Mathematical Biosciences 276 (2016) 44–58

Contents lists available at ScienceDirect

Mathematical Biosciences

journal homepage: www.elsevier.com/locate/mbs

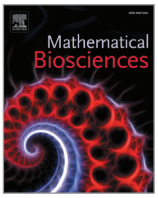

Efficient search and responsiveness trade-offs in a Markov chain model of evolution in dynamic environments

Amor A. Menezes[a,*], Pierre T. Kabamba[b,1]

[a] California Institute for Quantitative Biosciences, University of California, Berkeley, 2151 Berkeley Way, Berkeley, CA, 94704-5230, USA
[b] Department of Aerospace Engineering, University of Michigan, 1320 Beal Avenue, Ann Arbor, MI, 48109-2140, USA

ARTICLE INFO

Article history:
Received 14 May 2015
Revised 30 November 2015
Accepted 4 March 2016
Available online 11 March 2016

Keywords:
Optimization
Stochastic dynamical system
Information
Lyapunov function
Resilience

ABSTRACT

Motivated by the desire to study evolutionary responsiveness in fluctuating environments, and by the current interest in analyses of evolution that merge notions of fitness maximization with dynamical systems concepts such as Lyapunov functions, this paper models natural evolution with a simple stochastic dynamical system that can be represented as a Markov chain. The process maximizes fitness globally via search and has links to information and entropy. These links suggest that a possible rationale for evolution with the exponential fitness functions observed in nature is that of optimally-efficient search in a dynamic environment, which represents the quickest trade-off of prior information about the genotype search space for search effort savings after an environment perturbation. A Lyapunov function is also provided that relates the stochastic dynamical system model with search information, and the model shows that evolution is not gradient-based but dwells longer on more fit outcomes. The model further indicates that tuning the amount of selection trades off environment responsiveness with the time to reach fit outcomes, and that excessive selection causes a loss of responsiveness, a result that is validated by the literature and impacts efforts in directed evolution.

1. Introduction

1.1. Background

Evolution has long been studied as a natural optimization method [1–3] that may explain both observed phenotype (i.e., physical and behavioral trait) diversity [1,4] and optimal phenotype adaptation to an environment [5]. The method is considered by evolutionary ecologists to maximize fitness through natural selection [6] regardless of the definition of fitness [7] or the specifics of the map [8] from genotype (i.e., heritable genetic composition) to phenotype and thence to fitness. The method's power has been experimentally harnessed by the synthetic biology technique of directed evolution [9] to produce unnatural phenotypes [10–12] using tailored selective pressures and designed environment dynamics. As an optimization method, evolution has often been found to embody or cause trade-offs [13,14] that may be further enhanced by variations of the environment [15].

Evolution has also been viewed as a stochastic [3,16,17] search [18] process, with the results of stochasticity analyzed in fluctuating environments [19–21] and modeled by the related concepts of entropy [3,22,23] and information [24,25]. In its stochastic form, the process of evolution has been borrowed by computer science to perform optimization in the guise of genetic algorithms [26–28] (which can be modeled by Markov chains [29]) and evolution strategies [30–33] (which can be successfully deployed in dynamic environments [34] to also obtain biological insight [35] akin to genetic algorithms [36]).

But despite the many studies and models of evolution, there continue to be 'differing views about the efficiency, or optimality, of the adaptation model' of natural selection as a primary driver of evolution [37], and there is thus a still-open question about the optimality of the evolutionary process in addition to a question about whether natural selection maximizes fitness (the recent [38] has more on the latter question). Even the meaning of fitness maximization is unclear [38]: four varieties that are regarded by biologists for a given population are (1) the equivalence between a stable genetic equilibrium and mean fitness maximization, (2) the increase in mean fitness by natural selection when not at genetic equilibrium, (3) the equivalence between a stable genetic equilibrium and the adoption of a phenotype by all organisms that maximize individual fitness, and (4) the increase in the number of organisms adopting a phenotype that maximizes individual fitness by natural selection when not at genetic equilibrium [38].

Nevertheless, to account for, as [38] puts it, 'some cases [where] evolution by natural selection has led to traits that approximately

---

[*] Corresponding author. Tel.: +15106433722.
E-mail address: amenezes@berkeley.edu (A.A. Menezes).
[1] Deceased.

http://dx.doi.org/10.1016/j.mbs.2016.03.002
0025-5564/

maximize fitness within a set of feasible alternatives [6,37,39,40],' fresh analyses of evolution have been pursued, using inclusive fitness [41,42] for instance. The ongoing Formal Darwinism (FD) project [43–46] calls for 'a new kind of argument...to link equations of motion on the one hand to optimization programs on the other, and a major point is that the biologist's concept of fitness maximization is not represented by concepts from dynamical systems such as Lyapunov functions and gradient functions' [45]. The FD project seeks to axiomatize fitness and link equilibrium concepts of population genetics to solution concepts of optimization programs [46]. The FD project tries to vindicate the "individual as maximizing agent analogy" [38,45,47] and show that phenotypes present in an equilibrium state are optimal because no other phenotype does better in that equilibrium state.

