Journal of Bioinformatics and Computational Biology Vol. 8, No. 6 (2010) 1027–1040 © Imperial College Press DOI: 10.1142/S0219720010005075



QUANTIFYING SLOW EVOLUTIONARY DYNAMICS IN RNA FITNESS LANDSCAPES

PETR ŠULC*,^{†,||}, ANDREAS WAGNER^{‡,§,**} and OLIVIER C. MARTIN*,^{†,¶,††}

*Univ Paris-Sud, UMR8626, LPTMS Orsay, F-91405, France [†]CNRS, Orsay, F-91405, France [‡]University of Zürich, Department of Biochemistry Winterthurerstrasse 190 CH-8057 Zürich, Switzerland §The Santa Fe Institute 1399 Hyde Park Road Santa Fe, NM 87501, USA ¶UMR de Génétique Végétale INRA — Univ Paris-Sud — CNRS — AgroParisTech Ferme du Moulon, F-91190 Gif-sur-Yvette, France $\|petr.sulc@polytechnique.edu$ **aw@bioc.uzh.ch ^{††}olivier.martin@u-psud.fr

> Received 14 May 2010 Revised 23 July 2010 Accepted 4 August 2010

We re-examine the evolutionary dynamics of RNA secondary structures under directional selection towards an optimum RNA structure. We find that the punctuated equilibria lead to a very slow approach to the optimum, following on average an inverse power of the evolutionary time. In addition, our study of the trajectories shows that the out-of-equilibrium effects due to the evolutionary process are very weak. In particular, the distribution of genotypes is close to that arising during equilibrium stabilizing selection. As a consequence, the evolutionary dynamics leave almost no measurable outof-equilibrium trace, only the transition genotypes (close to the border between different periods of stasis) have atypical mutational properties.

Keywords: RNA neutral networks; genotype; phenotype; evolution; slow relaxation; classification description; systems biology.

1. Introduction

In many realistic systems, dynamics undergo a pronounced slow-down, a feature characteristic of "complex" landscapes.¹ Such phenomena arise for instance in physical systems (thermal relaxation²), in optimization problems (diminishing returns

on search efforts³), and in evolutionary dynamics (punctuated equilibria, i.e. long periods of stasis in the evolutionary record⁴). In this work, inspired by the theoretical framework of statistical physics for glassy systems, we reconsider a simple model of evolutionary dynamics, namely RNA secondary structure evolution.^{5–8} We find (i) slow evolutionary dynamics, whereby the time to find an advantageous phenotypic change has a distribution with a fat (power-law) tail; (ii) non-self-averaging behavior, i.e. even for long RNA molecules, directional selection for some targets will lead to significantly slower dynamics than for other targets; (iii) only weak out-of-equilibrium effects: the genotypes visited by an evolutionary trajectory are quite similar to those arising in equilibrium under stabilizing selection, so except near the ends of periods of stasis, the genotypes produced during an evolutionary trajectory do not have anomalously high or low mutational robustness.

The paper is organized as follows. In Sec. 2, we introduce the evolutionary model. In Sec. 3, we exhibit the non-exponential nature of the relaxation process; empirically, relaxation seems to follow an inverse power law. We also show that the relaxation curves remain sensitive to the target used for directional selection even in the limit of very long molecules. Finally in Sec. 4, we give evidence that at nearly all times the evolutionary trajectory is in quasi-equilibrium; more precisely, the "innovative" genotypes produced by a transition to a new period of stasis are only a bit different from random genotypes as measured by the phenotypic effects of mutations; furthermore, during the periods of stasis, the initial genotype seems to be quickly forgotten, so no significant trace of innovation seems to be maintained. Concluding remarks are given in the final section.

