



Fluctuation domains in adaptive evolution

Carl Boettiger^{a,*}, Jonathan Dushoff^{b,*}, Joshua S. Weitz^{c,d,*}

^a Center for Population Biology, University of California, Davis, United States

^b Department of Biology, McMaster University, Hamilton, ON, Canada

^c School of Biology, Georgia Institute of Technology, Atlanta, United States

^d School of Physics, Georgia Institute of Technology, Atlanta, United States

ARTICLE INFO

Article history:

Received 22 April 2009

Available online 18 October 2009

Keywords:

Evolution

Fitness landscapes

Fluctuation dissipation

Fluctuation enhancement

Canonical equation

ABSTRACT

We derive an expression for the variation between parallel trajectories in phenotypic evolution, extending the well known result that predicts the mean evolutionary path in adaptive dynamics or quantitative genetics. We show how this expression gives rise to the notion of fluctuation domains – parts of the fitness landscape where the rate of evolution is very predictable (due to *fluctuation dissipation*) and parts where it is highly variable (due to *fluctuation enhancement*). These fluctuation domains are determined by the curvature of the fitness landscape. Regions of the fitness landscape with positive curvature, such as adaptive valleys or branching points, experience enhancement. Regions with negative curvature, such as adaptive peaks, experience dissipation. We explore these dynamics in the ecological scenarios of implicit and explicit competition for a limiting resource.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

Fitness landscapes have long been an important metaphor in evolution. Wright (1931) originally introduced the concept to explain his result that the mean rate of evolution of a quantitative trait is proportional to gradient of its fitness. The result and its accompanying metaphor continue to arise in evolutionary theory, having been derived independently in quantitative genetics (Lande, 1979), game-theoretic dynamics (Hofbauer and Sigmund, 1998; Abrams, 1993), and adaptive dynamics (Dieckmann and Law, 1996). The metaphor creates a deterministic image of evolution as the slow and steady process of hill climbing. Other descriptions of evolution have focused on its more stochastic elements – the random chance events of mutations and the drift of births and deaths that underlie the process (Kimura, 1983, 1968; Ohta, 2002). In this manuscript, we seek to characterize the deviations or fluctuations of evolutionary trajectories around the expected evolutionary path.

Our main result is that the size of the deviations of evolutionary trajectories is determined largely by the *curvature* of the evolutionary landscape whose *gradient* determines the mean rate of evolution. The landscape metaphor can be used to understand the interplay of stochastic and deterministic forces by identifying regimes where the selection will counterbalance or enhance such

stochastic fluctuations. Because the curvature can be positive (concave up) or negative (concave down), the evolutionary landscape can be divided into domains where fluctuations are enhanced or dissipated. To make this more precise, we focus on a Markov model of evolution used in the theory of adaptive dynamics. This Markov model gives rise to the familiar gradient equation – the one which first inspired the fitness landscapes metaphor – and also a more precise statement of the fluctuation dynamics. The adaptive dynamics framework is general enough that we can consider how these results apply in a variety of ecological scenarios. We illustrate the surprising consequence of bimodal distributions of expected phenotypes among parallel trajectories emerging from a fixed starting point on a single-peak adaptive landscape due to fluctuation enhancement. We also provide a landscape interpretation that suggests how the ideas of fluctuation dynamics apply to other ecological and evolutionary models.

2. Theoretical construction

Various definitions of fitness landscape have been used in the study of evolutionary dynamics. The conventional idea postulates a mapping between trait values and fitness (Levins, 1962, 1964; Rueffler et al., 2004). The difficulty with this conception is that, in many cases, the fitness of a given population type is mediated largely through competition with other populations, and is thus dependent on the current distribution of populations and their respective phenotypes in the environment. This dependence on densities or frequencies turns the static fitness landscape into

* Corresponding author.

E-mail addresses: cboettig@ucdavis.edu (C. Boettiger), dushoff@mcmaster.ca (J. Dushoff), jsweitz@gatech.edu (J.S. Weitz).

a dynamic landscape, whose shape changes as the number of individuals, and their corresponding traits, changes.

