Evolvability: a Formal Approach

Alexis Gallagher
St Peter's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2009

ABSTRACT

Evolvability: a Formal Approach
D.Phil, Trinity 2009
St Peters College
Alexis Gallagher

This dissertation clarifies the concept of evolvability, the increased capacity of some organisms or systems to support evolution, especially the evolution of life-like complexity. I survey the literature, which is spread over the fields of population genetics, developmental biology, artificial life, and microbial and molecular evolution. Finding that researchers have often used the term vaguely and incompatibly I identify five distinct kinds or senses of evolvability. I also identify five key constituent ideas, which I discuss in the context of *organismic evolvability*, a sense of evolvability with deep roots in the traditional fields of animal development and macroevolution. In these fields research into evolvability has historically been hampered by an insufficiently detailed knowledge of development. Research in molecular evolution has produced a thorough knowledge of the folding of RNA into secondary structure, which can be regarded as a model of development. This has motivated new approaches to evolvability based on representing development via a single genotype-phenotype mapping function. I build on these approaches to invent new mathematical methods to formalise the traditional ideas. I create an exact model illustrating a classic example of evolvability, the capacity for repeated segmentation and simple modularity. I analyse this with two new formal approaches. First is the *genospace algebra*, a propositional calculus based on graph theory. It is a formal language for describing genotype-phenotype maps. It provides a system for making calculations, proofs, and diagrams about mutational structures in genotype space, and it is flexible enough to allow description at arbitrary degrees of resolution. Second is a pair of concepts, the *genetic leverage* and the *genetic fulcrum*. The leverage provides a crude numerical measure of evolvability, and the fulcrum provides a heuristic for identifying the genomic and developmental causes of evolvability. Besides its specific relevance to diversification and development, evolvability is also crucial to the fundamental question of how evolution produces ordinary biological life. Simulation systems that implement only a conventional textbook model of evolution – systems possessing only variation, inheritance, and selection – fail to evolve anything resembling the complexity of the biological world. Research into evolvability is our best bet to illuminate the "missing ingredient" for life-like evolution.

Preface

A question with half an answer

An animal – any animal – is a miraculous thing. Composed of many parts fitted together in an exquisite harmony, it eats, grows, reproduces, moves, and acts with every appearance of purpose. In all these regards an animal stands out from the non-living world of stuff. Although every organism is really a chemical machine at the microscopic level, the living world is so far beyond any man-made mechanism that it appears, still, quite different in kind.

So it is natural to wonder what created the living world, what explains it. This is the fundamental question of biology, so basic it is easily taken for granted. What explains the living world? That is, what explains the unique complexity that distinguishes the living world?

And the fundamental answer is evolution by means of natural selection. This process has transformed our planet, over billions of years, from an inanimate rock into a teeming greenhouse. Seen at the fundamental level, as a logical process, it requires only three conditions: variation, selection, and inheritance. This is our most basic explanation for the existence of the natural world.

However, a couple problems make this explanation alone unsatisfactory.

First, it omits details. This explanation describes a process of microscopic change but does not say how those changes add up to macroscopic changes. No doubt they do, but how? If the evolution of every complex organism is a unique story, then there is nothing to say beyond cataloguing these stories in all their particularity. However, it seems plausible that these stories share themes, that they fit into categories – in short, that there are principles that describe the evolution of complexity in general. What are they?

Macroevolutionary patterns suggest that other forces besides natural selection play an important role. Some organisms seem to survive longer, to diversify more widely, or to develop over the long run according to an internal logic. What are the features of an organism that facilitate its evolution? How much are these features also an explanation of complexity?

A second problem is that it is easy to create artificial systems that possess the conditions for evolution – variation, selection, and inheritance – but then this evolution does not produce anything like the complexity of the biological world. This shows that these conditions are necessary but not sufficient for life as we know it. Biological evolution depends on some extra condition that we do not understand but, presumably, that we could understand. This claim is either painfully obvious or bizarrely radical, depending on what kind of specialist you ask. In disciplines where it is not experimentally feasible to create artificial systems, the ideas those systems illustrate can come to seem unthinkable.