2. The RNA Evolutionary Model

2.1. RNA structure and its fitness landscape

For the purpose of this study, RNA molecules will be thought of as chains or strings of L nucleotides, taken from an alphabet of four possible nucleotides (A, C, G and U). Chemically, the bases can pair via hydrogen bonds. In addition to Watson-Crick pairings (A-U and G-C), the U-G pairing is also possible, though it is weaker. The pairings between bases give rise to an RNA secondary structure, as illustrated in Fig. 1. Apart from this graphical representation, the secondary structure of an RNA molecule can be specified by the more convenient dot-bracket notation where a non-paired base of the sequence is denoted by a dot "." and a paired base is denoted by a left or right parenthesis. This representation allows one to reconstruct the pairings as long as the secondary structure is "planar".⁹ (Planarity means that if the L bases are positioned on a line and the pairings are represented by arcs between the bases, these arcs can be drawn in the plane without any crossings.) An RNA molecule will spontaneously (i.e. via thermodynamic forces only) fold into the structures with lowest free energies. To simulate this folding *in silico*, we use the "Vienna RNA package"¹⁰ to find, for any given sequence, the pairings (secondary structure) which lead to the *minimum* free energy. In effect, this procedure produces



a map from sequences (genotypes) to secondary structures (phenotypes).^{5–8,11} The CPU time needed to determine a minimum free energy structure is $O(L^3)$ for a chain of L bases.

All genotypes that have the same phenotype form a neutral network.¹¹ More precisely, given a genotype to phenotype map, the neutral network associated with the phenotype P is a graph whose nodes are all the genotypes having that phenotype; the edges of that graph connect genotypes if and only if they are nearest neighbors. It is thus necessary to introduce a notion of neighborhood in genotype space. For the present work, two RNA genotypes (sequences) will be considered as nearest neighbors if and only if they differ by a single base. Other choices would have been possible¹¹ such as allowing a difference of two bases, meaning that two successive point mutations would transform one genotype to one of its neighbors.

One may also want to consider multiple *phenotypes*, leading one to also define a distance d_p between two phenotypes. In our case, we take the distance between two RNA secondary structures P_1 and P_2 to be 1/L times the Hamming distance between their dot-bracket string representations. Again, other choices of phenotypic distance functions have been used in the literature^{11, 13} often interchangeably,

but we use only the Hamming distance one for its simplicity. Given the neighborhood definitions both for genotypes and phenotypes, one can introduce a fitness landscape,¹⁴ where genotype and phenotype spaces both have a distance metric, and where each phenotype can be assigned a "fitness". We take the fitness of a phenotype to be a monotonically decreasing function of its distance to some "target" phenotype P_t . For specificity, the reader can think of a decreasing exponential of distance, but as will be clear later on, the choice of function makes no difference. Finally, we will also extend the notion of a neutral network to that of a *neutral ensemble*: we call neutral ensemble the union of all neutral networks whose corresponding phenotype has a given fitness, regardless of the specific phenotype. This definition reflects the fact that secondary structures are relevant in the evolutionary search mainly via their distance to P_t . RNA sequences with their associated fitness form a fitness landscape with hills, valleys and passes (saddles). Analogous landscapes also arise in other systems; for example spin glasses¹⁵ are physical materials in which energy is often identified with fitness; the landscape's ruggedness and many valley structure are important for the associated dynamical properties.

2.2. Dynamics under directional selection

We take the unit of time to be the expected waiting time Δt between two mutations. Thus if μ is the mutation rate per RNA molecule and per generation, our unit of time is $\Delta t = 1/\mu$. For our evolutionary dynamics, we simulate the process of RNA evolution towards a target structure by allowing a genotype to change by a point mutation at each unit of time. The directional selection associated with this evolutionary search then proceeds as follows. At each step, the genotype is mutated at one base taken at random; the corresponding phenotype (secondary structure for that genotype) is determined, and selection is applied: if the distance to the target structure has increased, the mutation is refused and the previous genotype is reinstated, otherwise the new genotype is accepted. This process is called "blind ant" dynamics¹⁷ but it also is referred to as an "adaptive walk".¹⁸ It can also be thought of as using Metropolis Monte Carlo¹⁹ where the temperature has been set to 0. (Note that if one were to consider only moves that were acceptable rather than trying and then testing, one would have instead a random "neutral" walk performing "myopic" ant dynamics.²⁰) In practice, we continue the random walk until the target structure is reached, or until a maximum number of trial mutations is reached.

