



Selective pressures on genomes in molecular evolution

Charles Ofria^{a,*}, Christoph Adami^{b,c}, Travis C. Collier^d

^a *Computer Science and Engineering, Michigan State University, East Lansing, MI 48824, USA*

^b *Digital Life Laboratory 123-93, California Institute of Technology, Pasadena, CA 91125, USA*

^c *Jet Propulsion Laboratory MS 126-347, California Institute of Technology, Pasadena, CA 91109, USA*

^d *Division of Organismic Biology, Ecology, and Evolution, University of California Los Angeles, Los Angeles, CA 90095, USA*

Received 13 June 2002; received in revised form 23 December 2002; accepted 15 January 2003

Abstract

We describe the evolution of macromolecules as an information transmission process and apply tools from Shannon information theory to it. This allows us to isolate three independent, competing selective pressures that we term compression, transmission, and neutrality selection. The first two affect genome length: the pressure to conserve resources by compressing the code, and the pressure to acquire additional information that improves the channel, increasing the rate of information transmission into each offspring. Noisy transmission channels (replication with mutations) give rise to a third pressure that acts on the actual encoding of information; it maximizes the fraction of mutations that are neutral with respect to the phenotype. This neutrality selection has important implications for the evolution of evolvability. We demonstrate each selective pressure in experiments with digital organisms.

Published by Elsevier Science Ltd.

Keywords: Avida; Digital Life; Information theory; Molecular evolution; Neutrality

1. Introduction

The techniques of Shannon information theory are steadily gaining application in the field of biology, especially bioinformatics. Information theory has been used to study the information content of biomolecules (Gatlin, 1972; Schneider et al., 1986; Pavesi et al., 1997; Adami and Cerf, 2000), as well as the mode and manner in which this information is acquired (Adami et al., 2000; Schneider, 2000), but not as a tool to isolate the selective pressures that adapting populations are subject to. This is surprising as information theory describes not only the difference between entropy and information, but also the transfer and storage of information.

Recently Adami and Cerf (2000) argued that a genomic complexity can be defined rigorously within standard information theory as the information the genome of an organism contains about its environment. Adami et al. (2000) then demonstrated how this complexity must generally increase in simple evolving

systems. Here, we focus on the selective pressures of molecular evolution and how they contribute to the evolution of robust and complex structures. Thus, information theory can be a tool that ties together evolutionary biology and molecular biology.

From the point of view of information theory, it is convenient to view Darwinian evolution on a molecular level as a collection of information transmission channels, subject to a number of constraints. In these channels, the organism's genomes code for the information (a message) to be transmitted from progenitor to offspring, and are subject to noise due to an imperfect replication process. Information theory is concerned with analysing the properties of such channels, how much information can be transmitted, and how the rate of perfect information transmission of such a channel can be maximized. This information transmission rate can be obtained by multiplying the channel's symbol transmission rate (the rate at which symbols can be sent from the source to the destination without taking into account the error rate) and its capacity (the maximal information that can be encoded per symbol while being able to guarantee faithful transmission at a given level of noise). A major assumption of the information-theoretic

*Corresponding author. Tel.: +1-517-432-1207; fax: +1-517-353-3955.

E-mail address: ofria@msu.edu (C. Ofria).

treatment of evolution is that selection acts to maximize the probability that the full genetic information is faithfully transmitted to the next generation. This maximization involves a number of factors that we isolate below, and guides us in elucidating the pressures evolution exerts on the genome itself.

We shall identify three pressures acting on evolving genomes and recast them in an information-theoretic language. This allows us to formulate the forces acting during selection from a unified point of view, while making unique predictions that go beyond the standard population genetics treatment of evolution. The first pressure results in the *compression* of a message that is not perfectly coded, and (in the absence of any other pressures) tends to reduce message length. The second pressure favors longer messages by increasing the amount of information transmitted per usage of the channel. The third pressure forces populations to code the information in a fault-tolerant manner (when mutation rates are high). While the first two pressures are well known within standard theory (Bürger, 2000; Baake and Gabriel, 2000; Baake and Wagner, 2001), the third one is novel, and its theoretical treatment has only been elucidated recently (van Nimwegen et al., 1999; Wilke et al., 2001; Wilke, 2001a, b). The pressure to code information in a fault-tolerant manner implies that codes should evolve that are robust to deleterious mutations, or, in other words, that the fraction of mutations that are neutral with respect to fitness should increase.

