \cdot e$. Due to the permutation invariance, the genotype-phenotype landscape is highly degenerate, such that the number of genotypes N_m in a mismatch class m is distributed according to the asymmetric binomial distribution:

$$N_m = (a - 1)^m \frac{L!}{(L - m)!m!}, \quad (4)$$

where a is the number of alleles per locus ($a = 4$ for transcription factor binding sites, because they are DNA sequences). Fig. 6A shows this distribution for transcription factor binding sites of length $L = 8$. Note that N_m is maximized at $m = 6$.

Fig. 6B shows the transformability of genotype-phenotype landscapes of transcription factor-DNA interactions constructed using the mismatch model under selection for an optimal mismatch class m_{opt} , reported in terms of the absolute change in the number of peaks in the fitness landscape, relative to the genotype-phenotype landscape. Based on our analysis of Mt. Fuji genotype-phenotype landscapes, we anticipated asymmetric transformability about the mismatch class $m_{\text{opt}} = 6$, because the distribution of the number of genotypes per mismatch class is asymmetric with a maximum at mismatch class $m = 6$ (Fig. 6A)

Figure 6. Global transformability: genotype-phenotype landscapes of transcription factor-DNA interactions. Landscapes constructed using (A,B) the mismatch model and (C,D) experimental measurements from protein binding microarrays for 1,137 eukaryotic transcription factors. (A) The number of genotypes, shown in relation to mismatch class. (B) The absolute change in the number of peaks in the fitness landscape, relative to the genotype-phenotype landscape, shown in relation to the optimal mismatch class m_{opt} . Labels indicate the number of genotypes per peak in the fitness landscape. Note the symmetry in the absolute change in the number of peaks around mismatch class $m_{opt} = 4$, as well as the tripling of the number of genotypes per peak for each increment in m_{opt} . The grey shaded circles are a schematic representation of the growing width of the peaks. (C) The number of genotypes per binding affinity class, where protein binding microarray E -scores are used as a proxy for relative binding affinity. Violin plots show the distribution, and box-and-whisker plots the 25-75% quartiles, across genotype-phenotype landscapes for the 1,137 transcription factors. Closed symbols and the dashed line denote the median of each distribution. (D) Violin plots of the distribution of the absolute change in the number of peaks in the fitness landscape, relative to the genotype-phenotype landscape, shown in relation to the optimal binding affinity w_{opt} for $\sigma = 0.15$. The inset shows the mean absolute change in the number of peaks, in relation to w_{opt} . The x-axes in (A,B) are arranged such that binding affinity increases when read from left to right, in qualitative agreement with the x-axes in (C,D). The results in panels A-C are independent of σ .

and we expected all genotypes in the optimal mismatch class m_{opt} to be peaks. However, we observe symmetric transformability about $m_{opt} = 4$ (Fig. 6B). This occurs because $a > 2$, which renders some genotypes in the same mismatch class mutational neighbours. Consequently, the peaks can be broad and include more than one genotype, thus resembling plateaus. Specifically, each genotype has $(a - 1)^m - 1$ mutational neighbours that are in the same mismatch class and are therefore part of the same peak. This leaves $\binom{L}{m}$ clusters of genotypes to be peaks in each mismatch class m , an expression that is maximized with $m = 4$ when $L = 8$, thus forcing symmetry in the absolute change in the number of peaks about $m_{opt} = 4$. However, the width of the peaks increases as $(a - 1)^m$, leading to a tripling of peak width for each increment in m (Fig. 6B).

Empirical landscapes

We now study the transformability of genotype-phenotype landscapes of transcription factor-DNA interactions generated using experimental data from protein binding microarrays⁸¹ (Methods). For all possible DNA sequences of length $L = 8$, these data include an enrichment score (E -score) ranging from -0.5 to 0.5 that serves as a proxy for relative binding affinity, such that higher E -scores correspond to higher binding affinities^{80,81}. We have previously used these data to construct genotype-phenotype landscapes for 1,137 eukaryotic transcription factors, in which the surface of each landscape was defined by the E -score²¹. Due to limitations in the reproducibility of E -scores across microarray designs for genotypes that are bound non-specifically or with very low affinity^{80,81}, each genotype-phenotype landscape only includes DNA sequences with an E -score exceeding a threshold of 0.35, which corresponds to a false discovery rate of 0.001⁸⁰. As shown in our previous work²¹, these landscapes tend to exhibit little, if any, reciprocal sign epistasis and therefore comprise few peaks. As such, they bear resemblance to the genotype-phenotype landscapes constructed using the mismatch model. An important difference, however, is that genotypes in the lower half of the E -score range are omitted from each landscape due to the reproducibility issues mentioned above. Fig. 6C shows the distributions of the number of genotypes across all 1,137 genotype-phenotype landscapes, grouped into six binding affinity classes. Whereas the lowest binding affinity class contains the most genotypes, we cannot determine if this is the true maximum, because we do not know what these distributions look like for lower binding affinity classes. However, assuming the energetic contribution of each binding site to be additive⁶⁶, we expect lower binding affinity classes to have fewer genotypes and the maximum to occur at an intermediate binding affinity class, as can be seen in the mismatch model and other models in literature⁴³.

Fig. 6D shows the transformability of these 1,137 empirical genotype-phenotype landscapes under selection for an optimal binding affinity w_{opt} , reported in terms of the absolute change in the number of peaks in the fitness landscape, relative to the genotype-phenotype landscape for $\sigma = 0.15$. Interestingly, the effect of changing σ is small (Supplementary Fig. 4B) because increasing σ not only decreases the range of variation of fitness values, but also decreases the uncertainty in these values (Methods), leaving the number of peaks largely unchanged. The mean and variance in the absolute change in the number of peaks decreases as w_{opt} increases to 0.5, in line with the intuition that selection for high w_{opt} generates fitness landscapes that are topographically similar to the underlying genotype-phenotype landscape. More precisely, the percentage of landscapes that show a change in the number of peaks decreases from around 99% to 40%, as we go from $w_{\text{opt}} = 0.35$ to $w_{\text{opt}} = 0.5$. As anticipated from our results with additive genotype-phenotype landscapes (Mt. Fuji and the mismatch model), selection for low or intermediate phenotypic values is more likely to increase than to decrease the number of peaks, although the relative fraction of increase depends upon w_{opt} – while around 99% of the landscapes show an increase in peaks for $w_{\text{opt}} = 0.35$, this value decreases to around 40% for $w_{\text{opt}} = 0.5$. Further, when we separately analyzed the single-peaked ($\approx 66.75\%$) and multi-peaked ($\approx 33.25\%$) genotype-phenotype landscapes, we found the single-peaked landscapes to have a higher transformability (Supplementary Fig. 4A), in line with our results from the previous section.

In contrast to genotype-phenotype landscapes constructed using the mismatch model, the transformability of genotype-

phenotype landscapes constructed using protein binding microarray data is highest for the lowest binding affinity class, rather than an intermediate class. There are three non-mutually exclusive explanations for this. First, these empirical landscapes are not purely additive²¹, unlike the landscapes constructed with the mismatch model, so we do not expect perfect symmetry about an intermediate w_{opt} . Second, as previously mentioned, protein binding microarray data do not capture the full range of binding affinity, so the lowest binding affinity class in our data (E -score = 0.35), which contains the most genotypes (Fig. 6C), is unlikely to be the lowest binding affinity class, but rather an intermediate binding affinity class. Third, while the binding affinity of a sequence is highly correlated with the binding affinities of its mutational neighbors⁸², this correlation is not perfect, so neighboring genotypes that are in the same mismatch class may not have sufficiently similar binding affinities to be considered part of the same peak, unlike in the mismatch model.

There are two additional differences between the transformability of the empirical genotype-phenotype landscapes and those constructed using the mismatch model that are worth highlighting. First, in the empirical landscapes, the average height of the peaks is maximised when selecting for intermediate binding affinities and is lowest when selecting for the highest affinity (Supplementary Fig. 5A), whereas peak height is independent of m_{opt} in the mismatch model. Second, in the empirical landscapes, peak width is maximized when selecting for the highest binding affinity ($w_{\text{opt}} = 0.5$) (Supplementary Fig. 5B), whereas it is maximized when selecting for the lowest binding affinity ($m_{\text{opt}} = 8$) in the mismatch model (Fig. 6B). These two differences in transformability are important for understanding evolutionary simulations on these landscapes, which are the focus of the next section.

Evolutionary consequences

A key finding of our analyses so far is that selection for low or intermediate phenotypic values is more likely to increase than to decrease the number of peaks in the fitness landscape, relative to the genotype-phenotype landscape. Because the ruggedness of a fitness landscape has important implications for a diversity of evolutionary phenomena^{59–62}, we now study the evolutionary consequences of this finding. We do so using two metrics: (1) the length $\langle l \rangle$ of a greedy adaptive walk, averaged over all possible genotypes as starting points, and (2) the mean population fitness at equilibrium $\langle f \rangle$ under deterministic mutation-selection dynamics (i.e., assuming an infinite population size). These metrics provide complementary information to measures of landscape ruggedness, such as the number of peaks. To see why, consider our results with the mismatch model, in which the fitness landscape had one or more global peaks of equal height (Fig. 6A,B). We observed single-peaked fitness landscapes when the optimal mismatch class was $m_{\text{opt}} = 0$ or $m_{\text{opt}} = 8$, and the most rugged fitness landscapes when the optimal mismatch class was $m_{\text{opt}} = 4$. We also observed a tripling of the number of genotypes per peak as m_{opt} increased from 0 to 8, such that the landscape with $m_{\text{opt}} = 0$ comprised a single peak with one genotype, the landscape with $m_{\text{opt}} = 4$ comprised 70 peaks with 81 genotypes per peak, and the landscape with $m_{\text{opt}} = 8$ comprised 1 peak with 6,561 genotypes. Which landscape topography is most conducive to adaptive evolutionary change?