All these problems point to the fledgling idea of evolvability, the capacity to evolve. A theory of evolvability would speak to a variety of questions. What are the conditions that promote evolution – especially evolution of the sort of complexity taken for granted in the biological world? Is it something special about certain organisms or their developmental systems? About certain substances, like carbon or water? What can we say about the evolution of complexity, in detail and in general? If systems vary in their evolvability, then why is that? How do we measure it? How do we predict it?

These questions have been studied for decades in a variety of disciplines, but usually in an ad hoc and informal manner and without much communication between researchers. This dissertation reviews past work on evolvability, sorts out different senses of the term, and then develops a mathematical framework for precisely modelling the most promising ideas.

All of this, it is hoped, will shed light on the question of evolvability. This question is certainly ambiguous and hard to specify. It is also, however, fundamental and important and worth tackling head-on.

The history of "evolvability"

The literature on evolvability is quite fragmented so before discussing the idea it is worth outlining the varied usage of the term.

The term "evolvability" is first used by Dawkins (1989) in a short paper presenting his Biomorphs model. This model illustrates how an organism's developmental system might increase its ability to evolve, how this ability itself could evolve, and how this might affect the organism's evolutionary history.

But this term merely brought focus to a cluster of ideas with an older and more diffuse history. For instance, Maynard-Smith et al. (1985) present a lucid review of these issues under the rubric of "developmental constraints". Even earlier, Riedl (1978) presented similar ideas in the *Order of Living Organisms* (discussed by Wagner and Laubichler (2004)). Conrad (1979) discussed "amenability to evolution" more narrowly in regards to chemical evolution, and Rechenberg (1973) promptly encountered related issues in the first work in genetic algorithms.

In the 1990s the term starts appearing too frequently to be traced in detail.[1] However, one can identify a distinct school of researchers all now using the term in roughly the original sense as they study the same cluster of issues. The key work in this area is the work of Gunter Wagner, Marc Kirschner, and John Gerhart. There have also been significant contributions by Yang, von Dassow, Poole, and others. The sense of evolvability which is studied in that work I will refer to as *organismic evolvability*, because it is about kinds of organisms. Chapter 1 will explain that concept by walking through its key ideas and the main publications of these authors.

During this period other researchers start using the term in new senses which, while related, are narrower. These researchers typically stand outside of evolutionary biology and they modify the idea as they translate it to different disciplines – population genetics, microbial biology, or evolutionary computation. These new senses often emerge from efforts to clarify the original idea so it can be applied to experiments or exact models. Chapter 2 will review this work and distinguish these alternative senses of evolvability.

---

[1] As of 4 February 2008 searches on the biological and medical references services PubMed and Biological Abstracts give over 400 and 200 hits respectively. The multi-disciplinary service Google Scholar gives over 8,000 hits, with 2,200 in biology and life sciences, and over 3,600 in computer science and other engineering disciplines. (Engineering disciplines may be overrepresented because of fuller computerisation of their literature, and because of an unrelated usage of "evolvability" in systems engineering referring to tolerance for changing technology and usage scenarios (e.g., Rowe et al., 1998).)

I call them *trait evolvability* (the most prominent), *individual fitness evolvability*, *individual mutational evolvability*, and *substrate evolvability*.

The word also falls into vogue for a few years, resulting in a scattering of publications which use the term casually and without a technical meaning (Cowell et al., 1999; Matsuura et al., 2002; Suzuki et al., 2003; Woodruff, 2001; Yamauchi et al., 2002). I do not review this work.

Plan of this dissertation

Chapter 1: Organismic evolvability

Chapter 1 reviews the literature on organismic evolvability, the original sense of the term, and the one with the longest research history and most immediate explanatory promise. The earliest work began with research on the body plan, the effect of development on the body plan, and the consequent effects of developmental bias on biological diversity. This work provides the concrete sense of evolvability that underlies nearly all later intuitions. Much of this work connects with older and more mainstream discussions about the evolution of complex features (e.g., the "problem" of half a wing, etc.).

