Evolvability: What Is It and How Do We Get It?

Matthew Andres Moreno

April 17, 2017

Contents

1 Introduction **4**
- 1.1 Background . . . 4
- 1.2 Defining Evolvability . . . 5
- 1.3 Introductory Glimpses of Evolvability for Biologists . . . 7
- 1.4 Introductory Glimpses of Evolvability for Computer Scientists . . . 8
- 1.5 Synopsis . . . 9

2 Making Evolvability Concrete **12**
- 2.1 Characterizing Evolvability . . . 12
- 2.2 Quantifying Evolvability . . . 15
- 2.3 Explaining Evolvability . . . 17

3 Proximate Causes of Evolvability **19**
- 3.1 Complexification . . . 20
- 3.2 Duplication and Divergence . . . 22
- 3.3 Developmental Constraint . . . 23
- 3.4 Hidden Genetic Variation . . . 24
- 3.5 Exploratory Growth . . . 25
- 3.6 Weak Linkage . . . 26
- 3.7 Baldwin Effect . . . 27

4 Intermediate Causes of Evolvability **30**
- 4.1 Modularity . . . 30
- 4.2 Robustness . . . 31
- 4.3 Canalization . . . 32
- 4.4 Direct Plasticity . . . 33
- 4.5 Indirect Plasticity . . . 34
- 4.6 Intraindividual Degeneracy . . . 35
- 4.7 Interindividual Degeneracy . . . 37
- 4.8 Regularity . . . 38

5 Ultimate Causes of Evolvability **40**
- 5.1 Indirect Encoding . . . 41
- 5.2 Temporally Varying Goals . . . 43
- 5.3 Environmental Influence on Phenotype . . . 46
- 5.4 Fitness Degeneracy . . . 46

6 Discussion **49**
- 6.1 Development . . . 50
- 6.2 Selection . . . 53
- 6.3 Synthesis . . . 55

7 Conclusion **57**

Glossary **59**

References **61**

Listing of Original Figures

1.3 Illustration of Canalization Against Bilateral Asymmetry in *Drosophila melangoster* 7

2.1 Cartoon Illustration of Individual Evolvability 13
2.2 Cartoon Illustration Contrasting Individual and Population Evolvability 13

3.1 Illustrated Example of Complexification 22
3.2 Schematic Illustration of Local Phenotypic Search 28

4.1 Illustration of Indirect Plasticity in *Spea multiplicata* 35
4.2 Illustration of Intraindividual Degeneracy Among Mammalian Deoxyribonuclease Kinases . . 37
4.3 Illustrated Example of Degeneracy 38
4.4 Illustration of Phenotypic Regularity in *Aloe polyphylla* 39

5.2 Illustration of Indirect Genetic Encoding in an Elephant 42
5.3 Hypothetical Illustration of a Hummingbird Individual Evolvability 45
5.4 Illustration of Fitness Degeneracy Among *Uta stansburiana* 48

Chapter 1

Introduction

1.1 Background

The impressive matching of form to function in biological systems has long been admired by engineers, giving rise to the field of biomimicry, where design elements generated by the evolutionary process are employed in technological applications. Examples of biomimicry include legged locomotion in robotics that provides both efficiency and maneuverability [Grimes and Hurst, 2012], nanotextures mimicking shark skin on boats that discourage barnacle growth while simultaneously decreasing water drag on the vessel [Stenzel et al., 2011], and tire treads inspired by the wet-adhesive properties of tree frog toe pads [Persson, 2007]. Soon after the advent of modern computing, researchers began experimenting with biomimicry at a higher level of abstraction. Instead of mimicking the particular phenotypic forms generated through evolution, they harnessed the evolutionary process — repeated cycles of selection on random variation — to generate novel solutions to a wide array of problems. This approach has since blossomed into the field of evolutionary algorithm (EA) design [Mitchell, 1996]. Language used to discuss EA reflects the biological metaphor on which the algorithm is predicated. A glossary, reviewing evolutionary concepts from both a biological and a digital perspective, is provided. Terms that appear in the glossary are italicized upon their introduction.

Evolutionary algorithms operate on candidate solutions to a problem, which in the biological metaphor are equivalent to *individuals*. The aptitude of candidate solutions to solving a target problem is used to determine the candidate solution's *fitness*, the amount of offspring it generates. Evolutionary algorithms (EAs) traditionally begin with a *population* of randomly-generated candidate solutions. Then, through a series of successive generations, the population is regenerated through *recombination* of fit candidate solutions, so *selection* is performed for candidate solutions that better satisfy the target problem. In biological evolution, a distinction is drawn between the *phenotype* of an individual — the physical characteristics which govern its interaction with the environment, its morphological, physiological, chemical, and molecular

characteristics — and the *genotype* of an individual — the heritable information that influences the phenotype displayed by the individual, i.e. the ordered sequence of base pairs in its DNA. This distinction can become blurred in the realm of evolutionary algorithms, where the phenotypic characteristics of an individual might be directly encoded in the genotype.

The desired outcome of the evolutionary algorithm is, as generations elapse, to observe candidate solutions that provide an increasingly satisfactory solution to the target problem that was used to determine their fitness. Once a predefined stopping criterion is met, usually after a specific number of generations or at a threshold fitness score, the evolutionary algorithm halts. Researchers and engineers have widely demonstrated the ability of EAs to attack labor-intensive optimization problems and to discover novel solutions beyond the reach of human ingenuity [Poli et al., 2008]. For example, evolutionary methods have been successfully applied to evolve communication antenna designs to satisfy the demanding specifications necessary for use in miniaturized spacecraft [Hornby et al., 2006]. Figure 1.1 depicts an evolved antenna design from

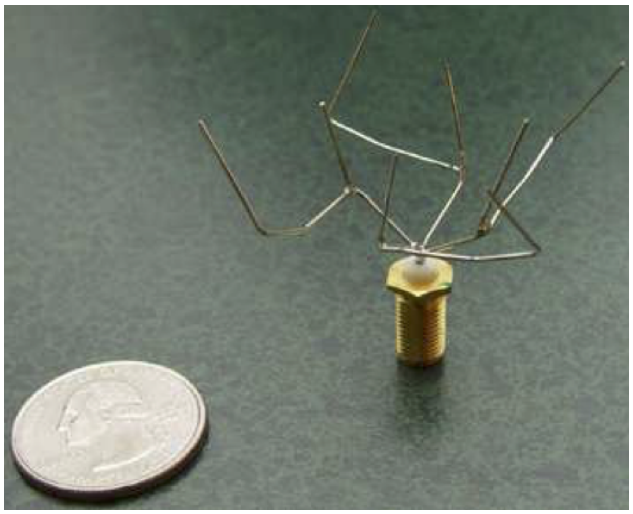

Figure 1.1: A spacecraft antenna design generated using evolutionary methods [Hornby et al., 2006, Figure 2(a)].

that project. Although its form appears alien to traditional human approaches to design, it is nonetheless effective.

1.2 Defining Evolvability

While biological phenotypic adaptation is indeed spectacular, another marvel of biology lurks just below the parade of phenotypes well-suited to their respective environments. It is hypothesized that biological organisms exhibit adaptation to the evolutionary process itself, not just to their environment over the course of their lifespans. That is, biological organisms are thought to possess traits that facilitate successful evolutionary search. The term evolvability was coined to describe such traits. A general consensus exists in the literature that evolvability stems from traits that facilitate the generation of *novel* heritable phenotypic

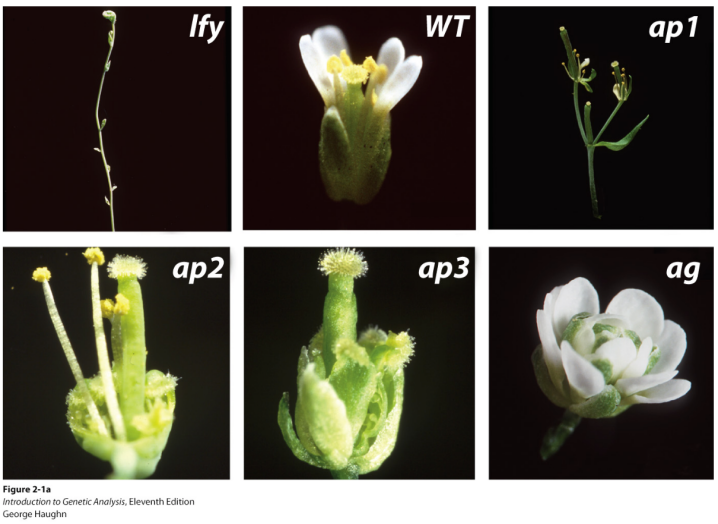

Figure 1.2: Wild-type and mutant strains of *Arabidopsis thaliana* [Griffiths et al., 2015]

variation that is *viable*.[1] Evolvability can be conceived of concretely by imagining a gallery of offspring as depicted in Figure 1.2. Evolvability is determined by the composition of this gallery, the degree to which variation introduced by mutation is deleterious and the amount of phenotypic diversity observed among the offspring in the gallery. An organism among whose potential offspring exist a nontrivial number of individuals that have relatively fit phenotypic forms that exhibit significant structural diversity among themselves and in relation to their parent is highly evolvable.

Breaking the concept down, evolvability stems from:

1. the amount of *novel*, heritable phenotypic variation among offspring,
2. the degree to which heritable phenotypic variation among offspring is *viable*,[2]

The dependence of evolution on these capacities is straightforward. Without any heritable variation, evolution would have no raw material to select from and would stagnate. Without any viable variation, evolution would select against all novelty and again stagnate. Hence, systematic evolutionary change depends the production of heritable, novel phenotypic variation, some of which must not be severely deleterious. We have established plausible traits that might facilitate evolution, but several important questions remain unanswered. How does evolvability manifest in biological organisms (i.e. what traits of biological organisms provide explanations for the presence of viable heritable variation among offspring)? Why does evolvability manifest (i.e. what ultimate mechanistic forces endow biological organisms with traits that promote evolvability)? Addressing these two questions gives us a shot at tackling a third: how can evolvability be

[1]This statement does not suggest that mutation is nonrandom, a controversial and widely discredited theory referred to biologists as adaptive mutation. Instead, it is predicated on the notion that the internal configuration of a biological system (i.e. the developmental process, modularity, degeneracy, etc.) constrains the outcomes of arbitrary perturbations to that system. It is hypothesized that biological organisms possess traits that influence the distribution of phenotypic effects of random mutation.

[2]This can be thought of in terms of the frequency at which lethal or otherwise severely harmful mutational outcomes are observed.

promoted in evolutionary algorithms? We will proceed to explore these questions, but let's begin by priming our intuition for evolvability by considering an artificial selection experiment performed on *Drosophila melangoster*, common fruit flies.

1.3 Introductory Glimpses of Evolvability for Biologists

Experiments performed by Tuinstra et al. (1990) and Coyne (1987) revealed that bilaterally asymmetric phenotypic traits, such as different-sized eyes, could not be induced through artificial selection. In contrast, other artificial selection criteria, such as overall smaller eyes, yielded observable phenotypic changes over the course of a number of generations. A cartoon summarizing these results is provided in Figure 1.3. The success

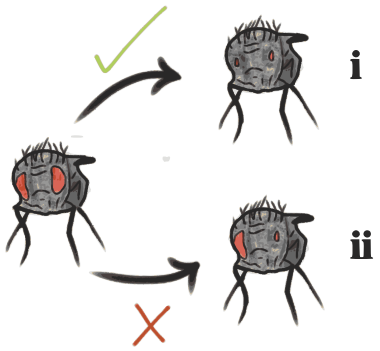

Figure 1.3: This cartoon summarizes an artificial selection experiment performed on *Drosophila melangoster*. In treatment *i*, selection is made for a bilaterally symmetric trait — overall decreased eye size in this illustration. However, in treatment *ii* artificial selection for a non-bilaterally symmetric trait — relatively smaller right eye size — fails [Tuinstra et al., 1990].

of artificial selection for most traits on *Drosophila* demonstrates the existence of a good amount of heritable phenotypic variation for those traits. It is hypothesized that the negative result in artificial selection for bilaterally asymmetric phenotypic traits is due to a lack of bilateral symmetry-breaking information during the embryological development of *Drosophila*. In other words, the very nature of the developmental process constrains the nature of phenotypic variation that can be observed in offspring, in this case curtailing the abundance of offspring that lack bilateral symmetry. As Tuinstra et al. (1990) phrase it, "the developmental system does not seem to allow this type of variation." In the life of a fly, buzzing about in search of food and sex, bilateral symmetry is usually more fit than asymmetry. In this way, the distribution of phenotypic diversity in offspring is biased away from a particular type of deleterious variation, asymmetry. The results from these artificial selection experiments can be cast in terms of evolvability: the distribution of phenotypic outcomes of mutation is not entirely arbitrary. *Drosophila melangoster* more readily exhibits heritable phenotypic variation for certain traits — overall eye size, for example — than for other traits, such as bilateral asymmetry.

In addition to qualities that constrain against non-viable mutational outcomes, discussion of evolvability is predicated on the notion that biological organisms can possess qualities that facilitate significant herita-

ble variation for some phenotypic trait. The regulatory action of hormonal signals such as somatotropin exemplify such a quality. This compound, also known as growth hormone, is well known for its widespread anabolic effects on tissues throughout the body. Mutations affecting the regulatory pathways that regulate somatotropin production and release, receptors and cell signaling components that mediate cellular response to somatotropin, and the protein itself all provide avenues for significant heritable variation in body size [Devesa et al., 2016].[3] The presence of such hormonal signaling pathways could be viewed as making a broad range of heritable phenotypic variation more readily realizable via mutation, increasing individual evolvabiltiy. Dog breeds, which exhibit a range of body weights nearly spanning an entire order of magnitude, evidence the accessibility of heritable variation for body size in animals. Among certain groups of dogs, much of this variation can be explained by just six genes, several of which are associated with pathways somatotropin participates in [Rimbault et al., 2013].

1.4 Introductory Glimpses of Evolvability for Computer Scientists

Computer scientists who have worked on software understand that two pieces of code that meet identical specifications — return identical output for any input given — can differ vastly in difficulty to extend, modify, or maintain. Software implementation, internal structures largely invisible from the perspective of an external interface, accounts for this discrepancy. Computer scientists use the derogatory phrase "spaghetti code" to describe software that is perhaps functional but implemented with such a convoluted control structure that making changes that result in a desired functional outcome becomes very difficult. For a moment, imagine that you are a corporate executive overseeing a large collection of junior developers sourced from the local primate house to develop a word processing application. Imagine that you prompt your army of monkeys at keyboards to begin making arbitrary changes to copies of your code base. You might quantify the functional outcomes of these arbitrary changes. Some changes might have no effect on the functionality of your software product; the software meets the same specifications before and after the code modification. Some changes might fundamentally break the software product, causing it to fail to compile or crash on load. Other changes might cause slight changes to its behavior that significantly degrade its efficacy such as fixing the text cursor at the top of the document. Yet other changes might cause significant changes to its behavior that do not significantly degrade its efficacy such as a complete shuffling of the user menu. It is not inconceivable that the internal configuration of your code base — the extent to which functionality is modularized, the extent to which constants are hard coded versus declared globally, etc. — would affect the outcomes of arbitrary changes made by your junior developers. If the code were structured as a single source file without exception handling, arbitrary changes to the code might be expected to fundamentally break

[3]Recent research implicates somatotropin in a number of processes unrelated to its classical association with metabolism and growth. Although the phenotypic consequences of mutations affecting somatotropin pathways are not exclusively limited to body size, somatotropin response nonetheless provides an avenue for heritable phenotypic variation in that regard.

the software more frequently. If styling information were factored out to a separate specification instead of provided individually for each element of the graphical user interface, arbitrary changes to the code might be expected more frequently to cause large non-lethal alterations to the software product by changing the styles of many aspects of the graphical user interface in one go.

The intent of this thought experiment is not to equate biological evolution and software design. These two processes differ fundamentally on several levels. For example, unlike biological mutation software modifications are not performed at random. The intention is instead to make concrete the notion that internal system configurations fundamentally constrain the outcomes of perturbation of the system, be it through source code changes or mutation. Computer scientists encountering difficulty concretely envisioning how evolvability might manifest in biological systems — or skeptical of the ability of internal system configuration to influence the outcomes of mutation — might take a few moments to recall their own experiences with "spaghetti code."

1.5 Synopsis

Before embarking on our exciting expedition through evolvability theory as it manifests in biological and computational domains, it would be wise to spend a moment consulting the travel agency's promotional brochure to preview the trip's itinerary and anticipate the souvenirs we might hope to take away. While the concept of evolvability, promotion of viable heritable variation, is straightforward, understanding how and why evolvability arises in biological organisms is much more complicated. As will be evident, evolvability can not be traced through a set of direct relationships back to a single point of origin; it appears instead to be an artifact of multiple causality. The causal roots of evolvability branch out in a dizzying multitude of directions.

The primary aim of this work is to survey, organize, and analyze factors that contribute to evolvability using evidence and theory that cuts across the fields of evolutionary computing and evolutionary biology. A novel conceptual framework for the causal factors related to evolvability, which are grouped into proximate, intermediate, and ultimate categories, is presented. This framework is inspired by the distinction between proximate causality, the immediate physical explanation for a phenomenon, and ultimate causality, the less immediate evolutionary explanation for a phenomenon, fundamental to traditional evolutionary thinking. Returning to the conceptual framework presented for evolvability, the label proximal is applied to causal factors of evolvability related to how the physical forms and processes of biological organisms dictate the types of phenotypic variation that can be realized. Discussion of intermediate causal factors related to evolvability steps back to a slightly greater level of abstraction. Intermediate causal factors might be succinctly described as design principles. More concretely, they are overarching patterns in the internal structure and processes of biological systems, manifesting in diverse contexts across and within different levels of biological organization.

