

Interim Report

IR-12-052

Adaptive dynamics

Johan A.J. (Hans) Metz (metz@iiasa.ac.at)

Approved by

Ulf Dieckmann
Director, Evolution and Ecology Program

February 2015

Adaptive Dynamics

J.A.J. Metz (Hans)

Evolution and Ecology Program, International Institute for Applied Systems Analysis,
A-2361 Laxenburg, Austria

Mathematical Institute & Institute of Biology, Leiden University
P.O. Box 9512, 2300AA Leiden, Netherlands

NCB Naturalis, Darwinweg 2, 2333 CR Leiden, Netherlands

j.a.j.metz@biology.leidenuniv.nl

p. 7 – 17 in
Alan Hastings & Louis J. Gross eds. (2012)
Encyclopedia of Theoretical Ecology
California University Press

Adaptive dynamics (AD) is a mathematical framework for dealing with eco-evolutionary problems, primarily based on the following simplifying assumptions: clonal reproduction, rare mutations, small mutational effects, smoothness of the demographic parameters in the traits, and well-behaved community attractors. However, often the results from AD models turn out to apply also under far less restrictive conditions. The main AD tools are its so-called canonical equation (CE) that captures how the trait value(s) currently present in the population should develop over evolutionary time, and graphical techniques for analyzing evolutionary progress for one-dimensional trait spaces like ‘pairwise invasibility plots’ (PIPs) and ‘trait evolution plots’ (TEPs). The equilibria of the CE, customarily referred to as evolutionarily singular strategies, or ess-es, comprise in addition to the evolutionary equilibria, or ESSes, also points in trait space where the population comes under a selective pressure to diversify. Such points mathematically capture the ecological conditions conducive to adaptive (Darwinian) speciation.

This research benefited from the support of the Chair Modélisation Mathématique et Biodiversité VEOLIA-École Polytechnique-MNHN-F.X.

1. Context

1.1. Micro-, meso- and macro-evolution

Adaptive dynamics was initiated as a simplified theoretical approach to *meso-evolution*, defined here as evolutionary changes in the values of traits of representative individuals and concomitant patterns of taxonomic diversification. This in contrast to *micro-evolution*, i.e., changes in gene frequencies on a population dynamical time scale, and *macro-evolution*, a term that then can be reserved for changes like anatomical innovations, where one cannot even speak in terms of a fixed set of traits.

Meso-evolution is more than micro-evolution writ large, and a similar statement applies to macro- versus meso-evolution. Each of these levels has its own emergent phenomena, and its own explanatory frameworks, which in the end should be based on idealised summaries of the outcomes of lower level mechanisms. Trait changes result from the micro-evolutionary process of mutant substitutions taking place against the backdrop of a genetic architecture and developmental system as deliverers of mutational variation, internal selection keeping the machinery of a body in concert, and ecological selection due to the interactions of individuals with their conspecifics, resources, predators, parasites and diseases. AD focuses on these encompassing mechanisms.

To get a clean story AD assumes a time scale separation between ecology and evolution. In reality this assumption holds good only rather rarely. The idea is that yet arguments based on it may often lead to outcomes that are fair approximations, provided one applies them selectively and takes a sufficiently gross look at reality. The time scale argument considerably eases the transition from population genetics to the perspective of ecologists, morphologists and taxonomists. AD aims at addressing that larger picture at the cost of being wrong in the details.

1.2. The fitness landscapes of meso-evolution

Meso-evolution proceeds by the selective filtering by the ecology of a continual stream of mutations. AD concentrates on the ecological side of this process, as at that end there are clearer a priori mathematical structures. The basic theory assumes clonal reproduction, and only a subset of the results extends to the Mendelian case, for monomorphic populations directly, for polymorphic populations after appropriate modification.

When approaching the evolutionary process from ecology, one sees immediately that fitnesses are not given quantities. They depend not only on the traits of individuals but also on their environment. The ecological feedback loop makes that in the monomorphic and clonal cases necessarily the fitnesses of all types present on the ecological time scale are zero. Only the fitnesses of potential mutants can be positive or negative. The signs and sizes of these mutant fitnesses determine the direction and speed of evolution. Evolution corresponds to uphill movement in fitness landscapes that keep changing so as to keep the resident types at zero (see Figure 1).

The main insight from the mathematical analysis of this picture has been the discovery of a potential mechanism for adaptive speciation that appears with a certain ubiquity in ecological models. Apart from that, the theory has produced effective tools for analysing special families of eco-evolutionary models.

Below it is every time assumed first that individuals reproduce clonally, when relevant followed by a discussion of the Mendelian case. In line with the landscape analogy, zero is referred to as sea level, etc.

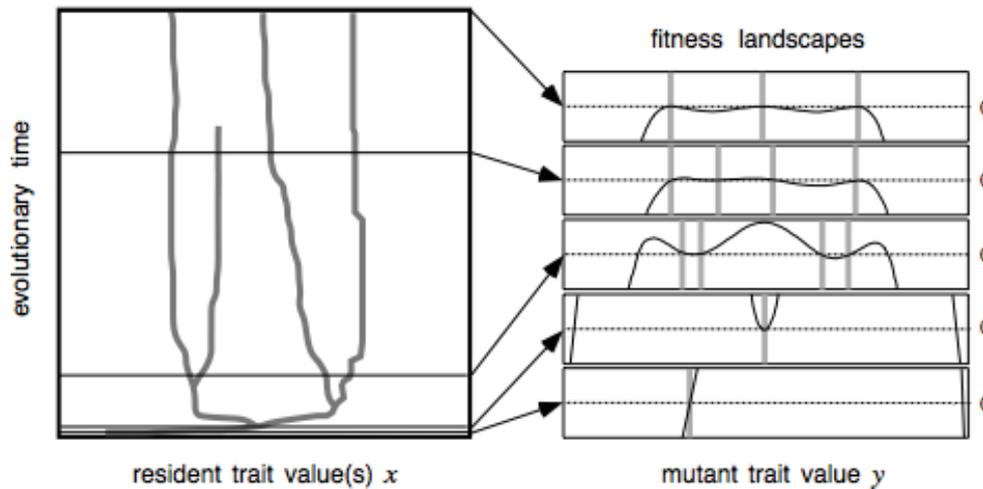


Figure 1. Left: Evolutionary path simulated on the basis of a population dynamical model, assuming clonal reproduction. Only the traits that are dominantly present in the population are shown. The second ascending branch finishes since its subpopulation went extinct. Right: The fitness landscapes for five population compositions as these occurred at the indicated times. The vertical bars indicate the types that at that moment were present in the population. At the second selected time the population resided at a branching point. At the final time the remaining three subpopulations reside at an evolutionarily steady coalition.