1.2. Goals, biological meaning of employed terms, and preview of results

This paper provides a simple stochastic dynamical system model of natural evolution that can be represented by a discrete-time homogeneous Markov chain (see Appendix A for a background on Markov chains), with the initial model motivation consisting of examining responsiveness in the immediate aftermath of an environment fluctuation. The goals of the model are to understand what happens to the evolution process in a dynamic environment rather than to determine exact phenotype outcomes, and, like the FD project, to specify the form of the fitness function as well as to investigate the effects of changes in process dynamics on understood equilibrium concepts within the context of optimization.

No constraints are imposed on the meaning of fitness in this paper, and the model is initially developed with an abstract fitness function that ascribes value to individual phenotypes. Because the population size is restricted to one at every time step with this model for analytical simplicity, and because the abstractness of the fitness function imposes no restrictions on the constituents of a fitness valuation, it is possible to interpret the results in this paper in the context of either population mean fitness (which is trivially equivalent here to the fitness of the population's sole phenotype at a time step) or individual fitness (which includes the fitness value effect of strategies pursued by a phenotype, a possibility that is allowed by the abstract fitness function). Since model extensions will lift the unity population size restriction so that the trivial equivalence described above no longer exists, an interpretation of individual fitness is a favored one. But for this paper at least, both population mean fitness and individual fitness are mathematically-plausible meanings of fitness.

The dynamics of the model, which assumes discrete generations and single reproductive events per individual per time step, embody changes caused by natural selection at each time step. Fitness maximization in this paper is also open to both dynamical interpretations debated by biologists: a process of increasing population mean fitness, or a process of increasing adoption of a phenotype that maximizes individual fitness. These interpretations will be elaborated upon with respect to the model where appropriate. Natural selection is taken to have the traditional meaning: local competition between phenotypes in a population, with phenotypic traits arising from genotypes in a way that includes environment effects [48].

Unsurprisingly, the local selection that is described by the model maximizes fitness globally via search, and the process also has links to the concepts of information and entropy, although these links are not imposed *a priori*. What is surprising is that, when determining the form of the fitness function, a possible rationale for the evolutionary process emerges. Borne out by the kinds of fitness functions that exist in nature, this rationale constitutes optimal search efficiency in a dynamic environment, and it represents the quickest trade-off of prior information about the search space for search effort savings after an environment fluctuation occurs. Moreover, a Lyapunov function exists to relate the dynamical representation of the fitness-maximizing evolution process with search information. It is further shown that this dynamical process dwells longer on more fit outcomes, whatever they may be, instead of being gradient-based.

Additionally, insights into the effects of varying levels of selection are obtainable from the model, such as a trade-off between responsiveness to a dynamic environment and the time to reach a fit outcome. A related detrimental result (which stipulates that excessive selection causes a loss of responsiveness) is shown to be corroborated by numerous literature examples. This result has implications for the fruits of directed evolution efforts. Lastly, the similarities and differences between the model's Markov chain optimization method and other methods that are used in computer science (genetic algorithms, evolution strategies, and simulated annealing [49]) are briefly described, with the model in this paper recapturing a known Markov chain Monte Carlo technique that was originally proposed to model physical phenomena [50].

2. Methods

2.1. Problem definition

As others have done before us (e.g., [18]), let us consider evolution to be a form of stochastic search, which looks for some genotype(s) that result(s) in some desirable phenotype(s). Let the set of genotypes, X, be a finite albeit large one, consisting of genotypes x_i, $1 \le i \le n$. As we shall see in Section 3.3, the fact that X is fixed does not preclude any "evolution of evolvability" [51] in a dynamic environment, a concept that we take to be a change in the ability to respond to a selection process. Let the set of phenotypes be called Z. Let the genotype-phenotype mapping be denoted by z, i.e., z: $X \to Z$, which is a function that is very important to specify when determining the phenotype outcomes of evolution, but since we are more interested in insights from a model of evolution, let us simply say that z is an unknown function that can change with time because it includes the effects of a changing environment. For notational simplicity, we will omit explicitly denoting a function's dependence on time.

Suppose that there exists some desirable phenotype z_{des} for a particular environment (possibly, but not necessarily, a phenotype that is "best" for its environment); of course, z_{des} can also change with time to reflect environment changes, and we may not know z_{des}. We will assume that it is possible to measure differences between any two phenotypes, i.e., Z is a metric space. Let the discrepancy between the phenotype that results from a genotype x, which is $z(x)$, and the desirable, possibly unknown, phenotype z_{des} be denoted by $\|z(x) - z_{des}\|$.

Since we model evolution as a stochastic search process, we consider it to produce a probability mass function ϕ_X over the set of genotypes, $\phi_X : X \to \mathbb{R}^+$, and our model will provide dynamic transition laws that cause X to be distributed according to ϕ_X. We have postulated that evolution searches for a desirable phenotype (whatever "desirable" means), so on average, the process results in

$$\mathrm{E}_{\phi_X}[\|z(x) - z_{des}\|] = 0. \tag{1}$$

Let $y(x) = \|z(x) - z_{des}\|$, so that we can rewrite the above as $\mathrm{E}_{\phi_X}[y(x)] = 0$, where y inherits the time-dependence of z. Because we do not know z, and we may not know z_{des}, $y(x)$ is effectively an unknown function for which we know an expectation.

Let us now consider phenotype fitness, however one chooses to define fitness. Let us simply say that there exists a function f