Observations of these overarching patterns are typically not limited to just their absence or presence, but often can describe the degree to which they are observed. Proximal and intermediate factors can be remarked upon by observing the properties of an individual in isolation from the larger evolutionary process. Ultimate causality focuses on the fundamental causal factors related to evolvability that give rise to intermediate and proximate causality. From an EA perspective, ultimate causal factors would reflect fundamental assumptions explicitly built into a model of biological evolution while intermediate and proximate causal factors would not be explicitly accounted for in the model but would instead be expected to emerge in the model under appropriate conditions. A review of causal factors related to evolvability, organized by the framework categories of proximate, intermediate, and ultimate, is laid out in Chapters 3, 4, and 5.

In Chapter 6, the proximate-intermediate-ultimate framework is leveraged to analyze possible strategies to promote evolvability. This analysis demonstrates how modifications to both major components of the evolutionary algorithm, the genotype-phenotype mapping and the phenotype-fitness mapping, can be achieved in a manner consistent or inconsistent with the biological metaphor. It is suggested that biologically implausible efforts to promote evolvability at the proximate and intermediate causal levels provide an avenue to achieve evolvability at a feasible computational cost. These strategies to promote evolvability in evolutionary computing discussed in this section are well trod, and lucratively so. This analysis aims not to unearth of novel strategies to promote evolvability in EA, but instead to provide an overarching theoretical explanation and contextualization of these strategies.

Discussion of biologically plausible versus biologically implausible intervention strategies nicely underlines the crux of evolutionary computing: what level of biological realism is necessary to realize the salient features of evolution? This train of thought arrives near the conclusion of our journey through evolvability. Chapter 7 reflects on the issue of evolvability to more broadly situate the fields of evolutionary computing and evolutionary biology in relation to one another. That crux of evolutionary computing can be recast as the question of what components of an evolutionary model are necessary in order to observe innovation on par biological evolution. It is argued that building and evaluating models shifts the task of evaluating the completeness of any particular confection of evolutionary theory from a qualitative endeavor towards a more quantitative endeavor. Evolvability itself crystallizes this point exactly. Without rigorous modeling, the imperative to appeal to evolvability for explanatory purposes is nebulous (although nonetheless recognized by some elements of the evolutionary biology community). Presenting evolution as a narrative tracing the interactions of mechanistic players, the structure of phenotypic variation can be plausibly dismissed as merely random or otherwise left unconsidered. In contrast, lackluster evolvability exhibited by naive models of evolution has loomed as a very concrete roadblock to evolutionary computing researchers; attempting to model evolution makes stark the problem of evolvability.

To summarize the road map compactly, Chapter 2 will equip us with vocabulary and metrics necessary to

precisely quantify and characterize evolvability. Chapters 3, 4, and 5 will review putative factors contributing to evolvability, illustrated with biologically-motivated examples and organized into a hierarchy from proximal to intermediate to ultimate. Chapter 6 will return focus to evolutionary computing, analyzing the factors related to evolvability to examine two overarching intervention strategies to promote evolvability in digital models of evolution. A handful of illustrative examples of these two strategies in action drawn from the evolutionary computing literature follows this exposition. Chapter 7 concludes the paper by reflecting on the relation between evolutionary biology and evolutionary computing. Evolvability specially exemplifies the necessity of modeling to rigorously evaluate the explanatory power of any synthesis of evolutionary theory. The overall goal of this work is to draw a fresh, rigorous connection between theory and major trends, and promising future directions, in evolutionary computing.

With the promotional pamphlet exhausted and the itinerary set, we are now well prepared to embark.

Chapter 2

Making Evolvability Concrete

2.1 Characterizing Evolvability

In relation to both aspects of evolvability, generation of heritable variation and bias towards heritable variation, researchers have developed further theoretical distinctions that allow evolvability to be discussed in a more nuanced and concrete manner. Focusing on evolvability as the generation of heritable variation aspect, discounting bias towards viable variation, Wilder et al. describe two measures of evolvability: individual evolvability and population evolvability [Wilder and Stanley, 2015]. Individual evolvability refers to the potential of an individual to yield a set of offspring that exhibit phenotypic diversity. Figure 2.1 contrasts high and low individual evolvability. In contrast, population evolvability refers to total amount of phenotypic diversity among potential offspring of a population as a whole [Wilder and Stanley, 2015]. These two measures of evolvability are contrasted in Figure 2.2. Although individual and population evolvability might be correlated to some extent, there is not a direct relationship between the two. As Wilder et. al admonish, "population-level evolvability is not equal to the sum over individual evolvability because the novel phenotypes contributed by different individuals may be redundant." [Wilder and Stanley, 2015] The difference between these two types of evolvability is more than semantic; it is thought that population-level evolvability is a much stronger indication of the ability of an evolutionary process to widely explore its search space, discover adaptive variability, and, ultimately, to generate highly-adapted individuals. Wilder et. al argue this point convincingly,

> "On the one hand, evolvable individuals are more likely than others to introduce phenotypic variation in their offspring. On the other hand, in evolvable populations a greater amount of phenotypic variation is accessible to the population as a whole, regardless of how evolvable any individual may be in isolation" [Wilder and Stanley, 2015]

Population evolvability and individual evolvability stem from a different set of proximal causes. An individual

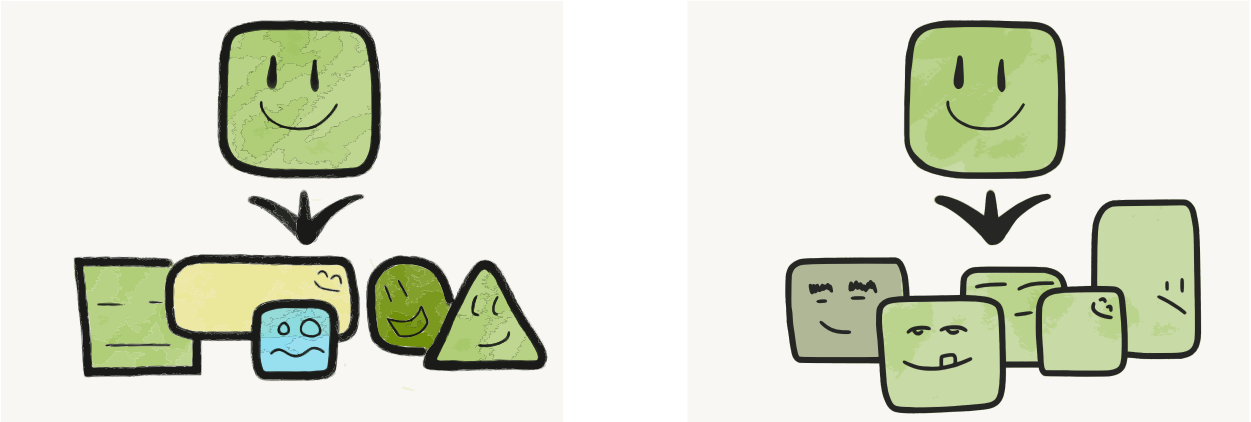

(a) high individual evolvability

(b) low individual evolvability

Figure 2.1: An illustration of individual evolvability, considering evolvability as heritable variation [Wilder and Stanley, 2015].

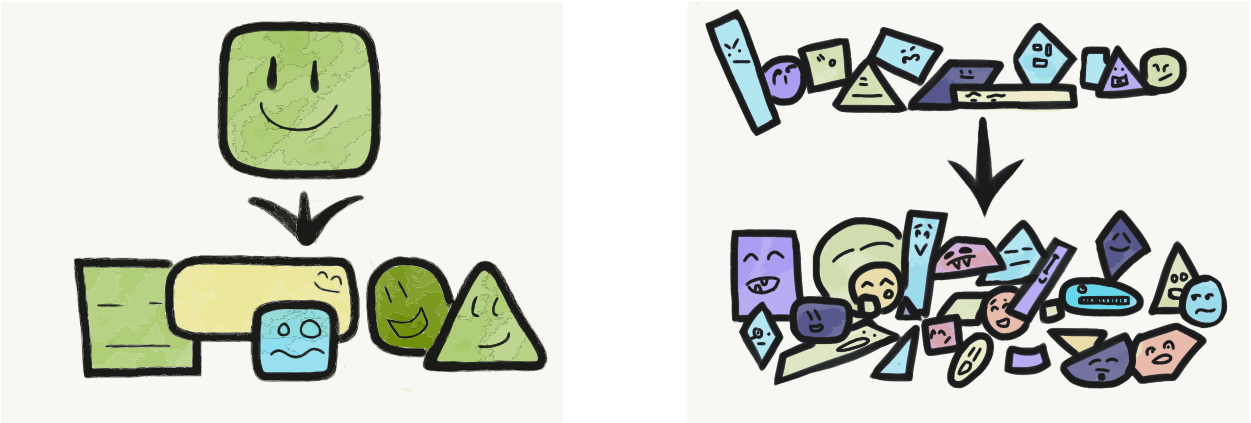

(a) individual evolvability

(b) population evolvability

Figure 2.2: An illustration contrasting individual and population evolvability [Wilder and Stanley, 2015].

with high individual evolvability occupies a region of the genotypic space that maps to a highly variable set of phenotypes; thus, a highly diverse set of phenotypes may be easily reached via small changes in the genotype space. In contrast, high population evolvability is likely due not just to the positioning of individuals in the genotype space, but also to the spreading of individuals from one another throughout the genotype space. In other words, a diverse set of parents will generate a diverse set of offspring.

Considering the other aspect of evolvability, focusing exclusively on bias towards viable variation, another theoretical distinction can be made between innate evolvability, latent evolvability and acquired evolvability. The terms latent evolvability and acquired evolvability were introduced by Reisinger et al. in [Reisinger et al., 2005] to discuss canalization, the ability of a population to control the variability generated among its offspring to in order to exploit fitness biases specific to its environment. Recalling the set of experiments from [Tuinstra et al., 1990, Coyne, 1987] reported in Section 1.1 and summarized in Figure 1.3, the observed bias that maintains bilateral symmetry among *Drosophilia melangoster* is a form of canalization. Distinguishing between innate, latent, and acquired evolvability, it is key to observe that canalization is a "learned" bias, developed over the course of evolution in response. In the case of *Drosophilia melangoster*, the canalization is due to the lack of symmetry breaking information in the developmental process, which itself is defined by the genome of *Drosophilia*. Thus, the information enabling canalization is stored in the genome. As Reisinger et. al put it, "evolvability emerges over the course of evolution with a specific fitness function, and is defined within the terms of that function" [Reisinger et al., 2005]. To better describe the learned nature of canalization, Reisinger et al. introduce the differentiation between latent evolvability and acquired evolvability. According to Reisinger et al., latent evolvability describes "the representations underlying capacity for becoming evolvable" while acquired evolvability describes "evolvability learned in response to a particular fitness function" [Reisinger et al., 2005]. In their experiments, acquired evolvability, which can be observed and quantified, is used as a proxy for latent evolvability. I introduce the term innate evolvability to describe bias towards viable variation that is inherent to a representational scheme. For example, Clune et al. identify bias towards phenotypic regularity, which in certain environments tends to be a useful trait, as an inherent trait of indirect genetic encoding [Clune et al., 2008]. (The relationship between indirect encoding and phenotypic regularity is discussed in detail in Section 4.8). To summarize, latent evolvability describes a representational scheme's potential to support canalization. Acquired evolvability describes actual canalization exhibited by an evolving population in response to a particular fitness environment. Innate evolvability refers to nonlearned bias towards viable variation. These distinctions emphasize the fact that bias towards viable variation can arise through canalization, which is a learned trait where learning is enabled by the representational scheme that relates genotype and phenotype, or can result from qualities innate to a representational scheme, such as a bias towards phenotypic regularity.

2.2 Quantifying Evolvability

Techniques to quantify evolvability are necessary to study it empirically. Reflecting the plurality of definitions for the term in the literature, several of measures of evolvability have been designed. These metrics are in no way mutually exclusive. In fact, they might be viewed as complementary as several measure distinct aspects of evolvability. Techniques to quantify evolvability relevant to evolutionary computing only, not biology, will be reviewed here.

The maximum fitness value achieved by an evolving system provides a convenient and straightforward, but blunt, assessment of evolvability. While evolvability is an important factor in successful evolutionary search, such fitness-based measures do not directly assess the the potential of an evolutionary system to generate viable variation. Thus, such a measure is not directly predictive of properties associated with evolutionary systems with high evolvability, such as the ability to adapt to new, untested environments [Tarapore and Mouret, 2015].

Reisinger et al. employ an evolvability measure that specifically targets acquired evolvability, bias towards viable variation learned in response to the fitness structure of a particular environment[Reisinger et al., 2005]. In their words, their measure targets "the representations ability to retain and generalize information learned about a changing domain." Essentially, Reisinger et al. assesses evolvability as the correlation between the bias towards viable variation exhibited by an evolving system and the amount of information on the fitness structure of an environment the representation receives. Reisinger et al. work with temporally varying fitness functions, where fitness criteria are adjusted from generation to generation. Importantly, these adjustments are made to preserve invariant properties across all fitness criteria. In the bit-string domain employed by Reisinger et al., for example, all fitness criteria manifest bilateral symmetry. The amount of information broadcast to the evolving system is controlled by varying the target drift rate, the rate at which changes to the fitness function are made. Maximal information on the invariant properties of a fitness environment is broadcast at a middling rate of fitness criteria adjustment; adjusting fitness criteria too quickly amounts to training the evolving system on "random noise" and adjusting them too slowly amounts to training the evolving system under a nearly static evaluation scheme where insufficient evolutionary pressure to learn the invariant properties of a fitness environment is exerted. Reisinger et al. measure the assess the amount of bias towards viable variation learned by an evolving system by measuring the efficiency with which an evolving system trained in a fitness environment with invariant properties (i.e. bilateral symmetry) adapts to entirely novel fitness criteria that maintain those invariant properties. It is contended that in a highly evolvable system, maximum bias towards viable variation should be observed at a middling rate of fitness function change. A numerical measure of evolvability is thus calculated as the variance of adaptation over a range of target drift rates encountered during training.

Mengistu et al. develop a measure for individual evolvability, relying upon the characterization of evolv-

ability as heritable variation, to power their evolvability search method [Mengistu et al., 2016]. To assess individual evolvability, a sample of an individual's offspring are generated. A Euclidean distance measure in the phenotype space is used to bin the generated offspring. Individuals exhibiting phenotypic distance from the representatives of each pre-existing bin exceeding a pre-defined, domain-specific threshold are used to found a new bin. The number of bins generated by this process, essentially a count of distinct behaviors present among a sample of the individual's offspring, serves as the metric of evolvability. Wilder et al. take a parallel approach to Mengistu et al., but sample offspring from an entire population instead of an individual in order to measure population evolvability instead of individual evolvability [Wilder and Stanley, 2015].

Clune et al. assess evolvability by examining the relationship between parental fitness and offspring fitness [Clune et al., 2011]. Their approach evolvability provides a window into the amounts of both viable variation and (to a lesser degree) overall heritable phenotypic variation present among offspring. The fitness impacts of variation, how much is how helpful or harmful, are assessed by examining the proportion of offspring with fitness greater than (or less than) their parents and, using reasonable threshold values, the proportion of offspring with fitness much greater than (or much less than) their parents. The spread of differences between parental and offspring fitness is taken as a proxy for the diversity of phenotypic form among offspring. In this scheme, a wide spread of fitness changes would indicate high individual evolvability. While phenotypic variability and fitness variability are certainly correlated to some extent, the distinction between the two should be noted; this measure provides at best an indirect measure of phenotypic variability among offspring.

Tarapore et al. recently introduced an evolvability measure that attempts to take a more clear-eyed view of both of the primary aspects of evolvability: the amount of heritable variation among offspring and the fitness effects of that variation [Tarapore and Mouret, 2015]. They forgo use of a scalar metric to describe evolvability, instead reporting evolvability using what they term a "signature." Essentially, the signature is a two-dimensional heatmap presenting the changes in phenotypic form and fitness observed in individual offspring from a single parent. Normalized mutual information between the phenotypic states of parent and offspring is used to quantify the amount of change in phenotypic form observed in an offspring. Proportion decrease in fitness is used to quantify the fitness difference between parent and offspring. For a highly evolvable individual, we would expect to see offspring occurring with significant frequency in the corner of the heatmap indicating significant change in phenotypic form with slight or no loss of fitness. The evolvability signature provides a nuanced snapshot of evolvability, allowing for interaction between the two primary components of evolvability to be visualized. Such information can be highly diagnostic, for example alerting researchers to phenomena that might appear falsely promising using other metrics, such as genetic changes that alter phenotypic form significantly but at great cost to fitness or genetic changes that are beneficial to fitness but fail to uncover novel phenotypic form.