2. Preliminaries

2.1. Fitness

2.1.1. History

The concept of fitness as a quantitative measure of competitive prowess is a recent invention. Darwin never used the term in this meaning, and neither did the founding fathers of population genetics. (Except in one of Fisher's early papers; elsewhere Fisher, Haldane and Wright use terms like selective advantage.) In modern population genetics, fitness is generally seen as the probability to survive to reproduction. However, this only works for relatively simple ecological scenarios with the different life phases neatly separated and synchronised. In ecology one has to account for a less simple world where populations have age, size, spatial or other structures, and where demographic properties vary with the weather and local conditions.

2.1.2 Population structure and evolutionary environments

In AD anything outside an individual that influences its population dynamical behavior, which by definition consists of impinging on the environment, giving birth, and dying, is called *environment*. It is then always possible in principle to find a representation of that behavior in terms of a state space, transition probabilities that depend on the course of the environment and outputs that depend on the state of the individual and the condition of the environment at the time. Given the course of the environment, individuals independently move through their state spaces, the population state is a frequency distribution over this space, and the mathematical expectation of this frequency distribution, which is again a frequency distribution, moves according to what mathematicians call a positive linear dynamics, linear due to the independence, positive

since we cannot have negative numbers of individuals. Mathematics then tells that generally the expected size of a population in an ergodic environment will in the long run on average grow or decline exponentially. (The technical term ‘ergodic’ means roughly that the environment may fluctuate but that these fluctuations have no persistent trend.) This growth rate ρ is what ecologists call *fitness*. It necessarily is a function of two variables, the type of the individuals, parametrized by their traits Y , and the environment E , to be written as $\rho(Y | E)$ (the vertical bar as separator of the arguments is a notation borrowed from probability theory; pronounce: the fitness of Y when the environment is E). The mathematical theory of branching processes (a mathematical discipline that deals with independently reproducing objects) moreover tells that a population starting from a single individual will, barring some technical conditions, either eventually go extinct or grow exponentially, with the probability of the latter being positive if and only if its fitness is so.

For constant environments ρ is usually written as r and referred to as intrinsic rate of natural increase or Malthusian parameter.

In the theory of longer-term adaptive evolution one is only interested in populations in which the number of individuals exposed to similar environments are sufficiently large that the internal workings of these populations can be modeled in a deterministic manner. In nature populations are necessarily bounded. A mathematical consequence is that the community dynamics will converge to an attractor, be it an equilibrium, a limit cycle, or something more complicated. The corresponding environments are not necessarily ergodic, but the exceptions tend to be contrived. So from now on it will be assumed that community attractors generate ergodic environments.

Let the environment generated by a coalition of clones $C = \{X_1, \dots, X_n\}$ be written as $E_{\text{attr}}(C)$. Then $\rho(Y | E_{\text{attr}}(C))$ is the *invasion fitness* of a new type Y in a C -community.

It follows immediately that all residents, i.e., types that are present in a community dynamical attractor, have zero fitness, since resident populations by definition do not in the long run grow or decline (see Figure 1). This fact is basic to the following considerations.

For ease of exposition, it is assumed throughout that $E_{\text{attr}}(C)$ is unique; the main results extend to the general case with small modifications.

2.1.3. Mendelian diploids

The extension of the previous framework to Mendelian populations is easier than perhaps expected (although implementing it can be horrible). For the community dynamics one has to distinguish individuals according to their genotypes, and incorporate their mating opportunities with different genotypes into the description of the environment (in the case of casual matings, with pair formation it becomes necessary to extend the state space of individuals to keep track of their marriage status). Alleles reproduce clonally and as such have invasion fitnesses. It is also possible to define a (mock) fitness of phenotypes by introducing a parallel clonal model with individuals passing through their lives like their Mendelian counterparts and having a reproduction equal to the average of the contributions through the micro- and macro-gametic routes (for humans semen and ova). With such a definition some essential, but not all, fitness-based deductions for the clonal

