UNIVERSITY OF SOUTHAMPTON

How evolution learns to evolve: Principles of induction in the evolution of adaptive potential

by

Kostas Kouvaris

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the

Faculty of Engineering, Science and Mathematics

School of Electronics and Computer Science

June 2018

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS
SCHOOL OF ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by Kostas Kouvaris

Explaining how organisms can exhibit suitable phenotypic variation to rapidly adapt to novel environmental conditions is central in evolutionary biology. Although such variability is crucial for the survival of a lineage and its adaptive potential, it remains poorly understood. Recent theory suggests that organisms can evolve designs that help them generate novel features that are more likely to be beneficial. This is possible when the environments that the organisms are exposed to share common regularities. Selection though cannot favour phenotypes for fitness benefits that have not yet been realised. Such capacity implies that natural selection has a form of foresight, which is inconsistent with the existing evolutionary theory. It is unclear why selection would favour flexible biological structures in the present environments that promote beneficial phenotypic variants in the future, previously unseen environments.

In this thesis, I demonstrate how organisms can systematically evolve designs that enhance their evolutionary potential for future adaptation relying on insights from learning theory. I investigate how organisms can predispose the production of useful phenotypic variation that helps them cope with environmental variability within and across generations, either through genetic mutation or environmental induction. I find that such adaptive capacity can arise as an epiphenomenon of past selection towards target optima in different selective environments without a need for a direct or lineage selection. Specifically, I resolve the tension between canalisation of past selected targets and anticipation of future environments by recognising that induction in learning systems merely requires the ability to represent structural regularities in previously seen situations that are also true in the yet-unseen ones. In learning systems, such generalisation ability is neither mysterious, nor taken for granted. Understanding the evolution of developmental biases as a form of model learning and adaptive plasticity as task learning can provide valuable insights into the mechanistic nature of the evolution of adaptive potential and the evolutionary conditions promoting it.

Contents

List of Figures

Acknowledgements

First of all, I would like to thank my supervisors, Richard A. Watson and Markus Brede, for giving me the opportunity to do this PhD, but also for their continuous support, guidance and valuable feedback. I would also like to thank everyone in evolution group for creating such a pleasant work atmosphere. I am really grateful to Alfredo Rago, Louis Kounios, Tobias Uller and Jeff Clune for their invaluable help, inspiration, fruitful discussions and feedback. Katsia Pabortsava, I cannot thank you enough for your immense support, motivation and always being there for me. Special thanks to Melis Cevatoglu, Andrew Said, Billy Hurndall, Ilse Bloom, Ruvi Kaur, David Lang and Aude Mazuel for the wonderful times we shared, and all the necessary, and unnecessary, distractions from my PhD. Last but by no means least, I would like to thank my family and friends in Greece for all their support and encouragement.

This work was funded by the Institute for Complex Systems Simulation, Southampton (EP/G03690X/1) and by the Defense Science and Technology Laboratory (DSTLX-1000074615).

List of Publications

- Kouvaris, K., Clune, J., Kounios, L., Brede, M. and Watson, R.A., 2017. How evolution learns to generalise: Using the principles of learning theory to understand the evolution of developmental organisation. PLoS computational biology, 13(4), p.e1005358.
- Watson, R.A., Mills, R., Buckley, C.L., Kouvaris, K., Jackson, A., Powers, S.T., Cox, C., Tudge, S., Davies, A., Kounios, L. and Power, D., 2016. Evolutionary connectionism: algorithmic principles underlying the evolution of biological organisation in evo-devo, evo-eco and evolutionary transitions. Evolutionary biology, 43(4), pp.553-581.

Glossary

Adult Phenotype A set of phenotypic characteristics of a fully developed organism. This is analogous to the actual output of learning model.

Batch Learning In batch learning, all of the samples in the training set are presented at a time during the learning phase. The performance of the learning system is thus evaluated over all of training samples at each learning step.

Bias Variance Dilemma The prediction errors of a learning model can be decomposed into two main parts: one error type due to bias, and one error type due to variance. The bias-variance dilemma refers to the trade-off between the two sources of error when a predictive model minimises the total error. Increasing the bias will decrease the variance, and vice versa. Understanding the bias-variance tradeoff can help us alleviate the problem of under- or over-fitting.

Directional Selection Selection favours an extreme phenotype in a population, resulting in adaptive change in the direction of that phenotype. This form of selection often occurs under environmental change. For example, when populations migrate in different habitats characterised by different environmental pressures.

Evolution by means of natural selection (ENS) ENS refers to Darwin's theory that all living organisms evolve through natural selection of small, heritable variations that enhance their ability to survive and reproduce. Accordingly, evolution can be defined as change in allele frequencies in a population. Evolution by means of natural selection requires three mechanisms: variation, selection and inheritance. These conditions are necessary to explain the living world as we know it.

Evolvability (Organismic) The propensity of an organism to acquire novel and useful functions (phenotypes) through genetic variation. Evolvability describes how random (undirected) genetic variation is channelled into partially non-random (directed) adaptive phenotypic variation. Development plays a central role in evolvability, since it is development that translates the genotype into the adult phenotype.

Fitness The fitness of an organism indicates how well-adapted the organism is to its current selective environment. It is estimated based on how close the adult developed phenotype is to the target phenotype of the current environment the organism is in.

Generalisation Error Generalisation error is a measure that indicates how well a learning model can predict previously-unseen data, i.e., how well the model can generalise to future data. This error is evaluated over the observations kept in the test set.

Inductive Biases A set of selection criteria that favours models of a desired behaviour. Inductive biases aid the learner to improve the accuracy of its predictions to new situations that have not been encountered before.

Jittering Training with noise. A techique used to alleviate the problem of over-fitting by adding a small amount of noise during the training phase. Adding noise makes it difficult for the optimisation process to precisely fit the training data, and thus tends to avoid fitting the noise in the data. Theory suggests that jittering has similar effects to regularisation.

Learning Learning is the process of discovering patterns and regularities in the data to gradually improve the performance of the learner over a given task. Learning involves the acquisition and the representation of knowledge about a specific problem domain. In addition, learning accounts for the storage and the organisation of such knowledge to perform a certain task.

Learning Rate The rate of change in the model parameters during the learning phase. Learning rates indicate the amount of correction that occurs at each learning step. They are analogous to the amount of adaptive genetic mutation accumulated by the evolving organisms in response to a set of selective environments.

Modularity A variational property of an organism that describes to what extend groups of morphological traits are independent. Modularity is usually determined empirically by the genetic or phenotypic correlations and is heavily dependent on the structure of genetic architecture.

Online Learning In online learning, one sample from the training set is presented at a time during the learning phase. The performance of the learning system is thus evaluated over one training sample at each learning step.

Overfitting Over-fitting occurs when a learning model captures the noise in the data and not the true underlying trend/regularity in the data. Intuitively, the model fits the training data so well that captures information that irrelevant to the given

learning problem. A model that over-fits is characterised by low training error and high test error. Although the learnt model performs really well in observations used during the training phase, it exhibits poor performance on new data.

Parsimony Pressure Techniques used to penalise the complexity of a learning model. Favouring simpler models tends to alleviate the problem of over-fitting and improve the generalisation performance of the learner.

Plasticity Phenotypic plasticity is the capacity of an individual genotype to express different phenotypes when exposed in different environmental conditions. For instance, *Daphnia longicephala* tends to develop tail spines and protective crest when its predator, Notonecta, is present in its environment. Plasticity is an important developmental property of the organisms to cope with unpredictable environmental change. Phenotypic plasticity can be neutral, adaptive or maladaptive with respect to an organism's fitness for a given environment. The environment pays two roles: it performs as an ''agent of development" affecting what phenotype is produced from a given genotype, and as an ''agent of selection" evaluating how well- or poorly-adapted the produced phenotypes are.

Plasticity First Hypothesis The traditional paradigm in evolutionary biology suggests that genes 'lead' the process of adaptive evolution and phenotypes 'follow'. In contrast, the plasticity first hypothesis suggests that novel traits can be initiated by environmental change and thus plasticity may precede and even guide genetic adaptation.

Reaction Norms Graphical representations that describe the association of phenotypic responses to environmental change by plotting phenotypic values across the range of possible selective environments.

Regularisation A way to alleviate the problem of over-fitting and improve the generalisation performance of a learning system on a task. Regularisation introduces an additional term in the original objective function that favours desirable solutions. The regularisation term embeds extra information in the criterion function about what is desired and what is not (e.g., smoothness or invariance) by introducing additional constraints and biases towards simpler hypotheses, limiting the complexity of the model and thus reducing the excessive variance of the model that can lead to over-fitting.

Robustness A fundamental property of an evolving organism that allows for stability and reliable functionality against genetic (internal) and/or environmental (external) perturbations.

Stabilising Selection Selection favours non-extreme values by penalising (selecting against) phenotypes that deviates from the phenotypic optimum.

Supervised Learning This form of learning aims to change the internal parameters of a learning model so that to improve its behaviour on a given task. The learning process involves an optimisation process that changes the parameters to reduces the error between the actual outputs (predictions of the model) and the desired outputs, i.e., training error. The error signals in the optimisation process imply the presence of an external teacher that informs the learning process on how to change the parameters in the right direction. This form of learning is analogous to how genes change over evolutionary time so that the phenotypes better match the target phenotypes in a given set of selective environments.

Target Phenotype A set of desired (optimal) phenotypic traits an organism should have for a given selective environments. This is analogous to the desired outputs (labeled data) in supervised learning.

Test Set A set of observations (data) used to evaluate the generalisation performance of a model after the training phase. The test set comprises of new observations that have not been used during training to estimate how well the model generalises to new situations. This is analogous to future previsouly-unseen phenotypic targets.

Training Set A set of observations (data) used during the model building phase (learning). This is analogous to the past selected phenotypic targets.

Training Error This type of error is a measure that indicates how well a learning model can fit the data used during the learning phase, i.e., how well the model can memorise the past experience. This error is evaluated by how close the model predictions are to the observations in the training set.

Under-fitting Under-fitting occurs when a learning model is incapable of capturing the underlying trend/regularity in the training data. Intuitively, the model does not fit the data very well. A model that under-fits is characterised by both high training and test error, and exhibits poor performance on new data. Under-fitting contrasts the problem of over-fitting and can be explained by the bias-variance dilemma.

Unsupervised Learning This form of learning aims to perform statistical inference over input data. The aim of an unsupervised learning algorithm is to optimise a criterion that is independent of a specific task in the absence of labeled data. The objective function in this type of learning is a function of the actual output produced by the learning system. The desired (target) outputs are not known, and thus learning does not involve external teacher signals to correct the behaviour of the learner. The learner can discover and extract regularities in the training samples, such as correlational patterns over input features, or different categories and clusters. Unsupervised learning is often used to reduce the dimensionality of the features space.

Chapter 1

Introduction

1.1 Evolution of adaptive potential

1.1.1 Overview

The current paradigm in evolutionary biology is that of the Modern Synthesis (Mayr and Provine, 1998); an extension of Darwin's theory that all living organisms evolve through natural selection of small, heritable variations that enhance their ability to survive and reproduce. Explaining how organisms can develop useful physiological, morphological or behavioural characteristics (i.e., phenotypes) that help them cope with new environments is central in evolutionary biology (Bedau et al., 2000; Adami et al., 2000; Lenski et al., 2003; Bedau, 2009; Moczek et al., 2011). Such capacity is often identified as a prerequisite for their ability to exhibit adaptive evolution (i.e., evolvability) (Pigliucci, 2008). Biological research has focused on how different forces of selection can determine the paths of the evolution of such adaptive phenotypic characteristics (Lynch, 2007). It is, however, equally important to understand the mechanisms that enable the production of adaptive phenotypes in the first place (Riedl and Jefferies, 1978; West-Eberhard, 2003; Toussaint, 2002; Brakefield, 2006; Gerhart and Kirschner, 2007; Toussaint and von Seelen, 2007; Braendle et al., 2010; Pigliucci, 2008). In evolutionary biology, it is generally assumed that such adaptive phenotypes are available, either as expressed phenotypes in a population or as potential but yet unrealised phenotypic outcomes (Kirschner and Gerhart, 1998; Pigliucci, 2008). Yet, work in evolutionary computation questions the validity of this assumption (Lipson et al., 2002; Kashtan and Alon, 2005). Understanding how new and useful phenotypic traits arise is not clear.