2.3 Explaining Evolvability

Considering the evolutionary process through the lens of evolvability, one realizes that the value of a solution in terms of the ultimate outcome of evolutionary search isn't completely described by its fitness scores. Although high-scoring individuals are the ultimate goal of the EA process, during the EA process individuals that have the potential to lead to high fitness offspring being discovered by the evolutionary process are the most desirable (although they themselves might not have excellent fitness scores). The fitness score is simply a rough proxy for the quality of future offspring. It does not explicitly take into account an individual's potential (or lack of it) for innovation and adaption in evolutionary time. This observation raises a quandary: how can natural selection "favor properties that may prove useful to a given lineage in the future, but have no present adaptive function" [Pigliucci, 2008]? Researchers are searching out explanations in two main areas:

- evolutionary selection mechanisms, positing that perhaps forces beyond traditional static selection such as divergent selection or a fluctuating fitness function might encourage evolvability, and
- developmental mechanisms, positing that indirect encoding of the phenotype adds inherent bias towards regular, modular phenotypes or allows for the introduction of learned biases that canalize mutational effects towards selectively-advantageous ends.

In biological parlance, these hypotheses fall under the umbrella of ultimate causality, factors that are considered root causes of a phenomenon. Biologists also pursue proximal explanations, which appeal to immediate physical properties and processes to account for observations, to understand phenomena they encounter. In investigating why an flower is purple, for example, a biologist may discuss refractive properties of biochemical compounds present in the flower's skin, proximate causation, as well as the fact that the flower's color scheme has been evolutionarily selected for to attract pollinators, ultimate causation. Our discussion of factors that promote evolvability will trace a similar trajectory. Chapter 3 will discuss proximate factors that promote evolvability in biological systems and Chapter 5 will discuss ultimate explanations for biological evolvability. A third section, Chapter 4, presents a collection of properties commonly associated with biological systems thought to be related to evolvability, which occupy an intermediate ground between ultimate and intermediate. Recalling the superficial analogy drawn between software and biology in Section 1.4, one might liken proximate causal factors to specific design patterns in software that promote maintainability such as the exclusive use of return statements at the end of a function, adherence to don't-repeat-yourself, or avoidance of hard-coded values. Intermediate factors could be likened to high-level attributes of software such as modularity and robustness. Finally, ultimate factors could be related to mechanisms or structures external to a software package that influence its structure such as SCRUM or Agile processes, firm culture, or the operating system it is designed to run on. In Chapter 6, we will conclude with a graphical summary

and discussion of how high-level factors that promote evolvability (i.e. intermediate and ultimate causes) interact.

Chapter 3

Proximate Causes of Evolvability

The robustness of biological organisms, the ability to persist under a range of environmental conditions, is well understood by scientists, and envied by engineers. Students of human physiology drown in the plethora of mechanisms through which the body acts to counter any number of disruptive influences in its relentless pursuit of homeostasis. When blood sugar spikes, pancreatic cells respond by releasing insulin to prompt absorption of glucose from the bloodstream. When ambient temperature decreases, capillaries below the surface of the skin close to prevent loss of heat through the skin, evaporative cooling through sweating ceases, and a shivering reflex in skeletal muscles is triggered to spur heat production [Guyton, 1971, Benzinger, 1969]. Alter most any aspect of the body or its environment within reasonable bounds, and an array of biological systems spring into reaction.

Similar processes are at play in embryological development, cushioning that process against internal hiccups and external interference. The plasma membranes of embryological cells are known to be resilient to strain and perforation. Disruption of the barrier between the cell and its surroundings leads to an influx of calcium ions, which induce rapid incorporation of waiting intracellular vesicles at the site of the leak. These vesicles patch over the damage, allowing the cell to survive and continue along its proper developmental trajectory [Hamdoun and Epel, 2007]. Embryological development exhibits robustness at a broader scope, as well. Consider the processes that lay out the body plan of *Drosophila melanogaster*. In this organism, anterior-posterior directional information is broadcast to developing *Drosophila* cells through a concentration gradient of the morphogen bicoid. However, proper organization of body segments are maintained over a range of temperatures, across which bicoid distribution is known to be disrupted. The mechanisms that account for this surprising maintenance of fidelity to the proper body plan, as of 2007, have not been definitively established [Hamdoun and Epel, 2007].

However, not all biological components and subsystems — in embryos and beyond — are configured to suppress variability in phenotypic form. In response to seasonal cues, butterflies are known to alternatively

suppress or activate expression of a pronounced eye spot. *Daphia*, better known as water fleas, are known to modulate the fortitude of their immune system based on the amount of infectious bacteria mothers encounter during reproduction [Hamdoun and Epel, 2007]. Metabolic processes in adult rats are contingent on nutritive conditions encountered during embryological development. Rats whose mothers were kept on a restricted diet during pregnancy tend to gain more weight when exposed to high-calorie diets than rats whose mothers had better access to food [Wilson, 2007, p 57]. A similar embryological decision-making approach is at play in male dung beetles, which develop along alternate trajectories — one displaying a horned forehead in adulthood and the other displaying a naked forehead in adulthood — depending on their body size at a certain juncture of infancy. Horns are thought to be most useful to beetles capable of levying significant bulk against romantic competitors and are otherwise superfluous [Wilson, 2007, p 47].

Biological organisms are predisposed to minimize some types of phenotypic variability and to allow, or even encourage, other variability in other phenotypic dimensions. The systems and subsystems that regulate variability often persist under mild mutative perturbation. Thus, these systems and subsystems act upon the mutative perturbation and mold the phenotypic effects that result from the mutation. Some phenotypic changes are suppressed while others are tolerated.

Proximal causality of evolvability results from the relationship between the phenotypic form of biological organisms and the phenotypic changes that can manifest as a result from mutation. Kirschner et al. zero in on this idea, emphasizing that — although genotypic mutation is essentially an arbitrarily random process — "phenotypic variation cannot be random because it involves modification of what already exists" [Kirschner and Gerhart, 2005, p 220]. I agree with the sentiment of this observation, however the precision of the language Kirschner et al. use could be improved. They use the word random to express an idea that could be better termed as unconstrained or arbitrary randomness. Ultimately, phenotypic variation results from genetic perturbation, which is random and arbitrary. The key observation that Kirschner et al. have made, and which will be presented in this section, is that while phenotypic variation is stochastic, it is not arbitrary. This section aims to highlight how design patterns common in biological organisms transform random mutation into structured phenotypic variation. A good amount of material presented in this section is framed by or drawn from the Theory of Facilitated Variation, originally presented in [Kirschner and Gerhart, 2005], as reported in [Downing, 2015].

3.1 Complexification

3.1.1 Definition

Complexification refers to the development of sophisticated phenotypic functionality through the incremental refinement and elaboration over evolutionary time [Downing, 2015, pg 202]. In this process, a phenotypic

structure results not from a single mutation event but instead from a sequential series of modifications, each building upon the last. Many intermediate forms of the phenotypic feature are observed, with a trend towards increasing sophistication, are observed. It should be noted that the stepping stones of complexification, intermediate phenotypic forms, need not share the same functionality.

3.1.2 Relation to Evolvability

Complexification biases evolutionary search towards viable variation. In short, variants of already-existing phenotypic structures are more likely to be useful than randomly-generated phenotypic structures If existing phenotypic structures could not readily be elaborated upon, the chance of evolution happening upon viable variation would be vanishingly small.

The capacity of a genotype-phenotype mapping to accommodate complexification in evolutionary computing cannot be taken for granted. Consider, for example, an evolving artificial neural network. In this context, a network that starts out with very few nodes and adds nodes over the course of evolution — thus, developing a more complex phenotype — exhibits complexification [Clune et al., 2011]. Genetic encodings that encode phenotypes with a fixed number of nodes would be incapable of exhibiting this form of complexification. The capacity genotype-phenotype mapping to accommodate complexification — its capacity to translate additions of genetic information into refinement of existing phenotypic features — will be key to achieving evolvability in evolutionary computing.

3.1.3 Example

The evolution of the vertebrate eye epitomizes complexification. It is thought that this structure evolved through a series of intermediates, beginning with a simple region of enervated photosensitive cells. A folded-in, photosensitive pouch-like structure, which provided directional sensitivity, is thought to have arisen next. Pinhole and lens structures, which provide greater visual acuity, are thought to have descended from the pouch structures [Gregory, 2008]. These intermediate phenotypic structures, each elaborating on an existing phenotypic form, can be observed in extant organisms as illustrated in Figure 3.1.

The complexification of vertebrate eyes was enabled by the nature of the biological genotype-phenotype mapping. The biological genotype-phenotype mapping mapping is capable of accommodating increased amounts of genetic information describing an existing phenotypic trait and translating that new information into refinement of the existing phenotypic trait. For example, in the transition from light sensitive patches to photosensitive pouches the genotype-phenotype mapping accommodated new information altering the shape of the surface in which light sensitive cells are embedded while preserving the general arrangement and connectivity of the light sensitive cells.

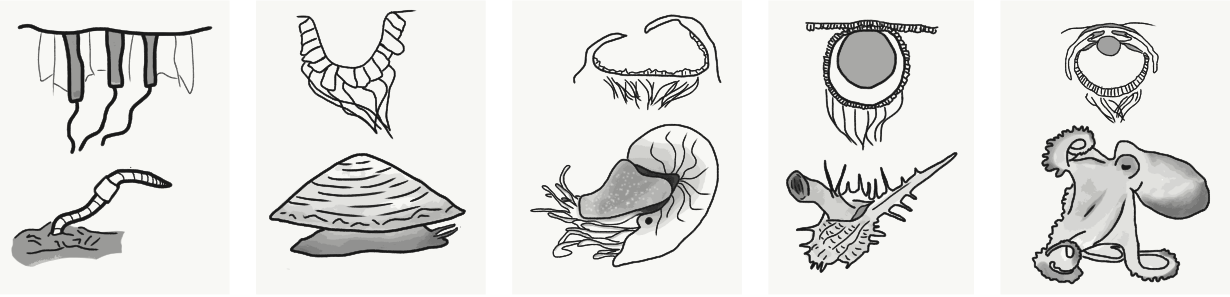

(a) light sensitive cells in the worm

(b) pit-type eye in the limpet

(c) pinhole-type eye in the nautilus

(d) eye with lens in the murex

(e) camera-type eye in the octopus

Figure 3.1: Extant organisms that illustrate several different types of eyes, which might have provided a potential route of complexification of the ocular organ by evolution [Gregory, 2008].

3.2 Duplication and Divergence

3.2.1 Definition

Duplication and divergence refers to a two step process thought to be essential to the evolution of new genes in biology. In the first step, duplication, the genetic information coding for a phenotypic feature is duplicated in the genome. In the second step, divergence, one of the gene copies is modified by mutation.

3.2.2 Example

Homeobox genes provide a canonical example of duplication and differentiation in action. These genes, which lay out the body plan of a developing embryo, have been duplicated and differentiated many times, at each step adding novel body plan features [Downing, 2015, p 203]. In particular, Homeobox genes have seen re-use in determining brain segmentation. Fascinating parallels can be drawn between the the topology of the brain and an animal body plan [Downing, 2015, p 201].

3.2.3 Relation to Evolvability

Duplication and divergence allows for re-use of evolved functionality. It is much more tractable to achieve sophisticated functionality by modification of existing biological components, for which evolution has already worked out most of the kinks, than starting from scratch. Further, duplication and divergence can jump start the evolution of new functionality without disrupting the functionality originally encoded in the genome. The original gene is preserved. By allowing gene evolution to fork out to achieve a widening array of functionality, duplication and divergence paves a tractable route to gradual phenotypic complexification [Downing, 2015, p 202].

3.3 Developmental Constraint

3.3.1 Definition

Developmental constraint refers to the influence of the development process on the distribution of phenotypic forms that can be generated in the offspring of an individual [Smith et al., 1985]. The concept of developmental constraint rests upon the idea that certain phenotypic forms are more likely to arise because they are more readily generated by physicochemical processes that underpin embryological development [Smith et al., 1985]. Because the genotype is interpreted into phenotypic form in large part through the developmental system, the phenotypic outcomes of mutation are closely linked to the configuration of the developmental system.

3.3.2 Relation to Evolvability

Developmental constraint might restrict nonadaptive variation, biasing evolutionary search towards viable phenotypes [Downing, 2015, pg 40]. Developmental constraint that funnels phenotypic variation towards viability contributes to canalization (Section 4.3). However, developmental constraints may also be arbitrary. For example, centipedes exhibit an apparent constraint towards developing an odd number of body segments [Arthur and Malcolm Farrow, 1999]. Thus, developmental constraint can, but need not necessarily, contribute to canalization.

3.3.3 Example

Artificial selection experiments performed by several laboratories on *Drosophila melangoster* revealed developmental constraint towards bilateral symmetry. Artificial selection experiments on symmetrically distributed traits (i.e. overall eye size, overall bristle count, etc.) *Drosophila melangoster* have a high rate of success. However, artificial selection for left-right traits generally fails. For example, artificial selection experiments were unable to introduce an asymmetric bias in thorasic bristle counts despite many generations of selection. Although phenotypic variation in left-right thorasic bristle number existed in the population, that variation was not heritable. Similar results have been reported in relation to artificial selection for asymmetric distribution of eye size [Coyne, 1987], eye facet number [Manyard Smith and Sondhi, 1960], and wing-folding behavior [Purnell and Thompson, 1973] in *Drosophila melanogaster*. Tuinstra et al. hypothesize that this canalization results from the developmental process, "as no evidence is available for an independent left-right gradient in the embryo, quantitative traits can only be expressed variably along an existing gradient of positional information or a morphagen" [Tuinstra et al., 1990]. It seems that the nature of the developmental process makes heritable variation for bilaterally asymmetric traits in *Drosophila* more difficult to come by. A cartoon summarizing the outcome of these experiments is provided in Figure 1.3.

3.4 Hidden Genetic Variation

3.4.1 Definition

Hidden genetic variation, also known as cryptic variation or neutral variation, refers to genetic diversity in a population that does not manifest as phenotypic diversity. In biological systems, many genotypes map to a identical or nearly identical phenotypes. This many-to-one relationship may be due to the presence of variation in non-coding DNA, variation in genes for which phenotypic effects are suppressed by regulatory mechanisms, or degeneracy in the genetic code (i.e. several codons encoding the same amino acid residue). It is thought that the homeostatic mechanisms that promote robustness facilitate hidden genetic variation by counteracting the phenotypic changes that might be induced by some forms of genetic variation [Moczek et al., 2011]. Environmental influence on the phenotype might contribute to an intermediate form of cryptic variation where a trait is expressed phenotypically in only a subset of the population due to differences in environmental conditions. Such a trait would therefore be somewhat hidden from phenotypic expression [Moczek et al., 2011].

3.4.2 Relation to Evolvability

Because cryptic genetic variation is not selected upon by evolution, it can accumulate in a population. This accumulated genetic variation is thought to promote evolutionary innovation [Wilder and Stanley, 2015] several ways. By allowing for a broader distribution of a population through a genotype space, cryptic variation increases the phenotypic diversity that can be realized in offspring from the population because individuals in the population can "access radically different phenotypes in their immediate mutational neighborhood" [Wilder and Stanley, 2015]. Cryptic variation is also thought to allow for larger steps to be taken in the mutational search space during evolutionary search. Significant accumulated cryptic variation can rapidly switch to being expressed through a sensitizing mutation or environmental change [Moczek et al., 2011].

3.4.3 Example

Work with strains of the model organism *Drosophila melanogaster* possessing the mutation *Scute* has unmasked elements of hidden genetic variation in wild type *Drosophila* [Wagner, 2003]. The mutation *Scute* increases the number of bristles observed on *Drosophila* and, more intriguingly, also boosts the variability in bristle count observed between individuals. *Scute* mutants exhibit exhibit more variability in bristle counts than wild type *Drosophila* by several orders of magnitude. Artificial selection experiments revealed a stronger response in the mutant populations to selection for bristle counts than in wild type populations. Thus, some of the variation of the *Scute* mutant phenotype seems to have a heritable genetic basis. It is hypothesized that the mutation *Scute* sensitized *Drosophila melanogaster* to existing genetic variation. The experiments

demonstrated "genetic variation for the character that is not expressed in the wild type, but becomes visible in the mutant background" [Wagner, 2003]. In other words, the mutation *Scute* revealed previously hidden genetic variation for bristle count.

3.5 Exploratory Growth

3.5.1 Definition

Exploratory growth refers to the incorporation of search into the developmental process. Instead of having developmental components grow to hard-coded proportions and in hard-coded locations, developmental processes incorporate information about the the current state of the organism into the developmental trajectories of system components [Downing, 2015, p 214].

3.5.2 Relation to Evolvability

Because other systems in a developing organism can change to adapt to changes in one system, exploratory growth reduces the probability of mutations leading to catastrophic fitness decline or outright mortality. Thus, exploratory growth promotes both canalization and robustness: translation of genetic changes into phenotypic effects are more likely to be resisted and, if they do occur, they are more likely to be viable because the development of the organism adapts to compensate for those changes. As Downing argues, due to exploratory growth "the production of novel phenotypes does not require concerted change to many parts of the genome, a very low-probability combination of events, but rather a single change to a factor affecting an early phase of development. The rest just grow to fit the altered context..." [Downing, 2015, p 214].