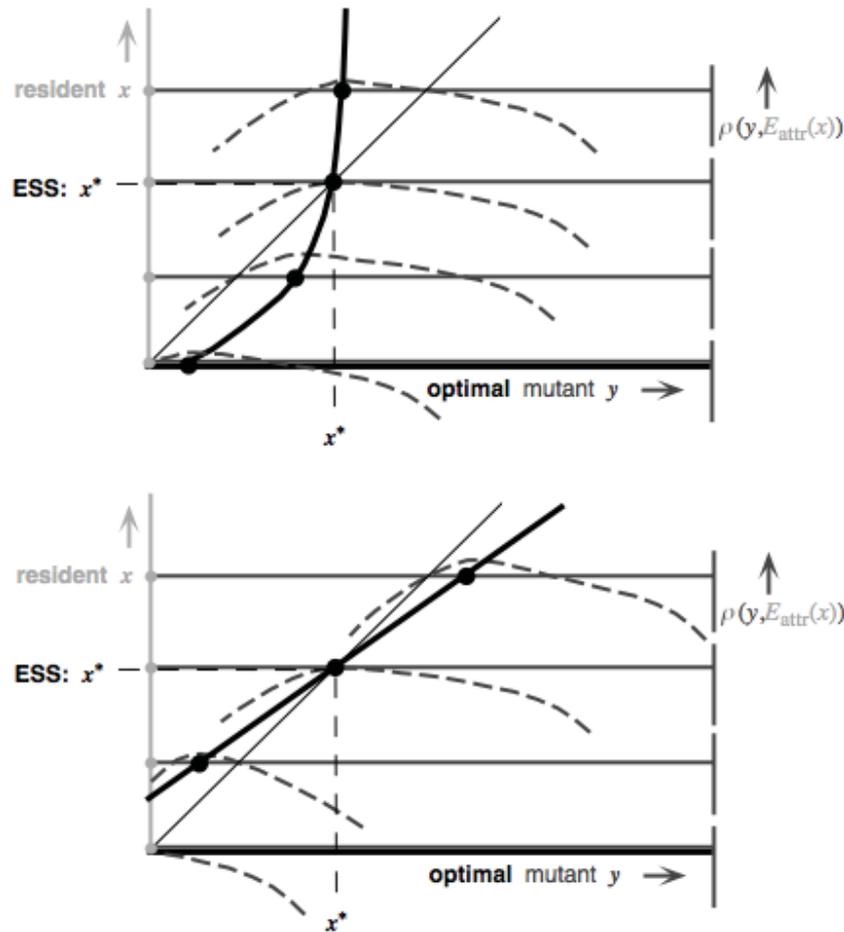


Figure 2. Scheme for calculating ESSes: For each of the possible resident populations, here characterised by a scalar trait, the invasion fitness of all potential mutants is calculated (interrupted curves). The mutant axis is drawn on the same scale as the resident axis. From these fitness curves the optimal strategy for the corresponding resident environment is calculated (fat curve). The ESS is the optimal reply to itself, to be calculated by intersecting the fat curve with the 45° line. In the upper panel the ESS attracts evolutionarily as can be seen from the fact that for other trait values the fitness landscape increases in the direction of the ESS. In the lower panel the fitness landscapes decrease in the direction of the ESS. Hence it repels.

case go through for Mendelian inheritance. In particular, for genetically homogeneous populations a resident also has fitness zero (since genetically homogeneous populations breed true). Moreover, the invasion of a new mutant into a homogeneous population is correctly predicted, as that mutant initially only occurs in heterozygotes that breed true by backcrossing with the homogeneous resident.

2.2. Meso-evolutionarily statics: ESSes

Evolution stays put whenever the community produces an environment such that all mutants that differ from any of the residents have negative fitness. In the special case where we have a single resident type, we speak of an *evolutionarily steady strategy* (ESS). (The old name evolutionarily stable strategy is a bit of a misnomer since ESSes

need not be evolutionarily attractive; see Figure 2 lower panel and Figure 7 lower left panel) The general case when there may be more than one resident type is called an *evolutionarily steady coalition* (ESC). ESCs are the equilibria of evolution.

One way of calculating ESSes is depicted in Figure 2. For each environment as generated by a possible resident the maximum of the invasion fitness landscape $\rho(Y | E_{\text{attr}}(X))$ is calculated. Next one intersects the resulting curve (or surface) $Y = Y_{\text{opt}}(X)$ with the line (or surface) $Y = X$ to get the ESS $X^* = Y^*$. As monomorphic residents have fitness zero, all potential mutants $Y \neq Y^*$ have negative fitness.

The situation for ESCs is more complicated, as there may be so-called genetic constraints, with heterozygote superiority as generic example. However, in the so-called *ideal free* (IF) case all phenotypes comprising an ESC have fitness zero (at least when there is a single birth state and the ESC engenders a community dynamical equilibrium). This IF case is defined by the requirement that there are no genetic constraints whatsoever, that is, mutants can occur that produce any feasible type as heterozygotes in the genetic background of the resident population.

2.3. Fitness proxies

The existence of a well-defined fitness forms the basis for the calculation of ESSes and AD. However, given its existence it is often possible to replace ρ by some more easily determined quantity that leads to the same outcome for the calculations of interest. For example, in optimization calculations ρ can be replaced with any quantity monotonely related to it, and for the graphical methods of AD one may replace ρ with any sign-equivalent quantity. Such quantities are referred to as *fitness proxies*. An example of a fitness proxy of the first type is the average rate of energy intake. Being a fitness proxy is always predicated on additional assumptions. For instance, it may help a forager little to increase its energy intake in an environment where this drastically increases its exposure to predation. A fitness proxy of the second type, restricted to non-fluctuating environments, is the logarithm of the average lifetime offspring number $\ln(R_0)$. If individuals may be born in different birth states, as is the case for spatial models, where birth position is a component of the birth state, R_0 is defined as the dominant eigenvalue of the next generation matrix (or operator in the case of a continuum of birth states). This matrix is constructed by calculating from a model for the behavior of individuals how many offspring are born on average in different birth states dependent on the birth state of the parent.

2.4. Optimization principles

Often biologists try to predict evolutionary outcomes by pure optimization. Of course, if one just measures the environment one may predict the evolutionarily steady trait values that go with that environment by maximizing fitness in that environment. This is why predictions from optimization work. In general, optimization procedures do not predict the outcome of evolution, for that entails also predicting the environment that goes with the ESS, but they often satisfactorily predict the strategies that go with that environment. However, such limited predictions are of little use when considering the consequences of environmental change like increasing fishing mortality or global warming.

Many papers on ESS theory also derive optimization principles by which ESSes may

be calculated. Hence, it is of interest to know when there exist properties of phenotypes that are maximized at an ESS. It turns out that this is the case if and only if one of two equivalent conditions holds good: the effect of the trait (environment) can be summarized in a single variable such that for each environment (trait) there exist a single threshold above which fitness is positive. These statements can be paraphrased as: the trait (environment) should act in an *effectively one-dimensional monotone* manner. (Think e.g. of the efficiency of exploiting a sole limiting resource.) Given such a one-dimensional summary $\psi(E)$ of the environment (which is often more easily found) it is possible to construct a matching summary of the traits $\phi(X)$, and vice versa, through

$$\phi(X) = -\psi(E_{\text{attr}}(X)) \quad (1)$$

Of course, the previous statements hinge on the interpretation of the term *optimization principle*. The latter should be interpreted as a function that attaches to each trait value a real number such that for any constraint on the traits the ESSes can be calculated by maximizing this function. The proviso mirrors the practice of combining an optimization principle derived from the ecology with a discussion of the dependence of the evolutionary outcome on the possible constraints.