Explaining the origins of novel complex traits is a long-standing problem in evolutionary biology (Muller and Wagner, 1991; Schlichting et al., 1998; Gould, 2002; Mayr, 1960; Wagner and Lynch, 2010; Müller, 2007; West-Eberhard, 2003; Moczek et al., 2011; Brigandt and Love, 2012; Wagner, 2014; Hall and Kerney, 2012). New traits or combination

of traits can originate through i) genetic mutation, ii) recombination of the genetic material, or iii) environmental induction; the ability of the organisms to conditionally express different phenotypes based on the environment they are in. All these three processes are potent sources of evolutionary novelty.

Evolution is often defined as change in the allele frequencies (Mayr and Provine, 1998). Traditional theory thus tends to emphasise on mutation as the prime mover and the only legitimate source of evolutionary change, and thus novelty (Carroll, 2008; Ghalambor et al., 2007; Pigliucci and Murren, 2003). The process of adaptive evolution through mutation only is conceptually clear when a new mutation triggers reliable and advantageous phenotypic effects across genetic and environmental contexts. Natural selection acts a sieve checking which mutations are beneficial and which are not. As a result, novel adaptive traits may arise as a mere accident. Similar considerations have been made for situations when variation arises due to environmental induction (De Jong, 2005; Ghalambor et al., 2007; Lande, 2009). Organisms experiencing a different environment, even if it has not been encountered before, then produce traits that might indeed be beneficial. However, such adaptive variation might appear, but not necessarily so. Novelty might again arise by accident. Consequently, such theories provide inadequate explanations on the capacity of the organisms for phenotypic innovation and subsequent adaptive evolution.

1.1.2 Brief History of Evolvability

The term evolvability is attributed to Dawkins (2003), where he introduced his Bimorphs model, demonstrating how the developmental system of an organism could enhance its ability to evolve, how this ability could evolve itself and how this ability could affect the evolutionary history of the organism. However, the actual concept of evolvability is much older. The concept of evolvability brought many related ideas scattered and diffused in the biological literature into focus. For instance, similar issues have been addressed before by Riedl and Jefferies (1978) and Smith et al. (1985), under the term of *developmental constraints*, and Conrad (1990) who argued about the amenability to evolution in chemical networks. Since its original proposition the term of evolvability has received a lot of attention and the term 'evolvability' has been included in hundreds of papers (at least 364 by 2008) since then -– mostly theoretical and conceptual (Pigliucci, 2008). Although the term has been frequently used throughout a large body of literature consistently with respect to its original meaning. Researchers from different disciplines, outside evolutionary biology, such as micro-biology, population genetics and evolutionary computing, usually use the term in new ways to serve their own cluster of issues. These new meanings are related but narrower in an effort to elucidate the original idea and apply it in specific experiments or models.

1.1.3 The Many Definitions of Evolvability

Although the concept of evolvability is central in evolutionary biology, there is no unambiguous definition for evolvability. This generates incoherence throughout the literature and often conceptual confusion (Pigliucci, 2008). Recently, several attempts have been made to formalise the concept of evolvability, leading to a plethora of definitions dependent on the context under consideration. Essentially, this pluralism stems from the fact that evolvability actually deals with a family of distinct, yet overlapping, issues and concepts. To better understand evolvability, it is thus important to review and distinguish between the multiple meanings of evolvability. According to Gallagher (2009), the definitions of evolvability can be categorised as follows:

- *Organismic Evolvability*: the propensity of an organism to acquire *novel* and *useful* functions (phenotypes) through genetic variation. Understanding of organismic evolvability entails understanding of how random (undirected) genetic variation is channelled into partially non-random (directed) adaptive phenotypic variation. Thus, the genotype-phenotype (G-P) map plays a central role, as it is the G-P map that translates genotypes into phenotypes. This type of evolvability is in focus in this work and it will be further described in the subsequent sections.

- *Trait Evolvability*: the propensity of a specific trait to rapidly respond to selection. Thus, unlike organismic evolvability, trait evolvability is a property of the trait and not the organism as a whole. The actual definition was originally proposed in the area of population genetics by Houle (1992), generally as 'the ability of a population to respond to natural or artificial selection'. In his analysis, however, the statistical measures of evolvability used (i.e., Houle's *genetic coefficient of variation*) concerned a single quantitative trait. Similar definitions have been also proposed by Flatt (2005); 'the ability of a population to respond to selection' and Griswold (2006) who addressed evolvability in terms of the speed of evolution of a single phenotypic character. Although they treated evolvability by explicitly measuring heritability, they all merely focused on the *standing* genetic variance within a population. Undoubtedly, the speed at which a trait evolutionarily responds to selection is an importance aspect of evolvability. Nevertheless, this narrow sense of evolvability cannot capture all macro-evolutionary facets of evolvability, such as the evolution of phenotypic innovations, new species and complexity.

- *Individual Fitness Evolvability*: evolvability can be defined by a number of fitness-related measures (Gallagher, 2009):

  1. "the expected variance of the fitness of its offspring, prior to selection"
  2. "the expected maximum fitness of its offspring, prior to selection"
  3. "the expected fitness of its offspring after selection (or of some future generation)" (Turney, 1999)

4. "the likelihood that the offspring generation includes an individual of higher fitness" (Altenberg, 1994; Smith et al., 2002)

The different measures above are in essence the same and correlated with each other. They are all defined in terms of fitness of individuals and implicitly motivated by the same mindset regarding the evolution of evolvability. The value of this definition stems from the conceptual clarity of its description using the familiar concepts of individuals and fitness — a fundamental notion in evolutionary biology. In short, evolvability, in that sense, is all about the expected offspring fitness distribution. Evolvability becomes important as natural selection works over generations and acts upon such distributions so as to increase absolute fitness in the future.

Despite the appealing advantages of its conceptual simplicity and compatibility with the existing evolutionary theory, the notion of individual fitness evolvability cannot capture two vital aspects of evolvability: i) the evolution of phenotypic development, e.g., how and what developmental architectures can facilitate evolvability and how they were evolved, and ii) the evolution of complexity, e.g., what mechanisms can enhance the evolution of complexity and how.

There is a plethora of other definitions that bear many resemblances with the aforementioned definitions, but do not strictly fall in any of these categories. For instance, Quayle and Bullock (2006), from a computational science perspective, measured evolvability as the time a population need to reach a certain target phenotype. Alternatively, in the area of computational learning, Valiant (2013, 2009) takes a different angle and equates evolvability with learnability in the PAC sense. The aim is to discover the ideal function that generates the most beneficial behaviour in every possible scenario. Mathematically, Valiant defines a measure of performance of a genome as to how close its evolving function is to the ideal one over the distribution of all possible conditions, i.e., the higher the performance, the higher the chance of the organism to produce proper behaviour, and thus survive and reproduce. Yet, evolvability in that sense is a target and not a by-product of evolution. Consequently, it might be susceptible to the criticism that the evolution of evolvability is un-Darwinian (see below).

1.1.4 Phenotypic variation and its nature

For organisms to evolve by means of natural selection, there should be a non-negligible probability to generate beneficial phenotypic variants, either through genetic mutation or environmental induction (i.e., plasticity). Variation is the raw material that selection works with. We thus expect that generating more phenotypic variation could potentially increase the chance of discovering fit variants. In contrast to the typical gene-centric view of the Modern Synthesis, the focal point here is the phenotype.

However, merely increasing the amount of phenotypic variation would not necessarily entail good exploration over the space of possible phenotypic variants (Kirschner and Gerhart, 1998, 2006; Moczek et al., 2011; Kashtan and Alon, 2005; Toussaint, 2002; West-Eberhard, 2003; Watson et al., 2014). Instead, we should also emphasise on how fit the produced variants are. We may, for instance, have higher adaptability with lower phenotypic variation being expressed by favouring the production of potentially more useful variants. In other words, selecting amongst a few but good phenotypes may be preferred. Interest lies in how biological systems are capable of responding to past and new environments and readily adapt to them. This means that the biological systems are predisposed to produce phenotypic traits that are better than random Kirschner and Gerhart (1998, 2006).

1.1.5 Potential phenotypic variation: Variability

To explain how new and useful phenotypes can reliably arise, we need to focus on the potential and the tendency of the organisms or populations to produce various phenotypes, rather than the phenotypes that are already expressed in a population of individuals (Schlichting and Murren, 2004; West-Eberhard, 1998; Dichtel-Danjoy and Félix, 2004; Pigliucci, 2008). We refer to the former as potential, yet-unrealised, phenotypic variation and to the latter as standing genetic variation within a population. To better understand the difference between the two, it is crucial to distinguish between variation and variability (Wagner and Altenberg, 1996).

Variation constitutes a basic element in population genetics models. Also, it is relatively easy to observe variation in natural populations, since it represents the actual morphological differences among individuals. The realised variation is the outcome of both the sampling effects of selection acting on different traits and the underlying *variability*. Variability, on the other hand, occurs *prior* to selection. It describes the range of *potential* phenotypic traits that can be generated by the organisms through development given the genetic variation that is introduced through genetic mutation and recombination. The study of variability diverts the focus from the fitness of phenotypes to the processes of mutation and development that produce the adult phenotypes upon which selection acts. Variability concerns the potential and the propensity of phenotypic traits to vary.

1.1.6 Adaptive developmental biases

The phenotypic variability of an organism is determined by development. Development thus is responsible for the flexibility of an organism to produce different phenotypes, and its potential for adaptive change (West-Eberhard, 2003; Hallgrímsson and Hall, 2005; Hendrikse et al., 2007). Indeed, the amount and the type of phenotypic variation

are heavily determined by the genetic and developmental structure (Smith et al., 1985; Yampolsky and Stoltzfus, 2001; Braendle et al., 2010). Development is the machinery that translates a genotype into a phenotype and hence is determined both genetically and environmentally (Waddington, 2014). The traditional paradigm of evolutionary biology treats the genetic mutations as the dominant driver of all phenotypic innovations (West-Eberhard, 2003). Yet, there is not an one-to-one mapping between genotypes and phenotypes.

Some early theories suggested that the production of adaptive traits can be promoted by the robustness of developmental organisations (Baldwin, 1896; Simpson, 1953; Waddington, 1959). Robustness refers to the property of a biological system to maintain functionality against mutations or environmental change. These theories however were limited as they treated development as a black box (Hall, 2003), not accounting for the mechanisms driving it (Kirschner and Gerhart, 1998; Carroll, 2000; Pigliucci, 2007; Müller, 2007). Over the past few decades, studies in evolutionary developmental biology (evo-devo) have managed to shed light on key molecular mechanisms involved in development. This enabled mechanistic explanations of how the structure of development can shape the potential phenotypic variation upon which selection can act. Featuring development is crucial in explaining how adaptive phenotypic variation can arise and more importantly in identifying the conditions that promote it.