NK landscapes

Fig. 7A,C shows the average length of a greedy adaptive walk $\langle l \rangle$ in NK landscapes with $L = 5$ and $L = 8$. For all $K < L - 1$, $\langle l \rangle$ is minimised when $w_{\text{opt}} = 0.5$, whereas for $K = L - 1$ (House-of-Cards landscapes), $\langle l \rangle$ is independent of w_{opt} . This can be understood by first considering additive genotype-phenotype landscapes ($K = 0$), wherein genotypes with a phenotype $w \approx 0.5$ are by definition in the “center” of the landscape. Selecting for $w_{\text{opt}} = 0.5$ thus minimizes the average length of a greedy adaptive walk to these peak genotypes. As K increases, correlations among the phenotypes of mutationally-neighboring genotypes are gradually broken and the landscape becomes more rugged. Yet, genotypes with a phenotype $w \approx 0.5$ tend to remain near the center of the landscape. When K is increased to $L - 1$, the phenotypes of mutationally-neighboring genotypes become completely uncorrelated and monotonic increases in phenotype along mutational paths to peak genotypes become exceedingly rare, rendering $\langle l \rangle$ independent of w_{opt} . Thus the “U”-shape of the curve of $\langle l \rangle$ versus w_{opt} flattens as K increases from 0 to $L - 1$. Interestingly, so long as $K > 0$, the average height of the peak reached by a greedy adaptive walk is also maximized at $w_{\text{opt}} = 0.5$, and these local peaks are $>97\%$ the height of the global peak for $\sigma = 1$ (Supplementary Fig. 6), an observation we explain below.

We next simulated deterministic mutation-selection dynamics. Fig. 7B,D show the change in mean fitness at equilibrium when selecting for w_{opt} , relative to fitness at equilibrium when selecting for $w_{\text{opt}} = 1$, for $\sigma = 0.5$ and $\mu = 0.1$ (see Supplementary Fig. 7 for other values of σ and μ). This change is always positive, increases with L and K , and tends to be maximized at $w_{\text{opt}} = 0.5$ when $K > 0$. The latter occurs because the optimal phenotype $w_{\text{opt}} = 0.5$ is, on average, closer to all other phenotypic values in the landscape. For instance for $K = L - 1$, the sum $\sum_{i=1}^{2^L} (w_i - w_{\text{opt}})^2$ is minimized when $w_{\text{opt}} = \sum_{i=1}^{2^L} w_i / 2^L = \langle w \rangle = 0.5$. Thus, selecting for $w_{\text{opt}} = 0.5$ creates a fitness landscape with higher average fitness than fitness landscapes under selection for any other w_{opt} . This also explains why the average height of local adaptive peaks is maximized at $w_{\text{opt}} = 0.5$ (Supplementary Fig. 6). In sum, these results show that selection for an intermediate phenotypic value decreases the length of a greedy adaptive walk and increases mean fitness at equilibrium, despite increasing the overall ruggedness of the fitness landscape, relative to the genotype-phenotype landscape.

Mismatch model

In the fitness landscapes generated using the mismatch model, the length of the greedy adaptive walk averaged over all starting genotypes is given by

$$\langle l \rangle = \frac{\sum_{m=0}^L (a-1)^m \cdot \frac{L!}{(L-m)!m!} \cdot |m_{\text{opt}} - m|}{a^L}. \quad (5)$$

Fig. 8A shows this expression for $L = 8$ and $a = 4$. Due to the additive nature and degeneracy of this genotype-phenotype landscape, the fitness landscapes have monotonically decreasing fitness values as the mutational distance from any peak increases. Therefore, the length of any individual greedy adaptive walk is simply the absolute difference of the mismatch class of the starting genotype and that of the optimal mismatch class (i.e., $|m_{\text{opt}} - m|$). As such, the average length of the greedy

Figure 7. Evolutionary consequences of landscape ruggedness in NK landscapes. (A,C) The average length of a greedy adaptive walk and (B,D) the change in mean fitness at equilibrium, relative to fitness at equilibrium when selecting for $w_{\text{opt}} = 1$, are shown for (A,B) $L = 5$ and (C,D) $L = 8$. In (B,D), $\mu = 0.1$ and $\sigma = 0.5$.

adaptive walk is minimized when selecting for $m_{\text{opt}} = 6$ (Fig. 8A), because this maximizes the number of genotypes in adaptive peaks (Fig. 6B; 28 global peaks $\times$ 729 genotypes per peak). If instead, we chose to start the greedy walk from only non-peak genotypes, the walk would still be minimised for $m_{\text{opt}} = 6$, because this class has the maximum number of genotypes that are Hamming distance one away from the peak genotypes. Recall that in each fitness landscape generated upon selection for m_{opt} , all the peaks are of the same height and therefore, the greedy walk always terminates on a global peak. This is in contrast with the NK landscapes, and as we will see below, the empirical landscapes.

For each m_{opt} , we next calculated the change in mean fitness at equilibrium under deterministic mutation-selection dynamics, relative to the fitness landscape generated for $m_{\text{opt}} = 0$. To do so, we exploit the permutation-invariance of these landscapes to group genotypes into a lower-dimensional state space defined by mismatch class (Methods). Specifically, we construct a transition matrix that defines the probability that a genotype from one mismatch class mutates into another, based on the frequency and fitness of genotypes in each mismatch class. We iterate this matrix until we reach steady state, which is guaranteed by the Frobenius-Perron theorem to be independent of initial conditions, because the matrix is irreducible⁸³. Fig. 8C shows the change in mean fitness at equilibrium in relation to the optimal mismatch class $m_{\text{opt}} = 0$. Our first observation is that the change in mean fitness is always positive for $m_{\text{opt}} > 0$ when $\mu = 0.1$ and $\sigma = 1$, even though selection for such mismatch classes always causes an increase in the number of peaks in the fitness landscape, relative to the underlying genotype-phenotype landscape (Fig. 6B). Our second observation is that the change in mean fitness relative to m_{opt} is always unimodal. For example,

selection for $m_{\text{opt}} = 6$ maximizes the change in mean fitness when $\mu = 0.1$ and $\sigma = 1$ (Fig. 8C). More generally, the m_{opt} that maximizes the change in mean fitness depends on the interplay between the mutation rate μ and the strength of selection σ . See Supplementary Fig. 9 for the phase diagram. When selection is strong or mutation is weak, mismatch class $m_{\text{opt}} = 8$ maximizes the change in mean fitness. As the strength of selection decreases or the mutation rate increases, the mismatch class that maximizes the change in mean fitness decreases from $m_{\text{opt}} = 8$ to $m_{\text{opt}} = 7$ and then to $m_{\text{opt}} = 6$, where it remains for weak selection or high mutation rates. This is because when σ is small (i.e., selection is strong), it is costly to step down from a peak and therefore, selecting for the class with the broadest peak ($m_{\text{opt}} = 8$) leads to the highest equilibrium mean population fitness. As σ increases and selection becomes weaker, it is no longer as costly to step down from a peak and therefore, selecting for the class with the maximum number of genotypes in peaks ($m_{\text{opt}} = 6$) maximizes mean fitness at equilibrium. While this effect is reminiscent of the “survival of the flattest”⁸⁴ phenomenon, the crucial difference is that in the mismatch model, the peaks always have the same height and thus, there is no trade-off between peak height and width. Another way of altering the strength of selection is by changing e , the energetic cost of a mismatch. Larger e corresponds to stronger selection and the phase diagram changes accordingly (Supplementary Fig. 10).

Empirical landscapes

The empirical genotype-phenotype landscapes are topologically more complex^{82,85} than the genotype-phenotype landscapes generated with the mismatch model, which are regular graphs (i.e., every genotype has $(a - 1) \cdot L$ mutational neighbors). We therefore used simulations to calculate the average length of a greedy adaptive walk $\langle l \rangle$ in these empirical landscapes⁸⁶, initiating the walks from all non-peak genotypes in the fitness landscape. Moreover, each binding affinity measurement (E -score) in the empirical landscapes is associated with a noise threshold that is used to determine whether two genotypes truly differ in phenotype (Methods). This noise threshold can cause landscapes to have large non-peak plateaus, in which many genotypes have indistinguishable fitness. We therefore modified the greedy adaptive walk such that when a non-peak plateau was encountered, we chose a random mutational neighbor of indistinguishable fitness for the next step in the walk, disallowing reversion mutations. We repeated this process until the plateau was traversed and a sequence with higher fitness was reached. Finally, we terminated the walk when a peak sequence was reached. Therefore, the walk was primarily a deterministic greedy walk, with some stochasticity due to the non-peak plateaus.

Fig. 8B shows $\langle l \rangle$, averaged over 100 simulations of the adaptive walk from each initial condition, in relation to w_{opt} . In contrast to the mismatch model, $\langle l \rangle$ is shortest for $w_{\text{opt}} = 0.35$ and $w_{\text{opt}} = 0.5$ and slightly higher for intermediate w_{opt} . This is because, as in the mismatch model, $\langle l \rangle$ is correlated with the total number of genotypes in peaks, which depends upon both the number of peaks and their widths. Whereas selecting for $w_{\text{opt}} = 0.35$ leads to the largest number of peaks (Fig. 6D), selecting for $w_{\text{opt}} = 0.5$ leads to the broadest peaks (Supplementary Fig. 5B), thus explaining the minimisation of $\langle l \rangle$ when selecting for these extreme phenotypes. When selecting for low or intermediate phenotypic values, most walks terminate at a local, rather than the global fitness peak. However, these local peaks tend to be nearly as high as the global peak, especially when selecting for intermediate phenotypic values (Supplementary Fig. 8).

Figure 8. Rugged fitness landscapes need not impede adaptation. The average length of the shortest adaptive walk to a fitness landscape's global peak $\langle l \rangle$ and the mean population fitness at equilibrium $\langle f \rangle$ is shown for landscapes of transcription factor-DNA interactions generated using (A,C) the mismatch model for $e = 0.05$ and (B,D) protein binding microarray data. In (B,D), violin plots show the distribution, and box-and-whisker plots the 25-75% quartiles, across the 1,137 empirical landscapes for each optimal binding affinity w_{opt} . Closed symbols and dashed line represent the mean. In (C), $\mu = 0.1$, $\sigma = 1$ while in (D) $\mu = 0.1$, $\sigma = 0.15$.

Next, we simulated deterministic mutation selection dynamics (Methods). For each w_{opt} , Fig. 8D shows the change in mean fitness at equilibrium for $\mu = 0.1$, $\sigma = 0.15$, relative to the fitness landscape generated after selecting for $w_{\text{opt}} = 0.5$. As in the mismatch model, the change in mean fitness tends to be positive, despite the increase in the number of peaks caused by selection for low or intermediate phenotypic values (Fig. 6D). Moreover, the mean change in fitness at equilibrium is maximized when selection favors an intermediate phenotypic value ($w_{\text{opt}} = 0.425$ for $\mu = 0.1$, $\sigma = 0.15$). This can be explained by the average peak heights and widths that occur when selecting for intermediate phenotypic values (Supplementary Fig. 5A and B for $\sigma = 0.15$). Exactly which w_{opt} maximizes the change in mean fitness depends on the interplay between the mutation rate μ and the strength of selection σ (Supplementary Fig. 11), converging on 0.425 for large μ and σ , similar to the convergence seen at $m_{\text{opt}} = 6$ for the mismatch model.

