Robustness and evolvability: a paradox resolved

Andreas Wagner^{1,2,3,*}

¹*Department of Biochemistry, University of Zurich, 8057 Zurich, Switzerland*

²*The Santa Fe Institute, Santa Fe, NM 87501, USA*

³*Swiss Institute of Bioinformatics, 4056 Basel, Switzerland*

Understanding the relationship between robustness and evolvability is key to understand how living things can withstand mutations, while producing ample variation that leads to evolutionary innovations. Mutational robustness and evolvability, a system's ability to produce heritable variation, harbour a paradoxical tension. On one hand, high robustness implies low production of heritable phenotypic variation. On the other hand, both experimental and computational analyses of neutral networks indicate that robustness enhances evolvability. I here resolve this tension using RNA genotypes and their secondary structure phenotypes as a study system. To resolve the tension, one must distinguish between robustness of a genotype and a phenotype. I confirm that genotype (sequence) robustness and evolvability share an antagonistic relationship. In stark contrast, phenotype (structure) robustness promotes structure evolvability. A consequence is that finite populations of sequences with a robust phenotype can access large amounts of phenotypic variation while spreading through a neutral network. Population-level processes and phenotypes rather than individual sequences are key to understand the relationship between robustness and evolvability. My observations may apply to other genetic systems where many connected genotypes produce the same phenotypes.

Keywords: evolvability; molecular evolution; neutral networks

1. INTRODUCTION

Robustness and evolvability are fundamental properties of biological systems. They determine a system's persistence and its potential for future evolutionary change. Multiple definitions of robustness and evolvability are in use (Houle 1992; Kirschner & Gerhart 1998; Bedau & Packard 2003; Hansen 2003; Masel & Bergman 2003; Michod *et al.* 2003; Nijhout *et al.* 2003; Earl & Deem 2004; Aharoni *et al.* 2005; Wagner 2005; Sniegowski & Murphy 2006). For my purpose, I call a biological system mutationally robust if its function or structure persists after mutations in its parts. The system is evolvable if mutations in it can produce heritable phenotypic variation. Both robustness and evolvability are quantitative properties. For instance, a system can be more or less robust in response to a given number of mutations. Both definitions apply to systems on different levels of organization, including RNA and protein molecules, small genetic circuits, genome-scale networks such as metabolic networks and even whole organisms. The appropriate nature of a mutation depends on the level of organization one focuses on. It may correspond to amino acid change for proteins, regulatory mutations in genetic circuits or changes in enzymatic reactions for a metabolic network. The appropriate notion of phenotype also depends on this level of organization. Examples include the spatial conformation or catalytic activity of a macromolecule, and a gene expression pattern for a genetic circuit.

At first sight, the above definitions imply an antagonistic relationship between robustness and evolvability: the more

robust a system is, the less phenotypic variation a given number of mutations generate and hence the less evolvable the system is (Ancel & Fontana 2000; Sumedha *et al.* 2007). However, from different perspectives, robustness and evolvability may go hand in hand. First, multiple experiments show that many organisms harbour cryptic genetic variation—reflecting robustness—which can become visible in certain environments or genetic backgrounds, and thus lead to enhanced evolvability (Waddington 1953; Dun & Fraser 1959; Waddington 1959; Rutherford & Lindquist 1998; Queitsch *et al.* 2002). Next, a biological system with a given phenotype typically has many alternative genotypes that can produce this phenotype (Wagner 2005). These genotypes are often connected; that is, they can be reached from each other through series of single mutations, as was shown first for RNA (Schuster *et al.* 1994). Importantly, this connectedness implies at least some degree of robustness, because for a typical genotype some mutations must leave the phenotype unchanged. This connectedness also suggests evolvability, because many *new* phenotypes might be produced by single mutations if a phenotype has many different genotypes that adopt it (Huynen 1996; Ciliberti *et al.* 2007b; Sumedha *et al.* 2007).

Thus from one perspective, robustness hinders evolvability. From another perspective, robustness promotes evolvability. I here show how this apparent paradoxical tension can be resolved. To this end, I use RNA secondary structure as a study system. RNA secondary structure is an important phenotype in its own right, because it is required for the biological function of many RNA molecules (Dayton *et al.* 1992; Baudin *et al.* 1993; Powell *et al.* 1995). Because computationally efficient algorithms exist to predict RNA secondary structure from an RNA sequence (Zuker & Sankoff 1984; Hofacker *et al.* 1994; Tacker *et al.* 1996), RNA secondary structure has proven

* Author and address for correspondence: Department of Biochemistry, University of Zurich, Building Y27-J-54, Winterthurerstrasse 190, 8057 Zurich, Switzerland (aw@bioc.uzh.ch).

Electronic supplementary material is available at <http://dx.doi.org/10.1098/rspb.2007.1137> or via <http://journals.royalsociety.org>.

to be an important computational model to understand how genetic variation maps into phenotypic (structural) variation (Schuster *et al.* 1994; Reidys *et al.* 1997; Fontana & Schuster 1998; Fontana 2002). The biological relevance and computational tractability of RNA structure allow one to explore the relationship between robustness and evolvability. The insights thus obtained may inform our thinking for many other biological systems that are less tractable.

One might argue that the above notion of evolvability, focusing only on variation is not sufficiently ambitious. A more expansive definition might involve the ability of a phenotype or genotype to produce new and desirable phenotypes, evolutionary innovations. In reality, evolutionary innovations are obvious only in hindsight. However, one can specify a closely related measure of evolvability in a computational analysis by asking how easily a blind random walk starting from a given phenotype can find a pre-defined but otherwise arbitrary 'target' phenotype. I show that this notion of evolvability is independent of robustness, and may depend very little on properties of the starting phenotype.