In our work, we allow only neutral or improving moves in the landscape, which corresponds to zero temperature dynamics in physical systems, or hill climbing in optimization theory. The trajectories are stochastic and are influenced both by the initial genotype chosen (e.g. an arbitrary RNA *sequence*), and by the target (an arbitrary *secondary structure*). We thus need to average over many trajectories, and consider the dependence of our results on these choices.

In a more general context, one can consider the evolution of a *population* in the fitness landscape. If μ is the mutation rate and N the population size, then the effective number of genotypes in a population scales approximately linearly with $N\mu$. When $N\mu \ll 1$, the population remains essentially monomorphic, and this is the case we focus on in this work. If instead $N\mu$ were large, the evolutionary dynamics would occur in a polymorphic population with many different genotypes.

2.3. Equilibrium sampling under stabilizing selection

To test whether the out-of-equilibrium evolutionary dynamics generates atypical genotypes, one must define the "null hypothesis"; clearly we want to compare to *random* genotypes in the fitness landscape. This means that we need to sample *uniformly* the fitness landscape for each given distance to the target phenotype; thus we focus on the genotypes that have a given fitness. The algorithmic procedure to do so is to start with any genotype in the fitness landscape with the specified fitness and then produce a long random walk with importance sampling using the Metropolis Monte Carlo algorithm.¹⁹ In our RNA neutral ensemble context, one accepts only the mutations that do not change the distance to the target; this is then identical to dynamics under a particular stabilizing selection. A long Monte Carlo then provides genotypes satisfying that constraint and their probability distribution is uniform.

From this sampling, we shall obtain equilibrium averages in this space. We will be particularly interested in the mean mutational robustness, where the mutational robustness of a genotype is defined as the fraction of the single-base mutations that do not change the fitness; this is the same thing as the number of neighbors of this genotypye that belong to the neutral ensemble divided by its total number of neighbors.

3. Slow Evolutionary Dynamics

3.1. Typical trajectories have long stasis times

Consider a typical evolutionary trajectory starting from a random initial genotype. At the beginning, there is a high frequency of advantageous mutations, so the phenotypic distance d_p to the target structure initially decreases fast. But at long times, as first realized in Ref. 5, the frequency of favorable mutations becomes small and long periods of "stasis" appear where the fitness remains constant.^{5–8} This is illustrated for two typical evolutionary trajectories in Fig. 2. Successive plateaus in fitness are separated by small changes in the Hamming distance to the target: the steps decrease this distance by 1 or 2 units typically, rarely more than that. Also, the time of a stasis period typically increases as the distance to the target decreases: one can speak of diminishing returns for the effects of mutations in approaching the target.



Fig. 2. Plot of the phenotypic distance to a target structure as a function of the number of mutation steps for chains of length 100. Shown are two typical trajectories, displaying periods of stasis.

3.2. Anomalously slow dynamics

To claim that the convergence to the target is particularly slow, it is appropriate to have a comparison benchmark. For this, consider the situation where directional selection is for a target genotype instead of a target phenotype. Just as for phenotypic distances, we define the genotypic distance of two genotypes G_1 and G_2 as $d_G = d_H(G_1, G_2)/L$ where $d_H(G_1, G_2)$ is the Hamming distance between the two strings defining the sequences for G_1 and G_2 . If d_G is then the distance between the *current* genotype and the target genotype, a mutation has a probability $d_G/3$ to produce a strictly better genotype. It follows that the convergence to the target is exponential in time:

$$\langle d_G(t) \rangle \approx d_G(0) e^{-t/3L}.$$
 (1)

Consider now the times τ when a favorable mutation arises; a transition between one stasis period and the next generates such an event which diminishes d_G by 1/L. If $\rho(\tau)$ is the density of these times, we have

$$-\frac{1}{L}\rho(\tau) = \langle d_G(\tau) - d_G(\tau-1) \rangle \approx \frac{d}{d\tau} \langle d_G(\tau) \rangle.$$
(2)

We then obtain $\rho(\tau) = d_G(0) \exp(-\tau/3L)/3$. This shows that selection for a target genotype leads to fast (exponential) relaxation.