3.5.3 Example

The structural components of the mammalian body develop in a "grow to fit" pattern. The development of bones and muscles is defined by a process in which they seek out attachment sites defined by other components of the developing system. These components "grow to fit" in the sense that they grow (or shrink) to fit an altered developmental context. The final form of these structures is determined through exploration of the developmental environment [Downing, 2015, pg 214]. Mesychme cells, a cell type observed in animal embryos that is a precursor to several tissue types, provide another striking example of exploratory growth. During embryological cellular migration, mesenchyme cells extend filopodia to explore their environment; several filopodia attach to different sites on the blastocoel wall, and the final position of the cell is determined from a tug-of-war between the filopodia — the mesenchyme cell will choose the site where it found most stable attachment. In this way, exploration performed by mesychyme cells contributes to the development process [Downing, 2015, pg 214].

The development of the nervous system is heavily marked by exploratory growth. In the brain, an excess of candidate neural networks are created during the development process. Fledgling networks with insufficient connectivity to the rest of the network are then culled; essentially, the final components of the brain are derived of a much larger number of competing trial subnetworks, many of which are failures [Downing, 2015, p 214]. Ennervation of the body during development proceeds in a similar trial-and-error fashion. Excess neurons migrate and extend processes to seek out and compete for targets; many are unsuccessful and, ultimately, die [Edelman and Gally, 2001].

3.6 Weak Linkage

3.6.1 Definition

The concept of weak linkage is predicated on a distinction between instructive and enabling signals [Downing, 2015, p 210]. Instructive signals contain significant amounts of information about the process to be performed, not just information that a process should be performed. In contrast, enabling signals are concise; the only information they contain is that the process should be performed. In an enabling signal scheme, information characterizing a response to the signal is stored in the system that receives the signal and not in the signal itself [Kirschner and Gerhart, 2005, p 283]. Weak linkage refers to interaction of subsystems of a biological organism coordinated by enabling, rather than instructive, signals.

Let us examine an economic analogy based around a hypothetical bakery to tease apart the difference between instructive and enabling signaling. Suppose the bakery has a telephone. Detailed instructions delivered over the telephone about the recipes, ingredient sources, and a schedule for production that should be used would be considered an instructive signal. Consider if, instead, the baker were to adjust his production based solely on the frequency with which the telephone rang throughout the day. Under this scheme, bakery output can still be influenced by external signals. However, the amount of information required to signal the bakery is reduced. Instead of needing to communicate recipes, ingredient sources, and a schedule for production, external sources just need to know the telephone number for the bakery. Thus, such a scheme would be considered enabling signaling. A collection of industries in small town coordinated by such enabling, rather than instructive, signals would be said to exhibit weak linkage.

3.6.2 Relation to Evolvability

With simple signaling protocols, the probability of mutation establishing interaction between two systems via signaling is increased. Of particular interest is the role of weak linkage in allowing an externally triggered signal to become innate. On a cellular level, many environmental signals, such as the concentration of a certain chemical compound in an organism's environment, are of an enabling nature. Environmental

enabling signals can often be mimicked by the cell itself and are therefore accessible to becoming innately triggered [Downing, 2015, p 210]. This scheme provides a plausible for phenotypic traits that are originally environmentally-induced (i.e. indirect plasticity, see Section 4.5) to be incorporated on a permanent, heritable basis.

Additionally, under a weak signaling regime information required to perform a process is more tightly contained within the context of the subsystem that performs the process. Thus, weak signaling promotes modularity.

3.6.3 Example

In the grossest terms, a neuron consists of an input component, which is responsive to the presence of a specific subset of neurotransmitters, and output, which releases a different specific subset of neurotransmitters. The specific subsets of neurotransmitters a neuron's input and output is sensitive to varies between neuron types. These two components of the neuron interface via electrical voltage. An action potential conducts information from the input component of the neuron to its output component. The action potential is an enabling signal. Information contained in the signal is minimized — binary on/off information is signified by the presence or absence of an action potential. This arrangement allows input and output components to be freely mixed and matched in a single neuron. The freedom to mix and match is enabled by the simple nature of the signal that interfaces the input and output components. Thus, weak linkage makes a large number of viable neural configurations readily accessible to evolution [Kirschner and Gerhart, 2005, p 139].

3.7 Baldwin Effect

3.7.1 Definition

The Baldwin effect postulates that local search in the phenotype space biases evolutionary search to allow advantageous phenotypic features originally acquired via phenotypic plasticity to be encoded into the genetic representation [Downing, 2010]. In the first phase of the Baldwin effect, advantageous phenotypic features discovered by local search in the phenotype space increase the fitness of individuals proximal to that phenotype. In the second phase of the Baldwin effect, continued evolutionary search encodes phenotypic features originally discovered by phenotypic plasticity into the genetic representation.

The first phase of the Baldwin effect is driven by local phenotypic search. This concept is illustrated in Figure 3.2, where the phenotype space is depicted as a three dimensional surface with points on the surface representing different phenotypes and the height of the surface denoting the fitness of the phenotypes at those points. In the illustration, the dark blue square represents the phenotype originally mapped to by the genetic representation of an individual, the blue shaded region represents the region of the phenotype space

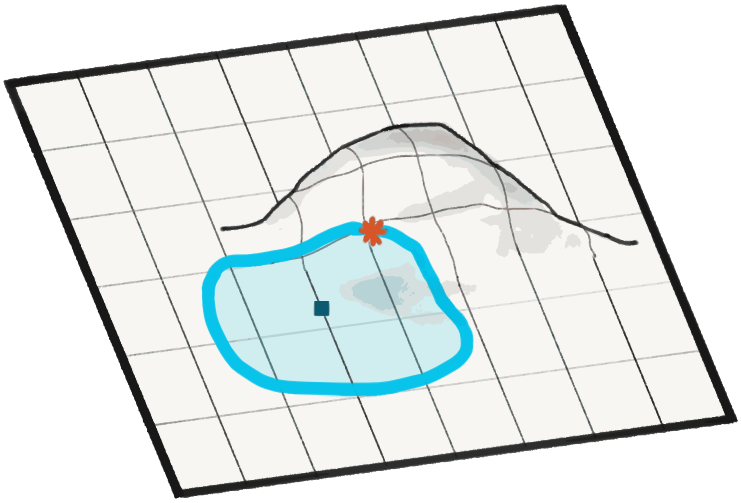

Figure 3.2: A schematic illustrating local search in a phenotypic landscape.

searched via phenotypic plasticity, and the red star represents a high fitness phenotypic variant reached via phenotypic plasticity. The fitness boost gained from phenotypic proximity to a high fitness solution biases evolutionary search to continue exploring that region instead of proceeding in other directions.

Cost or unreliability of generating the advantageous phenotypic feature via phenotypic plasticity provides evolutionary pressure that favors genetic encoding of that phenotypic feature, driving the second phase of the Baldwin effect. In the case of indirect genotype to phenotype mappings, a direct genetic encoding of the feature discovered via phenotypic plasticity might not exist; however, genomes that map to phenotypes closer to the phenotypic feature discovered by plasticity — which support attainment of the phenotypic feature by reducing the cost or increasing the reliability of acquiring that feature via plasticity, providing "scaffolding" for the local phenotypic search — may still arise and will be selected for [Downing, 2012].

3.7.2 Relation to Evolvability

By allowing a candidate solution to assume proximal phenotypic forms, local phenotypic search allows evolutionary selection to act on information about a candidate solution's local phenotypic neighborhood. Through selective pressure for phenotypes proximal to high-fitness phenotypic forms, local phenotypic search "buys evolutionary time" until heritable scaffolding arises to support phenotypic adaptation originally attained via plasticity [Downing, 2010].

3.7.3 Example

The Baldwin Effect has been hypothesized to play a role in mammalian brain evolution. Learning is thought to play the role of local phenotypic search [Downing, 2010]. In the first phase of the Baldwin Effect, advantageous brain organizational traits would have been achieved by learning. Individuals capable of achieving these traits by learning would exhibit greater fitness. This fitness reward allowed individuals with a baseline capacity to achieve advantageous brain organizational traits by learning to persist. Downing claims that over evolutionary time the processes of neurogenesis, synapogenesis, and synaptic tuning have shifted away from the postnatal life phase, becoming increasingly concentrated in the prenatal developmental developmental phase [Downing, 2012]. More simply put, brain organization activity retracted into the initial embryological phase of life. The phenotype is generally considered to be more strongly governed by genetic influences (as opposed to environmental influences) during the embryological phase relative to later phases of life. The retraction of brain organization activity brain into the initial embryological phase of life thus corresponds to the second phase of the Baldwin Effect, a shift towards stronger genetic influence on traits related to brain organization originally discovered by local phenotypic search. Under this hypothesis, brain structures that were originally obtained by learning (local phenotypic search) became increasingly encoded in the genotype.

Chapter 4

Intermediate Causes of Evolvability

Discussion of intermediate causal factors related to evolvability steps back to a slightly greater level of abstraction than proximal causal factors. These causal factors describe design motifs common to biological systems rather than specific processes or structures as was the case in discussion of proximal causal factors. The overarching patterns in the internal structure and processes of biological systems classified as intermediate causal factors manifest in diverse contexts across and within different levels of biological organization. Observations of intermediate causal factors related to evolvability are typically not limited to just their absence or presence, but often can describe the degree to which they are observed. That is, these intermediate causal factors related are often quantifiable, especially in a digital context.

These causal factors are termed intermediate because meaningful levels of causality exist above them. These higher level factors are termed ultimate causal factors related to evolvability and are presented in the following chapter. For the most part, intermediate causal factors describe qualities of a biological system as a whole instead of a describing a specific subcomponent or subsystem of the system as tended to be the case for proximate factors related to evolvability. By providing convenient language to describe causal factors related to evolvability at the level of entire organisms rather than their constituent subsystems and internal processes, intermediate causal factors yield a useful conceptual interface between proximate and ultimate causality.

4.1 Modularity

4.1.1 Definition

Modularity refers to the organization of a system into distinct subsystems where interaction between components of the subsystem outstrips interaction between components inside of a subsystem and components outside of it [Downing, 2015, p 207]. By distilling the system into compact units of independent func-

tionality, modularity allows reuse of that functionality in different parts of the system through duplication [Reisinger et al., 2005].

4.1.2 Relation to Evolvability

Downing argues convincingly for the utility of modularity,

> Once structures or mechanisms are consolidated and isolated (to some degree) from external influences, their probability of disruption declines and their potential for self-modification without global repercussions increases: they enhance both robustness and adaptability [Downing, 2015, p 208]

The compartmentalization of functionality makes independent changes to one aspect of functionality possible while other aspects of functionality are preserved. By increasing the robustness of other aspects of functionality in an organism to changes in one aspect of functionality, modularity ultimately biases evolutionary search space towards viable offspring. Further, compartmentalization of functionality allows for re-use of that functionality in another context in the system. Because complete units of phenotypic functionality are more likely to be useful in new contexts than randomly generated structures, modularity allows for the generation of divergent phenotypes (i.e. diverse offspring) that are likely to be viable. In this way, modularity promotes complexification.

4.1.3 Example

The segregated cell roles in multicellular organisms provides an excellent example of modularity. Consider, in particular, simple squamous epithelial tissue. This tissue cells are typically constitutes regions of membrane requiring rapid diffusion or filtration. The tissue's self-contained structure facilitates the reuse of its functionality across many different parts of the body. Simple squamous epithelial tissue appears in a diverse set of anatomical contexts including the air sacs of lungs, the walls of capillaries, and the kidneys [Owens and Lane, 2003].

4.2 Robustness

4.2.1 Definition

Robustness generally refers to the ability of a system's function to persist under perturbation. Biological systems exhibit robustness on three levels, as identified in [Richter et al., 2015]:

1. system functionality is not degraded by stochastic fluctuations in the system,
2. phenotypic traits are generally not degraded by genetic variation, and

3. the phenotype is able to function if environment changes.

4.2.2 Relation to Evolvability

By reducing outright mortality and other phenotypic variations with catastrophic fitness implications, robustness biases evolutionary search towards viable variation.

4.2.3 Example

The heart provides a prime example of level 1 robustness. The waves of electric potential that govern the function of the heart are generally robust to perturbation; the heart can usually recover from momentary disruptions to these waves and resume normal function after a short time. However, the heart is not a totally robust organ: it can be susceptible to ventricular fibrillation, where normal heart function is essentially halted by the emergence of a spiral wave pattern of electrical activation [cite]. In many cases, the consequences of the heart not being able to recover from the perturbation that induced the spiral wave pattern of activation are fatal.

Duplicate genes exemplify robustness at level 2 [Gu, 2003]. It has been shown that deletion duplicate genes have a significantly lower proportion of fatal outcomes and a significantly higher proportion of weak or no effect outcomes in *Saccharomyces cerevisiae* (yeast) [Gu, 2003]. Robustness is also provided by alternate metabolic pathways and regulatory networks that can compensate for the absence of a gene [Gu, 2003].

Finally, robustness at level 3 is exhibited by the brown rat (*Rattus norvegicus*). This creature has been wildly successful in a wide range of environments — today, its range spans nearly the entire globe — that differ vastly from the environment where it originated [Wikipedia contributors, 2016]. In fact, *Rattus norvegicus* is particularly notorious for thriving in urban areas, a completely novel environment that only came into existence in their current form over the last millennium or so. The brown rat displays impressive robustness to environmental variation.

4.3 Canalization

4.3.1 Definition

Canalization refers to "organizing the effects of mutations such that some features become more resistant to change" [Reisinger et al., 2005]. Canalization can stem from many aspects of phenotypic form and the phenotype-genotype mapping. For example, developmental constraint is associated with canalization. Importantly, canalization can be selective — some phenotypic features become more resistant to change than others.

4.3.2 Relation to Evolvability

Canalization can provide selective bias against inviable phenotypic outcomes while still allowing for viable phenotypic variability. Thus, canalization contributes directly to one of the major components of evolvability: bias towards viable variation. Through this bias towards viable variation, canalization contributes to an increase in the overall quality of offspring to be selected from [Downing, 2015, p 40].

4.3.3 Example

Developmental constraint is a form of canalization. The developmental constraint revealed by artificial selection experiments performed on *Drosophila melanogaster* illustrate canalization well. These experiments, presented in Section 3.3.3, revealed heritable bilaterally symmetric phenotypic variation is more readily accessible than bilaterally asymmetric variation. Hence, *Drosophila* exhibits canalization against systematic asymmetry.

4.4 Direct Plasticity

4.4.1 Definition

Plasticity, in general, refers to environmental influence on the phenotype. In biology, environmental and genetic influences, together, shape the phenotype. Environmental influences may alter the trajectory of the developmental process or may otherwise induce phenotype changes in response to environmental stimulus [Fusco and Minelli, 2010].

Direct plasticity plasticity is specifically related to environmental influence that is exerted directly and coercively on developmental or physiological processes. An organism that exhibits direct plasticity is robust to these environmental perturbations. At a fundamental level, successful direct plasticity entails *resistance* to environmental influence on the phenotype.

4.4.2 Relation to Evolvability

It is thought that the homeostatic mechanisms that mediate an organism's interactions with its environment (i.e. the mechanisms that support to direct plasticity) might also promote robustness to mutation [Moczek et al., 2011]. That is, the same traits that help an individual maintain functionality under environmental perturbation may also help protect the individual against catastrophically deleterious mutational outcomes.

4.4.3 Example

Osmotic pressure, generated via the diffusion of water into a cell due to a disparate solute concentrations inside and outside a cell, constitutes a grave threat to biological cells. Unchecked, this coercive physical influence may indeed burst a cell, such as an animal erythrocyte [Lodish et al., 2000]. Such cells, unable to withstand disruptive environmental physical influences, would be said to exhibit poor direct plasticity in that regard. Many biological organisms do, however, exhibit strong direct plasticity with respect to osmotic pressure. Plant, algal, fungal, and bacterial cells use a rigid cell wall to withstand osmotic pressures. Many protozoa cope with osmotic influx of water by capturing excess fluid in the cytosol in a contractile vacuole that is occasionally discharged into the extracellular environment [Lodish et al., 2000].

4.5 Indirect Plasticity

4.5.1 Definition

Indirect plasticity is specifically related to environmental influences more akin to informational signals than coercive physical influences. An organism in which these environmental signals prompt responses that are mediated by physiological or developmental systems exhibit indirect plasticity. This term describes the ability of an organism to receive, process, and respond to environmental information through adjustments to its own phenotype [Fusco and Minelli, 2010]. At a fundamental level, successful indirect plasticity entails strategic *amplification* of certain aspects of environmental influence on the phenotype.

4.5.2 Relation to Evolvability

Researchers have suggested that phenotypic modularity might be a key contributor to indirect plasticity, especially in plants [Schlichting, 1986, De Kroon et al., 2005]. Thus, selection for indirect plasticity, which is known to play an important role in adaptation to an unpredictable or variable environment [Fusco and Minelli, 2010], might promote modularity.

Conditional expression of phenotypic traits through plasticity allows for relaxed selection on the genotypic locus determining those traits. Thus, significant genetic variation can accumulate at that locus in a population. In a process known as genetic accommodation, the environmental influence on when rarely-expressed phenotypic traits are expressed can be diminished or erased through sensitizing mutation; what once was induced via environmental signal can become constitutive. Such processes have been observed experimentally via artificial selection [Moczek et al., 2011].