The existence of this neutrality selection has implications for the evolution of evolvability, an effect otherwise known as *canalization* (Waddington, 1942; Mather, 1953; Gibson and Wagner, 2000). It appears to be active in viral (Burch and Chao, 2000), as well as digital (Wilke et al., 2001) populations, and may even be responsible for the evolution of “robustness genes” such as Hsp90 (Rutherford and Lindquist, 1998; Wagner et al., 1999; Queitsch et al., 2002).

2. Selective pressures

Standard population genetics views the process of adaptive evolution (in the absence of frequency-dependent selection) as the maximization of a suitably defined fitness. Under simple conditions (a population of self-replicating molecules in a single niche, without co-evolution) and ignoring the effects of mutation, a genotype will dominate if its basic Malthusian parameter—the growth-rate of a population seeded by this genotype—is highest. When analysed from an information-theoretic point of view, we see that the growth rate is in fact affected by three distinct parameters: the *length* of the message, the maximum *rate* (in units of bits/s) at which that message can be transmitted, and the

message’s robustness to mutations. We shall take a look at each one of them separately, and demonstrate their effect on an evolving population of digital organisms.

2.1. Compression selection

No matter how information is transmitted, there is a cost for passing along each symbol. This general truth is obvious in biology where transmitting a symbol involves copying a nucleotide, and this cost can be minimized simply by transmitting fewer symbols (a shorter message). This results in a selective pressure that we call *compression selection*, and illustrate it in a simple experimental system: a population of digital organisms adapting to a virtual world whose complexity and structure is determined by the experimenter (Adami, 1998; Lenski et al., 1999; Adami et al., 2000; Wilke and Adami, 2002). The experiment demonstrating compression selection is identical in spirit to those of Spiegelman (Mills et al., 1967) involving RNA replicating in vitro.

Among a fixed-size population of self-replicators (see Section 3), a single limiting resource is allocated evenly to the organisms regardless of phenotype or genotype, such that the only trait selected for is self-replication. In this scenario, compression selection will dominate the other evolutionary pressures on the population. Since it takes time to copy each symbol, there is a pressure to shrink the genome to the minimal length that can still self-replicate. Taking organisms evolved in a complex environment (see Section 3) and introducing them into this simple one, we can demonstrate this selective pressure in action. Fig. 1(A) shows the adaptive drop in average replication time (CPU time itself being the single resource in the system) as sections of the genome that are meaningless in the simple environment are stripped away. In Fig. 1(B) we witness the corresponding drop in genome length. Of course, this is completely analogous to the seven-fold decrease in sequence length observed in the evolution of RNA in a simple in vitro environment (Mills et al., 1967). Naturally, the other experiments described below would be much more difficult (if not impossible) to carry out with RNA.

Although compression selection is fundamental and unavoidable, it is also expected to be weak in systems where the cost associated with transmitting the message on a per-symbol basis is low. However, in simple molecular systems such as the one described above, as well as in periods following the invasion of a new niche or an environmental change, the effects on the genome can be significant due to the removal of genes that have lost their usefulness.

2.2. Transmission selection

Most resources in extant biological systems are spent building “survival machines” as Dawkins (1976)

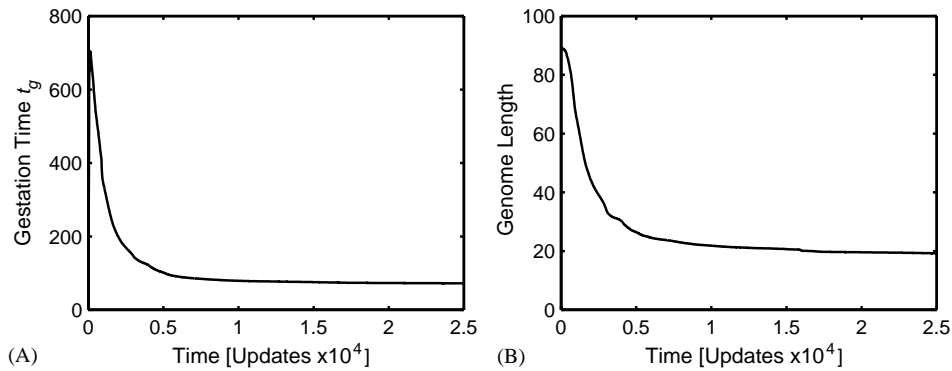


Fig. 1. Average replication time (A) and average genome length ℓ (B) for complex organisms introduced into a simple world (set I).