Fig. 3. Probability density for the times τ from the beginning of directional selection to transition points (separating two periods of stasis). (L = 100, averaged over 1,000 trajectories for each of 30 randomly chosen targets.) Superposed is the analytical form following from Eq. (2).

The situation is completely different in our system where selection is instead for a target phenotype. Guided by the above analysis, we have determined the corresponding distribution $\rho(\tau)$ in this case. Figure 3 shows the result when averaging over target and initial genotypes at L = 100. This distribution has a very clear fat tail which seems compatible with a power law and not at all with the analytic form derived for selection according to genotypic distance. This same behavior is observed for the other values of L tested (data not shown). One can thus say that convergence to the target under phenotypic selection is "slow", much slower than when compared to genotypic selection.

This difference between the selection types can also be seen by looking at the average distance to the target as a function of evolutionary time. In landscapes associated with disordered systems such as glasses or spin glasses, the relaxation processes encountered typically have non-exponential dynamics.² Empirically, two families of functions have been used to perform fits: the stretched exponential family which for our purposes corresponds to $\langle d_P \rangle \approx A \exp[-(t/T)^{\beta}]$, and the shifted power law family for which

$$\langle d_P(t) \rangle \approx \frac{A}{(B+t)^{\nu}}.$$
 (3)

For small L, both types of fits lead to satisfactory results, but at larger L, the data favor a shifted power law behavior. In Fig. 4, we show that $\langle d_P(t) \rangle$ decreases rather slowly, roughly as an inverse power of time. This leads to a nearly straight line on a log-log plot; also shown in that figure are the fits to shifted power laws. This overall behavior should be contrasted with the law [cf. Eq. (1)] found when using selection for a target genotype: there the approach to the target was much faster.



Fig. 4. Log-log plot of mean phenotypic distance to target structure as a function of the number of mutational steps for chains of length 100 and 200, when averaged over many initial genotypes and target structures. Also shown are the fits to a shifted power law (see text).

3.3. The slow dynamics is sensitive to the target

To further analyse the slow approach towards a target structure, we investigated how the evolutionary dynamics changes with the target. First, we ran simulations for genotypes of length 100 for different targets. In these simulations, for each target, we averaged the relaxation curves for different randomly generated initial genotypes. Figure 5 illustrates how the relaxation is different for different target structures. We conclude that there is slow dynamics whose speed depends on the target structure, a conclusion that holds for all the values of the chain length L we have investigated.

Although we cannot deal with arbitrary L because of computational limitations, it is nevertheless relevant to ask whether this dependence on targets survives for arbitrarily large L. It seems possible that the relaxation behavior is not selfaveraging, i.e. that fluctuations associated with different targets do not become negligible when L grows. To test this, we carried out simulations with multiple targets for different lengths of RNA chains, to determine whether the relaxation curves have smaller dispersion for the different targets when the chain length L increases. We carried out simulations for 30 different targets, and averaged the relaxation curve for each target over 1,000 evolutionary trajectories with random initial genotypes. We then measured the standard deviation σ of the relative distance to the



Fig. 5. Log-log plot of relative phenotypic distance from a target as a function of time for simulations with different targets (all of them of length 60). For each target, the plotted curve is an average over 1,000 different trajectories.

target phenotype at times where the *mean* distance (averaged over all curves) to the target was 10 percent of the chain length. Data are summarized in the following table.