Kirschner and Gerhart (2006) gathered evidence from cellular biology and metazoan evo-devo and identified the *core processes* which allow developmental systems enhance their adaptive potential. Based on their structural principles, these core processes can be categorised into i) weak regulatory linkage, ii) compartmentation, iii) exploratory mechanisms and iv) redundancy (or degeneracy). For instance, the bodies of *Drosophila melanogaster* are separated into predefined compartments. The type of body segments (from the *bauplan*) fitted in these compartments is solely determined by the value of the *hox genes*. Consequently, changes in the values of hox genes result in changes in the *ordering* of the segment structure (i.e., different combinations of body segments) and not distorted unrecognisable body parts. Such core components are fundamental for the organisms, since they affect the operation of many other components. The core processes are considerably difficult to emerge but when they do they persist over long periods of evolutionary time. They can, therefore, induce a complex machinery that remains constant and down-regulate the process of development, facilitating significant phenotypic change in new directions.

1.1.7 Modularity

Another keystone of adaptive potential is the ubiquitous modular nature of the biological systems (Hansen, 2003). If we want to understand the potential of the organisms to produce new and useful phenotypes, we have to understand the variational properties

of organisms, and modularity is one such property. In short, modularity describes to what extend groups of morphological traits are independent and is usually determined empirically by the genetic or phenotypic correlations. Modularity is a property heavily dependent on the structure of genetic architecture; i.e., the amount and patterns of the epistatic dependencies and pleiotropy.

Consider the case where a genetic mutation simultaneously affects every morphological trait, i.e., pleiotropic effects. Then, a mutation that was advantageous for a specific trait would be disadvantageous for others, and hence almost no mutation would be favourable overall. Modularity, however, can, in principle, facilitate adaptive potential by removing such interferences and de-coupling clusters of traits, enhancing their ability to evolve independently (Wagner and Altenberg, 1996). A good illustration of phenotypic modularity is the repetition of certain identical body segments (Dawkins, 2003).

Yet, this notion of modularity is too simplistic and cannot capture ideas such as segment ordering, namely, what types of segments occupy each (body) slot. (Watson et al., 2014) recently proposed a model, which not only allows for such interpretation of different modular orderings, but also enable us to understand the evolution of developmental modularity in a more profound sense, providing clear insights on how intra– and inter-modular phenotypic correlations can arise. Specifically, when selective environments systematically change in a modular fashion, the evolved developmental structure internalises models that reflect these modular regularities of the past environments, exhibiting a form of developmental memory. This modular flexibility of development also allows for generating suitable phenotypic variation for future unseen environments when they share similar structural regularities (see also (Parter et al., 2008)).

1.1.8 Robustness vs Flexibility

Robustness is a fundamental property of natural systems (Kirschner and Gerhart, 1998; Wagner, 2013). Living organisms are ubiquitously observed to exhibit remarkable robustness against a great range of genetic and environmental perturbations. The former correspond to genetic changes in the genetic makeup of the organism (e.g., point mutations, gene deletions and gene duplications), whereas the latter correspond to random environmental fluctuations (e.g., physical and chemical changes).

In developmental biology, the notion of robustness is widely known with the term *developmental canalisation* (Waddington, 1942). Accordingly, developmental canalisation is defined as the propensity of a developmental process to produce a certain phenotype over a range of genetic and environmental conditions. If, for instance, a phenotype is quite robust, then small perturbations will be ignored and development will return down in the same 'canal' generating (almost) the same phenotype. However, if the perturbation is strong enough development will move outside the 'canal' to unexplored areas. Then,

the conservative and myopic nature of natural selection would favour for the canalisation of previously selected phenotypes, enhancing the stability and reliability of the developmental system. For instance, gene regulatory networks (GRNs) and metabolic pathways are found to self-organise in such a way to accommodate proper functioning of living organisms in a wide spectrum of genetic and environmental conditions (Aldana et al., 2007).

At the same time, biological systems are also flexible and generate new phenotypic variants. Robustness and developmental flexibility occur throughout different levels of the biological complexity. However, it is not clear how these two concepts interact with each other and what mechanisms give rise to structures that facilitate both. On the surface, robustness and flexibility may seem to be two antagonistic forces. The more robust a developmental system is, the less the (phenotypic) variation it produces (upon which selection acts), and thus the less evolvable it is (Wagner, 2008). If a good phenotype, for example, is discovered, then no variation would be favoured by selection, as it comes with no immediate benefit. Then, how did living organisms evolve genetic structures that are at both flexible and robust the same time? Some view these notions as two-sides of the same coin (Kirschner and Gerhart, 1998; Brigandt, 2007; Draghi et al., 2010).

For instance, high genetic canalisation would allow for the accumulation of cryptic genetic variation within the population, since it does not produce any phenotypic effects. This mechanism serves as an evolutionary capacitor, releasing this cryptic variation under environmental stress or due to biological switches, accelerating adaptive evolution under new environmental conditions. This is a good example of robustness-induced flexibility (see also Wagner's discussion of mutational robustness and 'neutral spaces' in G-P maps (Wagner, 2013)). Alternatively, robustness could facilitate adaptive evolution by confining the phenotypic variability of an organism so as to decrease the probability of encountering deleterious genetic perturbations. Although in that sense the overall phenotypic variation is still decreased, the ability of the organisms to produce adaptive phenotypes is enhanced by constraining phenotypic variation in more promising regions in the phenotype space, and thus increasing the propensity of a beneficial phenotypic variant to arise (Kirschner and Gerhart, 1998; Watson et al., 2014).

More interestingly, the paradox of robustness versus flexibility can be explained in a more profound sense, when the developmental process that induces the variational structure of the organism evolves internal representations of the statistical correlation structure of past selective environments (i.e., time-invariant regularities). Then, the dimensionality of the phenotype space is decreased by facilitating the production of certain (potentially useful) combinations of phenotypic traits, without necessitating a reduction in individual traits' variability (Watson et al., 2014).

1.1.9 Evolution of adaptive development biases

Development is also subject to selection (e.g., exoskeleton versus endoskeleton). Hence, the potential phenotypic variation induced by development is also malleable and shaped by selection. Moreover, developmental evolution guides the direction of phenotypic variation and what the space organism's offsprings is that can be explored through genetic variation. This, for example, can be seen in how the evolution of insects' body size which is limited by the shape and size of other organs that is dictated by the given body plan (Nijhout and Emlen, 1998). We see that evolution accumulates information regarding its past behaviour and respond to selection differently in the future (Frank, 2009; Shalizi et al., 2009). It therefore exhibits a path-dependent behaviour, namely, future evolutionary trajectories are dependent on the whole history of all previous states (Riedl and Jefferies, 1978). Different developmental structures provide different evolutionary potentials. But how can selection act upon and differentiate them?

The main argument against the evolution of adaptive developmental biases is simple. Adaptive potentials exert fitness benefits concerning future situations of a given lineage in the long-term. Selection however cannot favour traits for fitness benefits that have not yet been realised. It is unclear why selection would favour biological structures in the present environments that promote good phenotypic variants in future previously unseen environments. Furthermore, the conservative nature of natural selection would favour for more stable and reliable developmental systems (Wagner et al., 2007; Pavlicev et al., 2010). Accordingly, we expect selection to make previously selected phenotypes more resistant against genetic or environmental variation. Such developmental robustness (i.e., canalisation) seems to be intrinsically opposed to an increase in phenotypic variability. Specifically, previous studies have shown that when selection has the control of variation, this nearly always reduces variation, favouring robustness over flexibility (Clune et al., 2013b). Then, how is it possible for natural selection to evolve developmental organisations that facilitate higher potential for adaptive phenotypic variation in the future? Such capacity appear to be impossible, as it implies that natural selection has a form of foresight, which is inconsistent with the existing theory.

1.2 Evolution as learning

An alternative theory suggests these mechanisms of phenotypic variation can act on information in the genome (Watson and Szathmáry, 2016; Livnat, 2017). Under considerations of parsimony, selection can improve adaptation of the organisms to their current environments, but also create new units that provide higher adaptability to new environments that were not encountered before. Novelty emerges at the system level from the interactions of such units.

Adaptive evolution can be seen and better understood as a form of learning from examples (Livnat, 2017; Watson and Szathmáry, 2016; Watson et al., 2015, 2014; Valiant, 2013, 2009). Evolution by means of natural selection requires three fundamental mechanisms: variation, selection and inheritance (Lewontin, 1970). Conventionally, the current dogma in evolutionary biology considers these mechanisms to be fixed. In nature, however, these mechanisms are malleable and shaped by natural selection (Shapiro, 2011). For instance, the structure of gene regulatory interactions that determines the distribution of potential phenotypic variants, and thus the direction of future evolutionary change, is also modified over evolutionary time by natural selection (Müller, 2007; Parter et al., 2008; Watson et al., 2014). Ecological interactions between species in an ecosystem are also shaped by past evolution altering the ecological attractors. As a result, the selective pressures on certain individuals traits also change (Metz et al., 1992). This is further supported by the major transitions in evolution which involve changes in the level of evolutionary unit following the evolution of new mechanisms of inheritance (Szathmary and Maynard Smith, 2004; Okasha, 2005).

We see that the product of evolution can therefore also modify its own underlying evolutionary processes; i.e., the Darwinian machine is *self-referential* (Watson and Szathmáry, 2016; Livnat, 2017). Such circular causality however creates many theoretical and conceptual roadblocks and controversies with the existing framework of evolutionary theory (Watson et al., 2015). How can we understand a machine that changes its own mechanisms as a result of past experience, including its own products; i.e., a program without designer? Yet, this idea not only is not strange in computational learning, but something extensively studied and well-understood. Learning theory can inspire new approaches to intriguing question in the area evolutionary developmental biology (Watson and Szathmáry, 2016), such as the evolution of adaptive potential that is addressed in this thesis.

1.2.1 What is learning?

Learning is an integral part of human cognition and intelligence. In general, learning involves the acquisition and the representation of knowledge about a specific problem domain and the storage and organisation of such knowledge to perform a certain task. A way of defining learning is as any change in the learning system that lead to higher performance the second time the system experience the same task (Simon, 1983). Alternatively, learning can be defined as improvement in performance for a given environment through information gained from experiencing that environment (Langley, 1996). 'Performance' refers to the evaluation, or measure, of behaviour of the learning system in a specific task, e.g., accuracy. 'Environment' refers to the external setup the system has to learn, i.e., problem setting. However, the environment should have an underlying structure for learning to take place. Finally, 'experience' refers to how received inputs

were converted into outputs, while 'improvement' refers to any beneficial change in term of performance. In short, learning refers to any *change* in the information processing of the system that is *adaptive* with respect to a given task.

1.2.2 Computational learning

In computational learning, learning can be defined as learning a target function, or concept, from a limited number of instances, or examples, drawn from that function, or concept; PAC learning (Valiant, 1984). The instances are presented to the learner in an input-output form. Let X be an instance space; a set of all possible input-output pairs over which a target function can be defined. Then, a concept c is a subset of X, $c \subset X$. For instance, if $X = \mathbb{R}$ is an instance space, then the set of all numbers between 0 and $\sqrt{10}$ could define a concept, e.g., an interval problem. A class of target concepts C can then be defined as a set of sub-sets over X. Consider a number of instances randomly drawn following an unknown distribution D over X, and that each instance is labelled as 1 or 0 respectively whether it belongs to a given concept c or not. If a randomly selected number belongs in the interval $[0, \sqrt{10}]$, then $c(x) = 1$, otherwise $c(x) = 0$. Subsequently, these observations along with their labels are presented to the learner. The set of these examplary observations is known as training set. The task of the learner then is to find an approximation of the given target function.