If an optimization principle ϕ exists, each successful mutant increases $\phi(X)$ and hence any ESS attracts. Moreover, $\psi(E_{\text{attr}}(X))$ decreases with each increase in $\phi(X)$. Since fitness increases with ψ where it counts, i.e., around zero, ψ may be dubbed *pessimization principle*. When a pessimization principle exists, in the end the worst attainable world remains, together with the type(s) that can just cope with it.

Optimization principles come closest to the textbook intuition for the meaning of fitness, which generally fails to account for the fact that the fitnesses of all possible types are bound to change with any change in the character of the residents. However, optimization principles, although frequently encountered in the literature, are exceptions rather than the rule.

In evolutionary ecology textbooks, the maximization of r or R_0 takes pride of place without any mention of the reference environment in which these quantities should be determined. Hence, all one can hope for is that the same outcome results for a sufficiently large collection of reference environments. It turns out that this is the case if and only if r can be written as $r(X|E) = f(r(X|E_0), E)$, respectively $\ln(R_0)$ can be written as $\ln(R_0(X|E)) = f(\ln(R_0(X|E_0)), E)$ for some function f that increases in its first argument, and E_0 some fixed, but otherwise arbitrary, environment. These two criteria are relatively easy to check in specific situations. A fair fraction of textbook statements, if taken literally, applies only when these special conditions are fulfilled. The latter happens for instance when the only influence of the environment is through an additional state- and type-independent death probability or rate, respectively when the life history can be decomposed into a number of stages that can only be entered through single states (so that no information about the past carries over), either X or E affects all stages after the onset of reproduction, and no earlier stage is affected by both X and E together.

3. Adaptive Dynamics (AD)

3.1. Traits, PIPs, MIPs and TEPs

Paleontologists and taxonomists are interested in the change of traits on an evolutionary time scale. What are traits to taxonomists are parameters to ecologists. So in AD one is after a dynamics in the parameter space of a community dynamics. The trick for arriving at such a simple picture is to assume that favorable mutants come along singly after a community has relaxed to an attractor. Another trick is to assume clonal reproduction, on the assumption that this way one can find out where the ecology would drive evolution if the latter were not hampered by the constraints of Mendelian genetics.

To get at a purely trait oriented picture, any reference to the environment should be removed from the expression for invasion fitness:

$$s(Y | C) = \rho(Y | E_{\text{attr}}(C)) \quad (2)$$

(often this is written as $s_C(Y)$ to emphasize the interpretation as a family of fitness landscapes).

This Subsection focuses on scalar traits, starting with the case where there is only a single clonally reproducing resident, $C = x$. The first step in the analysis is plotting a contour plot of $s(y | x)$. Usually this is simplified to plotting only the zero contours, as those matter by far the most. The result is called *pairwise invasibility plot* (PIP). See Figure 3.

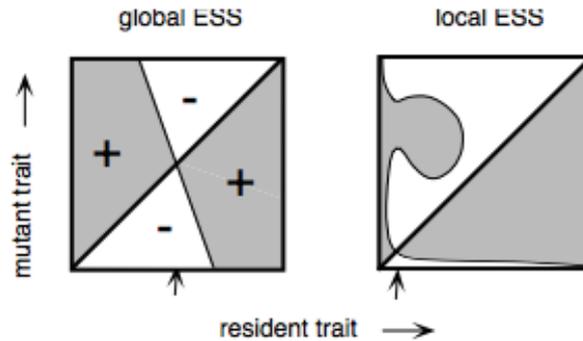


Figure 3. Pairwise invasibility plots: sign of the fitness of potential mutants as a function of the mutant and the resident traits.

Note that the diagonal is always a zero contour as residents have fitness zero. A point where some other contour crosses the diagonal is referred to as *evolutionarily singular strategy* (ess). The ESSes are a subset of the ess-es.

Now assume that mutational steps are small and that in the beginning there is only one resident trait value $x(0)$. Plot this value on the abscissa of the PIP, say in the left panel of Figure 3. After some random waiting time, mutation creates a new trait value y . This trait value can invade only when it has positive fitness, i.e., is in a plus area of the PIP. It can be proved that an invading type replaces its progenitor if the latter is not too close to an ess or a bifurcation point of the community dynamics, and the mutational step is not too large. If such a replacement has occurred we call the new trait value $x(1)$. In the considered PIP, if $x(0)$ lies to the left of the ESS then $x(1)$ lies to the right of $x(0)$, and

vice versa. Hence, repeating the process pushes the evolutionary path to a neighbourhood of the ESS. Upon reaching that neighbourhood it may become possible that ecologically the mutant and its progenitor persist together.

To see how such coexisting pairs of strategies fare evolutionarily it is necessary to consider the set of *protected dimorphisms*, i.e., pairs of strategies that can mutually invade, to be denoted as (x_1, x_2) . Its construction is shown in the Figure 4.

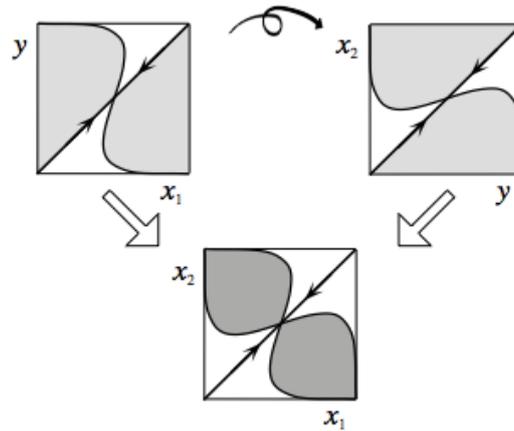


Figure 4. The construction of a *mutual invasibility plot* (MIP), depicting the set in (trait space)² harbouring protected dimorphisms. Not all polymorphisms occurring in AD are protected, but unprotected polymorphisms have the habit of never lying close to a diagonal.