Chain length	$T_{0.1}$	σ
40	499	0.0346
60	1439	0.02306
80	3349	0.02475
100	6933	0.02683
120	8953	0.02626

Here $T_{0.1}$ denotes the mean time at which the average distance to the target phenotype reaches the value 0.1 L, L being the chain length. The simulations were quite time-consuming which prevented us from testing lengths larger than 120. However, from the table we see that the dispersion σ decreases initially, but then remains practically unchanged. The relaxation process towards different targets therefore is compatible with a non-self-averaging behavior.

4. Weak Traces of Evolutionary Innovations

The transition from one period of stasis to the next is due to an evolutionary "innovation". It is appropriate to ask whether the associated transition genotypes can be considered to be atypical according to some measurable quantity. Here we shall study the *short term evolvability* of the entry genotype in a period of stasis and compare it to that of *random* genotypes at the same distance to the target. For a given genotype we define its short-term evolvability^{12, 23–26} as the fraction of single base mutations that lead to a phenotype with higher fitness. We found that the mean fraction of such beneficial mutations is approximately 40% *lower* for the entry genotype than for genotypes randomly sampled from the neutral ensemble. Thus with our definition, the entry genotype has an atypically low evolvability. How can this be qualitatively justified, at least naively? Suppose one thinks of the genotypes of given fitness as concentric spaces surrounding the target as Russian dolls. Suppose further that these spaces have some "thickness". If our "entry" genotype is in the space $S(d_P)$, at phenotypic distance d_P from the target, it will have fewer neighbors belonging to $S(d_P - 1)$ than the typical genotype in $S(d_P)$.

However, we also find that the entry genotype has a smaller fraction of deleterious mutations than the neutral ensemble average. (These mutations produce a phenotype with increased distance to the target.) This result should be contrasted with what is naively expected. Indeed, by construction the entry genotype has one particular mutation which is known to be deleterious (taking it back to the previous plateau). Neglecting all other effects, one would predict that on average the entry genotype would have its fraction of deleterious mutations be 1/(3L) above that of genotypes in the neutral ensemble. Instead, the effect is four times larger and in the opposite direction. These effects are summarized in Fig. 6.

During evolution on a plateau, one goes from the entry genotype (which we just saw is atypical according to some objective measure) to more random genotypes: to



Fig. 6. Comparison at L = 40 between entry ("first") genotypes and random (equilibrium) genotypes of the same fitness for the fraction of mutations that are beneficial (a) or deleterious (b). The x axis gives the phenotypic distance to the target.

some extent, one loses the memory of the entry genotype through successive mutations. Is it a slow change of genotypes that is responsible for the stasis periods? To address this question, we monitored the distance between a mutating genotype and the entry genotype of the corresponding period of stasis. Because L is moderately large in our simulations, the initial growth in distance is linear in the number of accepted mutations. At larger times, we see a saturation effect caused by multiple substitutions, as displayed in Fig. 7. Overall, the evolutionary dynamics on the neutral ensemble shows that even in the absence of phenotypic change, genotypes diffuse on a neutral ensemble at a rate not much slower than if their diffusion was not constrained by this set, until they discover a new phenotype closer to the target.

To understand the expected dynamics in the absence of selection, consider that a given site mutates with probability 1/L at each step, so the relaxation time scales as L. At long times one approaches the average distance 3/4L. This is the behavior, translated mathematically in the Jukes–Cantor formula,²² that the uppermost curve in Fig. 7 represents.



Fig. 7. Mean normalized genotypic distance to the genotype arising at the beginning of the current stasis period as a function of the number of accepted mutations divided by L, for chains of length L = 100 (shown for two different targets). Also shown is the Jukes–Cantor correction for multiple substitutions, which is the function $f(x) = \frac{3}{4}(1 - e^{-\frac{4}{3}x})$ used in evolutionary biology to estimate the normalized distance from the initial genotype after x mutations have taken place. The function f(x) approximates the amount of change to be expected if evolutionary change is unconstrained. Our curves lie somewhat below f(x) which is expected because genotype diffusion is constrained to remain on the neutral ensemble; nevertheless, the data show that diffusion in genotype space is rapid.