To do so, a hypothesis space is considered. Generally, in learning theory, the hypothesis space refers to a family of functions, or concepts, $H = \{h_\theta : \theta \in \Theta\}$, where Θ corresponds to the set of all possible parameter vectors of the respective class. After experiencing all of the training examples, the learner tries to find a hypothesis h that best approximate the target function. To evaluate how good a hypothesis is at approximating the target concept, a new set of instances is drawn following the distribution D used for the generation of training examples. The learner then has to guess the target values based on the hypothesis h. The general idea is that the closer the answers are to the target values for a given hypothesis h, the closer that hypothesis is to the target concept c. The probability of misclassifying new data in the testing phase is termed as true error. The set of the new observations the true error is evaluated upon during the testing phase is known as test set. The goal of learning is to minimise the true error and attain good generalisation performance. Yet, given that there is a limited number of training examples, there would potentially be multiple hypotheses that explain the given observations equally well, i.e., the problem is under-determined. It is quite difficult in practice to get the hypothesis with the minimum possible true error. In addition, different training sets provide different representativeness for the target concept. Since it is unrealistic, in practice, to provide the learner with the perfect set of training examples, the expectation of learning the true class is relaxed. Hence, the learner can probably

learn a hypothesis in the training phase that approximates the target function in the testing phase, i.e., Probably Approximately Correct (PAC) learning.

1.2.3 Types of learning

Learning can be broadly classified into the following three categories (Bishop et al., 2006; Hassoun, 1995; Luger, 2005)

- **Supervised Learning**: The learner is provided with a set of labeled training data consisting of input patterns received from the environment along with their desired outputs $\{x, z\}$ given by a teacher. The goal is to infer an appropriate input-output mapping from the training data to approximate the general target (or teacher) function. Ideally, the inferred function should map new unseen inputs to their correct output. Usually, the learning process involves a learning rule that gradually minimises the error between the actual output generated by the learning system and the desired output. This type of learning is also known as *learning with a teacher* and the desired outputs in the training set as supervisory signals.
- **Unsupervised Learning**: The learner is provided with a set of unlabeled training data $\{x\}$. The goal is to infer the hidden similarity structure of the given training data based on some criterion. The idea is that the learner has to discover good representations of the given input patterns capturing potential statistical regularities in the data. This may include a clustering process based on a distance metric, or discovering the statistical structure of the environment using Hebb's rule. The key point here is that there is not an external teacher or error signal during the learning phase.
- **Reinforcement Learning**: The learner dynamically interacts with its environment in order to perform a certain task. This may include learning driving a car, or learning how to play a game by actually playing the game against an opponent. Reinforcement learning is a process of trial and error. The external signal here is evaluative; it does not explicitly inform the learner whether it is close or not in achieving its goal. Accordingly, the external signals do not provide the learner with the correct answer, nor they explicitly correct learner's mistakes during the learning phase. The goal of the learner is to increase the probability of positive reinforcement.

1.2.4 Learning as a search problem

Learning can generally be viewed as a search problem in a multidimensional space for a state that optimises a pre-specified objective function, J (a.k.a. error, cost or criterion

function) (Hassoun, 1995). In neural networks, for example, the search space corresponds to all possible weight vectors. Accordingly, learning is viewed as searching for a vector of weight values that optimises an associated objective function. Note that the nature of learning (e.g., supervised or unsupervised learning) is encapsulated in the objective function J of the given search problem. For instance, in supervised learning problems the error function is designed in such way to measure the error between the actual output produced by the learning system and the respective desired output. Then, the goal of the learner is to minimise that error. Similarly, in Hebbian learning, which is an unsupervised learning problem, the criterion function is to maximisation of the variance of the outputs of a certain unit. In reinforcement learning problems, the goal is to increase the potential of receiving positive reward signals.

Consider $\theta \in \Theta$ represents the vector of the parameter values of the learning system (e.g., the set of weights in neural networks), where Θ corresponds to the set of all possible parameter states. We can then define the objective function as $J = J(\theta, x, z)$, where x indicates an input pattern and z a teacher or reward signal. As described above, the input signals are assumed to be generated from the environment following a distribution D with a probability density $p(x, z)$ in a supervised learning setting, or $p(x)$ in an unsupervised learning setting. Accordingly, the criterion function of the search problem determines the nature of learning based on the involvement of parameter z. If z is absent, then the learning process is unsupervised. The learning signals solely depend on the actual output of the learning system. For instance, consider the unsupervised Hebbian learning for a single linear neuronal unit y. The corresponding learning equation for continuous time is given by $\dot{w} = \rho yx$, and thus $J = -0.5y^2$. If now J incorporates external teacher signals z, then learning can be either supervised or reinforcement whether the signal is instructional or evaluative. The learning signals depend on both the actual output of the learning system and the desired output, or reward signal respectively.

1.2.5 Computational issues

There are two basic practical problems that learning has to deal with; firstly, the low quality of information and secondly, the information overload. The former problem is related to the well-known Plato's parable of the cave. Our sensory organs can capture only low-dimensional inputs, which can be seen as projections of the high-dimensional natural world. Thus, it is an ill-posed problem, since a great amount of information is missing, making it difficult to choose among a vast number of plausible internal models (Marr, 1982). Integrating across many separate experiences can help by representing aggregate statistics over a larger sample of sensory inputs, but how then would we know if anything useful is actually learned (a.k.a. empiricist's dilemma; (Luger, 2005))?. Instead, introducing prior biases that organise the way the incoming information is

processed, indicating what is important and what is not about the world, leads to more reliable and efficient ways of model learning (O'Reilly and Munakata, 2000).

Such designer's specifications, however, require top-down domain knowledge and may raise some more pragmatic issues regarding the inductive bias in learning. If the biases are too strong, the learners will not learn much from their experience. On the contrary, if the biases are too weak, learners will generate a great variability of models, reflecting the idiosyncrasies of their individual experience (high variance). This trade-off is well-described by the bias-variance dilemma (Geman et al., 1992) and the no-free-lunch theorem (Wolpert and Macready, 1995).

In order to deal with the second problem of learning, biases are usually introduced to favour simple parsimonious explanations. To put it simply, the most plausible hypotheses (least assumption) are selected — *Occam's razor*. In theory, regularisation can be justified as imposing Occam's razor on the produced solutions, by penalising the complexity of the models. As will be seen, the development of such parsimonious models is highly desirable, as it gives rise to enhanced generalisation.

1.2.6 Generalisation: The problem of induction

Generalisation from experience is a crucial part of learning (Luger, 2005). The problem of induction is directly related to the aforementioned computational issues. Ideally, the learner should not only improve its performance on the next encounter of an instance experienced before, but also on similar instances, or tasks, in the same task domain. Yet interesting domains tend to be enormous, and thus the learner is provided with a fraction of all possible instances or the class. The learner then has to generalise from that limited experience to new unseen instances drawn from the same class; the problem of induction. In most problem scenarios, the training examples presented are insufficient to guarantee perfect, or optimal, generalisation, i.e., the discovery of the hypothesis, h, with the minimum true error. The learning system must generalise by selecting those aspects and features of its past experience that can be directly relevant to and useful for the future. The corresponding selection criteria are generally known in machine learning as *inductive biases*.

1.2.7 The Problem of Overfitting: Fitting the noise

In statistical learning theory, overfitting is observed when the learning model is over-trained to memorise a particular set of exemplary input-outputs pairs (training set), at the expense of predictive performance on unseen data generated by the same teacher machine (test set) (Bishop et al., 2006). To put simply, overfitting occurs when the model starts fitting the noise (Abu-Mostafa et al., 2012).

Why is *overfitting* a problem? The problem arises when the derived statistical model describes the stochastic error of the data, and not the actual underlying problem structure. Generally in machine learning, the ultimate goal of the learner is to accurately predict desired outputs for *new* observations, i.e., not just the exemplary ones in the training set. To do so, the learning machine should be capable of inferring the underlying relationships between the data – in other words, to *generalise* to unseen observations. Consequently, overfitting opposes to the generalisation ability of the learner, and hence *generalisation can be enhanced by preventing overfitting.*

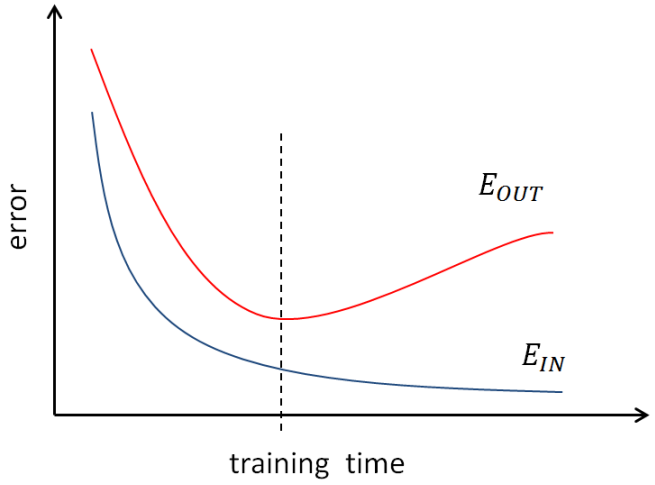

FIGURE 1.1: Overfitting in (supervised) learning. The blue line indicates the in-sample error (J_{IN}), while the red line the out-of-sample error (J_{OUT}) versus the training time. Overfitting occurs when the red line exhibits a positive slope. The best model can be obtained at the global minimum of the validation error (dashed line).

In learning theory, the hypothesis space refers to a family of functions, $H = \{h_\theta : \theta \in \Theta\}$, where Θ corresponds to the set of all possible parameter vectors of the respective family. Then, a hypothesis h_θ over-fits the training set, whenever another $h_{\theta'}$ exists, with lower empirical error in the test set. A practical way to detect overfitting is to simply plot the in-sample (J_{IN}) and out-of–sample errors (E_{OUT}) as training progresses. J_{IN} (training error) indicates how well the learner can describe the given exemplary data in the training set, and it monotonically decreases with training, since gradient descent to E_{IN} algorithms are typically used. J_{OUT} (validation error) initially exhibits similar behaviour to J_{IN}. Overfitting is detected whenever J_{OUT} starts increasing with training (Figure 1.1).

The problem of overfitting is a common issue in neural networks, not because of stochastic noise in the data, but because there too many parameters to be adjusted (*NB* that each weight is a parameter for the neural model) to a small amount of observed data. Again, we try to fit the noise, but a different type of noise this time – known as deterministic noise (Abu-Mostafa et al., 2012). Deterministic noise is present when the complexity of the model we try to fit is excessively high compared to the complexity of

our data (e.g., the number of available data). Choosing large hypothesis spaces, powerful enough to capture complex phenomena is convenient to accurately fit the model on the given data. However, the data values *per se* often cannot provide the information needed for the learner to generalise well to out-of-samples data, resulting in overfitting. In general, overfitting is reduced with the number of available data. Therefore, in the case of infinite number of training data, overfitting would not be a problem.

The problem of overfitting can be considered as a problem of model selection. There are two main approaches to deal with this problem; i) regularisation and ii) validation. The former concerns penalising the complexity of the model, introducing useful constraints (e.g., shrinkage methods), while the latter concerns testing the predictive accuracy of the derived statistical model in practice (e.g., cross-validation methods). More sophisticated techniques regarding model selection have been proposed by (O'Reilly, 1996), using a mixture of Hebbian (model) learning (conditional Hebbian-PCA; CPCA) and competitive learning (k-winner-takes-all; kWTA), known as self-organisation learning.

1.2.8 Regularisation: Putting the breaks

A way to control the generalisation performance of a learning system on a task is through conditioning the associated criterion function. As described above, the problem of overfitting is the problem of fitting the noise. Both stochastic and deterministic noise are generally characterised by rapid non-smooth fluctuations. Therefore, a practical technique to avoid this is to introduce biases towards simpler hypotheses, h_θ. Limiting the complexity of the model is a common technique used in the area of machine learning, reducing the excessive variance of the model that can lead to over-fitting at a cost of an increase in bias which may cause under-fitting when it is excessive (bias-variance trade-off). The regularisation term embeds extra information in the criterion function about what is desired and what is not (e.g., smoothness or invariance) by introducing additional constraints in order to prevent over-fitting. In situations of unsupervised learning, the regularisation term can be viewed as a constraint-satisfaction term. For example, regularisation in Hebbian learning can provide stability.