The evolutionary movement of the pair (x_1, x_2) is governed by $s(y | x_1, x_2)$. Under the assumption of small mutational steps a good deal of information can be extracted from the adaptive isoclines, calculated by setting the *selection gradient*

$$g_i(x_1, x_2) = ds(y | x_1, x_2)/dy|_{y=x_i} \quad (3)$$

equal to zero. As depicted in Figure 5, x_1 will move to the right when g_1 is positive and to the left when it is negative, and x_2 will move up when g_2 is positive and down when it is negative.

From the classification of the possible dynamics near an ess in Figure 7 it can be seen that the ess in the PIPs from Figure 3 also attract in the dimorphic regime.

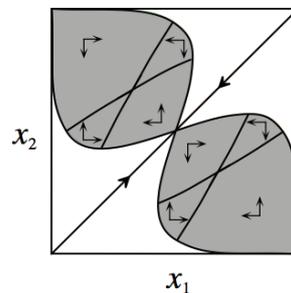


Figure 5. *Trait evolution plot* (TEP), i.e., MIP together with arrows that indicate the direction of the small evolutionary steps that result from the invasion by mutants that differ but little from their progenitor, and adaptive isoclines.

3.2. The canonical equation

3.2.1. Basics

The neat dependence of the dynamical outcomes on no more than the sign of the invasion fitness hinges on the ordering properties of the real line. For vectorial traits we have to proceed differently. The workhorse is the so-called *canonical equation* (CE) of AD, a differential equation that captures how the trait vector changes over evolutionary time, on the assumption of small mutational steps, accounting for the fact that favorable mutants do not always invade due to demographic fluctuations. For unbiased mutational steps the evolutionary speed and direction are given as the product of three terms: from left to right (1) the effective population size N_e (as in the diffusion equations of population genetics), (2) the probability of a mutation per birth event ε times a matrix C consisting of the variances and covariances of the resulting mutational step, (3) the selection gradient (the “curly d” notation stands for differentiating for that variable while treating the other variables as parameters)

$$G(X) = \left(\begin{array}{c} \partial s(Y | X) / \partial y_1 \\ \vdots \\ \partial s(Y | X) / \partial y_n \end{array} \right) \Big|_{Y=X} \quad (4)$$

which is a vector pointing from the position of the resident in the steepest uphill direction. The uphill pull of selection is thus modified by the differential directional availability of mutants expressed in the *mutational covariance matrix*. For diploid Mendelian populations an additional factor 2 appears, since the substitution of a mutant allele leads to a twice as far removed mutant homozygote (at the considered order of approximation).

The equilibrium points of the CE are the ess-es mentioned previously.

In reality many mutants attempt to substitute simultaneously. Luckily, for small mutational steps this appears to affect the environment only in the higher order terms that in the derivation of the CE disappear from sight. However, in the clonal case the effects of the mutants do not add up since a mutant may be supplanted while invading by a better mutant from the same parent type. In the Mendelian case the CE will do a better job as there substitutions occur in parallel on different loci, which to the required order of approximation act additively.

3.2.2. The link with Evo-Devo

From an AD perspective the link with Evo-Devo is through the mutational covariance matrices. Unfortunately, Evo-Devo has yet little to offer in this area. Therefore, at present the most AD researchers can do is work out how the outcomes of a specific eco-evolutionary model depend on the possible forms of the mutational covariance matrix. The answers from AD thus become Evo-Devo questions: is the mutational covariance matrix for these traits expected to fall within this or that class?

To show the importance of the missing Evo-Devo input in AD: mutational covariance matrices have an, often dominating, influence on the time scales of evolution (Figure 6, left), the basins of attraction of ess-es (Figure 6, right), even to the extent that they often determine whether an ess attracts or not.

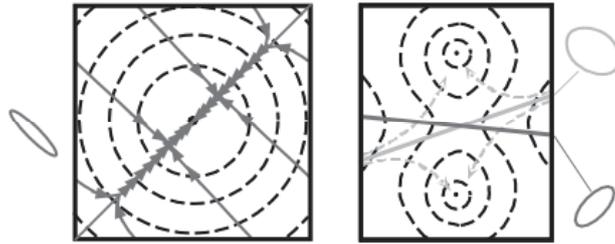


Figure 6. Two fitness landscapes that are supposed to keep their shape and only to sink when the adaptive trajectory moves uphill (as is the case if and only if the population regulation is through an additional state-independent death rate). Distributions of mutational steps are symbolised by ovals. left: The shape of the mutation distribution induces a time scale separation between the movement along the diagonal and anti-diagonal direction. Right: The difference in mutation distributions causes a difference in the domains of attraction of the two ESS-es.

On a more philosophical level it bears noting that the selection gradient points only in a single direction, while the components of the trait vector orthogonal to that gradient hitchhike with the selectively determined motion thanks to a developmental coupling as expressed in the mutational covariance matrix. The higher the dimension of the trait space the larger the contribution of development as a determinant of evolutionary motion. The dimensions of the trait spaces that are routinely considered thus make for the contrast in attitudes of behavioral ecologists and morphologists, with the former stressing selection and the latter developmental options for change.

3.3. Evolutionarily singular strategies

Evolutionarily singular strategies X^* can be calculated by setting the fitness gradient equal to zero. Figure 7 shows their classification according to dynamical type for scalar