Discussion

We systematically analyzed the transformability of a diversity of theoretical and empirical genotype-phenotype landscapes under selection for a low or intermediate phenotypic value. Such non-linear relationships between phenotype and fitness have

been observed in a diversity of biological systems^{28,31,33,35,87,88}, often reflecting a tradeoff between the costs and benefits of a particular phenotype, such as the expression level of a gene. These non-linearities, either in the genotype-phenotype or phenotype-fitness map are a cause of epistasis^{36–38,77}. The reason is that two or more mutations whose effects contribute additively to an intermediate phenotype, such as binding energy or gene expression, may contribute non-additively to a downstream phenotype, such as binding probability, or to fitness. Moreover, when mutations have epistatic interactions in their contribution to phenotype, a non-linear phenotype-fitness map can change the form of these interactions from negative to positive, or vice versa¹⁶, as well as introduce or remove sign epistasis⁸⁷. Our work complements these empirical observations and expands upon previous theoretical work⁷⁷, by systematically quantifying how and how often selection for a low or intermediate phenotypic value introduces or removes epistasis in the fitness landscape. Specifically, we show that the probability of changing the type of epistasis (e.g., positive to negative) is highest when selecting for low phenotypic values and the probability of introducing or removing sign epistasis is highest when selecting for intermediate phenotypic values.

Sign epistasis is a cause of landscape ruggedness⁵. We quantified the global transformability of genotype-phenotype landscapes under selection for low or intermediate phenotypic values in terms of the absolute change in the number of peaks in the fitness landscape, relative to the genotype-phenotype landscape. A key finding of this analysis is that selection for low or intermediate phenotypic values is more likely to increase than to decrease the number of peaks, with the probability of the two types of change being equal only in House-of-Cards genotype-phenotype landscapes. While increased landscape ruggedness is typically thought to frustrate the evolutionary process, because it limits the amount of adaptive phenotypic variation mutation can bring forth⁴, our evolutionary simulations show this need not be the case. Specifically, we find that the rugged fitness landscapes caused by selection for low or intermediate phenotypic values comprise local adaptive peaks that are nearly as tall as the global adaptive peak. Moreover, these local peaks tend to be accessible from throughout the landscape via a small number of sequential mutations that monotonically increase fitness. As a result, the mean population fitness at equilibrium is almost always higher when selecting for low or intermediate phenotypic values than when selecting for a high phenotypic value.

These topographical properties may help to explain the prevalence of low- and intermediate-affinity binding sites in the control of gene expression. Prior work has suggested an entropic argument⁸⁹: As with certain RNA secondary structures⁹⁰ or regulatory circuit motifs⁹¹, low- and intermediate-affinity binding sites appear more frequently simply because they are more “findable”. That is, because a transcription factor binds more distinct DNA sequences with low or intermediate affinity than with high affinity, low- and intermediate-affinity binding sites are more likely to evolve to control gene expression. Our work builds upon this “arrival of the frequent” argument by showing that low- and intermediate-affinity binding sites are not only more likely to evolve *de novo* due to their increased frequency, they are also easier to evolve via a short number of “uphill” mutational steps from throughout the landscape. This finding may also help to reconcile theoretical results suggesting long waiting times for the evolution of transcription factor binding sites⁴⁷ with the empirical observation that binding sites turn over rapidly in a diversity of species^{92–94}.

We made several modeling assumptions to simplify our analyses of transcription-factor DNA interactions, the relaxation of

which may open new avenues for future research. First, we assumed a single nonlinearity in the relationship between genotype, phenotype, and fitness. As noted by Domingo et al.³⁷, “from transcription to RNA processing, translation, and protein folding and all the way up to protein activity and cellular fitness, there are many layers of biological organization where the effects of a mutation can be transformed.” To date, our knowledge of how such nonlinearities combine to transform genotype-phenotype landscapes is based on a small number of experimental studies (e.g., ref¹⁶). A systematic theoretical analysis could build off the work presented here, for example by incorporating the sigmoidal relationship between binding site occupancy and binding affinity in modeling transcription factor-DNA interactions⁹⁵, such that fitness depends nonlinearly on occupancy, rather than affinity. Other nonlinearities, such as those caused by transcription factor cooperativity⁹⁵, could also be included.

Second, we assumed that selection acts directly on a single phenotype — binding affinity. While this assumption is common in models of the evolution of transcription factor binding sites^{21,43,46} and is supported by empirical data^{43,44}, the relationship between binding affinity and fitness is not so direct, because it is modulated by gene expression. Gene expression depends on a variety of factors, including the presence, arrangement, and affinities of binding sites for other competing or cooperating transcription factors in promoters and enhancers⁹⁶, as well as local sequence context⁹⁷, chromatin context⁹⁸, DNA methylation⁹⁹, and local transcription factor concentrations¹⁰⁰. Existing modeling frameworks that relate the architecture of entire regulatory regions to gene expression patterns may provide a path forward¹⁰¹, facilitating the study of landscape transformability when fitness depends upon the multitude of molecular phenotypes characteristic of eukaryotic gene regulation.

Third, we assumed a Gaussian phenotype-fitness map, which is commonly employed in a diversity of modeling frameworks^{29,73,75,102}, including those for transcription factor-DNA interactions⁴⁶. We expect alternative symmetric phenotype-fitness maps (e.g., ref.¹⁰³) will only affect our results quantitatively, because many of our findings, such as the changes in sign epistasis motifs and number of peaks, only depend on the rank ordering of fitness values. However, we expect asymmetric phenotype-fitness maps, such as those uncovered in experimental studies of biological systems as different as viruses³⁰ and yeast¹⁰⁴, to affect our results qualitatively. For example, in our analyses of NK landscapes, we often observed symmetries in transformability around an intermediate phenotypic value. These symmetries will almost certainly be lost. Understanding how asymmetric phenotype-fitness maps affect the transformability of genotype-phenotype landscape is therefore an outstanding challenge.

Finally, our study may open new lines of research on dynamic genotype-phenotype and fitness landscapes^{68,71,105–109}. For example, whereas we studied selection for a fixed phenotypic optimum, this optimum may in fact change in space or in time. Our results imply that even gradual changes in the phenotypic optimum may lead to abrupt changes in fitness landscape topography, which may have implications for an evolving population’s ability to track this optimum and thus for population persistence and extinction. Moreover, because our measures of transformability can be applied to any pair of landscapes so long as they are defined over the same set of genotypes, they may shed light on the kinds of topographical alterations induced by fluctuating environmental factors, such as DNA methylation⁹⁹ or the presence of protein partners¹¹⁰ on transcription factor-DNA interactions. As our ability to experimentally interrogate such complexities in the relationship between genotype,

phenotype, and fitness continues to improve, we anticipate a sharpened focus on landscape dynamics and their implications for the evolutionary process.

Methods

Epistasis

In a genotype-phenotype landscape, we classify the type of magnitude epistasis between a pair of loci using the following linear combination of phenotypic values:

$$\epsilon_{gp} = w_{00} + w_{11} - w_{01} - w_{10}, \quad (6)$$

where w_i represents the phenotype of genotype $i \in \{0, 1\}^2$. When $\epsilon_{gp} = 0$, there is no magnitude epistasis, because the phenotypic effects of the two alleles combine additively; when $\epsilon_{gp} > 0$, there is positive epistasis, because the phenotypic effects of the two alleles are greater than expected based on their individual phenotypic effects; when $\epsilon_{gp} < 0$, there is negative epistasis, because the phenotypic effects of the two alleles are less than expected based on their individual phenotypic effects.

Analogously, in the fitness landscape, we classify the type of magnitude epistasis between a pair of loci using the following linear combination of fitness values:

$$\epsilon_f = F_{00} + F_{11} - F_{01} - F_{10}, \quad (7)$$

where F_i represents the fitness of the corresponding genotype i . As in the genotype-phenotype landscape, $\epsilon_f = 0$, $\epsilon_f > 0$, and $\epsilon_f < 0$ indicate additive, positive, and negative epistatic interactions among loci, respectively.

We use ϵ_{gp} and ϵ_f to calculate the fraction of pairs of loci that have the same type of epistasis in the genotype-phenotype landscape and the fitness landscape, which we determine as the product of ϵ_{gp} and ϵ_f . This is because the type of epistasis is the same in the two landscapes when $\epsilon_{gp} \cdot \epsilon_f > 0$. While it is theoretically possible for both ϵ_{gp} and ϵ_f to be 0, in which case the type of epistasis would be the same in the two landscapes yet the condition $\epsilon_{gp} \cdot \epsilon_f > 0$ would not be satisfied, this never happens in practice.

NK Landscapes

We constructed the NK landscapes using the adjacency neighbourhood scheme, wherein each locus i of a genotype of length L , interacts with K adjacent loci to the right of locus i and $0 \leq K \leq L - 1$. We used periodic boundary conditions, such that the L -th locus interacts with the first K loci, and so on.

The phenotype $w(\tau)$ of genotype τ is computed as the sum of the individual contributions of all loci, each of which depends

on K other interacting loci:

$$w(\tau) = \sum_{i=1}^L f(\tau_i; \tau_i^1, \tau_i^2 \dots \tau_i^K), \quad (8)$$

where $f(\tau_i; \tau_i^1, \tau_i^2 \dots \tau_i^K)$ represents the contribution of the i -th locus, which depends on K other loci $\tau_i^1, \tau_i^2 \dots \tau_i^K$. We drew the contributions of each of the 2^{K+1} possible configurations from a uniform distribution between 0 and 1. Finally, we re-scaled the phenotypic values by subtracting the minimum and dividing by the maximum, such that they fell between 0 and 1.

Empirical genotype-phenotype landscapes of transcription factor-DNA interactions

We studied the transformability of 1,137 genotype-phenotype landscapes of transcription factor-DNA interactions. The procedure for constructing these landscapes has been described elsewhere²¹. In brief, each landscape corresponds to a single transcription factor, the genotypes it contains represent DNA sequences of length $L = 8$ that specifically bind the transcription factor, and the phenotype of each sequence is a quantitative proxy for relative binding affinity, which defines the surface of the landscape. These phenotypes are reported as enrichment scores (E -scores) derived from protein binding microarrays. In each landscape, two genotypes are considered mutational neighbors if they differ by a single small mutation, specifically a point mutation or a small indel that shifts an entire contiguous binding site by a single base⁸². We performed all analyses on the dominant genotype network (i.e., the largest connected component of the network), which comprises the vast majority of genotypes in each landscape²¹.