eloquently put it. If an organism obtains an increased share of resources (including expected lifetime), its opportunity for replication is correspondingly increased. This is often considered first by biologists when thinking of interesting examples of adaptation. In terms of information theory, these “survival machines” increase the symbol transmission rate of their channel. Indeed, in a population of self-replicators with a single limiting resource, an organism with twice the affinity to that resource compared to its competitors will produce twice as many offspring, and will soon dominate. In other words, by taking better advantage of its environment, an organism can increase the “bandwidth” of the channel through which it transmits its genetic message to future generations.

To demonstrate this selective pressure, we modified the previous environment that favored a short sequence length, by allowing the single limiting resource to be actively competed for. For our digital organisms, this meant that the resource (CPU cycles) was now distributed according to phenotype, in this case to those phenotypes that performed selected computations (see Section 3). Additionally, in these experiments we scaled the resource allocation with genome length to neutralize the effects of compression selection. This allocation method ensures that organisms with identical computational abilities are equally fit in a mutation-free environment, regardless of genome length. While a non-zero mutation rate creates an additional pressure to reduce the target area for mutations, in a complex environment we nevertheless observe a tendency for organisms to *increase* sequence length in order to store information about their environment. It is this information that modifies the organisms’ phenotype, allowing them to better exploit the environment (i.e. it increases the capacity of their channels, through which they pass their genomic information to their offspring.)

We performed trials (sets II–IV, see Table 1) starting from a simple, self-replicating ancestor in three environments of differing complexity: *simple* where no computational tasks are rewarded (as in our compression selection example), *medium* where the completions of 10

different logical operations are each rewarded by increased resource affinity (in turn, increasing the channel’s symbol transmission rate), and *complex* where a total of 78 distinct operations are rewarded. In this setup, we expect that transmission selection will cause populations in more complex environments to incorporate more information about those environments into their genomes, resulting in relatively more fit organisms.

Fig. 2(A) shows the trajectory of resource affinity divided by replication time (the fitness w) as function of time. This fitness is a Malthusian parameter measuring rate of replication. Replication time does increase as computational tasks are acquired, but this is more than balanced by the improved resource affinity, resulting in an increase of the overall fitness. Fig. 2(B) demonstrates that genome length increases in relation to increasing environmental complexity, as organisms learn to exploit their environment by acquiring information that increases the channel’s symbol transmission rate.

2.3. Neutrality selection

Up to this point, we have been careful to exclude the effects of mutation. Mutations are easily incorporated in an information-theoretical analysis as noise affecting the information transmission channel. This inclusion reveals a third selective pressure not easily apparent at the population level, but affecting the rate at which complexity grows in adapting populations with implications for the evolution of evolvability.

For illustrative purposes, suppose that all mutations are lethal, and as a consequence directly lower the number of expected offspring by the probability of a mutation occurring anywhere in the genome. In this case, the probability of faithfully transmitting genetic information to the next generation, which we have assumed is maximized by evolution, must account for the noise affecting the transmission process. However, just like there are many ways to express a concept in English words, there are many messages (sequences) that contain the same meaning (information), and thus in reality not all mutations are deleterious. If a mutation

Table 1
Sets of experiments with digital organisms to study evolutionary pressures

| Set | Complexity | Compression | Constraint | Ancestor | Rate (%) |
|-----|------------|-------------|------------|----------|----------|
| I | S | Yes | No | Complex | 0.75 |
| II | S | No | No | Simple | 0.75 |
| III | M | No | No | Simple | 0.75 |
| IV | C | No | No | Simple | 0.75 |
| V | C | No | Yes | Simple | 0.5 |
| VI | C | No | Yes | Simple | 1.0 |
| VII | C | No | Yes | Simple | 1.5 |

Each set consists of 100 trials that were averaged for statistical significance. Second column: complexity of environment (S; simple, M: medium, C: complex; see Section 3), third column: pressure for compression selection, fourth: constant length constraint, fifth: type of ancestor, sixth: mutation