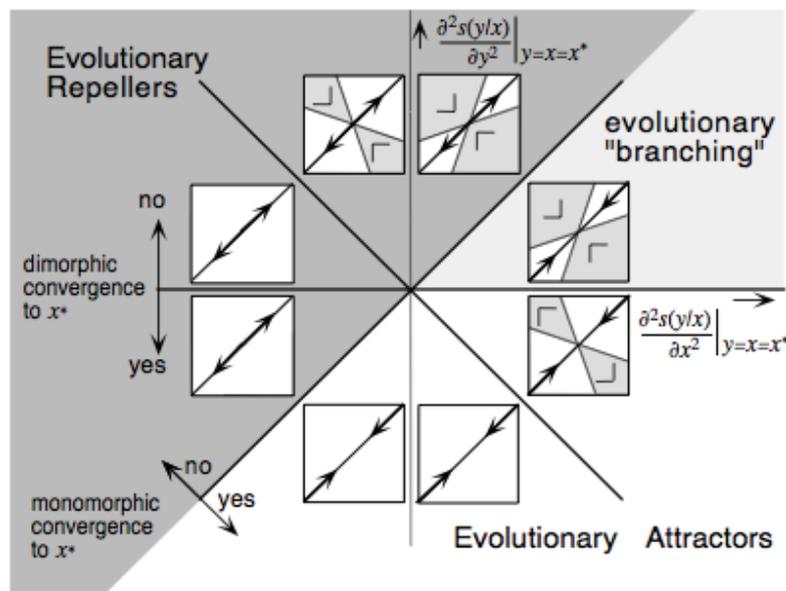


Figure 7. A classification of the ess-es for scalar traits. The cases in the lower half are all ESSes. The leftmost of these repels, the others attract. The latter ESSes are thus genuine evolutionary attractors. The branching points in the rightmost upper sector attract monomorphically but repel dimorphically.

traits. (Note that this classification was derived deductively from no more than some mathematical consistency properties shared by well-posed eco-evolutionary models.)

Devising a good classification for higher dimensional ess-es is an open problem. One reason is that in higher dimensions the attractivity of a singular point crucially depends on the mutational covariance matrix, except in very special cases.

3.4. Adaptive speciation

Perhaps the most interesting ess-es are *branching points*, where the eco-evolutionary process starts generating diversity. When approaching such points the evolutionary trajectory, although continually moving uphill, gets itself into a fitness minimum. More precisely, it is overtaken by a fitness minimum. An analogy may help intuiting this phenomenon. Somewhere gold has been found, attracting people to that spot. However, after the arrival of too many diggers it becomes more profitable to try one's luck at some distance.

The build up of diversity can take very different forms. In the clonal case the population just splits into two as depicted in Figure 1. In the Mendelian case the diversification starts with a broadening of the variation in the population. The fitness landscape locally has the shape of a parabola that increases away from x^* . This means that types more on the side have a higher fitness than those in the centre. It therefore pays not to get kids near the centre. The Mendelian mixer has the contrary tendency to produce intermediate kids out of dissimilar parents. Luckily, there are all sorts of mechanisms that may thwart this counterproductive mixing. The most interesting of these is the build up of some mechanism that lets like extremes mate only among themselves, thus ensuring that the branches become separate genetic units. A simple mechanism occurs in insects that diversify in their choice of host plants, with mating taking place on those hosts.

The author's conviction is that in cases where no automatic mating barrier puts itself in place a build up of other mechanisms engendering assortative mating is not unexpected. Present day organisms are the product of 3.5 billion years of evolution. During that time their sensory and signalling apparatus has been evolutionarily honed for finding the most advantageous mates. Hence, there will be an abundance of template mechanisms. These mechanisms, once recruited to the task, will probably have a tendency to enhance each other in their effect. One may thus expect that the available generalized machinery can often easily be adapted so as to genetically separate the branches whenever evolution brings the population to a branching point. However, most scientists working on the genetics of speciation do not appear to agree with this view.

3.5. Bifurcations

The bifurcations of AD encompass all the classical bifurcations found in ecological models. In additions there is a plethora of additional bifurcations. An example is the transition from an ESS to a branching point depicted in Figure 8.

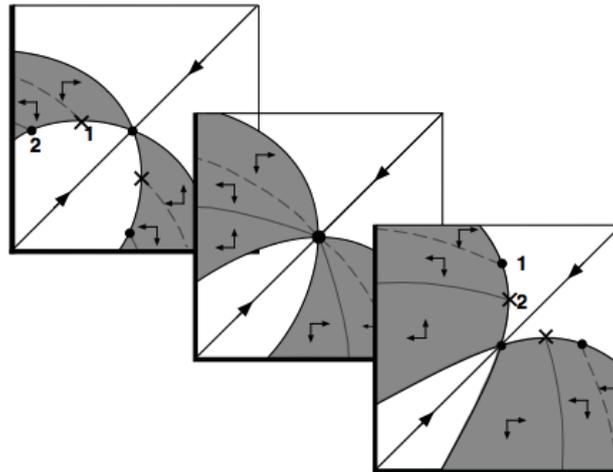


Figure 8. Three TEPs corresponding to a bifurcation of an ESS to a branching point. In this case the adaptive tracking of the ESS in the wake of slow environmental change stops due to a change in character of the ess. In the fossil record this scenario would correspond to a punctuation event that starts with speciation.

4. Justifying the AD approximation

4.1. Traits and genotype to phenotype maps

The real state space of the meso-evolutionary process is genotype space while the phenotypic trait spaces of AD are but convenient abstractions. Phenotypic mutational covariances reflect both the topology of genotype space, as generated by mutational distances, and the genotype to phenotype map generated by the developmental mechanics. This reflection can only be expected to be adequate locally in genotype space, and therefore locally in evolutionary time.

AD focuses on small mutational steps. A mechanistic partial justification comes from the expectation that the evolutionary changes under consideration are mostly regulatory. Coding regions of genes are in general preceded by a large number of short regions where regulatory material can dock. Changes in these regulatory regions lead to changes in the production rate of the gene product. The influence of a single regulatory region among many tends to be rather minor.

Note moreover that phenotypes should in principle be seen as reaction norms, i.e., maps from micro-environmental conditions to the characteristics of individuals (another term is conditional strategies). The phenotypes of AD are inherited parameters of these reaction norms. Only in the simplest cases a reaction norm is degenerate, taking only a single value.

4.2. Internal selection

Functional morphologists usually talk in terms of whether certain mechanisms work properly or not, and discuss evolution as a sequence of mechanisms all of which should work properly, with only slight changes at every transformational step. Translated into the language of fitness landscapes, this means that only properly working mechanisms give fitnesses in the ecologically relevant range, while the improperly working ones always give very low fitnesses. This leads to a picture of narrow, slightly sloping, ridges

in a very high dimensional fitness landscape. The slopes on top of the ridges are the domain of ecology, their overall location is largely ecology independent.