We used the Genonets Python package (version 0.31) to characterize the topographical properties of the empirical genotype-phenotype landscapes¹¹¹. Specifically, we used this package to compute the number of peaks per landscape and the number of genotypes in the peaks of each landscape. These calculations rely on a noise threshold δ , which is used to determine whether two genotypes actually differ in phenotype. For each transcription factor, we used the value of δ reported in ref.²¹, which was derived from a comparison of binding affinity measurements across two protein binding microarray designs.

While characterizing the topographies of the transformed fitness landscapes, we also had to transform δ following the rules of error propagation. Accordingly, the noise in fitness values, dF , depends on the noise in the phenotypic values $dw = \delta$ as follows:

$$dF = -2F \cdot \frac{w - w_{\text{opt}}}{\sigma^2} \cdot \delta. \quad (9)$$

To compute the number of peaks and their widths in the fitness landscapes, we adapted Genonets to specify different noise threshold values (dF) for each genotype.

Mutation-selection dynamics for the mismatch model

We grouped genotypes according to their mismatch class m and iterated a series of selection and mutation steps until the population reached an equilibrium distribution. In each selection step, the frequency of each mismatch class X_m was scaled by its fitness F_m and then normalized:

$$X_m^S(t+1) = \frac{F_m \cdot X_m(t)}{\sum_i F_i \cdot X_i}, \quad (10)$$

where $X_m^S(t+1)$ is the frequency of the mismatch class m after the selection step.

The selection step was followed by a mutation step, in which genotypes could either mutate within their mismatch class or mutate to an adjacent mismatch class. The total mutation probability is μ , so with probability $1 - \mu$, the genotype does not mutate. The frequency of each mismatch class was thus updated as

$$X_m^M(t+1) = \left[\mu \cdot \frac{L - (m-1)}{L} \right] \cdot X_{m-1}^S(t+1) + \left[(1 - \mu) + \mu \cdot \frac{(a-2) \cdot m}{(a-1) \cdot L} \right] \cdot X_m^S(t+1) + \left[\mu \cdot \frac{m+1}{(a-1) \cdot L} \right] \cdot X_{m+1}^S(t+1), \quad (11)$$

where $X_m^M(t+1)$ is the frequency of the mismatch class m after the mutation step and μ is the mutation probability. The mutation step implicitly accounts for the different number of genotypes per mismatch class. Finally, $X_m(t+1)$ was set to $X_m^M(t+1)$ and t was incremented. This process was repeated until equilibrium.

Mutation-selection dynamics for the empirical landscapes

The empirical landscapes comprise far fewer genotypes than the mismatch landscape. We therefore defined the state space of the empirical landscapes in terms of the individual genotypes in each landscape, merging each genotype with its reverse complement⁸². The mutational neighbors of each genotype were those that differed by a single small mutation, namely a point mutation or an indel that shifted an entire contiguous binding site by a single base⁸². In this analysis, we did not account for uncertainties in the fitness values, because these do not influence our results in this evolutionary regime (see below).

The recursion relation for mutation-selection dynamics in discrete time is

$$x_i(t+1) = \sum_j \mu^{d(\tau_i, \tau_j)} \cdot (1 - \mu)^{L - d(\tau_i, \tau_j)} \frac{f_j}{\bar{f}(\vec{x}, t)} x_j(t), \quad (12)$$

where $\vec{x}$ is the vector of genotype frequencies, x_i is the frequency of the i th genotype, μ is the probability of a single point mutation or a small indel mutation, L is the length of the genotypes, $d(\tau_i, \tau_j)$ is the minimum of the mutational distance between genotypes τ_i and τ_j and between genotypes τ_i and the reverse complement of τ_j , f_i is the fitness of the i th genotype, and $\bar{f}(\vec{x}, t)$ is the mean population fitness at time t .

We linearized the dynamics by substituting $z_i(t) = \frac{x_i(t)}{\prod_{\tau=1}^{t-1} \bar{f}(\vec{x}, \tau)}$, yielding

$$z_i(t+1) = \sum_j \mu^{d(\tau_i, \tau_j)} (1 - \mu)^{L-d(\tau_i, \tau_j)} f_j z_j(t), \quad (13)$$

from which we retrieved the normalized genotype frequencies with $\vec{x} = \vec{z} / (\sum_i z_i)$. In matrix form, the dynamics are

$$\vec{z}(t+1) = M \cdot S \cdot \vec{z}(t), \quad (14)$$

where the mutation matrix M has elements $M_{ij} = \mu^{d(\tau_i, \tau_j)} \cdot (1 - \mu)^{L-d(\tau_i, \tau_j)}$ and the selection matrix S is a diagonal matrix with $S_{ii} = f_i$, where f_i is the fitness of sequence τ_i . We calculated the eigenvector corresponding to the largest eigenvalue of $M \cdot S$ to determine the equilibrium state of the dynamics, which is guaranteed by the Frobenius-Perron theorem to be unique and stable.

Because the fitness values have some uncertainty around them, we additionally performed a sensitivity analysis by adding Gaussian noise to the fitness values, with standard deviation equal to $1/3 \cdot dF$, such that 99.7% of the sampled fitness values fell within $F \pm \delta$. Our results are robust to these perturbations, with the exception of two parameter combinations – $\sigma = 0.1, \mu = 0.001$ and $\sigma = 0.1, \mu = 0.0025$. For small σ , dF is large and therefore such sensitivity is expected. However, the quantitative changes in our results were small, leading us to conclude that fitness uncertainties do not significantly influence our results in this evolutionary regime.

Supplementary material

Derivations

Derivation 1: Derivation of equation 2

The epistasis in the fitness landscape is given by equation 7. We assume σ to be large, such that $(w_i - w_{\text{opt}})^2 / \sigma^2 \ll 1 \forall i \in \{0, 1\}^2$, an approximation that holds when selection is weak. Taylor expanding F_i and neglecting higher order terms, we find

$$F_i = \exp \left[- \left(\frac{w_i - w_{\text{opt}}}{\sigma} \right)^2 \right] \approx 1 - (w_i - w_{\text{opt}})^2 / \sigma^2, \quad (15)$$

where $i \in \{0, 1\}^2$. Let $\Delta_i = w_i - w_{\text{opt}}$, then

$$\begin{aligned} \varepsilon_f &\approx \frac{-1}{\sigma^2} (\Delta_{00}^2 + \Delta_{11}^2 - \Delta_{01}^2 - \Delta_{10}^2) = \frac{-1}{\sigma^2} [(\Delta_{00} + \Delta_{11})^2 - 2\Delta_{00}\Delta_{11} - (\Delta_{01} + \Delta_{10})^2 + 2\Delta_{10}\Delta_{01}] \\ &= \frac{-1}{\sigma^2} [(\Delta_{00} + \Delta_{11} + \Delta_{01} + \Delta_{10})(\Delta_{00} + \Delta_{11} - \Delta_{01} - \Delta_{10}) - 2(\Delta_{00}\Delta_{11} - \Delta_{10}\Delta_{01})] \\ &= \frac{-1}{\sigma^2} [(\sum_i \Delta_i) \varepsilon_{\text{gp}} - 2\varepsilon_{m,\Delta}] \end{aligned} \quad (16)$$

where $\varepsilon_{m,\Delta} = \Delta_{00}\Delta_{11} - \Delta_{10}\Delta_{01} = \varepsilon_{m,\Delta} = (w_{00} - w_{\text{opt}})(w_{11} - w_{\text{opt}}) - (w_{01} - w_{\text{opt}})(w_{10} - w_{\text{opt}}) \implies \varepsilon_{m,\Delta} = \varepsilon_m - w_{\text{opt}} \cdot \varepsilon_{\text{gp}}$, and

$\epsilon_m = w_{00}w_{11} - w_{01}w_{10}$ is the multiplicative epistasis. Substituting for $\epsilon_{m,\Delta}$, expanding the $\Sigma_i \Delta_i$ term and rearranging gives:

$$\epsilon_f = \frac{1}{\sigma^2} [2\epsilon_m + (2w_{\text{opt}} - \Sigma_i w_i) \cdot \epsilon_{\text{gp}}] \implies \epsilon_f \cdot \epsilon_{\text{gp}} = \frac{1}{\sigma^2} [2\epsilon_m \cdot \epsilon_{\text{gp}} + (2w_{\text{opt}} - \Sigma_i w_i) \cdot \epsilon_{\text{gp}}^2]. \quad (17)$$

Derivation 2: Derivation of equation 3

The absolute difference between the number of peaks in the genotype-phenotype landscape and the fitness landscape is

$$|p_f - p_{\text{gp}}| = \begin{cases} p_f - p_{\text{gp}}, & \text{if } p_f > p_{\text{gp}} \\ p_{\text{gp}} - p_f, & \text{if } p_f < p_{\text{gp}}. \end{cases} \quad (18)$$

We can rule out the case where $p_f = p_{\text{gp}}$, because the probability of change in the number of peaks tends towards 1 as L increases, so long as w_{opt} is sufficiently far from 1 (Fig. 4D). This implies

$$\langle |p_f - p_{\text{gp}}| \rangle = \langle p_f - p_{\text{gp}} \rangle_{p_f > p_{\text{gp}}} P(p_f > p_{\text{gp}}) + \langle p_{\text{gp}} - p_f \rangle_{p_{\text{gp}} > p_f} P(p_{\text{gp}} > p_f), \quad (19)$$

where $\langle p_f - p_{\text{gp}} \rangle_{p_f > p_{\text{gp}}}$ is the conditional mean of $p_f - p_{\text{gp}}$, given that $p_f > p_{\text{gp}}$.