In mathematics and statistical learning theory, regularisation is generally applied by introducing an additional term in the original objective function that favours desirable solutions. Therefore, the general form of this *augmented error* is given by (Abu-Mostafa et al., 2012):

$$J_{AUG} = J_{IN} + \lambda\Omega, \tag{1.1}$$

where J_{IN} is the initial criterion function and $\Omega : \theta \rightarrow R$ corresponds to a regulariser, rewarding or penalising the respective model, characterised by the parameter vector,

θ. The factor λ is a regularisation constant, which plays the role of control parameter. Usually, λ parameter is optimally tuned by validation techniques, such as cross-validation (Kohavi et al., 1995). Typical regularisers used in machine learning include L1– (LASSO method) and L2-norms (ridge regression method; (Tibshirani, 1996)), or a combination of the two (elastic net regularisation; (Zou and Hastie, 2005)). A prominent regulariser used in neural networks is the *weight decay* regulariser:

$$\Omega(\theta) = \frac{1}{2}\sum_{i=1}^{N}\theta_i^2, \tag{1.2}$$

where θ_i parameter correspond to the weights of the network. The regularisation parameter determines the *weight decay rate*. This additional term favours smaller absolute weights. To see that, consider that we want to minimise the augmented error, J_{AUG}, following a gradient descent method. Then, the weight update rule is given by:

$$w(t+1) = w(t) - \epsilon\nabla J_{IN}(w(t)) - \epsilon\lambda w(t). \tag{1.3}$$

Therefore, the weights decay towards zero proportionally to their size. All of the aforementioned techniques have the tendency to favour small values of parameters and are known as *shrinkage methods*. These methods are useful, as they penalise large weights which can lead to discontinuous input-output mappings and excessive variance in the generated outputs.

1.2.9 Early Stopping

An ad hoc solution to preventing overfitting is that of early stopping. The training set is typically divided into two sets; one used for training and one for validation. Accordingly, the performance of the network on the validation set is off-line tested, by evaluating the respective error over the validation set, J_{VAL}. The algorithm is halted, when J_{VAL} starts rising. To put simply, it stops when it detects that overfitting starts (Figure 1.1). The actual generalisation ability of learner is then tested on a separate set (test set). Although this is less principled technique, it is a form of regularisation and generally works well in practice (Bishop et al., 2006).

1.2.10 But is evolution the same as learning?

Valiant (2013) presents a simple and intuitive parable to address this question. Consider a species of monkeys having access to bananas, apples and berries. Suppose that a species of berries with bad taste invades the area the monkeys live in. When a monkey experiences eating this type of berries for the first time, the monkey will directly learn

to avoid eating them in the future by removing them from its dietary list. Suppose now the extreme case that the berries were lethal. Could evolution by natural selection do the same? The answer is *not exactly*. If a monkey dies prematurely because of these lethal berries, the genomes of its decedents will not be *immediately* amended to avoid eating berries. According to Darwin's theory, the produced offspring will have a chance to be mutated though and not having the lethal berries in their food preferences as their parents. Eventually, monkeys with this (beneficial) mutation will survive and persist through evolutionary time. Ultimately, the outcome in both of the cases is the same. The monkeys in later generations will learn to stop eating the lethal berries. This evolutionary kind of learning, however, is less direct as a consequence of the circular Darwinian feedback mechanism.

Valiant concludes that this form of learning is not as powerful as PAC learning, since there are classes of problems that are not learnable in the Darwinian sense. In Darwinian evolution, for instance, new variants are generated irrespectively of their parents' experiences during their life-time. Experience, however, affects the fitness of the current individual and thus its reproductive success. Another limitation of Darwinian's evolution as a process of learning is that, in contrast to traditional learning theory, it has to work and succeed with whatever hypothesis (genome) currently has. Undoubtedly, re-initialising is very convenient, but not biologically plausible. This is restrictive, but we also know from machine learning that there are classes of algorithms that successfully work from any starting point, e.g., the well-known perceptron learning (Rosenblatt, 1957).

Although the example with the monkeys makes some good and valid points about the differences between learning and evolution, it does not make a clear distinction between phenotypic and genetic learning. The former deals with the ability of an organism to learn and cognise throughout its lifespan (e.g., monkeys learning based on the taste of the food), whereas the latter deals with the ability of a population to adapt to varying selective environment through genetic mutation (e.g., change in its variational mechanisms).

1.2.11 Evolution of development as learning.

To see why evolution of development can be a form of learning, we view the part of the genome that describes the process of development as the corresponding hypothesis in computational learning. Then, the direct genetic effects on the embryonic phenotype along with the cues received from the environment play the role of the inputs, and the adult developed phenotype plays the role of the output. Likewise, the target phenotypes determined by the selective environments correspond to the desired outputs. Consequently, the evolution of the regulatory interactions in a fitness maximisation

manner by natural selection is directly analogous to a learning process by error minimisation, namely, searching the hypothesis space for better hypotheses that generate outputs closer to the desired ones.

1.3 Aims and objectives

The main aim of this thesis is to demonstrate how organisms can systematically evolve designs that enhance their evolutionary potential for future adaptation. Specifically, I study how organisms can *predispose* the production of suitable phenotypes which would help them cope with environmental change within and across generations. This could occur through either mutation or environmental induction. The origins and the conditions of such adaptive potential however remain elusive and is a long-standing quest in evolutionary biology. In this work I support the hypothesis that such adaptive capacity can arise as an epiphenomenon of past selection towards target optima in different selective environments. I attempt to demonstrate that the analogy between learning and evolution can provide valuable insights on the mechanistic nature of the evolution of adaptive potential and help us characterise general evolutionary conditions promoting it.

The key point here is that the selective environments do not change randomly, but in a systematic way (i.e., share common regularities that are invariant over time) – something which is ubiquitous in natural environments (e.g., natural laws or general characteristics) (Kashtan and Alon, 2005; Kashtan et al., 2007). Specifically, the environmental goal the organisms have to face when the selective environment changes shares similar sub-goals with the previous ones. For instance, goals such as digesting food, breathing and reproducing must be satisfied by the organisms in new selective environments but in different degrees or combinations. Thus, previous adaptations can potentially carry information from past selective environments that is directly relevant to future environment conditions, when they share a common 'language' as the past ones. I argue that the ability of evolution to discover and exploit such regularities is analogous to the ability of learning systems to generalise from limited past experience. Prediction in machine learning merely requires the ability to represent structural regularities in previously seen observations that are also true in the yet-unseen ones. In learning systems, such generalisation ability is neither mysterious nor for granted. It is not really about the past or the future, but about generalising from the data you have seen to the test cases you have not.

Firstly, I will characterise evolutionary conditions that enable and even enhance the evolution of evolvability relying on insights from learning theory. Specifically, I will assess how evolving developmental biases shape phenotypic variability and thus future

phenotypic exploration in yet-unseen environments through genetic change. By performing computer simulations of a gene regulatory network models I will investigate how evolution can find an appropriate context-dependent language to describe multiple varying selective environments. The simulations will specifically address the following objectives:

- show that the canalisation of evolved developmental processes to past or current selective environments and failure of natural selection to enhance evolvability in future selective environments is directly analogous to the problem of under- and over-fitting and failure to generalise in learning systems.
- show that the same conditions that alleviate under- and over-fitting in learning systems also enhance the evolution of evolvability.

Secondly, I aim to demonstrate how adaptive plasticity can systematically arise without the need of direct or lineage selection but rather as a by-product of selection within an intrinsically varying environment. Emphasis here will be given to how adaptive evolutionary potential can increase through plasticity. By performing computer simulations of a developmental network model, I intend the following:

- show that adaptive plasticity can systematically evolve in cases where plasticity is not expected to increase in situations when plasticity is costly and not needed in any individual.
- show that the evolution of adaptive plasticity under inter-generational environmental variability is favoured when the population experiences rapid environmental fluctuations or is characterised by low mutation rates.
- using learning theory, show that the plausibility of adaptive plasticity to arise depends on the expected adaptive change accumulated at each given environment.
- utilise learning theory to characterise general conditions under which natural selection favours developmental organisations which facilitate adaptive plastic responses under inter- and intra-generational environmental variability.

Thirdly, I intend to show that the plasticity-first hypothesis (West-Eberhard, 2003) can be better understood from a learning theory perspective. Specifically, I focus on the role of plasticity in shaping potential phenotypic variation and facilitating innovation in yet-unseen environments. According to the plasticity first hypothesis, novel traits can be initiated by environmental change and thus plasticity may precede and even guide adaptation. This will be addressed through the following:

- reviewing the theoretical and conceptual roadblocks of the 'plasticity-first' hypothesis within the existing framework of evolutionary theory.
- showing that the plasticity-first hypothesis and the evolution of evolvability face the same conceptual roadblocks, and as such the scientific problem domain for plasticity-first hypothesis is the evolution of evolvability.
- explaining how learning theory helps to overcome the conceptual issues of the plasticity-first hypothesis.
- characterise general conditions under which plasticity can potentiate and guide adaptive evolution.

1.4 Thesis structure

The thesis follows the three paper format. Each chapter is written as a standalone and self-contained manuscript. While the three chapters address different, on the surface, problems in the field of evolutionary biology, all three studies share common themes and methodologies. The research findings from the three chapters contribute towards a broader conceptual framework for the evolution of adaptive potentials. The first chapter focuses on genetic change, whereas the last two chapters focus on environmental change as potent sources of useful phenotypic variation.

The first two chapters involve a series of computational simulations to reproduce the behaviour of evolving systems under certain assumptions of interest. All the necessary information to reproduce the results in each chapter is included in the main text of the respective chapter. Additionally, the source code used for Chapter 1 and 2 can be found here: https://github.com/KostasKouvaris/.

1.4.1 Paper 1: How evolution learns to generalise: Using the principles of learning theory to understand the evolution of developmental organisation.

In this work, I study the evolution of adaptive developmental biases that enhance organismic evolvability. Specifically, I investigate how adaptive (directed) phenotypic change arises through (undirected) genetic variation.

Recent studies demonstrate that the evolution of developmental constraints might enhance evolvability in certain cases by facilitating the exploration and discovery of useful phenotypes (Gerhart and Kirschner, 2007; Kirschner and Gerhart, 2006; Parter et al., 2008; Kashtan and Alon, 2005; Kashtan et al., 2009; Clune et al., 2013b; Watson et al.,

2014). Yet, these examples currently lack a general theoretical framework that provides testable predictions of the conditions that enhance evolvability in novel environments. This is because evolvability seems to be concerned with future fitness benefits and thus its evolution by means of natural selection appears to be incompatible with the short-sighted concept of natural selection. This conceptual roadblock has prevented the development of a predictive theoretical framework for the evolution of evolvability.

A potential resolution is provided by learning theory. Evolving systems and learning systems share common underlying, conceptual and algorithmic, principles (Watson and Szathmáry, 2016). If this link is more than an analogy, then we expect that existing results from the learning domain can be transferred to the evolution domain; opening-up a well-developed theoretical framework and enabling it to be exploited in evolutionary theory. Here, I test this analogy between learning and evolution by verifying its predictions. Specifically, I show that the way natural selection fails to enhance evolvability is directly analogous to the way learning systems fail to generalise to new instances. If the hypothesis holds, then we expect that the same conditions that improve generalisation in learning systems will also enhance the evolution of generalised phenotypic distribution under natural selection, and thus the evolution of evolvability. Such functional equivalence between learning and evolution sheds light on how short-sighted natural selection can produce novel phenotypes fit for previously-unseen selective environments by equating evolvability with generalisation in learning systems.