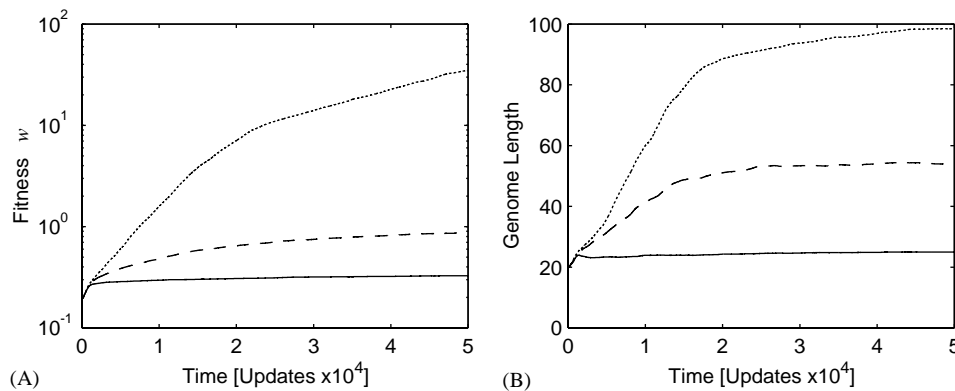


Fig. 2. Average fitness w (A) and average genome length ℓ (B) displayed for three experiments at distinct levels of complexity. Set II (solid, simplest environment), set III (dashed, intermediate environment), and set IV (dotted, most complex world).

modifies the genetic sequence without affecting the message, but does make it more robust to further mutations, such a sequence will have an enhanced probability of successful transmission, and therefore possesses a selective advantage. This is most easily thought of as a *neutrality selection* that acts to increase the probability that mutations are otherwise selectively neutral. While the possibility of such a pressure has been mentioned before (Clarke, 1970; Modiano et al., 1981), its theoretical basis has only been rigorously studied for the case that all mutations are either fully neutral or strongly deleterious (van Nimwegen et al., 1999; Wilke, 2001a). Here, we show that this pressure is more general, and exists even when the effects of mutations are distributed on a continuous scale.

As the replication process is not perfect, mutated genotypes often do not contribute to the strength of a species. Denote by F (the fidelity of replication) the probability that a birth process gives rise to a genetically identical offspring. If all mutations are deleterious, the true rate of growth (fitness) of a genotype is fidelity times raw replication rate,

$$w_F = Fw, \quad (1)$$

which is the number of genetically identical offspring produced per unit time. In the language of information

theory, this is the number of perfect copies of a message received per unit time.

In reality, of course, not all mutations are detrimental. From an information-theoretic perspective, a sequence whose information content is not altered is as good as an un-mutated copy: the neutrality of a code is akin to a limited error-protection. Therefore, we are interested in the probability that a replication process gives rise to either a true copy, or one that differs from the wild-type by neutral mutations only.

In our system (and macro-molecular evolution in general) the fidelity of the copy process can be written in terms of the error rate per monomer copied, R , and the length of the genome ℓ , as $F = (1 - R)^\ell$. Denote by ν the probability that a mutation does not affect the replication rate, w , of an organism's offspring. Then, the *neutral fidelity*

$$F_\nu = (1 - R + \nu R)^\ell \quad (2)$$

is the probability of either replicating accurately or else with only neutral mutations. Thus, the combination

$$w_\nu = F_\nu w \quad (3)$$

is the replication rate of the actual genetic information (in its original or a neutrally modified form). Because F_ν is an increasing function of ν , a pressure to increase

w_v implies that there is a selective pressure maximizing v . This is neutrality selection. Note, however, that evolution acts to increase the average w_v in the population (van Nimwegen et al., 1999; Wilke, 2001b). Thus, v -selection implies that selection at high mutation rates acts on groups of mutually neutral genotypes, rather than on the individuals themselves.

Before discussing evidence for neutrality selection in experiments with digital organisms, let us point out that the above formulae can be recast in terms of standard population biology of clonal or randomly mating sexual populations, if w is understood as the fitness contribution of a locus or gene.