As shown below, the probability of increase in the number of peaks is equal to the probability of decrease in the number of peaks (Proof 3); i.e., $P(p_f > p_{\text{gp}}) = P(p_{\text{gp}} > p_f) = 0.5$. This implies

$$\langle |p_f - p_{\text{gp}}| \rangle = 0.5(\langle p_f - p_{\text{gp}} \rangle_{p_f > p_{\text{gp}}} + \langle p_{\text{gp}} - p_f \rangle_{p_{\text{gp}} > p_f}). \quad (20)$$

In genotype-phenotype landscapes generated by the House-of-Cards model, the phenotypic values of mutational neighbors are completely uncorrelated, so selection for a low or intermediate phenotypic value generates a fitness landscape in which the fitness values of mutational neighbors are also completely uncorrelated. Thus, the distribution of the number of peaks is the same for the genotype-phenotype landscape and the fitness landscape. Moreover, p_f is independent of p_{gp} because p_f depends upon the number of genotypes with phenotypes close to the optimal phenotypic value, and this number is independent of the number of genotypes with phenotypes close to 1. Therefore, p_{gp} and p_f are independent and identically distributed (i.i.d) random variables. Now as $L \rightarrow \infty$, the distribution of the number of peaks in the House-of-Cards model tends to a normal distribution¹¹² with mean $\frac{a^L}{(a-1) \cdot L+1}$ and variance $\tau^2 = \frac{a^{L \cdot ((a-1) \cdot L - (a-1))}}{2((a-1) \cdot L+1)^2}$ ⁷⁹, where a is the number of alleles at each site and here $a = 2$. This implies the distribution of the random variable $p_f - p_{\text{gp}}$ is also normally distributed with mean zero and variance $2\tau^2$; i.e.,

$$p_f - p_{\text{gp}} := x \sim \mathcal{N}(0, 2\tau^2), \quad (21)$$

618 which is symmetric about zero, implying $\langle p_f - p_{gp} \rangle_{p_f > p_{gp}} = \langle p_{gp} - p_f \rangle_{p_{gp} > p_f}$. Thus,

$$\langle |p_f - p_{gp}| \rangle = \langle p_f - p_{gp} \rangle_{p_f > p_{gp}} = \frac{1}{\sqrt{2\pi\tau'}} \int_0^\infty x \cdot \exp \left[-\frac{x^2}{2\tau'^2} \right] dx = \frac{\tau}{\sqrt{\pi}} = \frac{1}{\sqrt{\pi}} \sqrt{\frac{a^L((a-1) \cdot L - (a-1))}{2((a-1) \cdot L + 1)^2}}, \quad (22)$$

619 where $\tau'^2 = 2\tau^2$.

620 So, for $a = 2$, as in our biallelic landscapes,

$$\lim_{L \rightarrow \infty} \langle |p_f - p_{gp}| \rangle = \sqrt{\frac{2^L \cdot (L-1)}{2\pi(L+1)^2}}. \quad (23)$$

Proofs

Proof 1: Selection for w_{opt} cannot transform a simple sign epistasis motif into a reciprocal sign epistasis motif

In order to have a reciprocal sign epistasis motif:

- No two adjacent arrows can point in the same direction.
- No two parallel arrows can point in the same direction.

Because sign epistasis motifs cannot be cyclic, the two conditions above imply one another. The simple sign epistasis motif determines the complete rank ordering of the phenotypic values, unlike the other two motifs, which only give a partial ordering. Therefore, upon selecting for any w_{opt} , at least two arrows in the genotype-phenotype landscape remain the same as in the fitness landscape or reverse in direction together. This results in a pattern of arrows in the fitness landscape that always contradicts both the above mentioned conditions for reciprocal sign epistasis (see Supplementary Fig. 1). ■

Proof 2: The phenotype-fitness map transforms a reciprocal sign epistasis motif into a no sign epistasis motif or a simple sign epistasis motif with equal probability

The reciprocal sign epistasis motif does not determine the complete rank ordering of the four phenotypic values. In particular, it does not determine the ordering of the two peaks and the two valleys. Let the phenotypic value of the higher of the two peaks be w_{p1} and that of the lower peak be w_{p2} . Similarly, let the labels w_{v1} and w_{v2} refer to the higher and lower valleys, respectively. By design, these values lie on the real number line between 0 and 1. We define the “neighbourhood” of each of the phenotypic values as the region on the number line around the phenotypic value that is closer to it than to any other phenotypic value (Supplementary Fig. 2). For transformation of the reciprocal sign epistasis motif into one of the other two motifs, w_{opt} must lie in the neighbourhood of either w_{p2} or w_{v1} and it must be second closest to the other (i.e., $F_{p2} > F_{v1} > \text{rest}$ or $F_{v1} > F_{p2} > \text{rest}$, where F_i represents the fitness corresponding to the phenotype w_i). If not, the reciprocal sign epistasis motif will be retained. So, for any w_{opt} , we only consider those configurations that satisfy this condition. Amongst these configurations, both phenotypic values are equally likely to be the closest to w_{opt} . Let’s call the probability of either event happening $P_0 = P(F_{v1} > F_{p2} > \text{rest}) = P(F_{p2} > F_{v1} > \text{rest})$. We denote the transition probabilities amongst motifs as $P(x|rs)$, where x represents either the no sign epistasis motif, ns , or the simple sign epistasis motif, ss , and rs represents the reciprocal sign epistasis motif. Now, for the no sign epistasis motif to emerge, the fitness values of the peaks or valleys need to be “separated”, while for the simple sign epistasis motif to emerge, they need to be “interspersed” (see Supplementary Fig. 2), giving us the following:

$$P(ns|rs) = P(F_{v1} > F_{p2} > \text{rest}) \cdot P(F_{p1} > F_{v2}) + P(F_{p2} > F_{v1} > \text{rest}) \cdot P(F_{p1} < F_{v2}), \quad (24)$$

649 and

$$P(ss|rs) = P(F_{v1} > F_{p2} > \text{rest}) \cdot P(F_{p1} < F_{v2}) + P(F_{p2} > F_{v1} > \text{rest}) \cdot P(F_{p1} > F_{v2}). \quad (25)$$

650 Finally, using $P_0 = P(F_{v1} > F_{p2} > \text{rest}) = P(F_{p2} > F_{v1} > \text{rest})$ gives:

$$P(ns|rs) = P_0 \cdot P(F_{p1} > F_{v2}) + P_0 \cdot P(F_{p1} < F_{v2}) \text{ and } P(ss|rs) = P_0 \cdot P(F_{p1} < F_{v2}) + P_0 \cdot P(F_{p1} > F_{v2}) \implies P(ns|rs) = P(ss|rs). \blacksquare \quad (26)$$

651 **Proof 3:** For House-of-Cards genotype-phenotype landscapes, selection for low or intermediate phenotypic values is equally
652 likely to increase or decrease the number of peaks in the fitness landscape, relative to the genotype-phenotype landscape

653 For any number of loci L , the probability that selection for a low or intermediate phenotypic value causes an increase in the
654 number of peaks in the fitness landscape, relative to the genotype-phenotype landscape, is

$$P_{\text{inc}} = \sum_{i>j} P(p_f = i | p_{\text{gp}} = j) \cdot P(p_{\text{gp}} = j), \quad (27)$$

655 whereas the probability of a decrease in the number of peaks is

$$P_{\text{dec}} = \sum_{i>j} P(p_f = j | p_{\text{gp}} = i) \cdot P(p_{\text{gp}} = i), \quad (28)$$

656 where $i, j \in \{1, 2, \dots, 2^{L-1}\}$, $P(p_f = j | p_{\text{gp}} = i)$ is the probability of having j peaks in the fitness landscape after selection for w_{opt}
657 given the genotype-phenotype landscape has i peaks, and $P(p_{\text{gp}} = i)$ is the probability of i peaks in the genotype-phenotype
658 landscape. The number of peaks in the fitness landscape, p_f , is independent of the number of peaks in the genotype-phenotype
659 landscape, p_{gp} , because p_f depends upon the number of genotypes with phenotypes close to the optimal phenotypic value, and
660 this number is independent of the number of genotypes with phenotypes close to 1. This implies:

$$P(p_f = i | p_{\text{gp}} = j) \cdot P(p_{\text{gp}} = j) = P(p_f = i) \cdot P(p_{\text{gp}} = j), \quad (29)$$

661 and

$$P(p_f = j | p_{\text{gp}} = i) \cdot P(p_{\text{gp}} = i) = P(p_f = j) \cdot P(p_{\text{gp}} = i). \quad (30)$$

662 Moreover, p_{gp} and p_f are identically distributed; i.e., $P(p_{\text{gp}} = i) = P(p_f = i)$

$$\implies P_{\text{inc}} = P_{\text{dec}} \blacksquare \quad (31)$$

Notes

Note 1: Number of maxima in the plot of $\langle |p_f - p_{gp}| \rangle$ versus w_{opt} for additive biallelic landscapes

Any biallelic, additive landscape can be relabelled such that the genotype with the lowest phenotypic value is represented by 000.....0 and the genotype with the highest phenotypic value by 111.....1. Moreover, the phenotypic values can be re-scaled such that genotype 000.....0 has phenotypic value zero and genotype 111.....1 has phenotypic value one. This yields a monotonically increasing genotype-phenotype landscape, wherein the phenotypic effects of every single mutant is larger than zero and they add up to 1; i.e.,

$$w_{100...0} + w_{010...0} + \dots + w_{000...1} = 1 \quad (32)$$

where $w_{100...0}$ is the phenotypic value of a single mutant with the mutation at the 1st site, and so on. Equation 32 implies that the vector of single mutant phenotypes belongs to a standard $L - 1$ -simplex[‡].

Two-locus case

A two-locus biallelic fitness landscape can have a maximum of two peaks. Because the additive genotype-phenotype landscape has only one peak, selection for a low or intermediate phenotypic value can increase the number of peaks in the fitness landscape, relative to the genotype-phenotype landscape, by at most one. Let the fitness of one of the single mutants be x , forcing the fitness of the other single-mutant to be $1 - x$. Due to the relabelling of the genotype-phenotype landscape, a two-peaked fitness landscape is only possible when the two single-mutants are peaks. This will occur when the two single-mutants have phenotypes that are closer to w_{opt} than the phenotypes of both genotypes 00 and 11. This leads to the following conditions on x :

$$1 - 2 \cdot w_{opt} < x < 2 \cdot w_{opt}, \forall w_{opt} < 0.5 \quad (33)$$

and

$$2 \cdot w_{opt} - 1 < x < 2 - 2 \cdot w_{opt}, \forall w_{opt} > 0.5. \quad (34)$$

From Eqs. 33 and 34, we can infer the shape of the $L = 2$ curves in Figures 4A,B. Specifically, for $w_{opt} < 0.25$ and $w_{opt} > 0.75$, we get the absurd condition $0.5 < x < 0.5$, which implies there is no possible value of x that will lead to two peaks. Further, for $w_{opt} = 0.5$, $0 < x < 1$, which implies there will always be two peaks. For $0.25 \leq w_{opt} < 0.5$ or $0.5 < w_{opt} \leq 0.75$, the fitness landscape can have either one or two peaks.