1.4.2 Paper 2: How can adaptive plasticity reliably evolve when plasticity is not selected for?

This work aims to understand how developmental systems can facilitate potential useful phenotypes that are conditionally-expressed in response to environmental change, i.e., adaptive plasticity. Specifically, I study how organisms can evolve appropriate environmental sensitivity which helps them cope with environmental fluctuations when plasticity is directly selected for.

Despite the plethora of theoretical studies, the role of plasticity in facilitating adaptive evolution remains a contentious issue. In nature, organisms encounter environmental changes across and within generations. The environments encountered by a population dictate which phenotypes are realised by the individuals along with their fitness consequences. When organisms experience multiple environments over their lifetime, i.e., fine-grained environmental variability, plasticity is expected to be evolutionarily advantageous, and thus selected for. Yet, when organisms experience one environment in their lifetime, i.e., coarse-grained environments, it is not clear how adaptive plasticity arises. Theoretical studies suggest that the selective pressures that maximise fitness to the local habitat of the population leads to adaptive genetic differentiation in when

within-generation environmental variation is absent. It is thus often expected for natural selection to favour plastic genotypes in fine-grained environmental variability and genotypic differentiation in coarse-grained environmental variability.

In this study, I investigate how adaptive plasticity can arise without a direct selection pressure. When retaining adaptive phenotypic plasticity in the short-term is unnecessary or even costly, then natural selection is expected to favour a decrease in plasticity. However I demonstrate that adaptive plasticity can reliably increase without lineage selection in situations where there is no direct selection *for* plasticity, or plasticity is even selected against, under certain circumstances. Specifically, I show that there is a systematic selection *of* plasticity in coarse-grained environment when the population experiences rapid environmental fluctuations or is characterised by low mutation rates. I argue that environmental sensitivity arises in such situation as a by-product of selection towards phenotypic targets in different selective environments across generations. The plausibility of adaptive plasticity to arise depends on the expected adaptive change accumulated at each given environment. I argue that such sensitivity of plasticity to the amount of adaptation is analogous to the sensitivity of prediction to learning rates in learning systems. Lastly, I outline how learning theory can provide useful insights on the evolution of plasticity and allows us to characterise general conditions that can facilitate the evolution of adaptive plasticity under natural selection.

1.4.3 Paper 3: Extending plasticity-first evolution with learning theory

In this chapter, I investigate whether plasticity can precede and even enhance the process of adaptive evolution, i.e., plasticity-first hypothesis. Firstly, I review the theoretical and conceptual roadblocks of the plasticity-first hypothesis within the existing framework of evolutionary theory. Then, I propose a new perspective on the plasticity-first hypothesis that help us overcome these issues transferring knowledge from the domain of learning theory.

Although the plasticity-first hypothesis has a long history in evolutionary biology, it still remains a contentious issue. Traditional views suggest that developmental plasticity buffers against environmentally induced variation and thus hinders adaptive evolution. Even if novel traits start as conditionally expressed alternative phenotypes, plasticity is of a little theoretical importance as the realised phenotypic variation tends to be undirected and does not provide any new insights into how evolution works. The existing models of plasticity tend to under-represent development and diminish the role of development in adaptive evolution. Instead, I argue that explaining whether plasticity can precede or even facilitate adaptive evolution requires an understanding on how development can reliably predispose the production of adaptive phenotypic variants in new environments. As such, the plausibility of plasticity-first hypothesis relies on the same

conceptual arguments as the evolution of evolvability. The discovery of new and useful phenotypes does not solely depends on the amount, but also the direction of the potential environmentally-induced phenotypic variation. Past selected developmental biases can constrain future evolutionary exploration into more promising phenotypic regions. Here I show how adaptive norms of reaction can be shaped by past selection and can be better understood from computational learning perspective. This allows us to characterise general conditions under which plasticity shapes and potentiate adaptation in new environments.

Chapter 2

How evolution learns to generalise: Using the principles of learning theory to understand the evolution of evolvable developmental organisation.

Abstract

One of the most intriguing questions in evolution is how organisms exhibit suitable phenotypic variation to rapidly adapt in novel selective environments. Such variability is crucial for evolvability, but poorly understood. In particular, how can natural selection favour developmental organisations that facilitate adaptive evolution in previously unseen environments? Such a capacity suggests foresight that is incompatible with the short-sighted concept of natural selection. A potential resolution is provided by the idea that evolution may discover and exploit information not only about the particular phenotypes selected in the past, but their underlying structural regularities: new phenotypes, with the same underlying regularities, but novel particulars, may then be useful in new environments. If true, we still need to understand the conditions in which natural selection will discover such deep regularities rather than exploiting 'quick fixes' (i.e. fixes that provide adaptive phenotypes in the short term, but limit future evolvability). Here we argue that the ability of evolution to discover such regularities is formally analogous to learning principles, familiar in humans and machines, that enable generalisation from past experience. Conversely, natural selection that fails to enhance evolvability is directly analogous to the learning problem of over-fitting and the subsequent failure to generalise. We support the conclusion that evolving systems and learning systems

are different instantiations of the same algorithmic principles by showing that existing results from the learning domain can be transferred to the evolution domain. Specifically, we show that conditions that alleviate over-fitting in learning systems successfully predict which biological conditions (e.g., environmental variation, regularity, noise or a pressure for developmental simplicity) enhance evolvability. This equivalence provides access to a well-developed theoretical framework from learning theory that enables a characterisation of the general conditions for the evolution of evolvability.

2.1 Linking the Evolution of Evolvability with Generalisation in Learning Systems

Explaining how organisms adapt in novel selective environments is central to evolutionary biology (Bedau et al., 2000; Adami et al., 2000; Lenski et al., 2003; Bedau, 2009; Moczek et al., 2011). Living organisms are both robust and capable of change. The former property allows for stability and reliable functionality against genetic and environmental perturbations, while the latter provides flexibility allowing for the evolutionary acquisition of new potentially adaptive traits (Wagner and Altenberg, 1996; Conrad, 1979; Kirschner and Gerhart, 1998; Schlichting and Murren, 2004; Moczek et al., 2011). This capacity of an organism to produce suitable phenotypic variation to adapt to new environments is often identified as a prerequisite for *evolvability*, i.e. the capacity for adaptive evolution (Conrad, 1972, 1979; Pigliucci, 2008). It is thus important to understand the underlying variational mechanisms that enable the production of adaptive phenotypic variation (Riedl and Jefferies, 1978; Wagner and Altenberg, 1996; Conrad, 1979; Altenberg, 1995; Toussaint, 2002; Brakefield, 2006; Gerhart and Kirschner, 2007; Toussaint and von Seelen, 2007; Braendle et al., 2010).

Phenotypic variations are heavily determined by intrinsic tendencies imposed by the genetic and the developmental architecture (Smith et al., 1985; Conrad, 1998; Yampolsky and Stoltzfus, 2001; Braendle et al., 2010). For instance, developmental biases may permit high variability for a particular phenotypic trait and limited variability for another, or cause certain phenotypic traits to co-vary (Wagner and Altenberg, 1996; Hansen, 2003; Brakefield, 2006; Pavlicev et al., 2010; Pavlicev and Hansen, 2011; Watson et al., 2014; Pavličev and Cheverud, 2015). Developmental processes are themselves also shaped by previous selection. As a result, we may expect that past evolution could adapt the distribution of phenotypes explored by future natural selection to amplify promising variations and avoid less useful ones by evolving developmental architectures that are predisposed to exhibit effective adaptation (Conrad, 1972; Altenberg, 1995). Selection though cannot favour traits for benefits that have not yet been realised. Moreover, in situations when selection can control phenotypic variation, it nearly always reduces such variation because it favours canalisation over flexibility (Clune et al., 2013b,a; Wagner et al., 2007; Pavlicev et al., 2010).

Developmental canalisation may seem to be intrinsically opposed to an increase in phenotypic variability. Some, however, view these notions as two sides of the same coin, i.e., a predisposition to evolve some phenotypes more readily goes hand in hand with a decrease in the propensity to produce other phenotypes (Kirschner and Gerhart, 1998; Brigandt, 2007; Draghi et al., 2010). Kirschner and Gerhart integrated findings that support these ideas under the unified framework of *facilitated variation* (Kirschner and Gerhart, 1998, 2006). Similar ideas and concepts include the *variational properties* of the organisms (Altenberg, 1995), the *self-facilitation* of evolution (Conrad, 1998) and evolution as *tinkering* (Jacob, 1977) and related notions (Conrad, 1972; Riedl and Jefferies, 1978; Conrad, 1979; Wagner and Altenberg, 1996). In facilitated variation, the key observation is that the intrinsic developmental structure of the organisms biases both the amount and the direction of the phenotypic variation. Recent work in the area of facilitated variation has shown that multiple selective environments were necessary to evolve evolvable structures (Parter et al., 2008; Kashtan and Alon, 2005; Kashtan et al., 2007; Watson et al., 2014; Clune et al., 2013b). When selective environments contain underlying structural regularities, it is possible that evolution learns to limit the phenotypic space to regions that are evolutionarily more advantageous, promoting the discovery of useful phenotypes in a single or a few mutations (Kashtan and Alon, 2005; Kashtan et al., 2007). But, as we will show, these conditions do not necessarily enhance evolvability in novel environments. Thus the general conditions which favour the emergence of adaptive developmental constraints that enhance evolvability are not well-understood.

To address this we study the conditions where evolution by natural selection can find developmental organisations that produce what we refer to here as *generalised phenotypic distributions* — i.e., not only are these distributions capable of producing multiple distinct phenotypes that have been selected in the past, but they can also produce novel phenotypes from the same family. Parter et al. have already shown that this is possible in specific cases studying models of RNA structures and logic gates (Parter et al., 2008). Here we wish to understand more general conditions under which, and to what extent, natural selection can enhance the capacity of developmental structures to produce suitable variation for selection in the future. We follow previous work on the evolution of development (Watson et al., 2014) through computer simulations based in gene-regulatory network (GRN) model. Many authors have noted that GRNs share common functionality to artificial neural networks (Wagner, 1996; Vohradský, 2001a,b; Watson et al., 2014; Fierst and Phillips, 2015). Watson et al. demonstrated a further result, more important to our purposes here; that the way regulatory interactions *evolve* under natural selection is mathematically equivalent to the way neural networks *learn* (Watson et al., 2014). During evolution a GRN is capable of learning a memory of multiple phenotypes that were fit in multiple past selective environments by internalising their statistical correlation structure into its ontogenetic interactions, in the same way that learning neural networks store and recall training patterns. Phenotypes that were

fit in the past can then be recreated by the network spontaneously (under genetic drift without selection) in the future or as a response to new selective environments that are partially similar to past environments (Watson et al., 2014). An important aspect of the evolved systems mentioned above is modularity. Modularity has been a key feature of work on evolvability (Wagner and Altenberg, 1996; Wagner et al., 2007; Lipson et al., 2002; Watson et al., 2015) aiming to facilitate variability that respects the natural decomposable structure of the selective environment, i.e., keep the things together that need to be kept together and separate the things that are independent (Wagner and Altenberg, 1996; Lipson et al., 2002; Conrad, 1998; Riedl and Jefferies, 1978). Accordingly, the system can perform a simple form of generalisation by separating knowledge from the context in which it was originally observed and re-deploying it in new situations.