It is possible to depict the fitness landscape that emerges even in the face of density- or frequency-dependent competition by returning to Wright's notion of a landscape. To do so, consider a resident population of individuals in an environment, each of which have an identical trait value, *i.e.*, a monomorphic population. Next, consider the per-capita fitness of a small number of mutant individuals with a similar, but different trait. The per-capita fitness of mutants in the environment may be larger, smaller or identical to the per-capita fitness of residents. The gradient in fitness around the trait represented by the resident is termed the selective derivative (Geritz et al., 1997) in adaptive dynamics (analogous to the selective gradient in quantitative genetics). The fitness landscape is defined by integrating over these derivatives in the trait space, to obtain a *local* picture of how fitness changes. In the limit of small, slow mutation, it can be shown that the population will evolve as predicted by the shape of the landscape: by climbing towards a peak most rapidly when the slope is steep (Rueffler et al., 2004), see Fig. 1 for three illustrative examples. The trait on the horizontal axis is that of the resident, not the mutant. The lower panels show the slope of the mutant fitness as a function of the resident trait (up to a multiplicative factor)– the slope of the mutant fitness changes as the resident trait changes. The trait of the resident population evolves in the direction given by that slope. Integrating the lower panel along the resident trait axis reveals the fitness landscape the resident trait climbs in an evolutionary process (top panels). Note that changes in the mutant fitness are captured in the changes in slope of the landscape; the fitness landscape itself remains fixed. We derive the expected evolutionary dynamics of the phenotypic trait of the resident and its variation among parallel trajectories in the following sections.

2.1. Model

Consider a population monomorphic for a particular phenotypic trait, x . The abundance of the population, $N(x, t)$ of trait x is governed by its ecological dynamics:

$$\frac{dN(x)}{dt} = f(x, N, E), \quad (1)$$

which may depend on the trait, population abundance, and the environmental conditions, E . The evolutionary dynamics proceed in three steps (Dieckmann and Law, 1996; Champagnat et al., 2006). First, the population assumes its equilibrium-level abundance $N^*(x)$, determined by its phenotypic trait x . Next, mutants occur in the population at a rate given by the individual mutation rate μ times the rate of births at equilibrium, $b(x)$. The mutant phenotype, y , is determined by a mutational kernel, $M(x, y)$. The success of the mutant strategy depends on the invasion fitness, $s(y, x)$, defined as the per-capita growth rate of a rare population. Mutants with negative invasion fitness die off, while mutants with positive invasion fitness will invade with a probability that depends on this fitness (Geritz et al., 1997).

The invasion fitness is calculated from the ecological dynamics, Eq. (1), and will in general depend on the trait x of the resident population as well as that of the mutant, y . We assume that a successful invasion results in replacement of the resident (Geritz et al., 1997, 2002), and the population becomes monomorphically type y . The details of this formulation can be found in Appendix A. The important observation is that such a model can be represented by a Markov process on the space of possible traits, x . The population can jump from any position x to any other position y at a transition rate $w(y|x)$ determined solely by the trait values x

and y . The probability that the process is at state x at time t then obeys the master equation for the Markov process,

$$\frac{d}{dt}P(x, t) = \int dy [w(x|y)P(y, t) - w(y|x)P(x, t)]. \quad (2)$$

No general solution to the master equation Eq. (2) exists. However, if the jumps from state x to state y are sufficiently small (*i.e.*, if mutants are always close to the resident), we can obtain approximate solutions for $P(x, t)$ by applying a method known as the Linear Noise Approximation (van Kampen, 2001), Appendix B. Doing so yields a general solution for the probability $P(x, t)$ in terms of the transition rates, $w(y|x)$, Eq. (23). The linear noise approximation *derives* a diffusion equation as an approximation to the original jump process, as is commonly postulated. While this diffusion can be written as a partial differential equation (PDE), following Kimura, or as a stochastic differential equation, the solution for the probability density can be proven to be Gaussian and hence it suffices to write down ordinary differential equations for the first two moments (Kurtz, 1971).