Finally, the role plastic processes such as learning might play in concert with phenotypic regularity is an open question. It is possible that phenotypic remodeling induced via environmental cues may provide a mechanism of irregular refinement of highly regular phenotypic structures generated via development. That

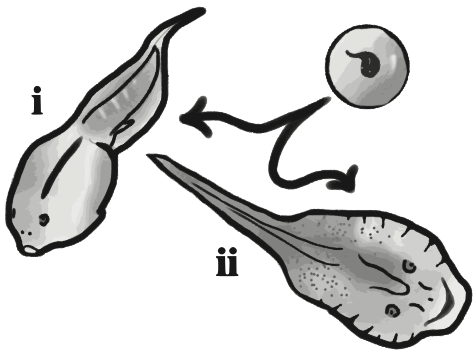

Figure 4.1: An illustration of the alternate phenotypic forms of *Spea multiplicata*. Form *i* is known as the omnivorous form while form *ii* is known as the carnivorous form. Developing tadpoles assume one or another of these forms on the basis of environmental signaling, thought to be the consumption of brine shrimp [Pfennig, 1992].

is, interaction of an individual with its environment may induce irregular phenotypic adjustments which increase fitness [Clune et al., 2011].

4.5.3 Example

Tadpoles of *Spea multiplicata*, which is also known as the Mexican spadefoot toad, develop along two alternate trajectories in response to environmental signals. The two tadpole morphs of *Spea multiplicata*, depicted in Figure 4.1 exhibit different jaw and digestive tract configurations and prefer different diets [Fusco and Minelli, 2010]. One morph is specially suited to an carnivorous diet while the other is suited to an omnivorous one[Pfennig, 1992].

While *Spea multiplicata* illustrate phenotypic plasticity in development, phenotypic plasticity — although often more subtle — also manifests on an ongoing basis throughout the lifespan of most creatures. *E. coli* provide a textbook example of phenotypic alteration in response to environmental stimulus. These bacteria selectively express an enzyme used in the digestion of lactose, β-galactosidase in the presence of lactose when alternate food sources, such as glucose, are unavailable. This phenotypic change is mutable; in the absence of lactose or the presence of glucose production of β-galactosidase halts [Griffiths et al., 2015].

4.6 Intraindividual Degeneracy

4.6.1 Definition

Degeneracy refers to the exhibition of similar functionality by structurally different systems [Edelman and Gally, 2001]. For example, one might refer to push doors, rotating doors, and sliding doors as degenerate; although they all accomplish similar tasks, they are structurally dissimilar.

In an organism exhibiting intraindividual degeneracy, distinct components with different structures perform the same function. Degenerate components are functionally, but not structurally, redundant.

4.6.2 Relation to Evolvability

Intraindividual degeneracy is thought to be key to the evolvability observed in biology. Intraindividual degeneracy provides a solution to the seemingly paradoxical relationship between robustness (Section 4.2) and evolutionary innovation, the ability to generate heritable novel variation [Whitacre and Bender, 2010]. How can a system have both robustness, which constrains change to the system, and innovation? Intraindividual degeneracy seems to be the answer. By having multiple independent components that perform the same task, the system is secured against the failure of one of the components and, therefore, more likely to persist under extreme environmental conditions due to the diverse response of the degenerate components to those conditions [Whitacre and Bender, 2010]. However, intraindividual degeneracy also provides a diverse set of phenotypic structures that can be repurposed or elaborated upon. Essentially, intraindividual degeneracy protects against loss of function mutations, which are often deleterious, while providing ample jumping off points to generate novel phenotypic function. In this way, intraindividual degeneracy promotes individual evolvability.

Intraindividual degeneracy is also thought to be related to direct plasticity. Under different circumstances — such as unusual environmental stress — the functionality of degenerate components might diverge due to their structural differences [Richter et al., 2015]. The presence of degenerate components, each of which might be able to maintain functionality under environmental stresses that would neutralize the other components, increases the range of environmental perturbations that a organism might be able to withstand.

4.6.3 Example

Mammalian deoxyribonucleoside kinases serve as a prime example of intraindividual degeneracy. Mammalian deoxyribonucleoside kinases are enzymes that process the precursors to DNA. As shown in Figure 4.2, these enzymes act on four substrates: dT, dC, dA, and DG. These kinases are said exemplify degeneracy because each unit of functionality, the chemical processing of each of the four precursors, is accomplished by two structurally unique components. Processing of the compound dT, in particular, is accomplished via both TK1 and TK2, enzymes from different protein families that fundamentally differ in structural composition. The ability of an organism to metabolize the dT DNA precursor is thus more robust to perturbations that disable one of Tk1 or Tk2. Additionally, Tk1 and Tk2 both serve as jumping-off points for evolutionary innovation via the repurposing of either enzyme.

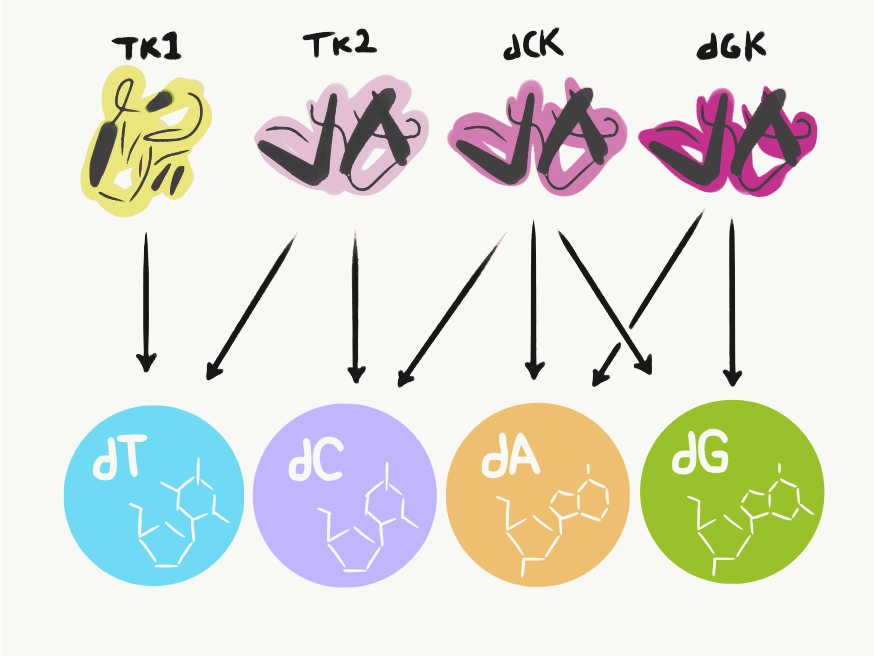

Figure 4.2: Mammalian deoxyribonucleoside kinases exhibit degeneracy [Sandrini and Piskur, 2005].

4.7 Interindividual Degeneracy

4.7.1 Definition

I define interindividual degeneracy as the presence of phenotypically dissimilar individuals in a population that, despite those dissimilarities, interact with their environment in identical or near-identical ways (i.e. are functionally equivalent). Interindividual degeneracy is related to the idea of diversity but specifically refers to phenotypic diversity that is functionally interchangeable. Although structurally divergent, individuals in a population that exhibits interindividual degeneracy are functionally similar.

Interindividual degeneracy can be seen as stemming from tolerance of the phenotype-fitness mapping for certain types of phenotypic variation. That is, certain aspects of phenotypic form are inconsequential to fitness.

Interindividual degeneracy can also be seen as stemming from robustness. Robustness enables a biological system to maintain function despite variations to some aspect of phenotypic form. (Without robustness, overall function of an organism would be more likely to be negatively impacted by variation in a particular aspect of phenotypic form).

Although distinct from intraindividual degeneracy, intraindividual degeneracy might bolster interindividual degeneracy by providing bountiful opportunities for neutral variation to manifest.

4.7.2 Relation to Evolvability

Similar to how intraindividual degeneracy promotes individual evolvability, population degeneracy promotes population evolvability. Succinctly put, by providing a wider array of jumping-off points, the presence of

(a) attached earlobe

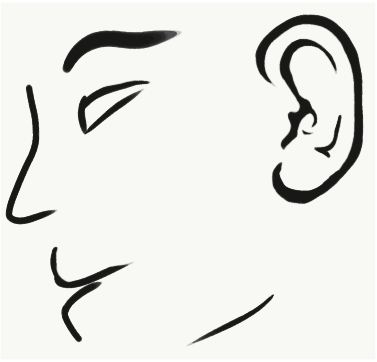

(b) detached earlobe

Figure 4.3: An illustration of alternate earlobe phenotypes in *Homo sapiens sapiens*. Although divergent, these alternate phenotypes do not impact phenotypic functionality. Thus, earlobe types are an instance of inter-individual degeneracy.

structurally divergent individuals in a population increases the variety of phenotypic forms that can be realized in the offspring of that population.

4.7.3 Example

Earlobe attachment in *Homo sapiens sapiens* provides an illustrative example of interindividual degeneracy. This phenotypic trait is distributed over a spectrum between attachment and detachment. The extremes of this spectrum are illustrated in Figure 4.3. This trait is widely understood to be genetically determined [Dutta, 1979]. This element of phenotypic diversity, which occurs in both forms within populations around the world, seems to be functionally neutral. That is, the trait does not affect how an individual interacts with his or her environment to determine his or her fitness. Hence, diversity in earlobe attachment represents interindividual degeneracy.

4.8 Regularity

4.8.1 Definition

Informally, regularity can be used to describe repetition phenotypic form. Repetitive form might manifest as symmetry and/or recurring modular substructures. Formally, regularity refers to how much information is required to describe a structure [Clune et al., 2011].

4.8.2 Relation to Evolvability

A bias towards regularity in phenotypic form tends to represent a bias towards viable variation. That is, in many situations, regular phenotypes tend to outperform highly irregular phenotypes. This conclusion that regular phenotypes tend to be more viable, of course, depends on the demands of the environment that the phenotype inhabits. Phenotypic regularity tends to be useful in more regular environments (that is, environments that exhibit regular characteristics) [Clune et al., 2011]. Many problem domains of interest to EANN researchers are highly regular [Clune et al., 2011]. The natural world, the realm of flesh-and-blood

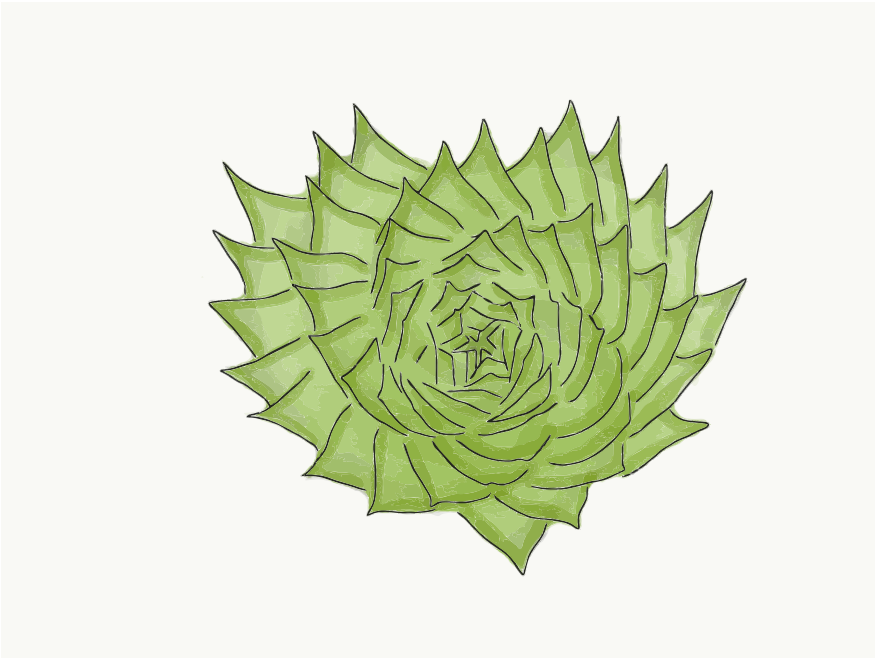

Figure 4.4: An illustration of Phenotypic Regularity in *Aloe polyphylla*. Notice the outward spiraling pattern formed by the leaves.

creatures, also exhibits significant regularity [Downing, 2015, pg 161]. Encodings biased towards regularity can also be thought of as promoting structures that were not directly selected for during evolution. These incidental phenotypic structures are called "spandrels," referencing an architectural phenomenon that arises in a similarly unplanned for way: the awkward space between a curved archway set in a rectangular wall. These regular structures might be more likely to perform well in fitness cases that were not tested. It has been found that a tendency for regularity in evolving artificial neural networks produces networks that are capable of more generalized learning. That is, EANN that are more likely to successfully perform tasks beyond those they were explicitly tested on during evolution [Tonelli and Mouret, 2013]. Viewing learning as post-developmental (plastic) irregular refinement of regular neural structures created by development, a connection between regularity, irregular refinement (or plastic complexification), and plasticity is apparent in promoting phenotypic fitness. In these ways, bias towards regularity increases evolvability because a higher proportion of individuals are viable.

4.8.3 Example

Aloe polyphylla, which is known for the striking spiral arrangement of its leaves (Figure 4.4), exemplifies regularity in nature [Royal Horticultural Society,]. The phenotypic regularity exhibited by *Aloe polyphylla* is an important adaptation. Phyllotaxis, the regular arrangement of leaves in plants, by minimizing the conflict between leaves for light, promotes efficient photosynthesis [Kappraff, 2004].

Chapter 5

Ultimate Causes of Evolvability

Establishing the ultimate cause of a phenomenon can easily devolve into a fraught and ill-defined endeavor. A bicycle rider noticing decreasing return on her pedaling effort would very well wonder what is wrong with the bike. Perhaps the tire is flat. Why is the tire flat? Perhaps a shard of glass punctured it. Why? Perhaps the biker's commute goes through a derelict neighborhood. Why? The answer to that question lies somewhere in urban history and zoning ordinance. This process could be continued. Often, there is no clearly defined top rung on the ladder of causality. Instead, causal hypotheses become increasingly obscure, eventually escaping the realm of science and slipping into that of philosophy or religion. Practical considerations can help to resolve this dilemma by informing a clear point at which to stop climbing and declare an ultimate cause. In terms of our biker, the route of her commute provides a clear avenue for intervention. Perhaps she should avoid biking through that seedy district.

Similar practical considerations inform organization of this section. High level causal factors that might be practically employed to promote evolvability in evolutionary computing were included. These causal factors frequently act by promoting the traits discussed in Chapter 4 that manifest in the proximal causes, phenotypic characteristics observed in biological organisms, presented in Section 3. While the contents of this section might hold the greatest direct practical relevance to researchers seeking to promote evolvability in artificial systems, it should not be viewed as a comprehensive listing of approaches to promoting evolvability in those systems. As will be overviewed in Chapter 6, some techniques devised by evolutionary computing researchers do not have a direct analog in nature, although many are inspired by the principles of nature's evolvability-promoting toolbox. Hopefully, this vantage on the roots evolvability, although inspired by the practical considerations of evolutionary computing, by considering how certain fundamental characteristics of biology promote evolvability will also prove satisfying to the biologically-motivated reader.

5.1 Indirect Encoding

5.1.1 Definition

The phrase indirect encoding describes the relationship between the genotype and the phenotype. In an indirect encoding, a one-to-one direct correspondence is not guaranteed between each phenotypic characteristic and a single entry in the genotype. The opposite of an indirect encoding is a direct encoding. In a direct encoding, each and every phenotypic characteristic is independently described by a single entry in the genotype.

Indirect encodings can be broken down into two categories: expanded and generative. In an expanded indirect encoding, a one-to-one relationship exists between information in the genotype and in the phenotype, but the encoding lacks independent control of each phenotypic characteristic by a single genotypic entry. That is, altering one piece of genetic information can affect multiple phenotypic characteristics. In a generative indirect encoding, the one-to-one relationship between phenotypic characteristics and genetic information is relaxed. Such an indirect encoding is deemed generative because, typically, a large number of phenotypic characteristics are generated from a smaller amount of genetic information via a developmental process [Downing, 2015, p 175].

5.1.2 Relation to Evolvability

Generative indirect encodings are generally biased towards phenotypic regularity (Section 4.8) because phenotypic information is generated from a smaller amount of genetic information. Developmental processes may allow for genetic information to be reused to describe different characteristics of the phenotype in a systematic manner [Clune et al., 2011]. Thus, phenotypic patterns — e.g. regularity — tend to be observed. Experiments by [Cheney et al., 2013] with soft bodied robots illustrate this point elegantly. Virtual soft-bodied robots evolved for locomotion using an indirect encoding based on Compositional Pattern Producing Networks display greater regularity than their directly encoded peers, manifesting in the robots as large contiguous patches of identical tissue type. Figure 5.1, which compares representative direct encoded and indirect encoded champions, illustrates the impact of indirect encodings on regularity.

5.1.3 Example

In biology, phenotypic information is indirectly encoded in DNA. The genotype of an organism is translated into its phenotype via the developmental process. The indirect relationship between phenotypic and genotypic information is universal in biology; heritable information in the genotype is translated into phenotypic characteristics through the production of proteins and other gene products.