Using Dawkins's model of biomorphs as a simplified illustration, this chapter describes five key ideas that have organised discussions of evolvability: the structure of variability, the developmental architecture, the effects on macroevolutionary patterns, the incompatibility with Darwinian selection, and the effects on the evolution of complexity. A central idea in work on organismic evolvability is the idea of a *core component*, a stable part of an organism's developmental architecture which modifies the expression of less stable parts in a way that increases the organism's phenotypic variability and biases it towards complexity. This chapter also describes the ways in which these discussions fit into the traditional research literature.

(The technical parts of this thesis – chapters 4 through 6 – are dedicated to trying to clarify organismic evolvability by expressing it in simple, heuristic formal models.)

Chapter 2: Taxonomy of evolvabilities

Organismic evolvability is about organisms and is rooted in the study of how their developmental systems shape their evolution. Chapter 2 names and distinguishes four other notions of evolvability which have emerged independently or as a result of efforts to formalise the original notion of organismic evolvability.

These are the notions of trait evolvability, substrate evolvability, individual fitness evolvability, and individual mutational evolvability. These notions come from research in population genetics, molecular evolution, microbial evolution, and evolutionary computation. Many discussions of evolvability have been thought-provoking but also somewhat vague or mutually contradictory. Clarifying the distinctions and relations between these various notions of evolvability is essential to evaluating them.

In addition, considering these other notions of evolvability and the research around them provides valuable insights for how to analyse organismic evolvability. Research using the notion of individual fitness evolvability has shown clearly how the evolution of evolvability is compatible with Darwinian selection, and sheds light on the key idea of variability. The notion of substrate evolvability, which underlies much work in evolutionary computation, highlights the connection between evolvability and deeper questions about the fundamental conditions necessary for evolution.

But most of all, research on evolvability in the molecular evolution of RNA shapes has stimulated the development of a new formal methods for analysing the genotype-phenotype map. This work is the foundation for the modelling approaches that I will develop in the following chapters.

Chapter 3: The topological phenotype

Chapter 3 recapitulates past efforts to use the mathematical construct of a topological space in order to represent and analyse the genotype-phenotype mapping. These efforts rely on a chain of intermediate formal constructs – from the mapping, to the accessibility measure, to the accessibility relation, to the set of neighbourhoods, and finally to the topological space. The topological space describes the mutational relationships between phenotypes, as implied by the mapping.

However, in reviewing these efforts, it is shown that the notion of a topological space has been applied incorrectly or only as a kind of inspiration. Applied strictly, a topological space destroys so much information that it renders very different biological systems identical, making it a poor basis for models. Furthermore, the intermediate construct of the *accessibility measure* does not accurately measure the probability it is supposed to measure and it conceals structural information which might be crucial to evolvability.

However, aspects of this approach seem very worthwhile. It develops other worthwhile ideas like the *phenotypic pre-image*, and the accessibility measures and relations are themselves powerful aides in thinking clearly about mutational relationships. This chapter offers a new accessibility measure, κ', which corrects the probability calculation. And considering the structural blind spots in the measure leads to a clearly definable contra-indication of evolvability, the idea of a *dead end genotype*, which is a genotype lacking the mutational possibilities of its phenotypic peers.

Chapter 4: Evolvable segmentation and the genospace algebra

Chapter 4 presents a model and a new kind of analysis. The model illustrates one of the most frequently discussed mechanisms supporting evolvability in the biological literature focused on development: the capacity for repetition of a phenotypic feature, such as a body segment. It compares two genotype-phenotype maps of differing evolvability, a naive mapping ϕ which expresses genes directly and an evolvable mapping ϕ_E where a regulatory mechanism provides an abstracted control over the number of body segments.

Analysis shows that the capacity for segmentation in the evolvable mapping eliminates all dead end genotypes. Analysis also shows that the evolvable mapping always allows certain direct mutational transitions between phenotypes, transitions which are not available in the naive mapping. Both of these results are consistent with the basic idea of organismic evolvability – the idea that the mapping with a capacity for abstracted control is indeed more evolvable.