2. RESULTS

(a) *Genotype and phenotype robustness*

To resolve the tension between robustness and evolvability, I first recast definitions of robustness and evolvability in terms of RNA genotypes and phenotypes. The relevant genotype space is the set of all 4^n possible RNA sequences of some given length n . Two sequences are 1-mutant neighbours or simply neighbours in this space if they differ in one nucleotide. The relevant phenotypes are the set of all possible RNA structures, whose number scales approximately as 1.8^n (Gruner *et al.* 1996; Schuster 2003). The set of all genotypes forming the same structure is often called a neutral set or, when connected in genotype space, a neutral network (Schuster *et al.* 1994). A neutral neighbour of a sequence is a neighbour that has the same structure. I define the 1-neighbourhood of a *structure* or, equivalently, the 1-neighbourhood of a neutral network as follows: it is the set of sequences that differ from the sequences that fold into the structure by exactly one nucleotide. With these notations, I now introduce two different definitions of robustness and evolvability.

- (i) *Genotype (sequence) robustness.* The number R_G (or fraction r_G) of neutral neighbours of a genotype G .
- (ii) *Phenotype (structure) robustness.* The number R_P (or fraction r_P) of neutral neighbours averaged over *all* genotypes G with a given phenotype. I will refer to this quantity also as the number of neutral neighbours of the structure.
- (iii) *Genotype (sequence) evolvability.* The number E_G of different structures found in the 1-neighbourhood of a sequence G .
- (iv) *Phenotype (structure) evolvability.* The number E_P of different structures found in the 1-neighbourhood of a structure P .

Below, I will show that the tension between robustness and evolvability disappears with the distinction just introduced. Specifically, genotypic robustness is negatively associated with genotypic evolvability, whereas phenotypic robustness is positively associated with phenotypic

evolvability. Note that sequence robustness and structure robustness cannot be meaningfully compared to one another and neither can sequence evolvability and structure evolvability. Also note that the above definitions of evolvability are special cases of the more general definition in the introduction. A sequence that is more evolvable has more structural variants in its 1-neighbourhood. Similarly, a structure that is more evolvable has more structural variants in its 1-neighbourhood.

(b) *High sequence robustness means low sequence evolvability*

For reasons detailed in the methods (electronic supplementary material), I focus throughout on random RNA sequences of length $n = 30$, which strikes a balance between structural richness and computational tractability. Biological RNAs are currently not suitable for this analysis, owing to the necessity of analysing many different structures of the same length and the unsolved problem of estimating the abundance or *frequency* f of a given structure. This frequency is defined as the number of sequences adopting the structure. In protein engineering, it is also known as a structure's designability.

It is well known that the distribution of structure frequencies is highly skewed (Hofacker *et al.* 1998; Schuster 2003). That is, there are relatively few structures adopted by many sequences, but many structures adopted by few sequences. Figure 1a shows the distribution of structure frequencies for a sample of 10^6 RNA sequences. Structure frequencies in this sample vary over a factor 1600, from less than 1.06×10^{-6} (structures found only once) to 1.7×10^{-3} .

A sequence G 's mutational robustness R_G and the number of sequences in the 1-neighbourhood of G whose structure is different from that of G are trivially and inversely related, i.e. this number is equal to $3n - R_G$. But what is the relationship between R_G and the number of structures in the 1-neighbourhood of a sequence that are *unique*, i.e. different from both the structure of G and from each other? To find out, I determined the number U of unique structures that are different from one another in the neighbourhood of G . To obtain this number, I counted all structures in the neighbourhood, but counted structures only once that occurred twice or more. By the above definition of sequence evolvability, $E_G = U$. To normalize for sequence length, one can also consider the normalized robustness $r_G = R_G/3n$, and e_G , an appropriately normalized analogue of U . Starting from U , however, there are two different ways of normalization to obtain the *proportion* of unique structures. The first approach is to divide U by the number $3n(1 - r_G)$, i.e. by the total number of sequences in the neighbourhood of G that have structures different from the structure of G . The statistical association between r_G and $U/3n(1 - r_G)$ is weak, although highly significant in a large sample of sequences (Spearman's $s = -0.06$, $p < 10^{-17}$; sample size $n_s = 7.5 \times 10^4$). It is perhaps remarkable that the association is negative despite the fact that robustness exerts a positive influence on e_G merely through this normalization: the higher r_G , the smaller $(1 - r_G)$ and therefore the greater $U/3n(1 - r_G)$.

The second approach to normalize and obtain the proportion of unique structures from U is to divide U simply by $3n$, the total number of sequences in the neighbourhood of G . Arguably, this approach is more sensible, because it reflects the likelihood that a new structure is encountered in a blind evolutionary

Figure 1. (a) The number of sequences folding into one structure has a highly skewed distribution. Structures found in a random sample of 10^6 sequences were ranked according to their frequency, defined as the number of sequences in the sample that adopt a structure divided by 10^6 . The plot shows structure rank (x -axis) plotted against structure frequency (y -axis). Note the logarithmic scale on the y -axis. Structure frequencies vary by more than a factor 10^3 in this sample. (b) High genotype robustness implies low genotype evolvability. Data shown are based on 7.5×10^4 different RNA structures ($n = 30$ nucleotides) whose frequencies span three orders of magnitude, and on one RNA sequence inversely folded for each structure. Genotype robustness (r_G) and evolvability (e_G) were calculated for these inversely folded RNA sequences. Lengths of error bars indicate one standard error of the mean, calculated for each of the 20 bins of data grouped according to r_G . Bars are too short to be visible for most of the data points. Standard deviations were below 0.13 for all 20 bins. The dashed line indicates points where $e_G = 1 - r_G$. Note that $1 - r_G$ is the fraction of sequences in the 1-neighbourhood of a sequence G , that have an MFE secondary structure different from that of G .

exploration of the neighbourhood of G . Figure 1b shows the association between this measure of evolvability $e_G = U/3n$ and r_G for 7.5×10^4 sequences whose structures span three orders of magnitude in frequency. The two quantities are highly negatively correlated with mutational robustness (Spearman's $s = -0.64$, $p < 10^{-17}$; $n_S = 7.5 \times 10^4$). The dashed line in the figure indicates $1 - r_G$, the fraction of sequences with structures different from that of G . The discrepancy between this dashed line and the data indicates that many structures different from that of G are not different from each other, otherwise the data would fall on the dashed line. The data shown in figure 1b is an average over structures with different frequencies.