2.2. The fluctuation equation

We assume the mutational kernel $M(y, x)$ is Gaussian in the difference between the resident and mutant traits, $y - x$, with width σ_μ and that mutations occur at a rate μ . A more thorough discussion of these quantities can be found in Appendix A, where they are developed in the process of deriving the transition probability $w(y|x)$ that specifies the underlying Markov process. Using Eq. (23) which results from the systematic expansion of the master equation (2), the mean trait \hat{x} obeys

$$\frac{d\hat{x}}{dt} = \frac{1}{2}\mu\sigma_\mu^2 N^*(\hat{x})\partial_y s(y, \hat{x})|_{y=\hat{x}} \equiv a_1(\hat{x}), \quad (3)$$

with an expected variance σ^2 that obeys

$$\frac{\partial\sigma^2}{\partial t} = 2\sigma^2\partial_{\hat{x}}a_1(\hat{x}) + 2\sqrt{\frac{2}{\pi}}\sigma_\mu |a_1(\hat{x})|. \quad (4)$$

Eq. (3) recovers the familiar canonical equation of adaptive dynamics (Dieckmann and Law, 1996) for the mean trait. Eq. (4), which describes the variance, we term the *fluctuation equation* of adaptive dynamics. The linear noise approximation (see Appendix B) also predicts that the probability distribution itself is Gaussian, so that Eqs. (3) and (4) determine the entire distribution of possible trajectories in this approximation. To illustrate the local behavior of the model, we define the fitness landscape as $L(x) = \int_0^x a_1(y) dy$.

3. Fluctuation domains

A closer look at Eq. (4) will help motivate our geometric interpretation of fluctuation domains. The approximation requires that the mutational step size σ_μ is very small, hence the second term will almost always be much smaller than the first (except when the fluctuations σ themselves are also very small). Note that the second term resembles the right-hand side of (3), differing only by a small scalar constant and the fact that it is always positive. Interestingly, this means that the second term vanishes at the singular point where $a_1(\hat{x}) = 0$ and $\sigma^2 = 0$: fluctuations cannot be introduced at an evolutionary equilibrium. That the first term is the gradient of the deterministic (mean) dynamics and the second term is smaller by a factor of the small parameter in the expansion are general features of the linear noise approximation.

The first term of (4) implies that the variance σ^2 will increase or decrease exponentially at a rate determined by $\partial_{\hat{x}}a_1(\hat{x})$; *i.e.*, the *curvature* of the fitness landscape. This landscape and its gradient




Fig. 1. Fluctuation domains. The lower panels show the rate of evolution given by Eq. (3) vs. the resident trait x for three ecological models: (a) implicit competition for a limiting resource (Section 4), (b) explicit resource competition (Appendix C), and (c) symmetric branching for dimorphic populations (with a resident population of trait x and another at $-x$, Appendix E). The horizontal dashed lines in the lower panel denote where $a_1(x) = 0$, i.e., the selective derivative is zero, and correspond to evolutionarily singular points, e.g., adaptive peaks and valleys. Shaded and unshaded regions in the lower panels correspond to trait values for which fluctuations are enhanced (shaded) and dissipated (unshaded). The adaptive landscape is defined as the integral of this expression (upper panels) – where populations climb the hills at a rate proportional to their steepness. Note that the adaptive landscape for the symmetric branching case has two peaks, and so evolutionary trajectories will lead to a stable dimorphism.

are depicted in Fig. 1 for several ecological scenarios. Plotting the gradient itself makes it easier to see when fluctuations will increase or decrease. On the gradient plot the singular point is found where the curve crosses the horizontal axis.

In the neighborhood of a singular point corresponding to an adaptive peak, the slope is negative, hence the first term in Eq. (4) (the coefficient of σ^2) is negative and fluctuations dissipate exponentially. This means that two populations starting with nearby trait values will converge to the same trajectory. In this region no path will stochastically drift far from the mean trajectory, hence the canonical equation will provide a good approximation of all observed paths. We term the part of the trait-space landscape where $\partial_x a_1(x) < 0$ the fluctuation-dissipation domain, analogous to the fluctuation-dissipation theorem found in other contexts (van Kampen, 2001).

Farther from the singular point corresponding to an adaptive peak, $\partial_x a_1(x)$ becomes positive (see Fig. 1(a) and (b)). While this part of trait-space still falls within the basin of attraction of the singular point, the variance σ^2 between evolutionary paths will grow exponentially. Initially identical populations starting with trait values in this region will experience divergent trajectories due to this enhancement. The variation can become quite large, with evolutionary trajectories that differ significantly from the mean. We call this region of trait-space where $\partial_x a_1(x) > 0$ the fluctuation enhancement domain. Eventually trajectories starting in this region will be carried into the fluctuation-dissipation domain, where they will once again converge.