Here we show that this functional equivalence between learning and evolution predicts the evolutionary conditions that enable the evolution of generalised phenotypic distributions. We test this analogy between learning and evolution by testing its predictions. Specifically, we resolve the tension between canalisation of phenotypes that have been successful in past environments and anticipation of phenotypes that are fit in future environments by recognising that this is equivalent to prediction in learning systems. Such predictive ability follows simply from the ability to represent structural regularities in previously seen observations (i.e., the training set) that are also true in the yet-unseen ones (i.e., the test set). In learning systems, such generalization is commonplace and not considered mysterious. But it is also understood that successful generalisation in learning systems is not for granted and requires certain well-understood conditions. We argue here that understanding the evolution of development is formally analogous to model learning and can provide useful insights and testable hypotheses about the conditions that enhance the evolution of evolvability under natural selection (Watson and Szathmáry, 2015a; Watson et al., 2015). Thus, in recognising that learning systems do not really 'see into the future' but can nonetheless make useful predictions by generalising past experience, we demystify the notion that short-sighted natural selection can produce novel phenotypes that are fit for previously-unseen selective environments and, more importantly, we can predict the general conditions where this is possible. This functional equivalence between learning and evolution produces many interesting, testable predictions (Table 2.1).

In particular, the following experiments show that techniques that enhance generalisation in machine learning correspond to evolutionary conditions that facilitate generalised phenotypic distributions and hence increased evolvability. Specifically, we describe how well-known machine learning techniques, such as learning with noise and penalising model complexity, that improve the generalisation ability of learning models have biological analogues and can help us understand how noisy selective environments and the direct selection pressure on the reproduction cost of the gene regulatory interactions can enhance evolvability in gene regulation networks. This is a much more sophisticated and

powerful form of generalisation than previous notions that simply extrapolate previous experience. The system does not merely extend its learned behaviour outside its past 'known' domain. Instead, we are interested in situations where the system can create new knowledge by discovering and systematising emerging patterns from past experience, and more notably, how the system separates that knowledge from the context in which it was originally observed, so that it can be re-deployed in new situations.

Some evolutionary mechanisms and conditions have been proposed as important factors for improved evolvability. Some concern the modification of genetic variability (e.g., Kashtan et al. (2007); Friedlander et al. (2013); Livnat (2013) and Livnat et al. (2008)), while others concern the nature of selective environments and the organisation of development including multiple selective environments (Kashtan et al., 2007), sparsity (Aldana et al., 2007), the direct selective pressure on the cost of connections (which can induce modularity (Clune et al., 2013b; Friedlander et al., 2013) and hierarchy (Mengistu et al., 2016)), low developmental biases and constraints (Arthur, 2006) and stochasticity in GRNs (MacNeil and Walhout, 2011). In this paper, we focus on mechanisms and conditions that can be unified and better understood in machine learning terms, and more notably, how we can utilise well-established theory in learning to characterise general conditions under which evolvability is enhanced. We thus provide the first theory to characterise the general conditions that enhance the evolution of developmental organisations that generalise information gained from past selection, as required to enhance evolvability in novel environments.

Table 2.1: **Predictions Made By Porting Key Lessons of Learning Theory to Evolutionary Theory.** Confirmed by experiment: † Conditions that Facilitate Generalised Phenotypic Distributions, ‡ How Generalisation Changes over Evolutionary Time, ◇ Rate of Environmental Switching (Learning Rates) and ⋆ Sensitivity Analysis to Parameters Affecting Phenotypic Generalisation.

	Learning Theory	**Evolutionary Theory**
(a)	Generalisation; ability to produce an appropriate response to novel situations by exploiting regularities observed in past experience (i.e., not rote learning).	Facilitated variation; predisposition to produce fit phenotypes in novel environments (i.e., not just canalisation of past selected targets).†
(b)	The performance of online learning algorithms (i.e., processing one training example at a time) are learning-rate dependent. Both high and low learning rates can lead to situations of under-fitting; failure of the learning system to capture the regularities of the training data (Bishop et al., 2006).	The evolution of generalised phenotypic distributions is dependent on the time-scale of environmental switching. Both high and low time-scales can lead to inflexible developmental structures that fail to capture the functional dependencies of the past phenotypic targets.◇
(c)	The problem of over-fitting: improved performance on the training set comes at the expense of generalisation performance on the test set. Over-fitting occurs when the model learns to focus on idiosyncrasies or noise in the training set (Abu-Mostafa et al., 2012). Accordingly, the model starts learning the particular irrelevant relationships existing in the training examples rather than the 'true' underlying relationships that are relevant to the general class. This leads to memorisation of specific training examples, which decreases the ability to generalize, and thus perform well, on new data.	Failure of natural selection to evolve generalised developmental organisations: improved average fitness gained by decreasing the phenotypic variation of descendants comes at the expense of potentially useful variability for future selective environments. Favouring immediate fitness benefits would lead to robust developmental structures that canalise the production of the selected phenotypes in the current selective environment. Yet, this sets up a trade-off between robustness and evolvability, since natural selection would always favour inflexible developmental organisations that reduce phenotypic variability and thus hinder the discovery of useful phenotypes that can have fitness benefits in the future.‡
(d)	Conditions that alleviate the problem of over-fitting: (1) training with noisy data, i.e., adding noise during the learning phase (jittering), (2) regularisation (parsimony pressure), i.e., introducing a connection cost term into the objective function that favours connections of small values (L_2-regularisation) or fewer connections (L_1-regularisation).	Evolutionary conditions that facilitate the evolution of generalised phenotypic distributions, and thus evolvability: (1) extrinsic noise in selective environments, (2) direct selection pressure on the cost of ontogenetic interactions, which favour simpler developmental processes and sparse network structures.†‡
(e)	L_2-regularisation results in similar behaviour as early stopping; an ad-hoc technique that prevents over-fitting by stopping learning when over-fitting begins (Bishop et al., 2006).	Favouring weak connectivity via connection costs results in similar behaviour as stopping adaptation at an early stage.†‡.
(f)	Training with noise results in similar behaviour to L_2-regularisation (Bishop et al., 2006).	Noisy environments can enhance the evolution of generalised developmental organisation in a similar manner as favouring weak connectivity.†‡.
(g)	Generalisation performance is dependent on the appropriate level of regularisation and the level of noise, i.e., it depends on the inductive biases, or prior assumptions about which models are more likely to be correct, such as a priori perference for simple models via parsimony pressures.	The evolution of generalised phenotypic distributions is dependent on the strength of selection pressure on the cost of connections and the level of environmental noise.⋆
(h)	L_1-regularisation results in better generalisation performance than L_2-regularisation in problems with simple modularity/independent features.	Favouring sparsity results in more evolvable developmental structures than favouring weak connectivity for modularly varying environments with weak or unimportant inter-modular dependencies.†‡

2.2 Experimental Setup

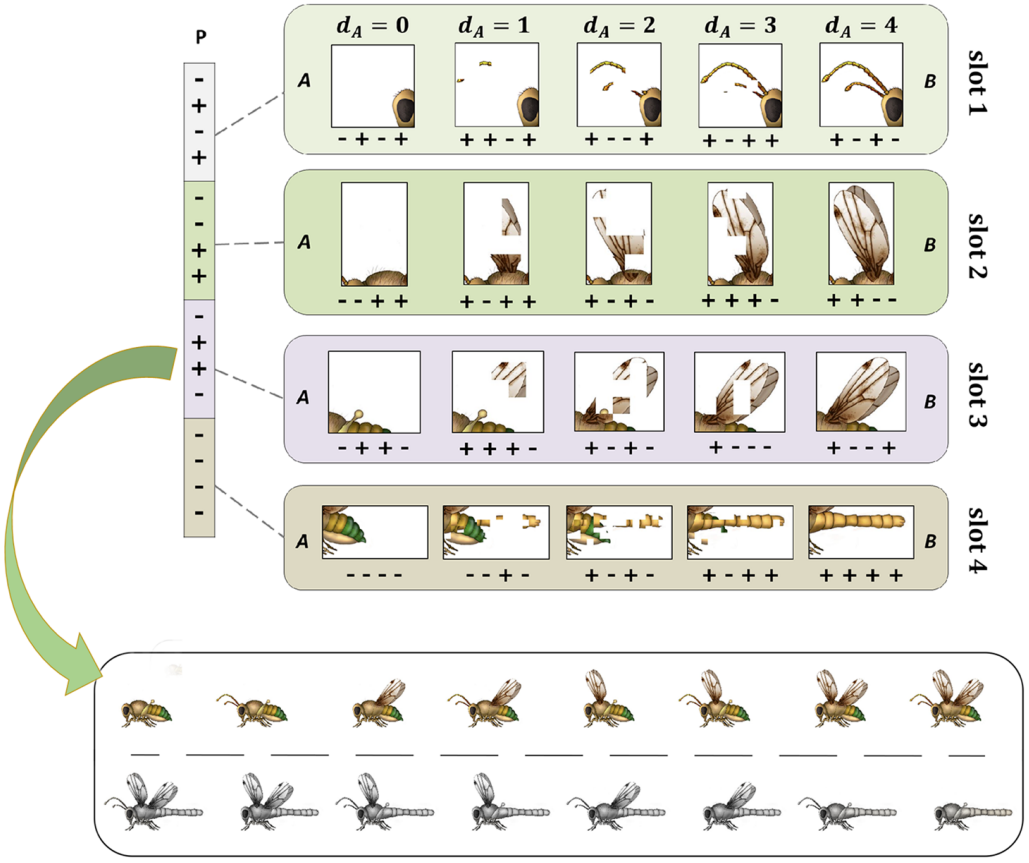

FIGURE 2.1: **Pictorial representation of phenotypes.** (Top) Schematic representation of mapping from phenotypic pattern sequences onto pictorial features. Each phenotypic 'slot' represents a set of features (here 4) controlling a certain aspect of the phenotype (e.g., front wings, halteres and antennae). Within the possible configurations in each slot (here 16), there are two particular configurations (state A and B) that are fit in some environment or another (see Developmental Model in S1 Appendix A). For example, '$++--$' in the second slot (from the top, green) of the phenotypic pattern encodes for a pair of front wings (state B), while '$--++$' encodes for their absence (state A). States A and B are the complement of one another, i.e., not neighbours in phenotype space. All of the other intermediate states (here 14) are represented by a random mosaic image of state A and B, based on their respective distance. d_A indicates the Hamming distance between a given state and state A. Accordingly, there exist $\binom{4}{d_A}$ potential intermediate states (i.e., 4 for $d_A = 1$, 6 for $d_A = 2$ and 4 for $d_A = 3$). (Bottom) Pictorial representation of all phenotypes that are perfectly adapted to each of eight different environments. Each target phenotype is analogous to an insect-like organism comprised of 4 functional features. The grey phenotypic targets correspond to bit-wise complementary patterns of the phenotypes on the top half of the space. For example, in the rightmost, top insect, the antennae, forewings, and hindwings are present, and the tail is not. In the rightmost, bottom insect (the bitwise complement of the insect above it), the antennae, forewings, and hindwings are absent, but the tail is present. We define the top row as 'the class' and we disregard the bottom complements as degenerate forms of generalisation.

The main experimental setup involves a non-linear recurrent GRN which develops an embryonic phenotypic pattern, G, into an adult phenotype, P_a, upon which selection can act (Watson et al., 2014). An adult phenotype represents the gene expression profile that results from the dynamics of the GRN. Those dynamics are determined by the gene regulatory interactions of the network, B (Kauffman, 1993; Vohradský, 2001a,b; Gu et al., 2005; Aldana et al., 2007) (see SI: Developmental Model). We evaluate the fitness of a given genetic structure based on how close the developed phenotype is to the target phenotypic pattern, S. S characterises the direction of selection for each phenotypic trait, i.e., element of gene expression profile, in the current environment. The dynamics of selective environments are modelled by switching from one target phenotype to another every K generations. K is chosen to be considerably smaller than the overall number of generations simulated. Below, we measure evolutionary time in *epochs*, where each epoch denotes $N_T \times K$ generations and N_T corresponds to the number of target phenotypes. (Note that *epoch* here is a term we are borrowing from machine learning and does not represent geological timescale.)