Earlier work (Ancel & Fontana 2000) showed—for a given structure—a negative association between mutational robustness of sequences folding into the structure and the structural repertoire found in a sequence neighbourhood. In contrast, the data shown in figure 1b are based on sequences that were inversely folded from many structures with widely varying frequencies. This raises the question of whether structure frequency has any influence on the relationship between r_G and e_G . I addressed this question by determining a partial correlation coefficient between r_G and e_G that holds the structure frequency constant. The strongly negative association persists when we control in this way for structure frequency (Partial product-moment correlation coefficient $r = -0.65$, $p < 0.01$; $n_S = 7.5 \times 10^4$). This means that structure frequency does not have a strong influence on the negative association between r_G and e_G . In sum, the greater a sequence's robustness, the lesser its evolvability.

(c) High structure robustness means high structure evolvability

By the above definition, the robustness R_P of a phenotype (structure) P is the average number of neutral neighbours of *all* sequences with phenotype P , or their proportion

$r_P = R_P/3n$. The corresponding measure of evolvability E_P is the total number of structures different from P in the 1-neighbourhood of P .

Because neutral networks are so vast, neither robustness R_P nor evolvability E_P can be determined exactly. For example, in a sample of 10^6 sequences of length $n = 30$ nucleotides, the rarest structures are those that are found only once in the sample. Their frequency in sequence space cannot be estimated accurately from a sample of this size. However, any structure whose unique occurrence in such a sample is a reflection of its true frequency, e.g. a structure where $f = 10^{-6}$, would still be adopted by approximately $10^{-6} \times 4^n \approx 10^{12}$ sequences. In other words, even the rarest phenotypes in a large sample of sequence may be adopted by many sequences. Although this prevents the determination of robustness R_P and evolvability E_P through exhaustive enumeration, a sampling approach can still be used to estimate these quantities. That is, one can inversely fold a sample of sequences with phenotype P and determine R_P in this sample. Similarly, one can determine for such a sample the number U of unique structures; that is, structures different from each other and from P . I emphasize that this is not simply the number of unique structures in the 1-neighbourhood of each sequence, added over all sequences in the sample. The reason is that the 1-neighbourhoods of two different sequences in the sample may contain sequences folding into identical structures. Instead, one needs to count structures that occur in the 1-neighbourhood of two or more sequences only once.

Because structure evolvability depends on the size of a neutral network, the resulting number U then still needs to be multiplied by an estimate of the total size of a neutral network, such as the structure frequency f . In other words, e_P will be monotonically increasing with Uf , although this relationship is not necessarily linear. Figure 2a shows the

Figure 2. (a) High phenotype robustness implies high phenotype evolvability. For any one structure, the estimate of evolvability (e_p) used here is the total number U of structures different from each other that were found in the 1-neighbourhood of $k=100$ inversely folded sequences, multiplied by the structure frequency f . (b) Histogram of the ratio Q (see main text) indicating how many structures in the 1-neighbourhoods of $k (=100)$ sequences are different from each other. Q ranges from $Q=1/k$, if the k neighbourhoods are identical in their structure content, to a value of $Q=1$, if no two structures in any two 1-neighbourhoods are identical. The median of Q is greater than $1/2$, indicating that the majority of structures in different 1-neighbourhoods are different. (c) The ratio Q increases with structure frequency, indicating that the neighbourhood of a sequence folding into a structure with a larger neutral network contains greater number of structures unique to this neighbourhood. Data shown are based on the 2.5×10^4 different RNA structures ($n=30$ nucleotides) with the highest ranking from figure 1, and on $k=100$ inversely folded RNA sequences for each structure. Error bars indicate one standard error. (d) Mutational robustness r_G varies among sequences inversely folded from different structures. Data shown are based on 7.5×10^4 different RNA structures ($n=30$ nucleotides) whose frequencies span three orders of magnitude, and on one inversely folded RNA sequence for each structure.

relationship between phenotype robustness r_p and Uf (a proxy for e_p) based on 2.5×10^4 structures, and on 100 randomly sampled sequences from each structure's neutral network. Structure robustness and evolvability show a strong positive association (Spearman's $s=0.55$, $p < 10^{-17}$; $n_S = 2.5 \times 10^4$).

To understand this positive association, several observations are germane. First, the higher a structure's frequency, the higher its robustness (Spearman's $s=0.64$, $p < 10^{-17}$; $n_S = 2.5 \times 10^4$; see figure S1 in the electronic supplementary material). When a neutral network is viewed as a graph, this observation simply states that the average number of neutral neighbours of a typical genotype G on the network increases with the size of the neutral network. It also means that the greater the structure frequency f , the more sequences in the 1-neighbourhood of any one sequence G will have the same structure as G . Conversely, the greater the structure frequency f is, the smaller the proportion of the neighbours of any one G that will have unique and different structures.

Next, the structures found in the neighbourhoods of two or more different sequences are typically very different from

each other. To see this, consider the *total* number of different structures found in the 1-neighbourhood of a set of sequences ($G_1, \dots, G_k$) sampled from a neutral network with structure frequency f . Denote the set of structures that are different from each other in the 1-neighbourhood of G_i as $\{U_i\}$ and the size of this set as $|U_i|$. When one compares two different sets, say $\{U_i\}$ and $\{U_j\}$, there are two extreme possibilities and a wide spectrum in between: all of these structures could be identical, i.e. $\{U_i\} \cap \{U_j\} = \{U_i\} = \{U_j\}$, or all of these sequences could be different, i.e. $\{U_i\} \cap \{U_j\} = \emptyset$. The truth is closer to the second extreme, as illustrated by the following analysis. Consider the quantity

$$\frac{1}{k} \sum_{i=1}^k \frac{|U_i|}{3n}, \quad (3.1)$$

where k is some number of sequences sampled from the neutral network of a given structure. This quantity is the proportion of unique structures in the 1-neighbourhood of G_i , averaged over all k sampled sequences. Next consider the

quantity

$$\frac{\left| \bigcup_{i=1}^k U_i \right|}{3nk}, \quad (3.2)$$

which is the *total* number of structures different from each other found in the neighbourhood of *all* k -sampled sequences divided by the total number of sequences examined, that is, k times $3n$, because each of the k sequences has $3n$ neighbours. Of interest is the ratio of (3.1) and (3.2), that is

$$Q := \frac{\left| \bigcup_{i=1}^k U_i \right|}{\sum_{i=1}^k |U_i|}.$$

If the structures found in the neighbourhood of two sequences are very similar, then $Q \ll 1$, because in that case $|\bigcup_{i=1}^k U_i| \approx |U_i|$ for every i , which does not depend on k , and Q would thus approach zero with increasing k . For example, for the $k=100$ sequences sampled here, one would then expect that $Q \approx 0.01$. If, however, all the structures found in the neighbourhood of two sequences are different from each other, then Q would be of the order 1.