Equilibrium (meaning mutation-selection balance) demands that the true rate of growth of the wild-type F_{w_0} equals the average growth rate \bar{w} , or, if there are neutral mutations, $F_{v w_0} = \bar{w}$. (Note that the neutrality v going into the calculation of F_v is the average neutrality of a population grown from the wild type (Wilke 2001b), which we approximate in the experiments below by the average neutrality of all one point mutants of the wild type, see Section 3).

We can then calculate the *mutational load* on the population (the fraction of the population that dies each generation because of deleterious mutations) as

$$L = 1 - \frac{\bar{w}}{w_0} = 1 - F_v. \quad (4)$$

Now, using Eq. (2) and the fact that $R \ll 1$, we obtain

$$L \approx 1 - e^{-R/(1-v)}. \quad (5)$$

We see that a population can reduce its mutational load, at constant mutation rate, by increasing the neutrality (fraction of all possible one-point mutations that are not deleterious to the organism) v of its genomes. This has been noted independently by van Nimwegen et al. (1999) and by Wilke (2001a) for the case of strictly neutral evolution.

A decrease of mutation rate (e.g. by the development of error correction) would provide another means of lowering the mutational load. However, a lower mutation rate would also result in reduced genetic variability. According to Fisher's Fundamental Theorem, a lower genetic variability correlates to lower adaptive potential. As a consequence, neutrality selection may be thought of as an indirect selection for evolvability.

We demonstrate the influence of neutrality selection experimentally by directly measuring the neutrality of the organisms, v , (see Section 3) as a function of mutation rate (i.e. for populations adapting under different fixed mutation rate). We examined three sets of trials subjected to different mutation rates (sets V–VII), all in identical complex environments (equalizing the pressure of transmission selection) and constrained to remain at a genome length of 100 instructions

(preventing length changes, and hence nullifying compression selection, see Table 1).

Higher mutation rates result in a stronger pressure to increase neutrality. This trend is observed in the measured neutralities of populations adapting in environments at different mutation rates, see Fig. 3. Additionally, the interplay between transmission and neutrality selection is apparent from the slow decrease in neutrality for all three sets due to information that increases resource affinity, and is incorporated into the genomes. Since length is fixed, the genome sequence is being “filled up” with information in the form of adaptive mutations that become fixed in the genome, and neutrality must decrease. In other words, as the population gains information about the environment, sites that were once neutral will now cause a loss of information were they to be mutated.

3. Methods

Evolutionary experiments were performed on digital organisms (Adami, 1998; Lenski et al., 1999; Adami et al., 2000) implemented with version 1.3 of the *Avida* software. This software and its documentation can be obtained from <http://dllib.caltech.edu/avida>. The *Avida* system controls populations of self-replicating computer programs in a complex and noisy environment within a computer's memory. Self-replicating programs may be viewed, within limits, as the computational analogs of catalytically active RNA sequences that serve as the template of their own self-reproduction. Populations are propagated for 50,000 updates (an arbitrary unit of time: every update represents the execution of an average of 30 instructions per individual in the population) from the same single ancestor genotype of length 20 instructions (except for set I, see Table 1) to the carrying capacity of 3600

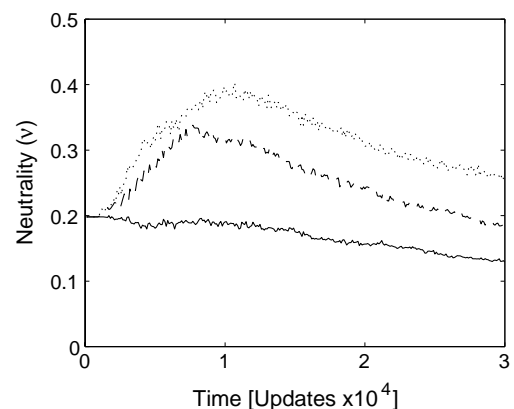


Fig. 3. Average neutrality for sets V (solid), VI (dashed), and VII (dotted), with fixed genome length in a complex environment (see Table 1). Their per-copy mutation rates are 0.5%, 1.0%, and 1.5% respectively.

individuals. Sequence length is variable except for trials V–VII. Mutations occur during replication, at the rate of R mutations per instruction copied per organism. Default settings were used in all experiments unless otherwise indicated.