Evolutionary branching points (Dieckmann and Doebeli, 1999; Geritz et al., 1998) provide another example of a fluctuation enhancement domain. Until now, we have assumed that successful mutants replace the resident population, a result known as “invasion implies substitution” (Geritz et al., 2002). However, at an evolutionary branching point, the mutant invader no longer replaces the resident population, and the population becomes dimorphic. In dealing with monomorphic populations, we have been able to describe the evolutionary dynamics in terms of the change of a single trait, describing a single resident population. For dimorphic populations, two resident populations coexist, and both influence the shape of the landscape. Thus, an invader’s fitness, $s(y; x_1, x_2)$, and its rate of evolution, $a_1(x_1, x_2)$, will depend on how it performs against both populations. The evolutionary dynamics after branching are two-dimensional and require a multivariate version of Eq. (4). Despite this, we can still gain qualitative insight using the intuition that connects fluctuation domains to curvature.

The existence of the stable dimorphism fundamentally distorts the landscape. While the population was monomorphic, being

closer to the singular point always ensured a higher fitness. Once a resident population sits on either side of a branching point though, the singular point becomes a fitness *minimum*. This description of evolution towards points which become fitness minima after branching has been addressed extensively elsewhere (Geritz et al., 1997, 1998; Geritz, 2004). Here, what interests us is the effect of branching on fluctuations. If the landscape is smooth, a minimum must have a positive curvature and therefore must be an enhancement domain. The farther a mutant gets from the branching point, the faster it can continue to move away, thus enhancing the initially small differences between mutational jumps. A rigorous description of this effect would require a multivariate version of Eq. (4) which is beyond the scope of this paper. Instead, we illustrate the enhancement effect by taking a slice of the two-dimensional landscape by assuming that the dimorphic populations have symmetric trait values about the singular strategy, $x_1 = -x_2 = x$. The fitness landscape for the case of symmetric branching is shown in Fig. 1(c), and the derivation provided in Appendix D. Note that the region around the branching point is a fluctuation enhancement regime and the two new fitness maxima far from the branching point are in fluctuation dissipation regimes.

To provide a more concrete illustration of the fluctuation dynamics, we will focus on the description of the ecological competition model and compare the theoretical predictions of Eq. (3) and Eq. (4) to point-process simulations of the Markov process (Gillespie, 1977).

4. An example of fluctuation domains in resource competition

4.1. An ecological model of implicit competition for a limiting resource

The logistic model of growth and competition in which populations compete for a limited resource is a standard model in ecological dynamics (Dieckmann and Doebeli, 1999). Here, we consider the population dynamics of $N(x, t)$ individuals each with trait x :

$$\frac{dN(x, t)}{dt} = rN(x, t) \left(1 - \frac{\sum_y N(y, t)C(x, y)}{K(x)} \right), \quad (5)$$

where r is the birth rate and $\sum_y rN(y)C(x, y)/K(x)$ is the density-dependent death rate. In the model, $K(x)$ is the equilibrium population density and $C(x, y)$ is a function which describes the relative change in death rate of individuals of type x due to

competition by individuals of type y . Given $C(x, x) = 1$, the equilibrium density of a monomorphic population is $N^*(x) = K(x)$. Following Dieckmann and Doebeli (1999) we assume the following Gaussian forms,

$$K(x) = K_0 e^{-x^2/(2\sigma_k^2)}, \quad (6)$$

and

$$C(x, y) = e^{-(x-y)^2/(2\sigma_c^2)}, \quad (7)$$

where σ_k and σ_c are scale factors for the resource distribution and competition kernels respectively. We focus on the case $\sigma_c > \sigma_k$, for which the model has a convergence-stable evolutionarily stable strategy (ESS) at $x = 0$ (Geritz et al., 1998). Having specified the ecological dynamics, we must also specify the evolutionary dynamics, which we assume occur on a much slower time scale. We assume that with probability μ an individual birth results in a mutant offspring, and that the mutant trait is chosen from a Gaussian distribution centered at the trait value of the parents and with a variance σ_μ^2 . From this we can assemble $w(y|x)$, and compute equations (3) and (4) (see Appendix B).