Figure 2*b* shows the distribution of Q for 2.5×10^4 structures. The median of this distribution is equal to 0.6, or 60-fold greater than the case where the structures found in the neighbourhoods of two sequences are identical. This means that most structures that occur in the neighbourhoods of different sequences along a neutral network are different from each other. For sequences that are not randomly sampled, but encountered along a random walk along a neutral network, this has been shown previously (Huynen 1996; Sumedha *et al.* 2007).

Because most structures in the 1-neighbourhoods of different sequences in a neutral network are different from each other, the size of a neutral network has an important influence on structure-based evolvability. Specifically, even though the number of structures found in the 1-neighbourhood of any one sequence decreases modestly with increasing neutral network size (see figure S1 in the electronic supplementary material), this decrease is much smaller than the increase in the number of different structures accessible from a much larger neutral network. In addition, the ratio Q increases modestly with structure frequency (figure 2*c*; Spearman's $s=0.11$, $p < 10^{-17}$; $n_s = 2.5 \times 10^4$). This means that the larger a neutral network is, the more distinct the structures found in the 1-neighbourhoods of sequences sampled from the neutral network.

A final observation useful to understand figure 2*a* emerges from a comparison between structure frequencies and sequence robustness. For any sequence of length n , the number of neighbours with the same MFE structure can vary between 0 and $3n$. In contrast, the number of sequences folding into a given structure can vary over a much broader range, from zero to a fraction of a per cent of the total number of sequences, i.e. of the order of the number 4^n of sequences itself. This discrepancy appears even in modestly sized samples of sequences. For example, figure 2*d* shows the mutational robustness, normalized to (0, 1), for a sample of 7.5×10^4 sequences whose structure frequencies vary over the same range as that in figure 1, i.e. by a factor 1600. In contrast, mutational robustness in this sample varies only by a factor 37. In other words, even in

this modest sample of sequences, structure frequencies are more than 40 times more variable than mutational robustness.

In sum, the positive association between phenotype robustness and evolvability can be explained as follows. First, the number of genotypes folding into any one structure can vary by many orders of magnitude, whereas mutational robustness among sequences of similar lengths varies more modestly. Next, most structures found near two or more sequences sampled from the same neutral network are different from each other. Thus, even though structure robustness increases modestly with structure frequency, this increase is much smaller than the vast increase in the number of different structures accessible found near a much larger neutral network.

(d) Populations evolving on large neutral networks can access greater amounts of variation

Neutral networks are vast in size and a finite population may take a very long time to explore such networks and all the structures in their neighbourhoods. It is thus important to show that robustness also affects evolvability on short time scales in finite populations. To determine whether this is the case, I first chose two different structures, one with high frequency ($f \approx 10^{-3}$) and thus high robustness, and the other with low frequency and low robustness ($f \approx 10^{-6}$). I then inversely folded 20 sequences for each of these structures. For each of the 40 ($= 2 \times 20$) sequences, I then established a population of $N=500$ identical sequences. Each population then underwent repeated rounds of mutation (one nucleotide per sequence generation) and selection that confined the population to the neutral network. That is, in each generation, mutants that were no longer on the neutral network were eliminated and replaced by randomly sampled mutants (with replacement) that still resided on the network. After each such round of mutation and selection, I determined the total number of unique structures found in the neighbourhood of the entire population. The results show that the more robust phenotypes can access much more variation in their evolution on a neutral network (figure 3*a*). For example, after a mere 10 generations, the neighbourhoods of the population on the large network contain 2118 (± 362 s.e.m.) unique structures. In contrast, the neighbourhood of the populations on the small network contains merely 874 different (± 148 s.e.m.) structures.

How can an evolving *population* with a robust phenotype access more variation, despite the fact that each *individual* typically has fewer unique structures in its neighbourhood? The answer is that the populations with the highly robust phenotype are more diverse, and this increased diversity is much greater than the decreased diversity around any one sequence. Figure 3*b* shows the number of different sequences for each population and figure 3*c* shows the mean Hamming distance of the sequences from each other. For both measures, the population with the robust phenotype rapidly accumulates greater diversity. But why is this increased sequence diversity observed in the first place? The reason is simply that in each generation, mutations kill fewer sequences with a more robust phenotype. Assume, for example, that in a population sequences have average robustness r_G . Then, a number of individuals proportional to $(1 - r_G)$ will