Three-locus case

[‡]An n dimensional simplex is the convex hull of $n + 1$ affinely independent points. E.g., a 0-simplex is a point, a 1-simplex is a line, a 2-simplex is a triangle, and a 3-simplex is a tetrahedron.

Here, the single-mutant phenotypic values belong to the 2-simplex formed by the points (0,0,1), (0,1,0) and (1,0,0). Let (w_1, w_2, w_3) be an arbitrary point in the simplex, where w_i is the phenotypic value of the single mutant with a mutation at the i th site and $i \in \{1, 2, 3\}$. Although, a three-locus, biallelic landscape can have at most four peaks, the transformed additive genotype-phenotype landscapes can have at most three, due to the correlations between the phenotypic values. So for instance, if the three single mutants become peaks, then the triple mutant cannot be a peak because the double mutants will be assigned a higher fitness than the triple mutant. This is caused by correlations in the additive landscape, which makes the double-mutants closer in phenotypic value to the single-mutants than the triple mutant. Thus, it is easy to see that the only three peak configuration possible is one in which either all the single mutants become peaks or all the double mutants become peaks. For the former to be the case, each single mutant phenotypic value w_i must satisfy the following conditions:

$$(w_i - w_{\text{opt}})^2 < (0 - w_{\text{opt}})^2 \quad (35)$$

and

$$(w_i - w_{\text{opt}})^2 < (1 - w_j - w_{\text{opt}})^2 \quad \forall j \neq i. \quad (36)$$

While Eq. 35 necessitates the single mutant fitness values to be higher than the wild type fitness, Eq. 36 requires them to be higher than the neighbouring double mutant fitness values. On average, $w_i = 1/3 \forall i$, so on average, the simplex area enclosed by the above conditions will be maximised when $w_{\text{opt}} = 1/3$. A similar analysis can be done for the case when the three double mutants become peaks and in this case, the area is maximised when $w_{\text{opt}} = 2/3$. This explains the presence of the two maxima in Fig. 4A at $w_{\text{opt}} = 1/3$ and $w_{\text{opt}} = 2/3$. Further, the dip at $w_{\text{opt}} = 0.5$ can be explained by the following argument: Assuming the single mutant phenotypic values are of the type $(x, y, 1 - x - y)$, the double mutant phenotypic value will be of the type $(x + y, 1 - x, 1 - y)$. At $w_{\text{opt}} = 0.5$, the wild type and triple mutant immediately become minima and leave an opportunity for the intermediate mutants to become peaks. However, if one of the single mutants (say with fitness x) is closest in fitness to 0.5, the other two cannot be peaks because the double mutant formed by combining these two mutations will have the same fitness i.e., $|x - 0.5| = |(1 - x) - 0.5|$. So at $w_{\text{opt}} = 0.5$, there are two peaks after selection for a low phenotypic value and therefore the net change is 1, which is smaller than what we see at $w_{\text{opt}} = 1/3, 2/3$.

Four and more locus cases

We could explicitly derive the shape of the curves for the $L = 2$ and $L = 3$ cases because it was easy to visualise the simplices. For $L \geq 4$, we can form an intuition about the expected result, by considering the average additive fitness landscape, which guides the behavior of the average change in the number of peaks. Because the only randomness in these landscapes arises from the effect of single mutations, which as described above, belong to the L -dimensional simplex, all we need to look at is the average L -dimensional vector of single mutation effects $\langle \vec{w}_1^L \rangle$ and the deviations around it. These random vectors belong to an L -dimensional Dirichlet distribution with all the concentration parameters equal to one. It can be shown by symmetry

714 that $\langle \vec{w}_1^L \rangle = (1/L, 1/L, \dots, 1/L)$. Therefore, on average, the fitness of a mutant with n mutations is n/L . As w_{opt} increases from
715 zero, the number of peaks will gradually start to increase as w_{opt} gets closer the value of the fitness of single mutants. Then as
716 w_{opt} approaches $1/L$, we encounter the first maxima because at this w_{opt} , all single mutants get to be peaks. Next, there is a
717 slight decline due to competition between the single and double mutants as peaks, followed by the emergence of the second
718 maxima at $2/L$ when all the double mutants become peaks, and so on. Naturally, because the number of genotypes with n
719 mutations increases as n reaches the value $L/2$, the height of the maxima increases accordingly. Thereafter, we observe a fall in
720 $\langle |p_f - p_{\text{gp}}| \rangle$. Therefore, for a genotype of length L , there are $L - 1$ maxima in the plot of $\langle |p_f - p_{\text{gp}}| \rangle$ vs. w_{opt} . However, the
721 maxima are close together and therefore difficult to visually discern.

Supplementary Figure 1. A simple sign epistasis motif (shown in brown) cannot be transformed into a reciprocal sign epistasis motif, for any w_{opt} . The no sign epistasis motif is shown in blue. The remaining colours represent the neighbourhood of the phenotypic values (w_i) corresponding to each genotype (g_i), where $i \in \{1, 2, 3, 4\}$ and the genotypes are labelled in descending order, such that the genotype with the highest phenotypic value is labelled 1, and so on. The four bottom arrows point to the transformed motifs when w_{opt} belongs to the neighbourhood from which the arrow emanates.

Supplementary Figure 2. Selection for w_{opt} transforms the reciprocal sign epistasis motif (shown in red) into the no sign epistasis motif (shown in blue) and the simple sign epistasis motif (shown in brown) with equal probability. The neighbourhood of each phenotypic value is shown in a different colour. For the no sign epistasis motif to emerge (top), the fitness values need to be "separated", while for the simple sign epistasis motif to emerge (bottom), the fitness values need to be "interspersed".

Supplementary Figure 3. Selection for $w_{\text{opt}} = 0$ does not change the type of epistasis motif for any “square” in the fitness landscape, relative to the genotype-phenotype landscape, yet it can change the number and location of peaks. To understand how, consider that any two adjacent faces of the hypercube (e.g., gray faces above) are sufficient to determine whether the genotypes on their common edge are peaks or not. After selecting for $w_{\text{opt}} = 0$, the type of epistasis motif does not change in the adjacent faces, yet the number and location of the peaks on their common edge does change. Peak genotypes are shown in red. Arrows point from lower to higher phenotypic or fitness values.

Supplementary Figure 4. Mean absolute change in the number of peaks in the fitness landscape relative to the genotype-phenotype landscape, shown in relation to w_{opt} , for (A) single-peaked vs. multi-peaked genotype-phenotype landscapes and (B) three values of σ .

Supplementary Figure 5. Average peak (A) height and (B) width for 1,137 empirical landscapes, shown in relation to the optimal binding affinity w_{opt} . Violin plots show the distribution across the landscapes for each w_{opt} . Data include both local and global peaks.

Supplementary Figure 6. The average height of peaks reached by greedy adaptive walks in NK landscapes with $\sigma = 0.5$, shown in relation to w_{opt} , for (A) $L = 5$ and (B) $L = 8$, with $K = 0 \dots L - 1$.

Supplementary Figure 7. Change in mean fitness at equilibrium for NK landscapes with (A,B) $L=5$ and (C,D) $L=8$, shown in relation to w_{opt} , for different values of σ and μ .

Supplementary Figure 8. (A) Local peaks on which the adaptive walks terminate tend to be nearly as tall as the global peaks in the 1,137 empirical landscapes. Violin plots show the distribution of the fractional height of local peaks reached by greedy adaptive walks, relative to the height of the global peak, for each optimal binding affinity w_{opt} . (B) Violin plots show the distribution of the fraction of walks terminating on the global peak, for each optimal binding affinity w_{opt} .

Supplementary Figure 9. The mismatch class m_{opt} that maximizes mean fitness at equilibrium is shown in relation to the strength of selection σ and the mutation rate μ . The three smaller panels show mean fitness at equilibrium in relation to m_{opt} for three combinations of σ and μ . The value of m_{opt} that maximizes mean fitness is indicated with a gray rectangle.

Supplementary Figure 10. The mismatch class m_{opt} that maximizes mean fitness at equilibrium is shown in relation to the strength of selection σ and the mutation rate μ for three different values of the mismatch penalty e : (A) $e = 0.025$, (B) $e = 0.05$ and (C) $e = 0.1$.

Supplementary Figure 11. The binding affinity w_{opt} that maximizes mean fitness at equilibrium is shown in relation to the strength of selection σ and the logarithm of the mutation rate ($\log \mu$) for the 1,137 empirical landscapes. The three smaller panels show the distributions of mean fitness at equilibrium as violin plots, in relation to w_{opt} , for three combinations of σ and μ . Box-and-whisker plots show the 25-75% quartiles. The value of w_{opt} that maximizes mean fitness is indicated. These results are robust to perturbations of the fitness values (Methods), with the exception of two parameter combinations – $\sigma = 0.1, \mu = 0.001$ and $\sigma = 0.1, \mu = 0.0025$. For these parameter combinations, the binding affinity w_{opt} that maximizes mean fitness at equilibrium changes from 0.35 to 0.385 (on average) and 0.375 respectively.

Acknowledgements

This work was supported by the Swiss National Science Foundation (Grant numbers PP00P3_170604 and 310030_192541). We thank members of the Computational Biology Group at ETH, Michael Manhart, Joachim Krug, and David M. McCandlish for discussions.