Figure 5.1.3 provides a cartoon example of this indirect encoding in biology, comparing the genotype and

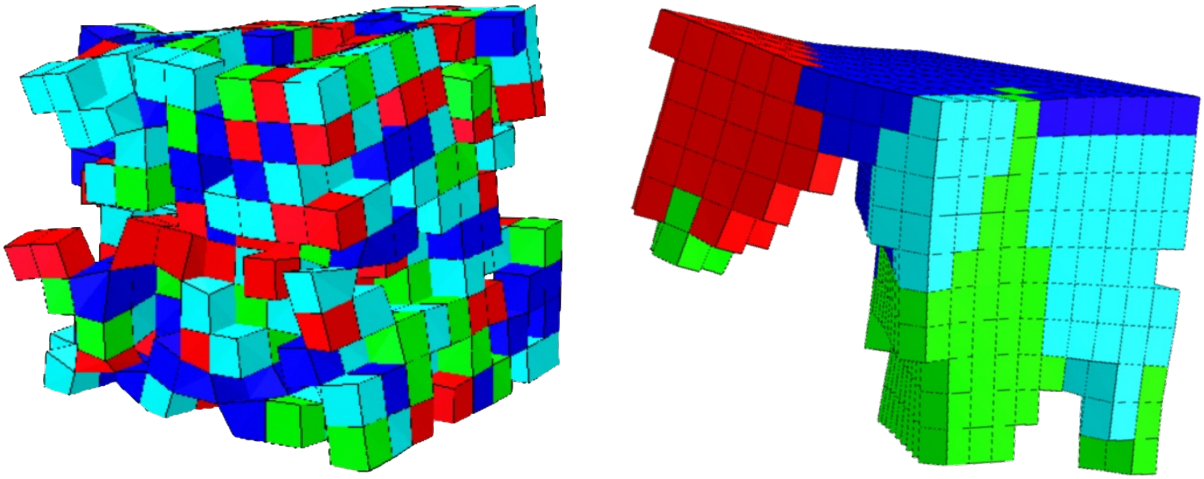

(a) direct encoding (low regularity)

(b) indirect encoding (high regularity)

Figure 5.1: Representative examples of soft robots evolved with direct and indirect representations [Cheney et al., 2013, Figures 6, 7]

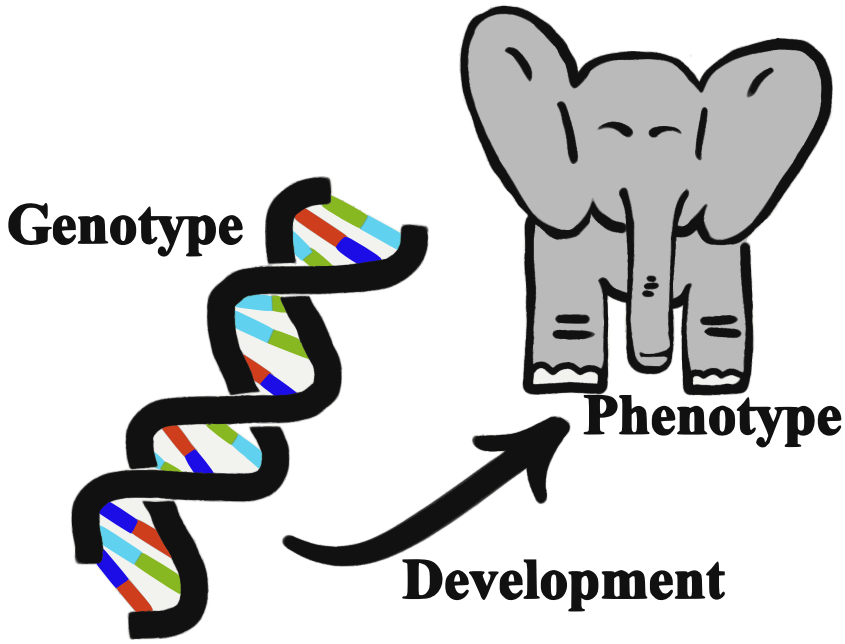

Figure 5.2: This cartoon illustrates principle of indirect encoding by juxtaposing the genotype and phenotype of the elephant. In the case of an elephant, the genotype is comprised of approximately 5×10^9 base pairs and the phenotype is comprised of approximately 1×10^{15} cells. Note how the much more information is required to describe the phenotype than the genotype [Kim, 2011, Hofreiter, 2008].

phenotype of an elephant.

Quantifying the amount of phenotypic information in an elephant is a nontrivial task. However, the claim that a cell constituting the elephant contains more information than a single base pair should be noncontroversial. As the number of cells constituting an elephant far outstrip the number of base pairs in its genome, the elephant also nicely illustrates the disequilibrium between phenotypic characteristics and genotypic information enabled by indirect encoding. As would be expected, regularity is observed on many phenotypic aspects of the elephant, from its bilateral symmetry to the repeated occurrence of many near-identical proteins and cellular structures in each of the trillions of cells throughout the elephant. [Clune et al., 2011] offer a similar comparison, perhaps a little closer to home, pointing out that in humans 25 000 genes describe a phenotype consisting of trillions of cells.

These two examples comparing genetic information to phenotypic characteristics in massively multicellular creatures might seem overly convenient. Unicellular creatures, for example, do not afford such a stark observation. Like elephants, quantifying the number of phenotypic characteristics of a protozoa or a bacteria is a nontrivial task. Despite their minuscule stature, these creatures do exhibit many phenotypic characteristics. However, a convincing argument that the amount of phenotypic information exceeds the amount of genetic information would be difficult to make. In any case, these creatures, which rely upon the same fundamental translation of genotypic information to phenotypic characteristics through gene products, nevertheless also lack a direct one-to-one correspondence between genetic information and phenotypic characteristics and, thus, are indirectly encoded. Like their multicellular counterparts, unicellular creatures exhibit phenotypic regularity. Such regularity might manifest in, for example, repeated occurrence of identical or near-identical functional subunits (i.e. many identical proteins, ribozymes, etc.).

5.2 Temporally Varying Goals

5.2.1 Definition

Most biological organisms do not evolve in a static environment. Instead, their environment changes over evolutionary time. These changes might be due to geophysical factors, such as changes in the climate or the geological composition of an area. These changes might also be due to biological factors, through evolutionary changes in the organisms with which they cooperate, compete, or otherwise interact. Thus the fitness-related goals of these organisms — that is, the specifications that they must meet in order to succeed — are temporally varying.

5.2.2 Relation to Evolvability

Evolutionary simulations have revealed that gradual changes to the environment — that is, a temporally varying fitness function — promote desirable characteristics related to evolvability including robustness, modularity, and individual evolvability [Kashtan et al., 2007, Wilder and Stanley, 2015]. It is thought that a gradually shifting fitness function might induce evolutionary pressure for these traits, essentially providing a means of selecting for them.

Robustness would be essential in a temporally varying goals scheme, allowing to persist under environmental conditions different from those of their ancestors. Similarly, modularity, which allows for the repurposing of evolved units of functionality and reduces the impact of changes in one aspect of functionality on other aspects of functionality, would also prove useful in responding to evolutionary pressure for constant, gradual adaptation to changing environmental conditions. Kashtan et al. point out that modularity is a particularly important trait because, throughout changes in environmental conditions, the necessity for many aspects of functionality is preserved.

> On the level of the organism, for example, the same subgoals, such as feeding, mating, and moving, must be fulfilled in each new environment but with different nuances and combinations. On the level of cells, the same subgoals such as adhesion and signaling must be fulfilled in each tissue type but with different input and output signals. On the level of proteins, the same subgoals, such as enzymatic activity, binding to other proteins, regulatory input domains, etc., are shared by many proteins but with different combinations in each case [Kashtan et al., 2007].

In fact, phenotypic modularity has been observed to emerge spontaneously in artificial evolution experiments performed with a temporally varying fitness function [Kashtan et al., 2007].

Finally, under a temporally varying goals regimen, individuals that are predisposed to yielding variable offspring are advantaged over individuals that do not. Although much of that variation is likely not to be useful, some of it is likely to be and will allow to exert dominance over individuals without fresh adaptation to the changing environmental conditions. As Wilder, et al. put it, "if selection sets a moving target, individuals will be more likely to introduce variation in their offspring to adapt to an uncertain future; mutations to the genotype will be more likely to result in phenotypic change" [Wilder and Stanley, 2015]. By inducing a selective pressure for individuals capable of generating relatively swift adaptive change to track changing environmental conditions, temporally varying goals promote a number of essential traits related to evolvability.

(a) low individual evolvability (b) high individual evolvability

Figure 5.3: A hypothetical illustration of how individual evolvability might be selected for under a modularly varying fitness function [Kashtan and Alon, 2005].

5.2.3 Example

Let us discuss a hypothetical population of hummingbirds that feed on purple flowers. Suppose that in order to successfully feed, hummingbird beak lengths must match the length of their food source — beaks can be neither too short nor too long or the hummingbird will be unable to feed effectively. Suppose also that the lengths of the purple flowers on which the hummingbirds depend were to be systematically manipulated over evolutionary time, proceeding through cycles of gradual increase and decrease. Under this scheme, individuals whose offspring exhibit greater variability in beak length would be favored. Figure 5.3 contrasts the offspring of an individual that exhibits low individual evolvability in relation to beak length and an individual that exhibits higher individual evolvability in relation to beak length. A greater diversity of beak lengths are observed among the offspring of the hummingbird with high evolvability. Although many beak length outcomes among offspring of a hummingbird that exhibits high individual evolvability will be deleterious, some fraction will be adaptive to changes in purple flower length. If selection is strong enough, the offspring will enjoy high fitness compared to other members of its generation. Hence, individual evolvability will be selected for.

As discussed in the previous subsection, temporally varying goals are thought to promote a slew of causal factors related to evolvability. In the context of our hummingbird example, one might expect to observe

- increased variability in beak length between siblings (i.e. individual evolvability),
- increased segregation of developmental processes that determine beak length from other developmental processes (i.e. modularity), and/or
- greater ability of the hummingbird to persist with limited nutritional resources (i.e. robustness)

among hummingbirds evolved in a temporally varying environment compared to a population of hummingbirds that exist in a static environment.

5.3 Environmental Influence on Phenotype

5.3.1 Definition

The mantra $P = G + E$, stating that phenotype is a product of both genotype and environment, has come to be universally accepted among biologists. On a basic level, it is clear that, because biological organisms do not exist in a vacuum, they are physically influenced by their environments. Temperature, chemical substances, electromagnetic radiation, and other external physical phenomena directly act upon biological organisms.

5.3.2 Relation to Evolvability

Environmental influence on the phenotype is included as an ultimate cause of evolvability because of an idea that is a bit more subtle: that these environmental factors might play an important role in the evolutionary process. Environmental influence on the phenotype can be divided into two major categories: disruptive noise and information. That is, certain aspects of the physical interaction of an organism with its environment can be considered as disruptive perturbations that an organism must be robust to. This type of environmental influence is related to direct plasticity, an intermediate causal factor presented in 4.4. Other aspects of the physical interaction of an organism with its environment might be considered as providing information to the organism that it can exploit in order to increase its fitness. This type of environmental influence is related to indirect plasticity, another intermediate causal factor presented in 4.5.

5.3.3 Example

Himalayan rabbits provide a charismatic example of environmental influence on the phenotype. These rabbits carry a gene that produces melanin, giving their fur a brown color. However, this gene is only active at temperatures significantly below body temperature so the brown coloration is typically only observed on the outer extremities of the rabbit. If a Himalayan rabbit is reared at a temperature above the threshold for melanin creation, however, the gene will be inactivated throughout the rabbits body and a pure white phenotype will be observed [Lobo, 2008].

5.4 Fitness Degeneracy

5.4.1 Definition

I define fitness degeneracy as the ability of divergent strategies for reproduction and survival to succeed in a particular environment. From an evolutionary computing perspective, a fitness function would be considered to exhibit fitness degeneracy if a set of functionally different phenotypes can all attain high fitness. This

concept roughly corresponds to the idea of biological niches, but considers the ability of different strategies for reproduction and survival to succeed within a population (as opposed to between species).

5.4.2 Relation to Evolvability

By allowing for functionally divergent phenotypes to succeed, fitness degeneracy promotes the accumulation of diversity within a population. This diversity broadens the range of phenotypic forms that are accessible to the offspring of a population, thus promoting population evolvability . The coexistence of speciated populations within the same ecosystem that exploit radically different strategies to earn their metabolic keep can also be seen as fitness degeneracy. This degeneracy, the ability of radically different fitness strategies to coexist within the biosphere, has vastly expanded the breadth of phenotypic form explored by evolution.

> A related approach is to set a vast number of separate, diverse evolutionary objectives, which approximates the infinite number of niches created by the ever-changing natural world. Such optimization leads to frequent goal-switching, as lineages fit on one objective invade other objectives, which rewards lineages that produce behaviorally diverse offspring and increases evolvability [Mengistu et al., 2016]

5.4.3 Example

The male mating patterns of *Uta stansburiana*, a lizard found in the coast range of California, illustrates fitness degeneracy — how different strategies stemming from phenotypic diversity can be successful in terms of fitness. Male *Uta stansburiana* are found in three morphs: orange, blue, and yellow [Sinervo and Lively, 1996]. The orange morph is characterized as "hyper-masculine." It staunchly defends its territory and will battle male trespassers. Lizards of the yellow morph are "sneakers." They pass themselves off as females to get past other males. The blue morph is described as the "mate-guarding" morph. Blue morph lizards are more circumspect to combat and not as easily deceived by interloping yellow morphs. Figure 5.4 provides a visual overview of how these morphs interact. Much like the game of rock, paper, scissors, each morph outcompetes one morph and is itself outcompeted by another morph. No single phenotypic strategy is favored by the environment. Although their strategies for survival and reproduction vary greatly, each lizard morph can be successful.

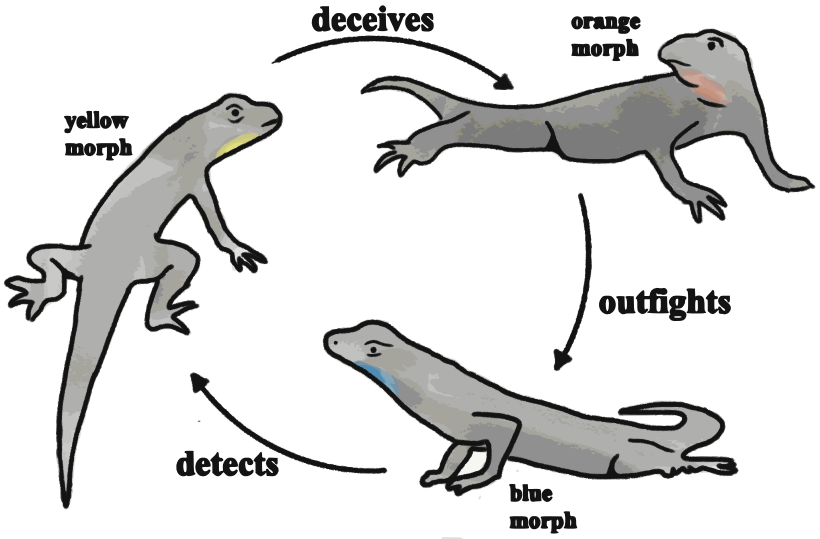

Figure 5.4: The three male morphs of *Uta stansburiana*, which exhibit different throat coloration and mating strategies, are shown to illustrate fitness degeneracy. Arrows indicate outcompetition for mates [Sinervo and Lively, 1996].

Chapter 6

Discussion

In Chapters 3, 4, and 5 we have toured through a sprawling menagerie of causal factors associated with evolvability. These factors included aspects of internal structures of biological organisms (e.g. modularity), the processes that give rise to the phenotypic structures observed in mature biological organisms (e.g. exploratory growth), the constitution of the environments biological organisms inhabit (e.g. fitness degeneracy), interplay between the structure of biological organisms and the environment (e.g. plasticity), and specific evolutionary processes through which complexity is thought to emerge (e.g. duplication and divergence). This expedition traveled through a wide range of explanatory scope — we discussed causal factors related to the immediate physical forms of biological organisms (Chapter 3) and causal factors related to overarching design patterns associated with biological organisms (Chapter 4) Finally, Chapter 5 reviewed causal factors rooted in the fundamental characteristics of organisms and their environments. Despite the broad nature of the survey presented in Chapters 3, 4, and 5, it does not constitute an exhaustive presentation of factors related to evolvabiltiy. However, this survey is wide-ranging enough to begin to appreciate broader patterns among the factors that contribute to evolvability. This background should also be sufficient to shed light on strategies to promote evolvability practiced in evolutionary computing.

This chapter aims to leverage overarching theoretical perspective on evolvability from the proximate-intermediate-ultimate organizational scheme to describe strategies to promote evolvability in an evolutionary computing setting. Selected efforts to promote evolvability from the evolutionary computing literature will be reported to illustrate the strategies described.

As we have seen in Chapters 3, 4, and 5, evolvability is a diffuse concept with multiple causality. Recall that at the most fundamental level, as introduced in Chapter 2, evolvability stems from the availability of heritable variation that is viable. Recall, also, the two major elements at play in relation to evolvability: an ability to readily produce *novel* heritable phenotypic variation and a bias towards *viable* variation. Chasing evolvability further down the rabbit hole, the chain of causality rapidly branches out to a diverse set of

interconnected factors; evolvability stems from a dense web of interactions between many concepts. The causal factors related to evolvability do not readily lend themselves to being cleanly teased apart. However, the proximate-intermediate-ultimate conceptual framework does provide some useful insight into the overall causal framework behind evolvability. We will consider separately the two major components of an evolutionary algorithm: development (Section 6.1) and selection mapping (Section 6.2). (Development refers to the genotype-phenotype mapping, and selection refers to the phenotype-fitness mapping; the terms development and genotype-phenotype mapping and the terms selection and phenotype-fitness mapping will be employed interchangeably). The lens of the proximate-intermediate-ultimate conceptual framework will help us distinguish between biologically-inspired and artificial strategies to promote evolvability and better understand both. We will see that biologically-inspired strategies to promote evolvability act at the ultimate level of causality while artificial strategies instead act at the proximate level, in relation to the genotype-phenotype mapping and at the intermediate level, in relation to the phenotype-fitness mapping. At the end of the discussion in Section 6.3, we will compare biologically-motivated and artificial strategies to promote evolvability and consider how they reflect broader themes in the field of evolutionary computing.