Figure 3. Populations evolving on a large neutral network (robust phenotype) have access to greater amounts of phenotypic variation. (a) Numbers of unique structures in the 1-neighbourhood of evolving populations, (b) numbers of different sequences in the population and (c) pairwise Hamming distance among sequences in the population, as a function of the number of generations of evolution (x -axes) on a neutral network. Open and filled circles in (a–c) correspond to populations with less and more robust phenotypes, respectively. Data are based on 20 inversely folded sequences per structure and on populations of size $N=500$ ($\mu=1$). (d) Mean (circles) and standard errors (bars) of numbers of unique structures (y -axis) in the 1-neighbourhood of populations that have evolved for 10 generations on a neutral network associated with structures whose frequency is shown on the x -axis. Data in (d) are based on 4000 different structures ranging in frequency from 3.3×10^{-5} to 1.7×10^{-3} , and on one inversely folded sequence per structure that is used to seed a population size of $N=100$. Circles and bars indicate means and one standard error. (e) Same as (a), but for populations with $N=10$ and $\mu=0.01$. Also, the y -axis in (e) shows the *cumulative* number of unique structure, i.e. unique structures that occurred in the 1-neighbourhood of all genotypes that the population encountered between generation zero and the time shown on the x -axis. Note that because $N\mu=0.1$, a new genotype arises that will come to dominate the population only once in every 10 generations (Kimura 1983), implying that the exploration of the network by such a population is necessarily slower. As one would expect, the 1-neighbourhoods of the starting genotypes contained fewer unique phenotypes for the robust starting phenotype than for the less robust starting phenotype (0.33 versus 0.30), and these genotypes were also more robust ($r_G=0.48$ versus $r_G=0.27$). (f) Mean (circles) and standard errors (bars) of the *cumulative* number of unique structures (y -axis) in the 1-neighbourhood of populations that have evolved for 10^4 generations on a neutral network associated with structures whose frequency is shown on the x -axis. Data in (f) is based on 1500 different structures ranging in frequency from 7.4×10^{-5} to 1.7×10^{-3} , and on one inversely folded sequence per structure that is used to seed a population size of $N=10$ and $\mu=0.01$.

be eliminated in every generation as a result of mutations. Populations of sequences with greater robustness (lower $1 - r_G$) can thus accumulate greater diversity.

These observations are not peculiarities of the two structures I used. I repeated this approach with inversely folded sequences derived from 4×10^3 different structures, each of which was used to seed an evolving population. Figure 3d shows the number of unique structures in the neighbourhood of these 4×10^3 populations after 10 generations of mutation and selection. There is a modest but highly significant positive association between the structure frequency and the amount of phenotypic variation accessible to these structures. Populations with more frequent and thus more robust phenotypes thus have access to more phenotypic variation.

I note that these results would be qualitatively the same if I had not used 'soft' selection, where population sizes are held constant but 'hard' selection, in which population sizes are allowed to fluctuate. The reason is that in this case, populations evolving on small neutral networks would simply shrink faster over time than those on large neutral networks and show lower sequence diversity for this reason.

(e) Even when populations or mutation rates are small, populations with robust phenotypes access more variation

In the above analysis, mutations occurred at a rate of $\mu = 1$ per sequence per generation, such that the population mutation rate is $N\mu \gg 1$. The evolutionary dynamics in populations with $N\mu \gg 1$ is fundamentally different from that in populations with $N\mu \ll 1$ (Kimura 1983; van Nimwegen *et al.* 1999; Wagner 2005). Specifically, large populations or populations with high mutation rates are likely to be polymorphic in any given generation, and their members accumulate over time in regions of a neutral network where genotypes show high genotypic robustness. In contrast, in populations where $N\mu \ll 1$, populations are monomorphic most of the time, and they perform a random walk on a neutral network that samples the network uniformly. Given these differences, it is necessary to show that the greater structural diversity accessible to populations with robust phenotypes does not occur only when $N\mu \gg 1$.

When $N\mu \ll 1$, two populations differing in the robustness of their phenotypes cannot have access to dramatically different structural variation in any given generation. The reason is that in most generations, both populations will be monomorphic; that is, all of their members have identical genotypes. However, the likelihood that a mutation is deleterious is smaller in populations with more robust phenotypes. This suggests that such populations can explore the neutral network faster also when $N\mu \ll 1$ and thus encounter—over time—more structural variation. Figure 3e shows that this is indeed the case. The figure is based on populations that were initialized with a single randomly chosen member from a neutral network, and that evolved via selection and mutation, just like in figure 3a, but with $N\mu \ll 1$ ($N = 10$; $\mu = 0.01$). The x -axis indicates the number of generations. The y -axis indicates the cumulative number of unique structures that the population encountered in its 1-neighbourhood since generation zero. This cumulative number is obtained by recording all the genotypes on the neutral

network that have occurred in the population since generation zero, and by determining the number of unique structures in their 1-neighbourhood (counting each structure that occurs in the 1-neighbourhood of two different genotypes only once). Clearly, populations with more robust phenotypes encounter greater structural diversity along their evolutionary path. Also, both the cumulative sequence diversity and the cumulative number of different sequences encountered during an evolutionary trajectory are larger in populations with robust phenotypes (results not shown). And again, these observations are not artefacts of the two phenotypes chosen: there is a positive association between the structure frequency and the structural diversity a small population encounters during its evolution (figure 3f).

In sum, in populations with both $N\mu \gg 1$ and $N\mu \ll 1$, an evolving population with a robust phenotype accesses more variation along its evolutionary path on a neutral network. The reasons are fundamentally the same, namely that the mutations are less likely to have deleterious effects for robust phenotypes, and thus allow more rapid exploration of a neutral network.

(f) An evolutionary search's ability to find a target structure is only weakly correlated with robustness

A definition of evolvability that focuses only on the variation directly accessible from a given genotype G or phenotype P may seem limited. It does not address how a blind evolutionary search driven by mutations would find a phenotype (structure) that is not in the immediate neighbourhood of G or P , but an arbitrary distance away from it. Some genotypes G or phenotypes P might be more amenable to finding such arbitrary target phenotypes, and thus be more evolvable in this sense. If so, how is this kind of evolvability related to genotypic or phenotypic robustness?

To address this question, I pursued an approach that started with a set of 7.5×10^4 random RNA structures that span three orders of magnitude in structure frequency.

I drew pairs of structures (S , T) at random from this set (with replacement). For each such pair, I inversely folded a sequence G with structure S . From the starting sequence, I then performed a random walk towards the target structure T . Specifically, each step of this random walk consisted in a random change of a single nucleotide. If the mutation had not increased the Hamming distance to T , then the random walk was continued with the mutated sequence; otherwise, the original sequence was mutated again. This process was repeated until a sequence with an MFE structure T was obtained. The number of mutational steps needed to get to the target structure T can be used as a measure of evolvability.