As a simple example you may think of leg length. The left and right leg are kept equal by a strong selection pressure, which keeps in place a developmental system that produces legs of equal length, notwithstanding the fact that during development there is no direct coupling between the two leg primordia. Hence in a trait space spanned by the two leg lengths ecologists concentrate on just the diagonal.

The trait spaces dealt with in morphology are very high dimensional so that the top of a ridge may be higher dimensional, while away from the ridge the fitness decreases steeply in a far larger number of orthogonal directions. A picture similar to that of functional morphologists emerges from Evo-Devo. The long term conservation of developmental units, think of the phylotypic stage or homology, can only be due to strong stabilising selection, since mutations causing large pattern changes have many side effects with dire consequences. As a result, ecological selection generally acts only on quantitative changes in the shapes and sizes of homologous body parts.

Since the fitness differences considered in functional morphology and Evo-Devo largely result from the requirement of an internal coherence of the body and of the developmental process, people speak of *internal selection*. Here the term is used to label features of the fitness landscape that are roughly the same for all the environments that figure in an argument.

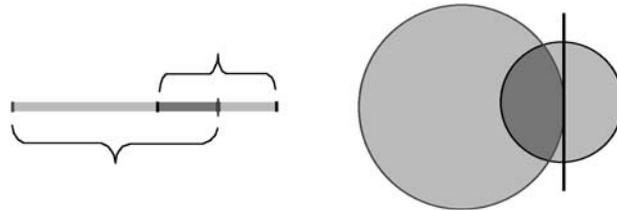


Figure 9. Left: Two balls in \mathbb{R}^1 , with the centre of the smaller ball on the boundary of the larger ball. The ratio of the volume of their intersection to the volume of the smaller ball is $1/2$. Right: A similar configuration in \mathbb{R}^2 . The volume of the intersection is now a smaller fraction of the volume of the smaller ball. For similar configurations in \mathbb{R}^n this fraction quickly decreases to zero for larger n . Now think of the larger ball as the part above sea level of a fitness hill and of the smaller ball as a mutation distribution. Clearly the fraction of favourable mutants will go to zero with n .

4.3. The assumptions of AD

4.3.1. Theory

Figure 9 illustrates an argument by Fisher showing that the higher the dimensionality of the trait space the more difficult becomes the final convergence to an adaptive top. This argument extends to the movement in a ridgy fitness landscape: the higher the number of orthogonal off-ridge directions, the more rare it is for a mutational step to end up above sea level, and small mutational steps have a far higher propensity to do so than large ones. Together these two arguments seem to underpin the requirements of AD that mutations in ecologically relevant directions are scarce and the corresponding mutational steps small.

Unfortunately the above arguments contain a biological flaw: the assumed rotational symmetry of the distribution of mutational steps. Real mutation distributions may be expected to show strong correlations between traits. Correlation structures can be represented in terms of principal components. The general experience with biological data is that almost always patterns are found like the ones shown in the lower left panel of Figure 10. The right panel shows that the existence of mutational correlations will in general enhance the rareness and smallness of the mutational steps that end up above sea level.

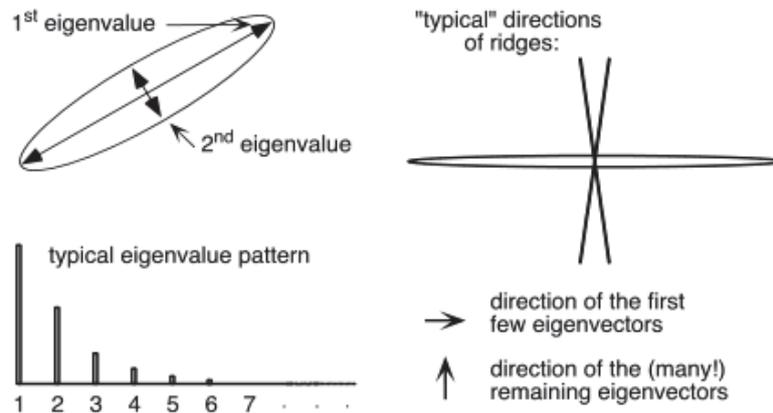


Figure 10. Upper left: Contour line of a bivariate distribution, supposedly of mutational steps. The lengths of the two axes of the ellipse, called principal components, are proportional to the square root of the eigenvalues of the mutational covariance matrix. Lower left: Typical eigenvalue pattern found for large empirical covariance matrices. Right: The mutation distribution will rarely be fully aligned with the fitness ridges. If in a high dimensional fitness landscape one takes one's perspective from the mutation distribution and looks at the orientation of the ridges relative to the first few principal axes of this the mutation distribution, then, when the number of the dimensions of the trait space is very large and the ridge has a relatively low dimensional top, the ridge will typically extend in a direction of relatively small mutational variation.

4.3.2. Data?

The conclusions from the previous Subsection seem to underpin nicely the assumptions of AD. Unfortunately, various empirical observations appear to contradict these conclusions. Populations brought into the lab always seem to harbor sufficient standing genetic variation to allow quick responses to selection, and a few loci with larger effect (quantitative trait loci or QTL) are often found to underlie the variation in a trait. However, these empirical observations may have less bearing on the issue than one might think.

First, given the speed of evolution relative to the changes in the overall conditions of life, populations in the wild are probably most often close to some ESS. Moreover, in noisy environments fitness maxima tend to be flat. This means that near neutral genetic variation will accumulate, which is exploited first when a population gets artificially selected on. At ESSes the mutation limitation question is largely moot. Beyond its statics, AD's main interest is in the larger scale features of evolutionary trajectories after the colonization of new territory or, even grander, a mass extinction. The scale of these

features may be expected to require a further, mutational, supply of variation. Second, directional selection on an ecological trait may be hampered by stabilizing internal selection on pleiotropically coupled traits. In the lab this stabilizing selection is relaxed. This means that far more variation comes available for directional selection than is available in the wild.