1038 P. Šulc, A. Wagner & O. C. Martin

In summary, because of the intricate relation between genotype and phenotype, there can be slow dynamics in the approach to the target phenotype despite rapid evolutionary change of genotypes. In fact, as displayed in Fig. 7, the rate of change of genotypes does not seem to be significantly different when comparing easy and difficult targets for the directional selection.

5. Discussion and Conclusions

Mappings from genotypes to phenotypes play a central role in biology, from the molecular scale up to whole organisms. Working at the level of RNA allowed us to use a framework for such mappings that are not biologically arbitrary, even though it is clearly idealized. One of this mapping's main advantages is that it is computationally tractable. Within this mapping, we showed a rich phenomenology of the evolutionary dynamics towards an optimum phenotype: (1) the "relaxation" towards the target undergoes severe slowing down as the target is approached; (2) this slowing down gives rise to stasis periods with fat tails, typical of what is expected in complex fitness landscapes; (3) the relaxation curves remain sensitive to the choice of the target, even in the limit of long RNA sequences; (4) the diffusion in genotype space during the periods of stasis is not slow, in obvious contrast with what happens at the level of phenotypes.

We observed that the probability of generating a favorable mutation decreases severely as one approaches the optimum, a property of the fitness landscape itself, i.e. the fraction of beneficial mutations goes down in this limit, much faster than in systems undergoing exponential relaxation. Why this leads to inverse power laws remains open, just as in many other landscape problems coming from other fields. Furthermore, because the stasis periods are long, the few innovative genotypes appearing in an evolutionary trajectory represent a tiny fraction of the whole, and thus most genotypes visited have little trace of the out-of-equilibrium dynamics.

Acknowledgments

PS and OCM acknowledge support by the Sixth European Research Framework (proposal number 034952, GENNETEC project). AW acknowledges support from the Swiss National Science Foundation as well as from the YeastX project of SystemsX.ch.

References

- 1. Gavrilets S, *Fitness Landscapes and the Origin of Species*, Princeton University Press, 2004.
- 2. Young AP, Spin Glasses and Random Fields, World Scientific Publishing, 1998.
- Hoos HH, Stützle T, Stochastic Local Search: Foundations and Applications, Elsevier, 2005.
- Eldredge N, Thompson JN, Brakefield PM, Gavrilets S, Jablonski D, Jackson JBC, Lenski RE, Lieberman BS, McPeek MA, Miller III W, The dynamics of evolutionary stasis, *Paleobiology* **31**, 2005.