Figure 4a shows that sequence robustness is only marginally associated with evolvability in this sense (Spearman's $s = 0.01$, $p = 0.016$; $n_S = 3.7 \times 10^4$). What other factors might influence the length of this random walk? One candidate factor is the frequency of the starting structure S . This frequency is associated with the size of the sequence space that can be explored while staying on a neutral network. However, it is not associated with the length of this random walk either (Spearman's $s = 0.006$, $p > 0.05$; $n_S = 3.7 \times 10^4$).

Figure 4. (a) Mutational robustness (x -axis) is not associated with the length of a random walk to a target structure (y -axis). (b) The frequency f of a target structure shows a weak negative association with the length of the random walk. (c) Distribution of the number of different secondary structures encountered during a random walk beginning from a sequence folding into a structure S to a sequence folding into a structure T . All data based on 3.7×10^4 random structure pairs (S , T).

Another candidate factor is the structure distance between the starting and the target structures. It might take longer to reach a given target structure if this structure is very dissimilar from the starting structure. However, two different structure distance measures are only weakly associated with the length of this random walk (Hamming distance: Spearman's $s = 0.03$, $p = 8 \times 10^{-7}$; $n_S = 3.7 \times 10^4$; base pair distance, the number of base pairs that need to be opened or closed to transform one structure into the other: Spearman's $s = -0.07$, $p < 10^{-17}$; $n_S = 3.7 \times 10^4$). The only

variable that is moderately associated with walk length is the frequency of the *target* structure (figure 4b; Spearman's $s = -0.21$, $p < 10^{-17}$; $n_S = 3.7 \times 10^4$). This means that regardless of the starting structure S , it is more difficult for a blind evolutionary search to get to a target structure T if this structure is rare.

Evolvability defined as the length of a random walk starting from a given sequence is a form of *sequence* evolvability. An analogous measure of *structure* evolvability can be defined as the *average* length of a random walk starting from a given structure S to a target structure T . This measure of evolvability, however, is also not associated with mutational robustness, when estimated for $k=100$ inversely folded sequences with structure S (Spearman's $s = -0.04$, $p > 0.05$; $n_S = 910$). Also it is not associated with the structure frequency of S ($s = -0.014$, $p > 0.05$; $n_S = 910$). The association between the length of this random walk and the distance between S and T is weak and depends on the distance measure used (Hamming: $s = -0.096$, $p = 0.004$; base pair distance: $s = -0.04$, $p > 0.05$; $n_S = 910$). Again, the only feature of some relevance is the frequency of the target structure T ($s = -0.34$, $p < 10^{-17}$; $n_S = 910$).

The reason why the starting sequence may be irrelevant for the length of this random walk becomes obvious if one asks how many different structures the random walk encounters between S and T . A histogram of this distribution is shown in figure 4c. The median (mean) of the distribution is 63 (121) with a 10th percentile at 19 structures. Thus, an evolutionary search starting at S traverses many other structures before arriving at T . Arguably, during this search, the properties of the starting structure may matter much less than the properties of the structures encountered during the search. Robustness and evolvability are also not associated if one restricts the analysis to random walks that traverse fewer than 10 (Spearman's $s = 0.04$, $p > 0.05$; $n_S = 745$) or fewer than 5 structures (Spearman's $s = 0.11$, $p > 0.05$; $n_S = 84$) before arriving at T . This implies that the properties of the starting point are rapidly 'forgotten' in an evolutionary search.

3. DISCUSSION

In sum, a highly robust RNA genotype has low evolvability. In contrast, a highly robust phenotype has high evolvability. This positive association is caused by (i) the vastly larger neutral networks (more sequences) associated with mutationally robust phenotypes and (ii) the different structures occurring in the neighbourhoods of two or more sequences sampled from a neutral network. This synergism between robustness and evolvability manifests itself in populations of genotypes spreading on a neutral network. On a large neutral network, such populations have on average higher robustness. They thus suffer lower losses through mutations, which allow them to accumulate greater genotypic diversity. In consequence, they can access greater phenotypic diversity in their neighbourhood.

I have explored robustness and evolvability for a specific genotype and phenotype, RNA and its secondary structure. However, at least the genotypic and phenotypic notions of robustness I use here can be applied to systems on all levels of organization, ranging from molecules to whole organisms. The reason is that for many systems, the same phenotypes can be achieved by vast numbers of

genotypes (Wagner 2005). Take the example of gene regulation circuits, like *Hox* genes guiding axial development in many animals (Carroll *et al.* 2001). They can be characterized according to a regulatory genotype that indicates which circuit genes interact. This genotype is encoded by DNA sequences that comprise both the gene coding regions and their regulatory regions on DNA, promoters and enhancers. The phenotype of such a circuit corresponds to a spatio-temporal expression pattern of circuit genes in response to some 'input' from genes upstream of the circuit. Clearly, any one phenotype can be adopted by many genotypes, partly because regulatory DNA is very flexible in its organization. Recent work suggests that in such circuits the same phenotypes can also be realized by vastly different numbers of genotypes (Ciliberti *et al.* 2007a,b). As in the case of RNA, genotypic robustness and evolvability are properties of one specific genotype, whereas phenotypic robustness and evolvability are properties of all genotypes with the same phenotype. To determine whether phenotypic robustness and evolvability go hand in hand also in such circuits is an important task for future work.

Whether to focus on genotype or phenotype when studying robustness and evolvability is to some extent a matter of taste. There is only one reason to prefer phenotypes. On evolutionary time scales, genotypes change constantly, and are thus a moving target for studies of robustness and evolvability. Phenotypes, however, can stay invariant over short and intermediate evolutionary time scales, if their conservation is important to the organism. To study evolvability of phenotypes—RNA and protein conformations, circuit gene expression patterns, etc.—may thus be of greater relevance for processes that take place on evolutionary time scales.