Lastly, AD style theory has shown that in the absence of assortative mating the initial increase of variability after the reaching of a branching point tends to get redistributed over a smaller number of loci with increasing relative effect. The end effect will be QTL, but produced through the cumulative effect of small genetic modifications.

Glossary

adaptive dynamics (AD): A mathematical framework for dealing with eco-evolutionary problems, primarily based on the following simplifying assumptions: clonal reproduction, rare mutations, small mutational effects, smoothness of the demographic parameters in the traits, and well-behaved community attractors.

adaptive (Darwinian) speciation: Speciation that results from a population evolutionarily getting to a branching point.

branching point: Phenotype that is approached from all directions by gradualistic evolution resulting from ecological interactions, at which the population then finds itself under disruptive selection.

canonical equation (CE): Differential equation that captures how the trait vector changes over evolutionary time on the assumption that mutations are sufficiently rare and mutational steps sufficiently small.

evolutionarily steady coalition (ESC): A combination of phenotypes (strategies) the bearers of which can live together on an ecological time scale and such that no mutant has positive fitness in the environment generated by such a community.

evolutionarily steady strategy (ESS): A phenotype such that no mutant has positive fitness in the environment generated by that strategy as resident. The acronym is often interpreted as evolutionarily stable strategy. However, as defined there is no need that an ESS is evolutionarily stable (i.e., attractive).

evolutionarily singular strategy (ess): A phenotype for which the fitness gradient is zero. ess-es comprise ESSes, branching points and various sorts of evolutionary repellors.

fitness: In population genetics: the probability to survive till reproduction. In evolutionary ecology: the asymptotic average rate of exponential growth of a (often hypothetical) clone of individuals in an a priori given (hypothetical or experimentally engineered) ergodic environment (i.e., the environment should at least be immune to the effects of this population growth), written as $\rho(Y|E)$ where Y typifies the sort of individuals and E stands for the environment. For constant environments ρ equals the intrinsic rate of natural increase, or Malthusian parameter r .

fitness landscape: Graph of the fitnesses of potential mutants as a function of their trait vector (for a given environment).

fitness proxy: Quantity that is not equal to fitness but can be substituted for it in specific calculations.

invasion fitness: Fitness of a new type in the environment generated by an existing community.

micro-, meso- and macro-evolution: Respectively, changes in gene frequencies on a population dynamical time scale, evolutionary changes in the values of traits of representative individuals and concomitant patterns of taxonomic diversification, and large scale evolutionary changes like anatomical innovations where one cannot even speak in terms of a fixed set of traits.

mutational covariance matrix: Matrix composed of the variances and covariances of the distribution of mutational steps.

mutual invasibility plot (MIP): Plot indicating for a scalar trait all combinations of two types such that each type can invade in a population consisting wholly of the other type.

optimization principle: Function defined on the trait space of an eco-evolutionary model such that under any constraint the ESSes for that model can be determined by maximizing that function.

pairwise invasibility plot (PIP): Plot indicating for a scalar trait the combinations of mutants and residents for which the mutant can invade.

pessimization principle: Function defined on the set of possible environments of an eco-evolutionary model such that under any constraint the ESSes for that model can be determined by minimizing that function.

selection gradient: Gradient of the fitness landscape at the position of a resident.

trait evolution plot (TEP): MIP together with arrows showing the direction of adaptive movement of each of the two types.

Suggestions for further reading

Dercole, F. & S. Rinaldi (2008) *Analysis of Evolutionary Processes: the Adaptive Dynamics Approach and its Applications*. Princeton University Press

Dieckmann, U., M. Doebeli, J.A.J. Metz, & D. Tautz eds. (2004) *Adaptive Speciation. Cambridge Studies in Adaptive Dynamics* 3. Cambridge University Press.

Dieckmann, U. & R. Law (1996) The dynamical theory of coevolution: a derivation from stochastic ecological processes. *J. Math. Biol.* **34**: 579–612.

Diekmann, O. (2004) A beginner's guide to adaptive dynamics p 47-86 in *Mathematical modelling of population dynamics*. Banach Center Publications, Volume 63. Institute of Mathematics. Polish Academy of Sciences, Warszawa

Durinx, M., J.A.J. Metz & G. Meszéna (2008) Adaptive dynamics for physiologically structured models. *J. Math. Biol.* **56**: 673-742

Geritz, S.A.H., É. Kisdi, G. Meszéna & J.A.J. Metz (1998) Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evol. Ecol.* **12**, 35-57.

Gyllenberg, M., J.A.J. Metz & R. Service (2011) When do optimisation arguments make evolutionary sense? In F. A. C. C. Chalub and J. F. Rodrigues (eds.) *The Mathematics*

- of Darwin's Legacy*. Birkhauser, Basel.
- Leimar, O. (2009) Multidimensional convergence stability. *Evol. Ecol. Res.* **11**: 191-208.
- Metz, J.A.J. (2008) Fitness. Pp. 1599-1612 in S.E. Jørgensen & B.D. Fath (eds.) *Evolutionary Ecology. Vol. 2 of Encyclopedia of Ecology*. Elsevier, Oxford.
- Metz, JAJ (2011) Thoughts on the geometry of meso-evolution: collecting mathematical elements for a post-modern synthesis. In F. A. C. C. Chalub and J. F. Rodrigues (eds.) *The Mathematics of Darwin's Legacy*. Birkhauser, Basel.
- Metz, J.A.J., S.A.H. Geritz, G. Meszéna, F.J.A. Jacobs & J.S. van Heerwaarden (1996) Adaptive dynamics, a geometrical study of the consequences of nearly faithful reproduction. Pp 183-231 in: S.J. van Strien & S.M. Verduyn Lunel (eds.) *Stochastic and spatial structures of dynamical systems*. North-Holland, Amsterdam.

See also the following articles: Behavioral Ecology, Apparent Competition, Invasion Biology, Stage Structure, Two-Species Competition, Adaptive Behavior, Adaptive Landscapes, Coevolution, Evolution of Dispersal, Evolutionarily Stable Strategies, Niche Construction, Bifurcations, Branching Processes, Markov Chains, Matrix Models, Pair Approximations, Stability Analysis.