4.2. Numerical simulation of stochastic evolutionary trajectories

There are many ways to represent evolutionary ecology models in numerical simulations, including finite simulations, fast-slow dynamics, and point processes (Dieckmann and Doebeli, 1999; Nowak and Sigmund, 2004; Champagnat et al., 2006; Champagnat and Lambert, 2007). We choose to model the evolutionary trajectories of the ecological model in Eq. (5) by explicitly assuming a separation of ecological and evolutionary time scales. First, given a trait x , the ecological equilibrium density $N^*(x)$ is determined. Then, the time of the next mutant introduction is calculated based on a Poisson arrival rate of $\mu r N^*(x)$. In essence, we are assuming that ecological dynamics occur instantaneously when compared to evolutionary dynamics. The trait of the mutant, y , is equal to x plus a normal deviate with variance σ_μ^2 . The mutant replaces the resident with probability given by the standard branching process result, $1 - d(y, x)/b(y, x)$ where d and b are the per-capita birth and death rate of the mutant, respectively, so long as $b > d$, and with probability 0 otherwise. The ecological steady state is recalculated and the process continues. Simulations are done using Gillespie's minimal process method (Gillespie, 1977). Estimates of the mean phenotypic trait \hat{x} and the variance σ^2 are calculated from ensemble averages of at least 10^5 replicates.

4.3. Comparison of theory and simulation

We consider simulations with three different initial conditions: starting within the fluctuation-dissipation domain ($x = 1$), just into the fluctuation-enhancement domain ($x = 2$), and deep into the fluctuation enhancement regime ($x = 3$) in successive rows in Fig. 2. For each condition, the simulation is repeated 10^5 times and we compare the resulting distribution of trajectories to that predicted by theory. In the first column we plot the numerically calculated mean path taken to the ESS (shown in circles), and compare to the theoretical prediction of Eq. (3). In the second column, we plot the variance between paths, and compare to the predictions of Eq. (4). As each replicate begins under identical conditions, the initial variance is always zero. In the third column, we present the distribution of traits among the replicates at a single instant in time. We chose the moment of maximum variance so that deviations from the theoretically predicted Gaussian kernel will be most evident.

4.3.1. Fluctuation-dissipation domain

Starting at the edge of the dissipation domain, fluctuations grow for only a short time before rapidly decaying, in the first row of Fig. 2. Since the simulation is sampled at regular time intervals, the vertical spacing of points indicate the rate of evolution. Theoretical predictions closely match simulation results for both the mean and variance, and the resulting distribution matches the theoretically predicted Gaussian. The variance represents an almost negligible deviation from the mean trajectory.

4.3.2. Fluctuation-enhancement domain

Our second ensemble at $x = 2$ begins clearly within the fluctuation-enhancement domain. In the second row of Fig. 2, the variance plot (center panel) begins by growing exponentially. Once the population crosses into the dissipation domain, at $x = 1$, at about $t = 1000$, the variance peaks and then dissipates exponentially. This transition appears as an inflection point in the mean path (left panel). Though the variation now represents significant deviations, the distribution still appears Gaussian (right panel).

4.3.3. Strong fluctuations

In the third row of Fig. 2, starting deep within the fluctuation-enhancement domain at $x = 3$, theory and simulation agree initially. Remaining long enough in this domain, fluctuations are eventually of the same order of magnitude as the macroscopic dynamics themselves, and the linear noise approximation underlying the theory begins to break down. The theory and simulation variance diverge, while the mean path is substantially slower than predicted by the canonical equation. In the right panel, the reason for this becomes clear. At this point in the divergence, the distribution is far from Gaussian, rather it has become bimodal, with some replicates having reached the stable strategy at $x = 0$ and others having hardly left the initial state.

The mechanism behind the emergence of bimodal probability distributions for trajectories can be seen in the final row of Fig. 2. The theoretical trajectory predicted by Eq. (3) is sigmoidal, comprised of slow change at the beginning and end and rapid evolution in the middle. Ecologically, this arises at the beginning from the small population size available to generate mutations, and at the end from the weakening selection gradient. Populations with intermediate trait values hence experience rapid trait evolution, giving rise to this transiently bimodal distribution. The theory provides insight into this case of strong fluctuations in two ways – first, it is driven by the existence of a fluctuation-enhancement domain, and second, it is anticipated by the theoretical prediction of fluctuations that are of the same order as the macroscopic dynamics. We observe that these macroscopic fluctuations can arise whenever the population remains long enough in the enhancement regime, independent of other model details. The fluctuation equation for the explicit resource competition (chemostat) model shown in Fig. 1(b) and described in Appendix C never predicts fluctuations that reach macroscopic values, and simulations of this system (not shown) do not produce bimodal trait distributions.