- Huynen M, Stadler P, Fontana W, Smoothness within Ruggendess: The role of neutrality, Proc Natl Acad Sci USA 93:397–401, 1996.
- Fontana W, Schuster P, Slow evolutionary dynamics of RNA structures and evolvability, *Science* 280(5368):1451–1455, 1998.
- Schuster P, Fontana W, Stadler PF, Hofacker IL, From sequences to shapes and back: A case study in RNA secondary structures, Proc R Soc Lond B 255:279–284, 1994.
- Fontana W, Schuster P, Shaping Space: The possible and the attainable in RNA Genotype-phenotype mapping, J Theor Biology 194(4):491–515, 1998.
- Gesteland RF, Atkins JF (ed.), *The RNA World*, Cold Spring Harbor Laboratory Press, 1993.
- 10. http://www.tbi.univie.ac.at/~ivo/RNA/.
- Fontana W, Stadler PF, Bornberg-Bauer EG, Griesmacher T, Hofacker IL, Tacker M, Tarazona P, Weinberger ED, Schuster P, RNA folding and combinatory landscapes, *Phys Rev E* 47, 1993.
- Wagner A, Robustness and evolvability: A paradox resolved, Proc Biol Sci 275: 91–100, 2008.
- Voss B, Meyer C, Giegerich R, Evaluating the predictability of conformational switching in RNA, *Bioinformatics* 20:1573–1582, 2004.
- 14. Stadler PF, Landscapes and their correlation functions, J Math Chem 20:1–45, 1996.
- Monthus C, Bouchaud J-P, Models of traps and glass phenomenology, J Phys A Gen 29:3847–3869, 1996.
- Bouchaud J-P, Weak ergodicity breaking and aging in disordered systems, J Phys I (France) 2:1705–1713, 1992.
- 17. Hughes BD, Random Walks and Random Environments, Oxford, Clarendon, 1996.
- Kauffman S, Levin S, Toward a general theory of adaptive walks on rugged landscapes, J Theor Biol 128:11–45, 1987.
- Binder K, Heermann DW, Monte Carlo Simulation in Statistical Physics: An Introduction, Springer, 2002.
- Van Nimwegen E, Crutchfield JP, Huynen M, Neutral evolution of mutational robustness, Proc Natl Acad Sci USA 96:9716–9720, 1999.
- Zuker M, Stiegler P, Optimal computer folding of large RNA sequences using thermodynamics and auxiliary information, *Nucleic Acids Res* 9(1):133–48, 1981.
- Kimura M, Ohta T, On the stochastic model for estimation of mutational distance between homologous proteins, J Mol Evol Z:87–90, 1972.
- Huynen MA, Stadler P, Fontana W, Smoothness within ruggedness: The role of neutrality in adaptation, Proc Natl Acad Sci USA 93(1):397–401, 1996.
- 24. Kirschner M, Gerhart J, Proc. Natl. Acad. Sci. USA, Evolvability 95:8420–8427, 1998.
- Meyers LA, Ancel FD, Lachmann M, Evolution of genetic potential, *PLoS Comput Biol* 1:236–243, 2005.
- Cowperthwaite MC, Meyers LA, How mutational networks shape evolution: Lessons from RNA models, Annu Rev Ecol Evol Syst 38:203–230, 2007.



Petr Šulc received his Master's degree in mathematical physics at the Faculty of Nuclear Sciences and Physical Engineering at Czech Technical University in Prague in 2009. He also holds Master's degree in quantum physics from Ecole Polytechnique, Palaiseau, France (2009). From 2008 to 2009, he worked as a research intern at the LPTMS, University Paris-Sud, Orsay, France. From August 2009 until August 2010, he was a graduate research assistant in CNLS, Los Alamos National Laboratory, New Mexico, USA.



Andreas Wagner is professor in the Department of Biochemistry at the University of Zurich in Switzerland, and External Professor at the Santa Fe Institute for the study of Complex Systems. His main research interest is the evolution of biological systems, from genes to complex biological networks with thousands of components. Wagner is the author of the award-winning book *Paradoxical Life* (Yale University Press, 2009), and of *Robustness and Evolvability in Living Systems* (Princeton University

Press, 2005). He received his Ph.D. in 1995 at Yale University, where his research won the J.S. Nicholas prize for best dissertation in his field. He has held research fellowships at several institutions, such as the Institute for Advanced Studies in Berlin, Germany, and the Institut des Hautes Etudes in Bures-sur-Yvette, France. He has authored more than 100 scientific publications, and has lectured widely in the U.S., Canada, Latin America, Europe, Asia, and Australia. He is a member of the Faculty of 1,000 Biology, as well as of the editorial boards of several scientific journals, including *Bioessays*, *BMC Evolutionary Biology*, *Molecular Genetics and Genomics*, and *Molecular and Developmental Evolution*. He lives in Zurich.



Olivier C. Martin is professor in the Departments of Physics and Biology at the University Paris-Sud, France; he is also researcher at INRA, the french national institute for agricultural research. He holds a Ph.D. in theoretical physics from the California Institute of Technology (1983). He is the author of over 100 scientific publications, mostly in physics. More recently his research interests have turned to the use of analytical and computational tools to solve complex problems in genetics and systems biology.