6.1 Development

6.1.1 Theoretical Analysis: Biologically-Motivated Strategies

Discussing genotype-phenotype mapping, it is important to recognize that the biological development process itself is largely defined in the genotype and is therefore subject to evolution. Thus, many aspects of the genotype-phenotype mapping manifest through the evolutionary process. For example, direct and indirect plasticity are evolved capacities to withstand and exploit environmental influence on the phenotype. In the context of evolutionary computing, these characteristics of the genotype-phenotype mapping would not be expected to be present in arbitrarily generated starter individuals. In our proximate-intermediate-ultimate scheme, these characteristics of the genotype-phenotype mapping would fall under proximate and intermediate causes of evolvability. In other words, the proximate and intermediate causal factors described in Chapters 3 and 4 manifest in biological organisms as a result of the evolutionary process, not as a precursor to it.

In contrast, aspects of the genotype-phenotype mapping that fall under the umbrella of ultimate causal factors related to evolvability are not encoded in the genotype. They can instead be seen as implicit to the model through which the genotype is interpreted. The genotype-phenotype map's indirect nature is an explicit assumption as is environmental influence on the phenotype. A maximally biologically-plausible model would explicitly account solely for these ultimate causal factors to promote evolvability. The other aspects of the genotype-phenotype map would be expected to emerge as a result of the evolutionary process

under the correct conditions.

6.1.2 Theoretical Analysis: Artificial Strategies

To begin to understand the role proximate factors play in relation to other causal factors, let us begin by developing an understanding of how these proximate factors fit together as a group. A cursory inspection reveals that this grouping of proximate factors falls into two major categories. Duplication and divergence, complexification, weak linkage, and the accumulation of hidden genetic variation are mechanisms that act on a population over evolutionary time (i.e. over the course of multiple generations). New phenotypic traits emerge from refinement of existing traits, duplication and modification of existing traits, the establishment of novel interaction between subsystems through enabling (rather than instructive) signaling, and sensitization to previously hidden accumulated genetic variation. Exploratory growth and developmental constraint are mechanisms that act on the phenotype during the lifetime of an individual (i.e. over the course of the developmental process). Components of the organism exhibit adaptivity during the developmental process, reacting to the state of other components and the developmental environment so otherwise independent phenotypic traits are synchronized. The nature of the developmental process governs how genotypic information manifests in the phenotype.

Although the collection of concepts deemed proximate contributors to evolvability are highly diverse, they might succinctly be described as patterns of how phenotypic features emerge through the developmental and evolutionary processes. All of these concepts ultimately boil down to the genotype-phenotype mapping, the manner in which phenotypic information is stored in the genome. The intimate intertwining of the genotype-phenotype mapping and both exploratory growth and developmental constraint, which act over developmental time, can be readily appreciated. The connection between the proximal traits that act over evolutionary time and the genotype-phenotype mapping is a bit subtler. All four of these proximate contributors to evolvability — duplication and divergence, complexification, weak linkage, and hidden genetic variation — are related to the addition of new phenotypic information to the genome. Duplication and divergence leverages existing information in the genome to jump start phenotypic innovation. Complexification is enabled by the ability of the genome to assimilate new information describing a phenotypic trait in increasing detail. Weak linkage is a phenotypic adaptation — propensity to use enabling over instructive signaling — that reduces the amount of information that must be added to the genome to establish regulatory links between separate systems. Hiding segments of the genome from phenotypic expression allows for the accumulation of novel genetic information in a population that can then be exposed through sensitizing mutation. The manner in which phenotypic information is stored in the genotype is inherently related to the genotype-phenotype mapping, hence the connection between these concepts and that mapping. The Baldwin Effect sits somewhat apart from other proximate causal factors related to evolvability. Although it

influences the genotype-phenotype mapping like other proximate causal factors, the Baldwin Effect is also tied to the phenotype-fitness mapping as it essentially performs local search in phenotypic space for fitness peaks. For the most part, with this slight exception, proximate causes of evolvability can be understood as acting on the genotype-phenotype mapping.

As we will see in the next section, manually-designed developmental processes can exploit this relationship between proximate causes of evolvability and the genotype-phenotype mapping. These proximate causes of evolvability can be explicitly accounted for in the artificial development process ("scaffolded" into the model).

6.1.3 Examples from Literature

Genetic representation design in evolutionary computing is typically a domain-specific endeavor. The manner in which a genotype-phenotype mapping is designed depends on the phenotype space it must map to. However, some design patterns do exist across domains and these are nicely exemplified by work being done in Evolving Artificial Neural Networks (EANNs). The goal of EANNs are to discover configurations of nodes and connective weights that cause an artificial neural network to exhibit a desired behavior. Scalability is a key hurdle for EANNs; evolving large artificial neural networks is difficult. EANNs have yet to rival neither the scale nor the intricacy of their biological counterparts and it is thought that poor genetic encoding is responsible for much of this shortcoming [Tonelli and Mouret, 2013].

Perhaps not surprisingly then, researchers have turned to biology to inspire new genetic encoding schemes for EANNs. On the more literal side of the biological inspiration coin lies Bongard and Pfeifer's artificial ontogeny (AO) system. In this system, the phenotype is constructed through successive duplication and differentiation of a single virtual cell modulated by simulated diffusion of chemical substances, which in turn stem from simulated genetic regulatory networks [Downing, 2015, p 345]. Mouret et al.'s map-based encoding occupies a middle ground in terms of literal interpretation of biological inspiration. Styled after models of the Basal Ganglia developed by neuroscientists, the map-based model jettisons the individual neuron as the fundamental unit of network composition. Instead, layers each consisting of an arbitrary number of neurons, which are strung together by several pre-defined regular connection schemes (i.e. one to one or one to all), are taken as the fundamental unit of network composition [Mouret et al., 2010].

HyperNEAT, a highly influential and ubiquitous encoding in evolutionary computing, lies on the reverse side of the biological inspiration coin — it exhibits highly abstracted biological inspiration. HyperNEAT is built off the NEAT encoding scheme developed by Stanley and Miikkulainen [Downing, 2015, p 324]. NEAT is a direct, variable-length encoding for neural networks that reduces omissions and redundancies introduced during crossover by tracking historical markers associated with each gene, in essence implementing a rough approximation of synapsis during meiosis. HyperNEAT repurposes NEAT to encode a Compositional Pattern

Producing Network (CPPN) instead of encoding a neural network.[1] The weight of the connection between two nodes in the neural network generated by HyperNEAT is determined by the output of the CPPN fed the coordinates of those two nodes. The CPPN in essence serves as the genotype in the HyperNEAT scheme. The activation patterns projected onto the substrate occupied by artificial neurons in the HyperNEAT scheme by the CPPN are analogous to patterns of chemical diffusion that play a key role in directing embryological development [Downing, 2015, p 339]. In the original HyperNEAT concept, artificial neurons were arranged in a simple grid pattern. Subsequent efforts have been made to implicitly define neuron placement based on CPPN output, allowing for the network to grow in numerical size over evolutionary time [Risi et al., 2010, Risi and Stanley, 2010].

Although great strides have been made in developing representational techniques, allowing effective EANNs of much larger scale than before, a lot of ground remains to be covered. EANNs still fall short of the scale exhibited by biological neural networks.

6.2 Selection

6.2.1 Theoretical Analysis: Biologically-Motivated Strategies

Under the proximate-intermediate-ultimate organizational framework, identifying biologically-motivated strategies to promote evolvability by acting on the phenotype-fitness mapping is straightforward. Any ultimate causal factors related to selection can serve as the basis for such a biologically-motivated strategy. Temporally varying goals and fitness degeneracy explicitly describe aspects of the fitness function — its fluctuation over evolutionary time and its tolerance of phenotypic diversity. Although environmental influence on the phenotype acts on the genotype-phenotype mapping it can also be seen as related to the phenotype-fitness mapping as it can relay information about the phenotype-fitness mapping — namely, which phenotypic outcome would be favored in terms of fitness.

6.2.2 Theoretical Analysis: Artificial Strategies

Although ultimate causal factors related to selection provide biologically plausible strategies to promote evolvability, intermediate causes of evolvability can provide means to promote evolvability that depart from strict adherence to the biological metaphor. Intermediate causal factors include modularity, canalization, robustness, individual evolvability, intraindividual degeneracy, and interindividual degeneracy. In nature, these intermediate contributors to evolvability are also qualities gained by evolving organisms through the evolutionary process, not built a priori into an evolving system. In concrete terms, this means that these

[1] A CPPN is a weighted network of nodes, each of which performs a transformation on its inputs to generate an output — essentially a neural network where nodes may perform transformations other than the logistic equation. The CPPN at its core is just a mathematical expression that accepts several inputs and generates several outputs.

traits would be essentially absent in a randomly generated population at generation zero but might develop as evolution proceeds under the right conditions. Intermediate causal factors tend to be less prone to direct manual instantiation compared to their proximate cousins because intermediate causal factors are largely not explicit components of the developmental process. However, that does not mean that intermediate causal factors are ill-defined or unobservable. Especially in digital organisms, these traits can be directly quantified (usually in the context of a single individual). Reisinger et al. quantify canalization by measuring the rate at which random mutation degrades fitness [Reisinger and Miikkulainen, 2007]. A domain-specific metric for modularity has been developed by [Kashtan and Alon, 2005]. The ability to quantify these traits associated with intermediate causality allows direct selection for them.[2] That is, individuals that exhibit a desirable trait related to evolvability, such as modularity or intraindividual degeneracy, could be explicitly chosen to populate the next generation. The hope would be to observe these traits become more prevalent over evolutionary time.

6.2.3 Examples from Literature

In biological evolution, selection criteria are neither monolithic nor static. Although single-objective, static selection is perhaps the most intuitive strategy to use in an evolutionary algorithm at first blush, such a selection scheme constitutes a rather simplistic interpretation of the biological metaphor. A number of recent efforts have leveraged inspiration from selection in biological evolution to great effect. Researchers have ventured away from static, objective-based selection by employing temporally varying goals, selection criteria that change over evolutionary time. This selection paradigm in evolutionary computing, reviewed in Section 5.2. Inspired by the abundance of niches in biological environments, Ngyuen et al. performed experiments considering the possibility of evolving towards several objectives at once. They found that, compared to evolving towards a single objective, defining a large set of fitness criteria and determining selection based on the maximum fitness score obtained under one of those criteria, yielded better quality solutions for each objective [Nguyen et al., 2015]. In their computational experiment, Ngyuen et al. observed objective switching, where offspring are better suited to a niche distinct from that occupied by their parent, occur at a significant rate. Mengistu et al. have hypothesized that this niche-inspired approach encourages individual evolvability by rewarding individuals with offspring that are phenotypically variable enough to jump between objectives [Mengistu et al., 2016].

As suggested in the analysis performed at the open of this section, researchers have also identified several successful selection strategies that depart from strict adherence to the biological metaphor. Highly notable among these is novelty search. This selection scheme discards traditional objective-based metrics in favor

[2]It should be noted that intraindividual degeneracy, which manifests in differences between individuals rather than within a single individual, is slightly different from the other listed intermediate contributors in this regard. However, selection to directly promote intraindividual degeneracy might nonetheless be possible by divergent selection to reward novel neutral phenotypic and/or genotypic variation.

of selecting for individuals that are the most novel, that are the most distinct from individuals that have already been encountered during evolutionary search [Lehman and Stanley, 2008]. Although not directly consistent with the biological metaphor, it has been suggested that such an approach captures aspects of the open-ended nature of biological evolution in ways that other selection schemes fail to do. This strategy occupies a gray area between biologically plausibility and artificiality. Surprisingly enough, it has been found that in certain scenarios jettisoning selection for an objective in favor of searching for novelty (and keeping track of the best-performing solutions encountered along the way) can actually lead to better performance in satisfying that objective [Lehman and Stanley, 2008]. Novelty search has been found to promote evolvability: selecting for novelty rewards individuals that create variable offspring [Lehman and Stanley, 2011].

More recently, Mengistu et al. demonstrated a scheme where individual evolvability is directly selected for [Mengistu et al., 2016]. This scheme selects on the characteristics of an individual — in this case, the ability to generate phenotypic variation among offspring). Thus, it exemplifies the avenue to promote evolvability provided by direct selection for intermediate causes of evolvability. Like novelty search, this selection scheme ignores objective performance when performing selection. Instead, individuals that create a set of offspring that exhibit the greatest phenotypic diversity amongst siblings are selected for. Also like novelty search, the best-performing individuals according to the objective function are simply logged along the way. Although subtle, it is an important point to that these two objective-bucking strategies are not random search. Individuals are instead selected for based upon criteria — either individual evolvability or novelty — that specifically aim to facilitate wide-ranging search. It is in the course of these wide-ranging journeys through phenotype space that these search strategies encounter solutions that well satisfy the objective.

6.3 Synthesis

Our analysis of strategies to promote evolvability has been divided between the two major components of the evolutionary algorithm: the genotype-phenotype mapping (development) and the phenotype-fitness mapping (selection). The proximate-intermediate-ultimate framework has shed light on strategies targeting both mappings, in particular how each can be accomplished both inside and outside of strict adherence to the biological metaphor. Possible strategies to promote evolvability can be summarized as:

- taking a broader view of selection
  - more open-ended environments that exhibit fitness degeneracy and temporal fluctuations (ultimate causes of evolvability)
  - non-objective-based selection mechanisms (such as selection for intermediate causes of evolvability)

- adopting a more nuanced view of the developmental process
  - allowing for a more biologically plausible developmental process that incorporates environmental influence on the phenotype and allows a high level of genetic control of the developmental process (relying on ultimate causes of evolvability)
  - manually building scaffolding for the developmental process to incorporate aspects of proximate contributors to evolvability at a feasible computational cost (scaffolding proximate causes of evolvability)

Since the conception of the field of evolutionary algorithm design efforts under each of these banners have already had transformative effects on the field. Such efforts are reviewed next in Sections 6.1 and 6.2.

At first blush, it very well might seem that these efforts to apply a broader, more nuanced view of evolution to the evolutionary algorithm amount to abandoning practical applications of the evolutionary algorithm. After all, these first components of the path forward for evolutionary computing entail moving away from exclusively selecting for the ability of a digital organism to fulfill a desired function. The second involves incorporation of more nuanced developmental processes, efforts that might not seem to directly relate to the ability of a mature phenotype to serve a desired function and instead merely appear to be an extravagant pretense to biological realism where "the complexity of the model dwarfs that of the task..." [Downing, 2015, p 354]. However, the exact opposite is the case — these efforts aim to unlock the power of evolutionary innovation, so that it may be better harnessed for applied ends. Indeed, these ideas have been adopted in applied scenarios where greater task performance is desired and have yielded it [Cheney et al., 2013, Mengistu et al., 2016, Reisinger and Miikkulainen, 2007, Lehman and Stanley, 2008]. The central question that evolutionary algorithms researchers confront seems to be: at what level of abstraction can a biologically-inspired route to profound and useful innovation via evolution be realized? At the heart of the matter, evolutionary algorithm researchers are challenged to sort out what biological factors are important to the evolutionary process. As our analysis with the proximate-intermediate-ultimate framework reveals, strict adherence to biological plausibility (i.e. ultimate causal factors) is not necessary to promote evolvability. It is possible to act on the genotype-phenotype mapping by scaffolding proximate causes of evolvability and to act on the phenotype-fitness mapping by selecting for intermediate causes of evolvability. Such clever adjustments to the evolutionary model allow for the acquisition of evolvability without undue computational cost of a naive, exclusively biologically plausible approach. The path forward for evolutionary computing will depend heavily on continued rigorous experimental work to suss out the constellation of causality behind evolvability in biological systems and digital models. It is hoped that this experimental work, guided and motivated by theoretical analysis, will ultimately yield stronger methodological techniques to promote evolvability in digital systems and, thus, continue to advance the practical utility of evolutionary algorithms in applied settings.

Chapter 7

Conclusion

Having thoroughly scrutinized the theoretical basis of evolvability, we will conclude by using evolvability as a lens to reflect on the relation between evolutionary biology and evolutionary computing to develop a broader perspective on the topic. It may not be surprising that evolvability is a topic of active discussion among elements of the evolutionary biology community [Pigliucci, 2008]. The concept falls under the umbrella of a broader effort to expand the theoretical framework of evolution called the extended evolutionary synthesis [Pigliucci, 2007]. Compared to the EA community, however, the concept of evolvability has been slower to gain traction among evolutionary biologists. As Kirschner and Gerhart, a pair of biologists known for their theory of facilitated variation, comment,

> "Many evolutionary biologists do not see a need to connect somatic adaptability to the generation of variation, and some see a need to keep them separate. For them, it is sufficient to say that random mutation is required and that the phenotypic variation arises haphazardly from it as random damage; the organism's current phenotype does not matter for the variation produced, and the output of variation is nearly random [Kirschner and Gerhart, 2005, p 219]."