Computations can be carried out by evolved programs if they develop sequences of code (“genes”) that perform logical, bitwise, computations on random numbers provided in their environment. Such genes evolve spontaneously if the performance is rewarded with bonus CPU cycles (their unit of energy). The complexity of the environment can be controlled by changing the number of logical operations whose performance is rewarded. The development of the necessary code to perform such computations is, in the digital world, the analog of the evolution of a sequence or gene enabling the organism to catalyse exothermic reactions leading to faster replication. Outside materials (in this case binary numbers) are taken in and processed for an energy advantage to the organism (see Wilke and Adami, 2002).

From an information-theoretic viewpoint, the rewards for these computations directly increase the symbol transmission rate of the channel, because longer messages can be sent per usage of the channel. In our simplest world (set II), no such computations are rewarded, and thus no genes beyond replication develop. In the intermediate world (set III), 10 computations are rewarded, and different combinations of these functions are acquired in the different replicas of that set. In the third and most complex set, 78 different operations are potentially rewarded.

Neutrality was measured by obtaining the fitness of all possible one-point mutants of the most abundant genotype, which allowed us to classify mutations as either deleterious, neutral, or beneficial. Neutral mutations are those for which the fitness has changed less than $1/N$, where N is the population size (3600 organisms, for all experiments presented here). The neutrality is the fraction of all mutations that are either neutral or beneficial (as these are those that do not decrease the information content).

4. Conclusion

In natural systems, neutral mutations often occur on a non-critical nucleotide (synonymous substitutions), or else in code that is never expressed. That most molecular evolutionary change is neutral with respect to the phenotype has been widely accepted since Kimura’s seminal work on the subject (Kimura, 1983). Explaining why so much evolutionary change is neutral has been difficult, and has complicated the discussion of increasing genomic complexity (even though genome length has generally been increasing over the course of

evolutionary history, Cavalier-Smith, 1985), since the correlation between genome length and organismic complexity appears weak—the so-called C-value paradox.

Here, we point out that we can think of the quantity maximized by the process of evolution as the probability of faithfully passing on the information encoded in the genome on to the next generation, which is simply the Malthusian parameter (w) multiplied by the neutral fidelity (2) of the organism. This implies that a population can minimize its mutational load by increasing the probability that a mutation is neutral, or, equivalently, decreasing the probability that a mutation is deleterious. Experiments with digital organisms confirm this view. In particular, an analysis of an aggregate of approximately 10 million generations of the evolution of digital organisms has never revealed a single case in which the effective fitness, Eq. (3), has decreased.

The selection of neutrality has a number of interesting consequences when translated to population biology. First, it might provide a convincing mechanism to explain the actual amount of neutrality in the genomes of higher organisms. Furthermore, because a higher neutrality implies both larger fitness variations and the potential for more information acquisition, neutrality selection implies a selective pressure for the evolution of evolvability (see, e.g. Wagner and Altenberg, 1996; Gerhart and Kirschner, 1997; and Ancel and Fontana, 2000) that is molecular at its origin, and provides added support for the idea that robustness is actively selected for in the evolution of development (Gibson and Wagner, 2000). That it may also have left traces in the genomes of higher organisms in the form of proteins that confer neutrality to other proteins (Rutherford and Lindquist, 1998; Wagner et al., 1999; Queitsch et al., 2002) is a tantalizing suggestion that awaits confirmation.

Acknowledgements

This work was supported by the National Science Foundation under grants PHY-9723972 and DEB-9981397. Access to a Beowulf system was provided by the Center for Advanced Computation Research at the California Institute of Technology. Part of this work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. We thank Richard Lenski and Claus Wilke for valuable suggestions.

References

- Adami, C., 1998. Introduction to Artificial Life. Telos Springer, New York, NY.
- Adami, C., Cerf, N.J., 2000. Physical complexity of symbolic sequences. *Physica D* 137, 62–69.