The observations made here suggest that a measure of phenotypic entropy, as defined below, plays a positive role in evolution. Common wisdom in biology holds that living things and the molecules therein are highly ordered and exceedingly unlikely constellations of matter. Cast in the metaphor of sequence space, nucleic acids and proteins that can perform any given function occupy exceedingly small regions in this space. From a thermodynamic perspective, such molecules show low information entropy (high information content) in sequence space: only a few strings of nucleotides or amino acids can adopt their conformation. More precisely, and expressed in terms of information content, consider that to specify any nucleotide sequence of length n , one needs $\log_2(4^n) = 2n$ bits. To specify a sequence that folds into a given structure, one needs $\log_2(4^n/f(4^n)) = -\log_2(f)$ bits, where f is the proportion of the 4^n sequences that fold into the structure, or an average of $-\log_2(f)/n$ bits per residue. A structure with low information entropy has a small f . Although it may be true that molecular phenotypes with any given function generally have small entropy, it is equally clear that the distribution of f among different phenotypes is very broad (figure 1a). The observations made here suggest that among all different phenotypes that can perform a given function, those with large f may show the highest evolvability. In other words, highly entropic phenotypes are more evolvable. This positive role of entropy in facilitating evolvability stands in contrast to the generally negative role entropy and disorder are held to play throughout biology.

This work leaves three important open questions. First, how robust and evolvable are biologically important *phenotypes*, such as RNA structures? To answer this question is currently impossible, because it requires the ability to estimate phenotype frequencies for many phenotypes. Partly owing to the vastness of genotype space, no reliable and tractable method to do this is currently available. Second, how general is the positive association between phenotypic robustness and evolvability that I showed for RNA structures? Does it occur in many other biological systems? Third, this work does not ask about the evolutionary forces that might cause high evolvability, of which there may be many (Gerhart & Kirschner 1998; Schlosser & Wagner 2004; Bloom *et al.* 2006). There are two principal possibilities. High evolvability might be an adaptation in its own right, or a by-product of other selective pressures. A good candidate for such a selective pressure is natural selection for mutational robustness or thermodynamic stability, which are positively associated (Ancel & Fontana 2000). A small number of studies suggest that biological evolution has produced RNA molecules with high genotypic or phenotypic robustness (Wagner & Stadler 1999; Meyers *et al.* 2004; Borenstein & Ruppén 2006; Sanjuan *et al.* 2006). Although high phenotypic evolvability could be a by-product of selection for high phenotypic robustness, the question whether this is generally the case remains to be resolved.

I am grateful for support through grant 315200-116814 from the Swiss National Foundation, as well as for support from the Santa Fe Institute.

REFERENCES

- Aharoni, A., Gaidukov, L., Khersonsky, O., Gould, S. M., Roodveldt, C. & Tawfik, D. S. 2005 The 'evolvability' of promiscuous protein functions. *Nat. Genet.* **37**, 73–76.
- Ancel, L. W. & Fontana, W. 2000 Plasticity, evolvability, and modularity in RNA. *J. Exp. Zool./Mol. Dev. Evol.* **288**, 242–283. (doi:10.1002/1097-010X(20001015)288:3<242::AID-JEZ5>3.0.CO;2-O)
- Baudin, F., Marquet, R., Isel, C., Darlix, J., Ehresmann, B. & Ehresmann, C. 1993 Functional sites in the 5' region of human-immunodeficiency-virus type-1 RNA form defined structural domains. *J. Mol. Biol.* **229**, 382–397. (doi:10.1006/jmbi.1993.1041)
- Bedau, M. A. & Packard, N. H. 2003 Evolution of evolvability via adaptation of mutation rates. *Biosystems* **69**, 143–162. (doi:10.1016/S0303-2647(02)00137-5)
- Bloom, J. D., Labthavikul, S. T., Otey, C. R. & Arnold, F. H. 2006 Protein stability promotes evolvability. *Proc. Natl Acad. Sci. USA* **103**, 5869–5874. (doi:10.1073/pnas.0510098103)
- Borenstein, E. & Ruppén, E. 2006 Direct evolution of genetic robustness in microRNA. *Proc. Natl Acad. Sci. USA* **103**, 6593–6598. (doi:10.1073/pnas.0510600103)
- Carroll, S. B., Grenier, J. K. & Weatherbee, S. D. 2001 *From DNA to diversity. Molecular genetics and the evolution of animal design*. Malden, MA: Blackwell.
- Ciliberti, S., Martin, O. C. & Wagner, A. 2007a Circuit topology and the evolution of robustness in complex regulatory gene networks. *PLoS Comput. Biol.* **3**, e15. (doi:10.1371/journal.pcbi.0030015)
- Ciliberti, S., Martin, O. C. & Wagner, A. 2007b Innovation and robustness in complex regulatory gene networks. *Proc. Natl Acad. Sci. USA* **104**, 13 591–13 596. (doi:10.1073/pnas.0705396104)
- Dayton, E., Konings, D., Powell, D., Shapiro, B., Butini, L., Maizel, J. & Dayton, A. 1992 Extensive sequence-specific