References

1. Alberch, P. From genes to phenotype: dynamical systems and evolvability. *Genetica* **84**, 5–11 (1991).
2. Lehner, B. Genotype to phenotype: lessons from model organisms for human genetics. *Nat. Rev. Genet.* **14**, 168–178 (2013).
3. Julien, P., Miñana, B., Baeza-Centurion, P., Valcárcel, J. & Lehner, B. The complete local genotype–phenotype landscape for the alternative splicing of a human exon. *Nat. Commun.* **7** (2016).
4. Payne, J. L. & Wagner, A. The causes of evolvability and their evolution. *Nat. Rev. Genet.* **20**, 24–38 (2019).
5. Poelwijk, F., Tănase-Nicola, S., Kiviet, D. & Tans, S. Reciprocal sign epistasis is a necessary condition for multi-peaked fitness landscapes. *J. Theor. Biol.* **272**, 141–4 (2011).
6. Østman, B., Hintze, A. & Adami, C. Impact of epistasis and pleiotropy on evolutionary adaptation. *Proc. Royal Soc. B: Biol. Sci.* **279**, 247–256 (2012).
7. Weinreich, D., Watson, R. & Chao, L. Perspective: sign epistasis and genetic constraint on evolutionary trajectories. *Evolution* **59**, 1165–1174 (2005).
8. Kinney, J. B. & McCandlish, D. M. Massively parallel assays and quantitative sequence–function relationships. *Annu. Rev. Genomics Hum. Genet.* **20**, 99–127 (2019).
9. Qiu, C. *et al.* High-resolution phenotypic landscape of the RNA polymerase II trigger loop. *PLoS Genet.* **12** (2016).
10. Wu, N., Dai, C. L., Olson, C. A., Lloyd-Smith, J. O. & Sun, R. Adaptation in protein fitness landscapes is facilitated by indirect paths. *eLife* **5**, e16965 (2016).
11. Tack, D. *et al.* The genotype-phenotype landscape of an allosteric protein. *Mol. Syst. Biol.* **17** (2021).
12. Sarkisyan, K. S. *et al.* Local fitness landscape of the green fluorescent protein. *Nature* **533**, 397–401 (2016).
13. de Boer, C. G. *et al.* Deciphering eukaryotic gene-regulatory logic with 100 million random promoters. *Nat. Biotechnol.* **38**, 56–65 (2019).
14. Lagator, M., Paixao, T., Barton, N. H., Bollback, J. P. & Guet, C. C. On the mechanistic nature of epistasis in a canonical cis-regulatory element. *eLife* **6**, e25192 (2017).

- 750 **15.** Schaerli, Y. *et al.* Synthetic circuits reveal how mechanisms of gene regulatory networks constrain evolution. *Mol. Syst.*
751 *Biol.* **14** (2018).
- 752 **16.** Li, X., Lalic, J., Baeza-Centurion, P., Dhar, R. & Lehner, B. Changes in gene expression predictably shift and switch
753 genetic interactions. *Nat. Commun.* **10** (2019).
- 754 **17.** Bassalo, M. C. *et al.* Deep scanning lysine metabolism in escherichia coli. *Mol. Syst. Biol.* **14** (2018).
- 755 **18.** Schenk, M. F., Szendro, M. L., Salverda, I. G., Krug, J. & de Visser, J. A. G. Patterns of epistasis between beneficial
756 mutations in an antibiotic resistance gene. *Mol. Biol. Evol.* **30**, 1779–87 (2013).
- 757 **19.** Gong, L. I., Suchard, M. A. & Bloom, J. D. Stability-mediated epistasis constrains the evolution of an influenza protein.
758 *eLife* **2**, e00631 (2013).
- 759 **20.** Jiménez, J. I., Xulvi-Brunet, R., Campbell, G. W., Turk-MacLeod, R. & Chen, I. A. Comprehensive fitness landscape for
760 RNA. *Proc. Natl. Acad. Sci.* **110**, 376–379 (2013).
- 761 **21.** Aguilar-Rodríguez, J., Payne, J. & Wagner, A. A thousand empirical adaptive landscapes and their navigability. *Nat. Ecol.*
762 *Evol.* **1** (2017).
- 763 **22.** Olson, C. A., Wu, N. C. & Sun, R. A comprehensive biophysical description of pairwise epistasis throughout an entire
764 protein domain. *Curr. Biol.* **24**, 2643–2651 (2014).
- 765 **23.** Hartman, E. C. *et al.* Experimental evaluation of coevolution in a self-assembling particle. *Biochemistry* **58**, 1527–1538
766 (2019).
- 767 **24.** Pitt, J. & Ferré-D'Amaré, A. Rapid construction of empirical RNA fitness landscapes. *Science* **330**, 376–379 (2010).
- 768 **25.** Wright, S. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. *Proc. 6th Int. Congr. Genet.*
769 356–366 (1932).
- 770 **26.** Romero, P. A. & Arnold, F. H. Exploring protein fitness landscapes by directed evolution. *Nat. Rev. Mol. Cell Biol.* **10**
771 (2009).
- 772 **27.** Kinsler, G., Geiler-Samerotte, K. & Petrov, D. Fitness variation across subtle environmental perturbations reveals local
773 modularity and global pleiotropy of adaptation. *eLife* **9**, e61271 (2020).
- 774 **28.** Dekel, E. & Alon, U. Optimality and evolutionary tuning of the expression level of a protein. *Nature* **436**, 588–592
775 (2005).
- 776 **29.** Lynch, M. & Hagner, K. Evolutionary meandering of intermolecular interactions along the drift barrier. *Proc. Natl. Acad.*
777 *Sci.* **112** (2015).
- 778 **30.** Rokyta, D. R. *et al.* Epistasis between beneficial mutations and the phenotype-to-fitness map for a ssDNA virus. *PLoS*
779 *Genet.* **7** (2011).

- 780 **31.** Duveau, F., Toubiana, W. & Wittkopp, P. J. Fitness effects of cis-regulatory variants in the *saccharomyces cerevisiae*
781 TDH3 promoter. *Mol. Biol. Evol.* **34**, 2908–2912 (2017).
- 782 **32.** Rest, J. S. *et al.* Nonlinear Fitness Consequences of Variation in Expression Level of a Eukaryotic Gene. *Mol. Biol. Evol.*
783 **30**, 448–456 (2012).
- 784 **33.** Keren, L. *et al.* Massively parallel interrogation of the effects of gene expression levels on fitness. *Cell* **166**, 1–13 (2016).
- 785 **34.** Labourel, F. & Rajon, E. Resource uptake and the evolution of moderately efficient enzymes. *Mol. Biol. Evol.* **38**,
786 3938–3952 (2021).
- 787 **35.** Perfeito, L., Ghozzi, S., Berg, J., Schnetz, K. & Lässig, M. Nonlinear fitness landscape of a molecular pathway. *PLoS*
788 *Genet.* **7** (2011).
- 789 **36.** Otwinowski, J., McCandlish, D. & Plotkin, J. Inferring the shape of global epistasis. *Proc. Natl. Acad. Sci.* **115**,
790 E7550–E7558 (2018).
- 791 **37.** Domingo Espinós, J., Baeza-Centurion, P. & Lehner, B. The causes and consequences of genetic interactions (epistasis).
792 *Annu. Rev. Genomics Hum. Genet.* **20**, 17.1–17.28 (2019).
- 793 **38.** Kryazhimskiy, S. Emergence and propagation of epistasis in metabolic networks. *eLife* **10**, e60200 (2021).
- 794 **39.** Crocker, J., Preger-Ben Noon, E. & Stern, D. The soft touch: Low-affinity transcription factor binding sites in development
795 and evolution. *Curr. Top. Dev. Biol.* **117**, 455–469 (2015).
- 796 **40.** Haberle, V. & Stark, A. Eukaryotic core promoters and the functional basis of transcription initiation. *Nat. Rev. Mol. Cell*
797 *Biol.* **19**, 621–637 (2018).
- 798 **41.** Shultzaberger, R., Malashock, D., Kirsch, J. & Eisen, M. The fitness landscapes of cis-acting binding sites in different
799 promoter and environmental contexts. *PLoS Genet.* **6** (2010).
- 800 **42.** Sharon, E. *et al.* Inferring gene regulatory logic from high-throughput measurements of thousands of systematically
801 designed promoters. *Nat. Biotechnol.* **30**, 521–530 (2012).
- 802 **43.** Mustonen, V., Kinney, C., J. and Callan & Lässig, M. Energy-dependent fitness: a quantitative model for the evolution of
803 yeast transcription factor binding sites. *Proc. Natl. Acad. Sci.* **105**, 12376–12381 (2008).
- 804 **44.** Haldane, A., Manhart, M. & Morozov, A. Biophysical fitness landscapes for transcription factor binding sites. *PLoS*
805 *Comput. Biol.* **10**, 1–16 (2014).
- 806 **45.** Gerland, U. & Hwa, T. On the selection and evolution of regulatory DNA motifs. *J. Mol. Evol.* **55**, 386–400 (2002).
- 807 **46.** Berg, J., Willmann, S. & Lässig, M. Adaptive evolution of transcription factor binding sites. *BMC Evol. Biol.* **4** (2004).
- 808 **47.** Tuğrul, M., Paixão, T., Barton, N., Tkačik, G. & Fay, J. Dynamics of transcription factor binding site evolution. *PLoS*
809 *Genet.* **11**, 273–283 (2015).

- 810 **48.** Weirauch, M. *et al.* Determination and inference of eukaryotic transcription factor sequence specificity. *Cell* **158**,
811 1431–1443 (2014).
- 812 **49.** Crocker, J. *et al.* Low affinity binding site clusters confer HOX specificity and regulatory robustness. *Cell* **160**, 191–203
813 (2015).
- 814 **50.** Kribelbauer, J. F., Rastogi, C., Bussemaker, H. J. & Mann, R. S. Low-affinity binding sites and the transcription factor
815 specificity paradox in eukaryotes. *Annu. Rev. Cell Dev. Biol.* **35**, 357–379 (2019).
- 816 **51.** Grönlund, A., Lötstedt, P. & Elf, J. Transcription factor binding kinetics constrain noise suppression via negative feedback.
817 *Nat. Commun.* **4** (2013).
- 818 **52.** Farley, E., Olson, K., Zhang, W. & et al. Suboptimization of developmental enhancers. *Science* **350**, 325–28 (2015).
- 819 **53.** Farley, E., Olson, K., Zhang, W., Rokhsar, D. & Levine, M. Syntax compensates for poor binding sites to encode tissue
820 specificity of developmental enhancers. *Proc. Natl. Acad. Sci.* **113**, 6508–6513 (2016).
- 821 **54.** Tanay, A. Extensive low-affinity transcriptional interactions in the yeast genome. *Genome Res.* 962–972 (2006).
- 822 **55.** Ramos, A. & Barolo, S. Low-affinity transcription factor binding sites shape morphogen responses and enhancer evolution.
823 *Philos. Transactions Royal Soc. B* **368** (2013).
- 824 **56.** Tsai, A. *et al.* Nuclear microenvironments modulate transcription from low-affinity enhancers. *eLife* **6**, e28975 (2017).
- 825 **57.** Cary, G. A., Jarvela, A. M. C., Francolini, R. D. & Hinman, V. F. Genome-wide use of high- and low-affinity Tbrain
826 transcription factor binding sites during echinoderm development. *Proc. Natl. Acad. Sci.* **114**, 5854–5861 (2017).
- 827 **58.** Wang, J., Malecka, A., Trøen, G. & Delabie, J. Comprehensive genome-wide transcription factor analysis reveals that
828 a combination of high affinity and low affinity DNA binding is needed for human gene regulation. *BMC genomics* **16**
829 (2015).
- 830 **59.** Clarke, B. The evolution of genetic diversity. *Proc. Royal Soc. London. Ser. B, Biol. Sci.* **205**, 453–474 (1979).
- 831 **60.** Gavrillets, S. Fitness landscapes and the origin of species. *Princet. Univ. Press.* (2004).
- 832 **61.** Misevic, D., Kouyos, R. D. & Bonhoeffer, S. Predicting the evolution of sex on complex fitness landscapes. *PLoS Comput.*
833 *Biol.* **5** (2009).
- 834 **62.** de Visser, J. A. G. M. & Krug, J. Empirical fitness landscapes and the predictability of evolution. *Nat. Rev. Genet.* **15**,
835 480–490 (2014).
- 836 **63.** Aita, Y., Takuyo anf Husimib. Fitness spectrum among random mutants on mt. fuji type fitness landscape. *J. Theor. Biol.*
837 **182**, 469–485 (1996).
- 838 **64.** Kingman, J. On the properties of bilinear models for the balance between genetic mutation and selection. *Math. Proc.*
839 *Camb. Philos. Soc.* **81**, 443–453 (1977).