To perform the analysis this chapter introduces a new mathematical formalism, the *genospace algebra*, a propositional calculus based on graph theory. It provides a system for making calculations, proofs, and diagrams about mutational structures in genotype space. This formalism is flexible enough to allow description at arbitrary degrees of resolution. It is a formal language for describing genotype-phenotype maps.

Chapter 5: Genetic leverage

Genetic leverage is another formal tool for measuring the evolvability of a genotype-phenotype mapping, cruder than the genospace algebra but simpler to apply. Roughly, it measures how much certain parts of the genome enjoy an amplified influence over phenotypic change, under certain genetic configurations. This kind of amplification is indicative of a high evolvability. The idea of genetic leverage works alongside the idea of the *genetic fulcrum*, a formalisation of the genetic configuration that enables this amplification. Chapter 5 discusses how to define the *natural* genetic fulcrum for a given genotype-phenotype mapping, and how to study a mapping by trying different fulcra and interpreting the results.

These constructs all formalise the basic idea that evolvability is enabled by the core features of an organism's developmental system, which remains relatively constant even as it influences the nature of evolutionary change outside of the developmental system.

Chapter 6: Evolvable modularity

Chapter 4 presented a model of evolvability produced through the abstracted control of the number of repetitions of a single phenotypic feature, a body segment. The model in chapter 6 incorporates abstracted control of which type of feature is expressed and of the ordering of the features in the body. This represents a richer kind of modularity than mere repetition.

The genetic leverage is calculated, confirming that the evolvable mapping is indeed more evolvable.

This model is only slightly more complex than the segmentation model. However, even this slight change introduces new facets to the genotype-phenotype mapping: module arrangement, gene discrimination, and regulatory addressing. These new facets enable new modes of abstracted control, by allowing new forms of abstraction beyond the simple abstraction of repetition number in the segmentation model. This illustrates the idea of abstracted control more fully.

In this way, the model suggests how evolvability can support the evolution of biologically realistic phenotypic modularity, just as the facets themselves better resemble biological development systems, such as hox gene systems.

Formal methods in this dissertation

At certain points this dissertation invents new mathematical constructs to analyse purely illustrative models. This is somewhat unusual and recalls a comment by Francis Crick :

> Elegance and a deep simplicity, often expressed in an abstract mathematical form, are useful guides in physics, but in biology such intellectual tools can be very misleading. For this reason a theorist in biology has to receive much more guidance from the experimental evidence.... (Crick, 1988, p 6)

This warning is quite sensible. Abstract mathematics has a certain rhetorical effect, and it can glorify what are merely more esoteric forms of confusion. Given this, it may be worth explaining how it was in fact a commitment to experimental evidence that led this work in such a mathematical direction.

Work on organismic evolvability springs from a long tradition of empirical research on development and macroevolutionary change. This tradition dates from the earliest efforts in paleontology, taxonomy, and morphology, starting with the first work on the body plan. But as development and genetics have been black boxes for years, this research could only describe an organism's phenotype (which is silent regarding the underlying genetics) and document its history (which is unavoidably unique and contingent). This limited basis of observation left open so many interpretations that arguments for evolvability have always remained a bit speculative, or vague to the point of being unfalsifiable.

But that situation is changing. Recent work on molecular and computational systems has produced much richer, finer-grained data relevant to evolvability, and it has spawned new methods for analysing these data. Also, the rapid progress in evolutionary developmental biology promises such data will become available for the original, traditional system of the organism itself.

In short, conceptual advances have made it possible and empirical advances will make it necessary to update older ideas on evolvability in order to accommodate a more detailed understanding of development. These ideas deserve to be integrated into the mainstream of thinking on development and evolutionary history. A basic first step is simply to spell these ideas out more clearly. This is where mathematics is useful, as this dissertation tries to frame these old ideas without ambiguity. If these ideas are wrong, a formal approach will at least make the error explicit. If these ideas are right, only a formal approach will let us measure how right they are.

Work on evolvability has been around long enough that there are interesting hypotheses leading in every direction. Only greater clarity can show the way forward.