5. Discussion

The metaphor of fitness landscapes has been a powerful tool for understanding evolutionary processes (Wright, 1932; Lande, 1979; Abrams, 1993; Dieckmann and Law, 1996). The concept of a landscape arises naturally from describing the rate of change of the mean trait as being proportional to the evolutionary gradient. We extend this metaphor to understand fluctuations away from this mean, where our analysis of a Markov process representation

Fig. 2. Simulations over fluctuation domains. The dynamics of the mean path, variance between paths, and snapshots of the trait distribution in time for three initial conditions: $x_0 = 1$, $x_0 = 2$, and $x_0 = 3$ (top to bottom successively). The three rows correspond to trajectories experiencing fluctuation dissipation (top) and fluctuation enhancement (bottom) along with an intermediate case (middle). The mean (first column) and variance (second column) among trait values are plotted over time. Circles are simulation averages from 10^5 replicates and lines are theoretical predictions. The final column shows the theoretical Gaussian distribution (solid line) at the point in time indicated by the dashed line and closed circle in the first two columns, and the actual histogram (circles) of positions across the replicates at that time. The bottom right panel shows a bimodal distribution among replicates in the case of fluctuation enhancement. **Parameters:** $\sigma_k^2 = 1$, $\sigma_c^2 = 1.01$, $r = 10$, $\sigma_\mu = 0.0005$, $\mu = 1$.

of evolution in the adaptive dynamics framework brings us back to the notion of a fitness landscape. In this case, it is the curvature of that landscape which informs the dynamics of these fluctuations away from the expected evolutionary path. Our result linking landscape curvature to fluctuation domains is consistent with similar results from quantitative genetics that relate the sign and magnitude of the quadratic selection coefficient in the Price equation to increases or decreases in trait variation (Lande and Arnold, 1983; Chevin and Hospital, 2008).

Though the mathematical analysis of the Markov process is somewhat technical, the result is generally intuitive. Graphs such as Fig. 1 provide a visualization of how fluctuations will behave on the landscape, while our fluctuation equation, Eq. (4), provides a more explicit description of those fluctuations. This result, like the gradient expression itself, relies on the underlying randomness being small relative to the scale of evolutionary change. We have seen how these approximations can break down in the case of very large fluctuations, resulting in a bimodal distribution of phenotypes. While Eq. (4) fails to describe this case accurately, it predicts the breakdown of the approximation by the explosion of fluctuations. Though we have focused on the adaptive dynamics model of evolution, we hope that the approach taken here can be extended to other landscape representations. Further, we hope that our model will be broadly applicable to assessing the repeatability of evolution in experimental studies of microbes

and viruses (Wichman et al., 1999; Cooper et al., 2003) and in ecological studies of rapid phenotypic change caused by biotic interactions (Duffy and Sivars-Becker, 2007) and/or anthropogenic factors (such as over-fishing) (Olsen et al., 2005).

In this work, we have compared theoretical predictions to numerical simulations. Even when a closed form expression such as (4) exists for the deviations, some of the most dramatic examples in Fig. 2 emerge only when the diffusion approximation breaks down and stochastic forces become macroscopic. In the spirit of scientifically reproducible research (Gentleman and Lang, 2007; Schwab et al., 2000; Stodden, 2009), we freely provide all the source code required to replicate the simulations and figures shown in the text. Though the numerical simulations are written in C for computational efficiency, we provide a user interface and documentation by releasing all the code, figures, text, and examples as a software package for the widely used and freely available R statistical computing language. This package can be downloaded from <http://ecothery.biology.gatech.edu/downloads/>.

Acknowledgments

We thank Sebastian Schreiber for many helpful discussions, and also thank Michael Turelli, Géza Meszéna and an anonymous reviewer for comments. This work is funded in part by the Defense