- 840 **65.** Kauffman, S. A. & Weinberger, E. D. The NK model of rugged fitness landscapes and its application to the maturation of
841 the immune response. *J. Theor. Biol.* **141**, 211–245 (1989).
- 842 **66.** Berg, O. G. & von Hippel, P. H. Selection of DNA binding sites by regulatory proteins. *J. Mol. Biol.* **193**, 723–750
843 (1987).
- 844 **67.** Wrenbeck, E., Azouz, L. & Whitehead, T. Single-mutation fitness landscapes for an enzyme on multiple substrates reveal
845 specificity is globally encoded. *Nat. Commun.* **8** (2017).
- 846 **68.** Li, C. & Zhang, J. Multi-environment fitness landscapes of a tRNA gene. *Nat. Ecol. Evol.* **2**, 1025–1032 (2018).
- 847 **69.** Chan, Y., Venev, S., Zeldovich, K. & Matthews, C. Correlation of fitness landscapes from three orthologous TIM barrels
848 originates from sequence and structure constraints. *Nat. Commun.* **8** (2017).
- 849 **70.** Somermeyer, L. *et al.* Heterogeneity of the gfp fitness landscape and data-driven protein design. *bioRxiv* (2021).
- 850 **71.** Bajić, D., Vila, J., Blount, Z. & Sanchez, A. On the deformability of an empirical fitness landscape by microbial evolution.
851 *Proc. Natl. Acad. Sci.* **115**, 201808485 (2018).
- 852 **72.** Poelwijk, F., Kiviet, D., Weinreich, D. & Tans, S. Empirical fitness landscapes reveal accessible evolutionary paths.
853 *Nature* **445**, 383–386 (2007).
- 854 **73.** Fisher, R. A. The genetical theory of natural selection. *Clarendon Press. Oxford.* (1930).
- 855 **74.** Tenaillon, O. The utility of fisher’s geometric model in evolutionary genetics. *Annu. Rev. Ecol. Evol. Syst.* **45**, 179–201
856 (2014).
- 857 **75.** Dieckmann, U. & Doebeli, M. On the origin of species by sympatric speciation. *Nature* **400**, 354–357 (1999).
- 858 **76.** Gao, H., Granka, J. M. & Feldman, M. W. On the classification of epistatic interactions. *Genetics* **184**, 827–837 (2010).
- 859 **77.** Hwang, S., Park, S.-C. & Krug, J. Genotypic complexity of fisher’s geometric model. *Genetics* **206**, 1049–1079 (2017).
- 860 **78.** Hordijk, W. & Stadler, P. Amplitude spectra of fitness landscapes. *Adv. Complex Syst.* **1**, 39–66 (1998).
- 861 **79.** Macken, C. A. & Perelson, A. S. Protein evolution on rugged landscapes. *Proc. Natl. Acad. Sci.* **86**, 6191–6195, (1989).
- 862 **80.** Badis-Breard, G. *et al.* Diversity and complexity in DNA recognition by transcription factors. *Science* **324**, 1720–1723
863 (2009).
- 864 **81.** Berger, M. *et al.* Compact, universal DNA microarrays to comprehensively determine transcription-factor binding site
865 specificities. *Nat. Biotechnol.* **24**, 1429–35 (2006).
- 866 **82.** Payne, J. & Wagner, A. The robustness and evolvability of transcription factor binding sites. *Science* **343**, 875–877
867 (2014).
- 868 **83.** Caswell, H. *Matrix population models* (Sinauer, Sunderland, Mass., 2001).

- 869 **84.** Wilke, C., Wang, J., Ofria, C., Lenski, R. & Adami, C. Evolution of digital organisms at high mutation rates leads to
870 survival of the flattest. *Nature* **412**, 331–333 (2001).
- 871 **85.** Aguilar-Rodriguez, J., Peel, L., Stella, M., Wagner, A. & Payne, J. The architecture of an empirical genotype-phenotype
872 map. *Evolution* **72**, 1242–1260 (2018).
- 873 **86.** Franke, J., Klozer, A., de Visser, J. & Krug, J. Evolutionary accessibility of mutational pathways. *PLoS Comput. Biol.* **7**
874 (2011).
- 875 **87.** Chou, H.-H., Delaney, N., Draghi, J. & Marx, C. Mapping the fitness landscape of gene expression uncovers the cause of
876 antagonism and sign epistasis between adaptive mutations. *PLoS genetics* **10** (2014).
- 877 **88.** Kemble, H. *et al.* Flux, toxicity, and expression costs generate complex genetic interactions in a metabolic pathway. *Sci.*
878 *Adv.* **6** (2020).
- 879 **89.** He, X., Duque, T. & Sinha, S. Evolutionary origins of transcription factor binding site clusters. *Mol. Biol. Evol.* **29**,
880 1059–1070 (2011).
- 881 **90.** Dingle, K., Ghaddar, F., Sulc, P. & Louis, A. Phenotype bias determines how natural RNA structures occupy the
882 morphospace of all possible shapes. *Mol. Biol. Evol.* **39** (2022).
- 883 **91.** Xiong, K., Gerstein, M. & Masel, J. Differences in evolutionary accessibility determine which equally effective regulatory
884 motif evolves to generate pulses. *Genetics* **219** (2021).
- 885 **92.** Moses, A. M. *et al.* Large-scale turnover of functional transcription factor binding sites in drosophila. *PLoS Comput. Biol.*
886 **2** (2006).
- 887 **93.** Villar, D., Flicek, P. & Odom, D. T. Evolution of transcription factor binding in metazoans — mechanisms and functional
888 implications. *Nat. Rev. Genet.* **15** (2014).
- 889 **94.** Weirauch, M. & Hughes, T. Conserved expression without conserved regulatory sequence: The more things change, the
890 more they stay the same. *Trends Genet.* **26**, 66–74 (2010).
- 891 **95.** Bintu, L. *et al.* Transcriptional regulation by the numbers: Models. *Curr. Opin. Genet. Dev.* **15**, 116–124 (2005).
- 892 **96.** Siggers, T. & Gordân, R. Protein–DNA binding: Complexities and multi-protein codes. *Nucleic Acids Res.* **42**, 1–13
893 (2013).
- 894 **97.** Dror, I., Golan, T., Levy, C., Rohs, R. & Mandel-gutfreund, Y. A widespread role of the motif environment on transcription
895 factor binding across diverse protein families. *Genome Res.* **25**, 1268–1280 (2015).
- 896 **98.** Xin, B. & Rohs, R. Relationship between histone modifications and transcription factor binding is protein family specific.
897 *Genome Res.* **28**, 321–333 (2018).
- 898 **99.** Yin, Y. *et al.* Impact of cytosine methylation on dna binding specificities of human transcription factors. *Science* **356**
899 (2017).

- 900 **100.** Tsai, A., Galupa, R. & Crocker, J. Robust and efficient gene regulation through localized nuclear microenvironments.
901 *Development* **147** (2020).
- 902 **101.** Duque, T. *et al.* Simulations of enhancer evolution provide mechanistic insights into gene regulation. *Mol. Biol. Evol.* **31**,
903 184–200 (2013).
- 904 **102.** Martin, G., Elena, S. & Lenormand, T. Distributions of epistasis in microbes fit predictions from a fitness landscape
905 model. *Nat. genetics* **39**, 555–60 (2007).
- 906 **103.** Roughgarden, J. Species packing and the competition function with illustrations from coral reef fish. *Theor. population*
907 *biology* **5**, 163–86 (1974).
- 908 **104.** Duveau, F. *et al.* Fitness effects of altering gene expression noise in *saccharomyces cerevisiae*. *eLife* **7**, e37272 (2018).
- 909 **105.** Vos, M. G. D., Dawid, A., Sunderlikova, V. & Tans, S. J. Breaking evolutionary constraint with a tradeoff ratchet. *Proc.*
910 *Natl. Acad. Sci.* **112**, 14906–14911 (2015).
- 911 **106.** Steinberg, B. & Ostermeier, M. Environmental changes bridge evolutionary valleys. *Sci. Adv.* **2** (2016).
- 912 **107.** Hayden, E. J. & Wagner, A. Environmental change exposes beneficial epistatic interactions in a catalytic RNA. *Proc.*
913 *Royal Soc. B: Biol. Sci.* **279**, 3418–3425 (2012).
- 914 **108.** Gorter, F. A., Aarts, M. G., Zwaan, B. J. & Visser, J. A. G. D. Local fitness landscapes predict yeast evolutionary dynamics
915 in directionally changing environments. *Genetics* **208**, 307–322 (2018).
- 916 **109.** Bank, C., Matuszewski, S., Hietpas, R. T. & Jensen, J. D. On the (un)predictability of a large intragenic fitness landscape.
917 *Proc. Natl. Acad. Sci.* **113**, 14085–14090 (2016).
- 918 **110.** Barrera, L. *et al.* Survey of variation in human transcription factors reveals prevalent DNA binding changes. *Science* **351**,
919 1450–1454 (2016).
- 920 **111.** Khalid, F., Aguilar-Rodriguez, J., Wagner, A. & Payne, J. Genonets server - A web server for the construction, analysis,
921 and visualization of genotype networks. *Nucleic Acids Res.* **44**, W70–W76 (2016).
- 922 **112.** Baldi, P. & Rinott, Y. Asymptotic normality of some graph-related statistics. *J. Appl. Probab.* **26**, 171–175 (1989).