Perhaps, in part, evolutionary biologists are less predisposed to interest in evolvability because they are not so directly stymied its absence. Success in attempts to emulate the evolutionary process to generate designs for sophisticated systems such as artificial neural networks or robotic bodies hinges on the ability of the evolutionary algorithm to generate viable, heritable variation. This development of this capability has been a major hurdle in EA research, especially in the field's early years. The intense and ubiquitous interest in evolvability among the EA community should therefore come as no surprise.

EAs have yielded interesting and useful results, but have not yet come close to replicating the intricacy or scale of biological systems [Tonelli and Mouret, 2011]. In classical EAs, a directly or trivially indirectly encoded population evolves against a static fitness function. Fitness gains are typically realized for a period of several hundred generations before innovation stagnates and the population settles out at

an equilibrium. This approach, predicated on a fundamentally accurate but extremely simplistic view of evolution, yields limited results. The stunted effectiveness of early EAs can be cast as a reflection of the limitations innate to the theory on which the algorithms are built. EA research provides uniquely direct and blunt evidence that it is *not* sufficient to say that "phenotypic variation arises haphazardly" [Kirschner and Gerhart, 2005, p 219]. EA can lend ammunition to biologists who posit that the ability of organisms to yield offspring more fit than themselves through "mutation, recombination and development" is "surprising and... demands an explanation" [Draghi and Wagner, 2008]. As skeptics in the evolutionary biology community point out, "the best way to elevate the prominence of genuinely interesting phenomena... is to strengthen the evidence for their importance" [Laland et al., 2014]. The evolutionary computing community have enjoyed a fruitful thrust to develop research algorithms that incorporate a broader array of theoretical factors that may influence evolution, such as varying fitness functions and phenotypic plasticity [Kashtan et al., 2007, Moczek et al., 2011, Downing, 2012]. Perhaps the evolutionary biology community would find widespread elevation of such theory beyond the status of circumstantial "add-ons to the basic processes that produce evolutionary change" similarly fruitful [Laland et al., 2014].

At present, it seems likely that evolvability stems from a large and diffuse web of cooperating mechanisms. The establishment — or rejection — of empirical evidence for causal links between factors such as plasticity or the developmental process and evolvability must be a key research goal in the field of evolutionary algorithm design. Such results will directly support efforts to refine the evolutionary algorithm and realize performance more closely akin to that of its biological counterpart. This line of inquiry raises and addresses questions of interest to the evolutionary biology community, especially in light of continuing controversy surrounding the extensions to the evolutionary synthesis. It will help determine which theoretical elaborations are necessary to account for evolution as observed in biology. It is hoped that further research in this vein — both *in silico* and *in vivo* — and, especially, continued exchange between EA and evolutionary biology researchers will yield both biological insight and more powerful digital engineering techniques.

Glossary

fitness Fitness refers to the success of an individual at passing its genetic information to the next generation. An individual with high fitness creates many offspring while an individual with low fitness does not. Success at surviving challenges posed by the environment is an important factor in determining fitness. In evolutionary algorithms, the concept of fitness is abstracted to the fitness function where an individual is scored based on its aptitude at performing a certain task.. 4

genotype Genotype refers to information that is used to determine the phenotype that is passed from generation to generation. In biology, a DNA sequence serves as the genotype. Although many different genotypic encodings are employed in evolutionary algorithms, in evolutionary algorithms the genotype ultimately boils down to a collection of digital information.. 5

individual Individuals are the object upon which evolution operates; evolution evaluates and selects on individuals and recombines individuals to form new individuals. In biology, and individual is an individual organism such as a single tree or a single bird. In evolutionary algorithms, an individual is abstracted as a candidate solution to a problem.. 4

phenotype Phenotype refers to the characteristics of an individual that interact with its environment to determine its fitness. In biology, the physical form of an organism (i.e. its body) is the phenotype. In evolutionary algorithms, the phenotype refers to the characteristics of an individual that are evaluated during selection.. 4

population A population is a collection of individuals that compete to transmit their genetic information to the next generation. These individuals are typically highly similar and, in many cases in both biology and evolutionary algorithms, recombine their genetic information to produce offspring.. 4

recombination Recombination refers to the generation of new genetic material from existing genetic material. This can involve combinations of two or more sets of genetic material, as in sexual reproduction, and/or random perturbation of genetic information (i.e. mutations).. 4

selection Selection refers to the determination of which individuals will pass genetic material on to the next generation by creating offspring (and how many offspring they will be able to generate) and which will not.. 4

Bibliography

[Arthur and Malcolm Farrow, 1999] Arthur, W. and Malcolm Farrow, a. (1999). The Pattern of Variation in Centipede Segment Number as an Example of Developmental Constraint in Evolution. *J. theor. Biol*, 200.

[Benzinger, 1969] Benzinger, T. H. (1969). Heat regulation: homeostasis of central temperature in man. *Physiological reviews*, 49(4):671–759.

[Cheney et al., 2013] Cheney, N., Maccurdy, R., Clune, J., and Lipson, H. (2013). Unshackling Evolution: Evolving Soft Robots with Multiple Materials and a Powerful Generative Encoding.

[Clune et al., 2008] Clune, J., Ofria, C., and Pennock, R. T. (2008). How a generative encoding fares as problem-regularity decreases. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*.

[Clune et al., 2011] Clune, J., Stanley, K. O., Pennock, R. T., and Ofria, C. (2011). On the performance of indirect encoding across the continuum of regularity. *IEEE Transactions on Evolutionary Computation*.

[Coyne, 1987] Coyne, J. A. (1987). Lack of response to selection for direction asymmetry in Drosophila melanogaster. *Journal of Heredity*, 78(119).

[De Kroon et al., 2005] De Kroon, H., Huber, H., Stuefer, J. F., and Van Groenendael, J. M. (2005). A modular concept of phenotypic plasticity in plants.

[Devesa et al., 2016] Devesa, J., Almengló, C., and Devesa, P. (2016). Multiple Effects of Growth Hormone in the Body: Is it Really the Hormone for Growth? *Clinical medicine insights. Endocrinology and diabetes*, 9:47–71.

[Downing, 2010] Downing, K. L. (2010). The Baldwin Effect in Developing Neural Networks. *GECCO Proceedings*.

[Downing, 2012] Downing, K. L. (2012). Heterochronous Neural Baldwinism. *Artificial Life*, 13.

[Downing, 2015] Downing, K. L. (2015). *Intelligence emerging : adaptivity and search in evolving neural systems*. MIT Press, Palatino.

[Draghi and Wagner, 2008] Draghi, J. and Wagner, G. P. (2008). EVOLUTION OF EVOLVABILITY IN A DEVELOPMENTAL MODEL. *Evolution*, 62(2):301–315.

[Dutta, 1979] Dutta, M. N. (1979). Earlobe Attachment Among the Ahom od Dibrugarh, Upper Assam. *Current Anthropology*, 20(2):399.

[Edelman and Gally, 2001] Edelman, G. M. and Gally, J. A. (2001). Degeneracy and complexity in biological systems. *PNAS*, 98(22):13736–13768.

[Fusco and Minelli, 2010] Fusco, G. and Minelli, A. (2010). Phenotypic plasticity in development and evolution: facts and concepts. *Philosophical Transactions of The Royal Society*, 365:547–556.

[Gregory, 2008] Gregory, T. R. (2008). The Evolution of Complex Organs. *Evolution: Education and Outreach*, 1(4):358–389.

[Griffiths et al., 2015] Griffiths, A. J. F., Wessler, S. R., Carroll, S. B., and Doebley, J. F. (2015). *Introduction to genetic analysis*.

[Grimes and Hurst, 2012] Grimes, J. A. and Hurst, J. W. (2012). THE DESIGN OF ATRIAS 1.0 A UNIQUE MONOPOD, HOPPING ROBOT *.

[Gu, 2003] Gu, X. (2003). Evolution of duplicate genes versus genetic robustness against null mutations.

[Guyton, 1971] Guyton, A. C. (1971). *Basic human physiology: normal function and mechanisms of disease*. Saunders.

[Hamdoun and Epel, 2007] Hamdoun, A. and Epel, D. (2007). Embryo stability and vulnerability in an always changing world. *Proceedings of the National Academy of Sciences of the United States of America*, 104(6):1745–50.

[Hofreiter, 2008] Hofreiter, M. (2008). Mammoth genetics. *Nature*, 456:330–331.

[Hornby et al., 2006] Hornby, G. S., Globus, A., Linden, D. S., and Lohn, J. D. (2006). Automated Antenna Design with Evolutionary Algorithms. *AIAA Space*, pages 19–21.

[Kappraff, 2004] Kappraff, J. (2004). Growth in Plants: A Study in Number. *Form*, 19:335–354.

[Kashtan and Alon, 2005] Kashtan, N. and Alon, U. (2005). Spontaneous evolution of modularity and network motifs. *PNAS*, 102(39):1377313778.

[Kashtan et al., 2007] Kashtan, N., Noor, E., and Alon, U. (2007). Varying environments can speed up evolution. *Proceedings of the National Academy of Sciences*, 104(34):13711–13716.

[Kim, 2011] Kim, O. (2011). How many cells are there in the human body? *MicrobeHunter Microscopy Magazine*.

[Kirschner and Gerhart, 2005] Kirschner, M. and Gerhart, J. (2005). *The plausibility of life : resolving Darwin's dilemma*. Yale University Press.

[Laland et al., 2014] Laland, K., Uller, T., Feldman, M., Sterelny, K., Müller, G. B., Moczek, A., Jablonka, E., Odling-Smee, J., Wray, G. A., Hoekstra, H. E., Futuyma, D. J., Lenski, R. E., Mackay, T. F. C., Schluter, D., and Strassmann, J. E. (2014). Does evolutionary theory need a rethink? *Nature*, 514(7521):161–164.

[Lehman and Stanley, 2008] Lehman, J. and Stanley, K. O. (2008). Exploiting Open-Endedness to Solve Problems Through the Search for Novelty. In *Proceedings of the Eleventh International Conference on Artificial Life (ALIFE XI)*, pages 329–336, Cambridge, MA. MIT Press.

[Lehman and Stanley, 2011] Lehman, J. and Stanley, K. O. (2011). Improving Evolvability through Novelty Search and Self-Adaptation.

[Lobo, 2008] Lobo, I. (2008). Environmental Influences on Gene Expresssion. *Nature Education*, 1(1):39.

[Lodish et al., 2000] Lodish, H., Berk, A., Zipursky, S. L., Matsudaira, P., Baltimore, D., and Darnell, J. (2000). Osmosis, Water Channels, and the Regulation of Cell Volume.

[Manyard Smith and Sondhi, 1960] Manyard Smith, J. and Sondhi, K. C. (1960). The genetics of a pattern. *Genetics*, 45:1039–1050.

[Mengistu et al., 2016] Mengistu, H., Lehman, J., and Clune, J. (2016). Evolvability Search: Directly Selecting for Evolvability in order to Study and Produce It. *GECCO Proceedings*.

[Mitchell, 1996] Mitchell, M. C. s. (1996). *An introduction to genetic algorithms*. MIT Press.

[Moczek et al., 2011] Moczek, A. P., Sultan, S., Foster, S., Ledó N-Rettig, C., Dworkin, I., Nijhout, H. F., Abouheif, E., and Pfennig, D. W. (2011). The role of developmental plasticity in evolutionary innovation. *Proc. R. Soc. B*.

[Mouret et al., 2010] Mouret, J.-B., Doncieux, S., and Girard, B. (2010). Importing the Computational Neuroscience Toolbox into Neuro-EvolutionApplication to Basal Ganglia. *Proceedings of GECCO10*.

[Nguyen et al., 2015] Nguyen, A., Yosinski, J., and Clune, J. (2015). Innovation Engines: Automated Creativity and Improved Stochastic Optimization via Deep Learning. In *Proceedings of the Genetic and Evolutionary Computation Conference*, Madrid.

[Owens and Lane, 2003] Owens, D. W. and Lane, E. B. (2003). The quest for the function of simple epithelial keratins. *BioEssays*, 25(8):748–758.

[Persson, 2007] Persson, B. N. J. (2007). Wet adhesion with application to tree frog adhesive toe pads and tires. *Journal of Physics: Condensed Matter*, 19(37):376110.

[Pfennig, 1992] Pfennig, D. W. (1992). Polyphenism In Spadefoot Toad Tadpoles as a Locally Adjusted Evolutionarily Stable Strategy. *Evolution*, 46(5):1408–1420.

[Pigliucci, 2007] Pigliucci, M. (2007). Do we need an extended evolutionary synthesis?

[Pigliucci, 2008] Pigliucci, M. (2008). Is evolvability evolvable? *Nature Reviews Genetics.*, 9(1):75–82.

[Poli et al., 2008] Poli, R., Langdon, W., McPhee, N., and Koza, J. (2008). *A field guide to genetic programming.*

[Purnell and Thompson, 1973] Purnell, D. J. and Thompson, J. N. J. (1973). Selection for asymmetrical bias in a behavioral character of Drosophila melanogaster. *Heredity*, 31:401–405.

[Reisinger and Miikkulainen, 2007] Reisinger, J. and Miikkulainen, R. (2007). Acquiring Evolvability through Adaptive Representations. *GECCO07 Proceedings.*

[Reisinger et al., 2005] Reisinger, J., Stanley, K. O., and Miikkulainen, R. (2005). Towards an Empirical Measure of Evolvability. *GECCO05 Proceedings.*

[Richter et al., 2015] Richter, A., Botsch, M., and Menzel, S. (2015). Evolvability of representations in complex system engineering: A survey. In *2015 IEEE Congress on Evolutionary Computation, CEC 2015 - Proceedings.*

[Rimbault et al., 2013] Rimbault, M., Beale, H. C., Schoenebeck, J. J., Hoopes, B. C., Allen, J. J., Kilroy-Glynn, P., Wayne, R. K., Sutter, N. B., and Ostrander, E. A. (2013). Derived variants at six genes explain nearly half of size reduction in dog breeds. *Genome research*, 23(12):1985–95.

[Risi et al., 2010] Risi, S., Lehman, J., and Stanley, K. O. (2010). Evolving the Placement and Density of Neurons in the HyperNEAT Substrate.

[Risi and Stanley, 2010] Risi, S. and Stanley, K. O. (2010). Indirectly encoding neural plasticity as a pattern of local rules. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics).*

[Royal Horticultural Society,] Royal Horticultural Society. Aloe polyphylla.

[Sandrini and Piskur, 2005] Sandrini, M. P. B. and Piskur, J. (2005). Deoxyribonucleoside kinases: two enzyme families catalyze the same reaction. *Trends in biochemical sciences*, 30(5):225–8.

[Schlichting, 1986] Schlichting, C. D. (1986). The Evolution of Phenotypic Plasticity in Plants. *Source: Annual Review of Ecology and Systematics Ann. Rev. Ecol. Syst*, 17(17).

[Sinervo and Lively, 1996] Sinervo, B. and Lively, C. (1996). The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 380(6571):240–243.

[Smith et al., 1985] Smith, J. M., Burian, R., Kauffman, S., Alberch, P., Campbell, J., Goodwin, B., Lande, R., Raup, D., and Wolpert, L. (1985). Developmental Constraints and Evolution: A Perspective from the Mountain Lake Conference on Development and Evolution The Quarterly Review of Biology. *The Quarterly Review of Biology*, 60(3):265–287.

[Stenzel et al., 2011] Stenzel, V., Wilke, Y., and Hage, W. (2011). Drag-reducing paints for the reduction of fuel consumption in aviation and shipping. *Progress in Organic Coatings*, 70(4):224–229.

[Tarapore and Mouret, 2015] Tarapore, D. and Mouret, J. B. (2015). Evolvability signatures of generative encodings: Beyond standard performance benchmarks. *Information Sciences*.

[Tonelli and Mouret, 2011] Tonelli, P. and Mouret, J.-B. (2011). On the Relationships between Synaptic Plasticity and Generative Systems. On the Relationships between Synaptic Plastic- ity and Generative Systems. 11:1531–1538.

[Tonelli and Mouret, 2013] Tonelli, P. and Mouret, J. B. (2013). On the relationships between generative encodings, regularity, and learning abilities when evolving plastic artificial neural networks. *PLoS ONE*.

[Tuinstra et al., 1990] Tuinstra, E., De Jong, G., and Scharloo, W. (1990). Lack of response to family selection for direction asymmetry in Drosophila melanogaster: left and right are not distinguished during development. *Proc. R. Soc. Lond. B*, 241(1301):146–152.

[Wagner, 2003] Wagner, G. P. (2003). Evolutionary Genetics: The Nature of Hidden Genetic Variation Unveiled. *Current Biology*, 13(24):R958–R960.

[Whitacre and Bender, 2010] Whitacre, J. and Bender, A. (2010). Degeneracy: A design principle for achieving robustness and evolvability. *Journal of Theoretical Biology*.

[Wikipedia contributors, 2016] Wikipedia contributors (2016). Brown rat.

[Wilder and Stanley, 2015] Wilder, B. and Stanley, K. (2015). Reconciling explanations for the evolution of evolvability. *Adaptive Behavior*, 23(3):171–179.

[Wilson, 2007] Wilson, D. S. (2007). *Evolution for everyone : How Darwin's theory can change the way we think about our lives*. Delta.