- Adami, C., Ofria, C., Collier, T.C., 2000. Evolution of biological complexity. *Proc. Natl Acad. Sci. USA* 97, 4463–4468.
- Ancel, L., Fontana, W., 2000. Plasticity, evolvability and modularity in RNA. *J. Exp. Zool. (Mol. Dev. Evol.)* 288, 242–283.
- Baake, E., Gabriel, W., 2000. Biological evolution through mutation, selection, and drift: an introductory review. *Ann. Rev. Comput. Phys.* VII, 204–264.
- Baake, E., Wagner, H., 2001. Mutation-selection models solved exactly with methods of statistical mechanics. *Genet. Res. (Cambridge)* 78, 93–117.
- Bürger, R., 2000. *The Mathematical Theory of Selection, Recombination, and Mutation*. Wiley, Chichester, UK.
- Burch, C.L., Chao, L., 2000. Evolvability of an RNA virus is determined by its mutational neighbourhood. *Nature* 406, 625–628.
- Cavalier-Smith, T., 1985. Eukaryote gene numbers, non-coding DNA and genome size. In: Cavalier-Smith, T (Ed.), *The Evolution of Genome Size*. Wiley, New York, NY, pp. 69–103.
- Clarke, B.C., 1970. Darwinian evolution of proteins. *Science* 168, 1009–1011.
- Dawkins, R., 1976. *The Selfish Gene*. Oxford University Press, Oxford.
- Gatlin, L.L., 1972. *Information Theory and the Living System*. Columbia University Press, New York, NY.
- Gerhart, J., Kirschner, M., 1997. *Cells, Embryos, and Evolution*. Blackwell Science, Malden, MA.
- Gibson, G., Wagner, G., 2000. Canalization in evolutionary genetics: a stabilizing theory? *Bioessays* 22, 372–380.
- Kimura, M., 1983. *The Neutral Theory of Molecular Evolution*. Cambridge University Press, Cambridge.
- Lenski, R.E., Ofria, C., Collier, T.C., Adami, C., 1999. Genome complexity, robustness, and genetic interactions in digital organisms. *Nature* 400, 661–664.
- Mather, K., 1953. Genetical control of stability in development. *Heredity* 17, 297–336.
- Mills, D.R., Peterson, R.L., Spiegelman, S., 1967. An extracellular experiment with a self-duplicating nucleic acid molecule. *Proc. Natl Acad. Sci. USA* 58, 217.
- Modiano, G., Battistuzzi, G., Motulsky, A.G., 1981. Nonrandom patterns of codon usage and of nucleotide substitutions in human α - and β -globin genes: an evolutionary strategy reducing the rate of mutations with drastic effects? *Proc. Natl Acad. Sci. USA* 78, 1110–1114.
- Pavesi, A., DeIaco, B., Granero, M.I., Porati, A., 1997. On the informational content of overlapping genes in prokaryotic and eukaryotic viruses. *J. Mol. Evol.* 44, 625–631.
- Queitsch, C., Sangster, T.A., Lindquist, S., 2002. Hsp90 as a capacitor of phenotypic variation. *Nature* 417, 618–624.
- Rutherford, S.L., Lindquist, S., 1998. Hsp90 as a capacitor for morphological evolution. *Nature* 396, 336–342.
- Schneider, T.D., 2000. Evolution of biological information. *Nucleic Acids Res.* 28, 2794–2799.
- Schneider, T.D., Stormo, G.D., Gold, L., Ehrenfeucht, A., 1986. Information content of binding sites on nucleotide sequences. *J. Mol. Biol.* 188, 415–431.
- van Nimwegen, E., Crutchfield, J.P., Huynen, M., 1999. Neutral evolution of mutational robustness. *Proc. Natl Acad. Sci. USA* 96, 9716–9720.
- Waddington, C.H., 1942. Canalization of development and the inheritance of acquired characters. *Nature* 150, 563–565.
- Wagner, G.P., Altenberg, L., 1996. Complex adaptations and the evolution of evolvability. *Evolution* 50, 967.
- Wagner, G.P., Chiu, C.-H., Hansen, T.F., 1999. Is Hsp90 a regulator of evolvability? *J. Exp. Zool. (Mol. Dev. Evol.)* 285, 116–118.
- Wilke, C.O., 2001a. Adaptive evolution on neutral networks. *Bull. Math. Biol.* 63, 715–730.
- Wilke, C.O., 2001b. Selection for fitness versus selection for robustness in RNA secondary structure folding. *Evolution* 55, 2412–2420.
- Wilke, C.O., Adami, C., 2002. The biology of digital organisms. *Trends Ecol. Evol.* 17, 528–532.
- Wilke, C.O., Wang, J.L., Ofria, C., Lenski, R.E., Adami, C., 2001. Evolution of digital organisms at high mutation rate leads to survival of the flattest. *Nature* 412, 331–333.