- information throughout the CAR RRE, the target sequence of the human-immunodeficiency-virus type-1 rev protein. *J. Virol.* **66**, 1139–1151.
- Dun, R. B. & Fraser, A. S. 1959 Selection for an invariant character, 'vibrissae number', in the house mouse. *Aust. J. Biol. Sci.* **12**, 506–523.
- Earl, D. J. & Deem, M. W. 2004 Evolvability is a selectable trait. *Proc. Natl Acad. Sci. USA* **101**, 11 531–11 536. (doi:10.1073/pnas.0404656101)
- Fontana, W. 2002 Modelling 'evo-devo' with RNA. *Bioessays* **24**, 1164–1177. (doi:10.1002/bies.10190)
- Fontana, W. & Schuster, P. 1998 Continuity in evolution: on the nature of transitions. *Science* **280**, 1451–1455. (doi:10.1126/science.280.5368.1451)
- Gerhart, J. & Kirschner, M. 1998 *Cells, embryos, and evolution*. Boston, MA: Blackwell.
- Gruner, W., Giegerich, R., Strothmann, D., Reidys, C., Weber, J., Hofacker, I., Stadler, P. & Schuster, P. 1996 Analysis of RNA sequence structure maps by exhaustive enumeration. II. Structures of neutral networks and shape space covering. *Monath. Chem.* **127**, 375–389. (doi:10.1007/BF00810882)
- Hansen, T. F. 2003 Is modularity necessary for evolvability? Remarks on the relationship between pleiotropy and evolvability. *Biosystems* **69**, 83–94. (doi:10.1016/S0303-2647(02)00132-6)
- Hofacker, I., Fontana, W., Stadler, P., Bonhoeffer, L., Tacker, M. & Schuster, P. 1994 Fast folding and comparison of RNA secondary structures. *Monath. Chem.* **125**, 167–188. (doi:10.1007/BF00818163)
- Hofacker, I., Schuster, P. & Stadler, P. 1998 Combinatorics of RNA secondary structures. *Discr. Appl. Math.* **88**, 207–237. (doi:10.1016/S0166-218X(98)00073-0)
- Houle, D. 1992 Comparing evolvability and variability of quantitative traits. *Genetics* **130**, 195–204.
- Huynen, M. A. 1996 Exploring phenotype space through neutral evolution. *J. Mol. Evol.* **43**, 165–169. (doi:10.1007/BF02338823)
- Kimura, M. 1983 *The neutral theory of molecular evolution*. Cambridge, UK: Cambridge University Press.
- Kirschner, M. & Gerhart, J. 1998 Evolvability. *Proc. Natl Acad. Sci. USA* **95**, 8420–8427. (doi:10.1073/pnas.95.15.8420)
- Masel, J. & Bergman, A. 2003 The evolution of the evolvability properties of the yeast prion [PSI⁺]. *Evolution* **57**, 1498–1512. (doi:10.1111/j.0014-3820.2003.tb00358.x)
- Meyers, L. A., Lee, J. F., Cowperthwaite, M. & Ellington, A. D. 2004 The robustness of naturally and artificially selected nucleic acid secondary structures. *J. Mol. Evol.* **58**, 681–691. (doi:10.1007/s00239-004-2590-2)
- Michod, R. E., Nedelcu, A. M. & Roze, D. 2003 Cooperation and conflict in the evolution of individuality IV. Conflict mediation and evolvability in *Volvox carteri*. *Biosystems* **69**, 95–114. (doi:10.1016/S0303-2647(02)00133-8)
- Nijhout, H. F., Berg, A. M. & Gibson, W. T. 2003 A mechanistic study of evolvability using the mitogen-activated protein kinase cascade. *Evol. Dev.* **5**, 281–294. (doi:10.1046/j.1525-142X.2003.03035.x)
- Powell, D., Zhang, M., Konings, D., Wingfield, P., Stahl, S., Dayton, E. & Dayton, A. 1995 Sequence specificity in the higher-order interaction of the rev protein of HIV with its target sequence, the RRE. *J. Acquir. Immune Def. Syndr. Hum. Retrovirol.* **10**, 317–323.
- Queitsch, C., Sangster, T. & Lindquist, S. 2002 Hsp90 as a capacitor of phenotypic variation. *Nature* **417**, 618–624. (doi:10.1038/nature749)
- Reidys, C., Stadler, P. & Schuster, P. 1997 Generic properties of combinatory maps: neutral networks of RNA secondary structures. *Bull. Math. Biol.* **59**, 339–397. (doi:10.1007/BF02462007)
- Rutherford, S. L. & Lindquist, S. 1998 Hsp90 as a capacitor for morphological evolution. *Nature* **396**, 336–342. (doi:10.1038/24550)
- Sanjuan, R., Forment, J. & Elena, S. F. 2006 *In silico* predicted robustness of viroids RNA secondary structures. I. The effect of single mutations. *Mol. Biol. Evol.* **23**, 1427–1436. (doi:10.1093/molbev/msl005)
- Schlosser, G. & Wagner, G. (eds) 2004 *Modularity in development and evolution*, Chicago, IL: University of Chicago Press.
- Schuster, P. 2003 Molecular insights into evolution of phenotypes. In *Evolutionary dynamics: exploring the interplay of selection, accident, neutrality, and function* (eds J. P. Crutchfield & P. Schuster), pp. 163–215. New York, NY: Oxford University Press.
- Schuster, P., Fontana, W., Stadler, P. & Hofacker, I. 1994 From sequences to shapes and back—a case-study in RNA secondary structures. *Proc. R. Soc. B* **255**, 279–284. (doi:10.1098/rspb.1994.0040)
- Sniegowski, P. D. & Murphy, H. A. 2006 *Evolvability*. *Curr. Biol.* **16**, R831–R834. (doi:10.1016/j.cub.2006.08.080)
- Sumedha, Martin, O. & Wagner, A. 2007 New structural variation in evolutionary searches of RNA neutral networks. *Biosystems* **90**, 475–485. (doi:10.1016/j.biosystems.2006.11.007)
- Tacker, M., Stadler, P., BornbergBauer, E., Hofacker, I. & Schuster, P. 1996 Algorithm independent properties of RNA secondary structure predictions. *Eur. Biophys. J. Biophys. Lett.* **25**, 115–130.
- van Nimwegen, E., Crutchfield, J. & Huynen, M. 1999 Neutral evolution of mutational robustness. *Proc. Natl Acad. Sci. USA* **96**, 9716–9720. (doi:10.1073/pnas.96.17.9716)
- Waddington, C. H. 1953 The genetic assimilation of an acquired character. *Evolution* **7**, 118–126. (doi:10.2307/2405747)
- Waddington, C. H. 1959 *The strategy of the genes*. New York, NY: Macmillan.
- Wagner, A. 2005 *Robustness and evolvability in living systems*. Princeton, NJ: Princeton University Press.
- Wagner, A. & Stadler, P. 1999 Viral RNA and evolved mutational robustness. *J. Exp. Zool./Mol. Dev. Evol.* **285**, 119–127. (doi:10.1002/(SICI)1097-010X(19990815)285:2<119::AID-JEZ4>3.0.CO;2-D)
- Zuker, M. & Sankoff, D. 1984 RNA secondary structures and their prediction. *Bull. Math. Biol.* **46**, 591–621.