In the following experiments all phenotypic targets are chosen from the same class (as in Parter et al. (2008); Watson et al. (2014)). This class consists of 8 different modular patterns that correspond to different combinations of sub-patterns. Each sub-pattern serves as a different function as pictorialised in Fig 2.1. This modular structure ensures that the environments (and thus the phenotypes that are fittest in those environments) share common regularities, i.e., they are all built from different combinations from the same set of modules. We can then examine whether the system can actually 'learn' these systematicities from a limited set of examples and thereby generalise from these to produce novel phenotypes within the same class. Our experiments are carried out as follows. The population is evolved by exposure to a limited number of selective environments (training). We then analyse conditions under which new phenotypes from the same family are produced (test). As an exemplary problem we choose a training set comprised of three phenotypic patterns from the class (see Fig 2.2 a).

One way to evaluate the generalisation ability of developmental organisations is to evolve a population to new selective environments and evaluate the evolved predisposition of the development system to produce suitable phenotypes for those environments (as per Parter et al. (2008)). We do this at the end of experimental section. We also use a more stringent test and examine the spontaneous production of such phenotypes induced by development from random genetic variation. Specifically, we examine what phenotypes the evolved developmental constraints and biases B are predisposed to create starting from random initial gene expression levels, G. For this purpose, we perform a post-hoc analysis. First, we estimate the phenotypic distributions induced by the evolved developmental architecture under drift. Since mutation on the direct effects on the embryonic phenotypes (G) in this model is much greater than mutation on regulatory interactions (B) (see Methods), we estimate drift with a uniformly random distribution

over G (keeping B constant). Then we assess how successful the evolved system is at producing high-fitness phenotypes, by seeing if the phenotypes produced by the evolved correlations, B, tend to be members of the general class (see Methods).

2.3 Results and Discussion

2.3.1 Conditions that Facilitate Generalised Phenotypic Distributions

In this section, we focus on the conditions that promote the evolution of adaptive developmental biases that facilitate generalised variational structures. To address this, we examine the distributions of potential phenotypic variants induced by the evolved developmental structure in a series of different evolutionary scenarios: 1) different time-scales of environmental switching, 2) environmental noise and 3) direct selection pressure for simple developmental processes applied via a the cost of ontogenetic interactions favouring i) weak and ii) sparse connectivity.

2.3.1.1 Rate of Environmental Switching (Learning Rates)

In this scenario, we assess the impact of the rate at which selective environments switch on the evolution of generalised developmental organisations. This demonstrates prediction (b) from Table 2.1. The total number of generations was kept fixed at 24×10^6, while the switching intervals, K, varied. In all reproductive events, G is mutated by adding a uniformly distributed random value drawn in $[-0.1, 0.1]$. Additionally, in half the reproduction events, all interaction coefficients are mutated slightly by adding a uniformly distributed value drawn from $[-0.1/(15N^2), 0.1/(15N^2)]$, where N corresponds to the number of phenotypic traits.

Prior work on facilitated variation has shown that the evolution of evolvability in varying selective environments is dependent on the time-scale of environmental change (Kashtan and Alon, 2005; Kashtan et al., 2007; Parter et al., 2008). This is analogous to the sensitivity of generalisation to learning rate in learning systems. The longer a population is exposed to a selective environment, the higher the expected adaptation accumulated to that environment would be. Accordingly, the rate of change in a given environment (learning rate) can be controlled by the rate of environmental change (sample rate). Slow and fast environmental changes thus correspond to fast and slow learning rates respectively.

We find that when the environments rapidly alternated from one to another (e.g., $K = 2$), natural selection canalised a single phenotypic pattern (Fig 2.2 b). This phenotype however did not correspond to any of the previously selected ones (Fig 2.2 a). Rather, this corresponds to the combination of phenotypic characters that occurs most in each

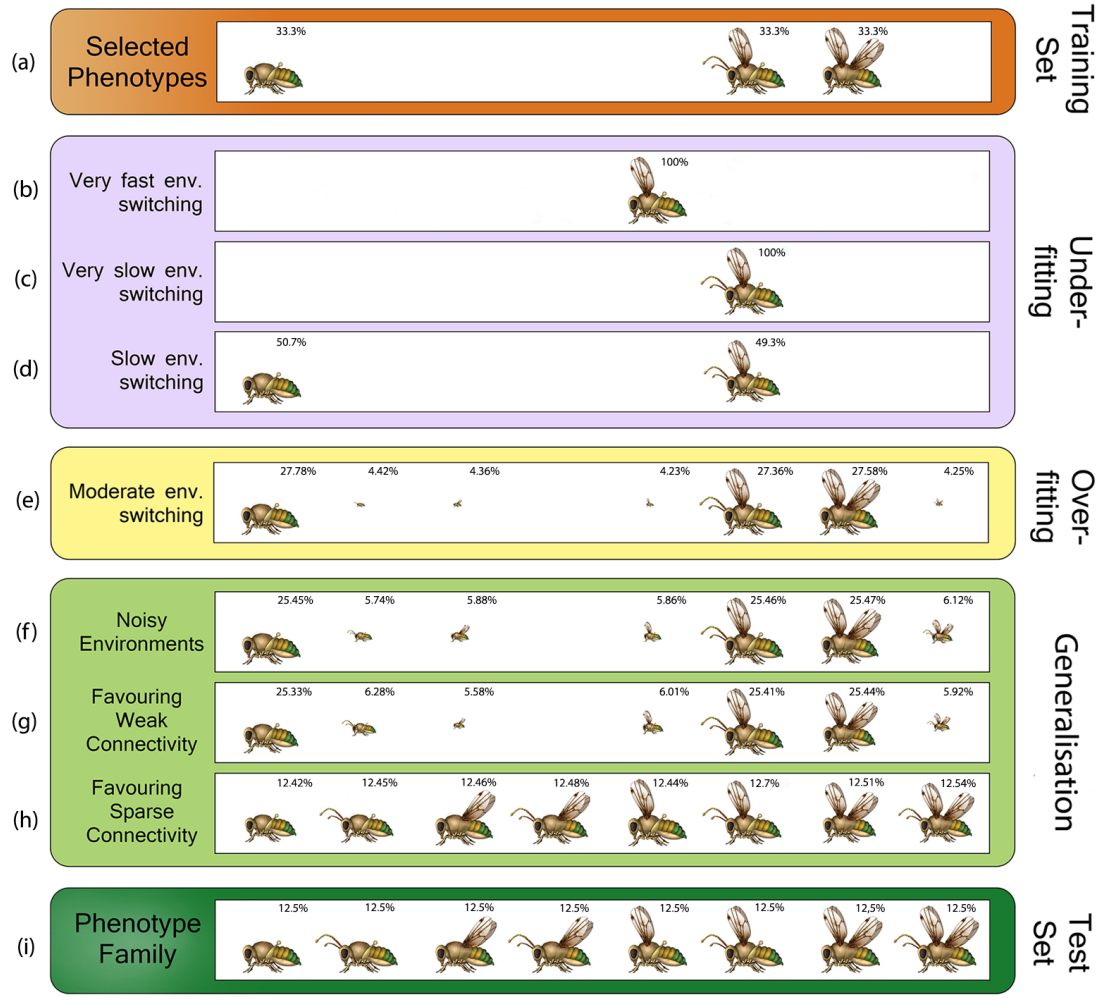

FIGURE 2.2: **Conditions that facilitate generalised phenotypic distributions.** Potential phenotypic distributions induced by the evolved developmental process under 1) different time-scales of environmental switching, 2) environmental noise ($\kappa = 35 \times 10^{-4}$) and 3) direct selection pressure for weak ($\lambda = 38$) and sparse connectivity ($\lambda = 0.22$). The organisms were exposed to three selective environments (a) from the general class (i). Developmental memorisation of past phenotypic targets clearly depends on the time-scale of environmental change. Noisy environments and parsimony pressures enhance the generalisation ability of development predisposing the production of previously unseen targets from the class. The size of the insect-like creatures describes relative frequencies and indicates the propensity of development to express the respective phenotype (phenotypes with frequency less than 0.01 were ignored). Note that the initial developmental structure represented all possible phenotypic patterns equally (here 2^{12} possible phenotypes).

of the seen target phenotypes. Hence, it does best on average over the past selective environments. For example, over the three patterns selected in the past it is more common that halteres are selected than a pair of back wings, or a pair of front wings is present more often than not and so on.

When environments changed very slowly (e.g., $K = 4 \times 10^6$), development canalised the

first selective environment experienced, prohibiting the acquisition of any useful information regarding other selective environments (Fig 2.2 c). The situation was improved for a range of slightly faster environmental switching times (e.g., $K = 2 \times 10^6$), where natural selection also canalised the second target phenotype experienced, but not all three (Fig 2.2 d). Canalisation can therefore be opposed to evolvability, resulting in very inflexible models that fail to capture any or some of the relevant regularities in the past or current environments, i.e., *under-fitting*. Such developmental organisations could provide some limited immediate fitness benefits in the short-term, but are not good representatives of either the past, or the general class.

When the rate of environmental switching was intermediate (e.g., $K = 4 \times 10^4$), the organisms exhibited developmental memory (Watson et al., 2014). Although initially all possible phenotypic patterns (here 2^{12}) were equally represented by development, the variational structure of development was adapted over evolutionary time to fit the problem structure of that past, by canalising the production of previously seen targets (Figure 2.2 e, see also Fig B in Supporting Figures in S1 Appendix A). This holds for a wide range of intermediate switching intervals (see Fig C in Supporting Figures in S1 Appendix A). This observations illustrates the ability of evolution to genetically acquire and utilise information regarding the statistical structure of previously experienced environments.

The evolved developmental constraints also exhibited generalised behaviour by allowing the production of three additional phenotypes that were not directly selected in the past, but share the same structural regularities with the target phenotypes. These new phenotypic patterns correspond to novel combinations of previously-seen phenotypic features. Yet, the propensity to express these extra phenotypes was still limited. The evolved variational mechanism over-represented past targets, failing to properly generalise to all potential, but yet-unseen selective environments from the same class as the past ones, i.e., over-fitted (see below). We find no rate of environmental variation capable of causing evolution by natural selection to evolve a developmental organisation that produces the entire class. Consequently, the rate of environmental change can facilitate the evolution of developmental memory, but does not always produce good developmental generalisation.

Here we argue that the problem of natural selection failing to evolve generalised phenotypic distributions in certain cases is formally analogous to the problem of learning systems failing to generalise due to either under- or over-fitting. In learning, under-fitting is observed when a learning system is incapable of capturing a set of exemplary observations. On the other hand, over-fitting is observed when a model is over-trained and memorises a particular set of exemplary observations, at the expense of predictive performance on previously unseen data from the class (Bishop et al., 2006). Over-fitting occurs when the model learns to focus on idiosyncrasies or noise in the training set (Abu-Mostafa et al., 2012). Similarly, canalisation to past selective environments can

be opposed to evolvability if canalised phenotypes from past environments are not fit in future environments. Specifically, canalisation can be opposed to evolvability by either 1) (first type of underfitting, from high learning rates) reducing the production of all phenotypic characters except those that are fit in the selective environments that happen to come early (Fig 2.2 c), 2) (second type of under-fitting, from low learning rates) reducing the production of all characters except those that are fit on average over the past selective environments (Fig 2.2 b), or 3) (over-fitting) successfully producing a sub-set of or all phenotypes that were fit in the past selective environments, but inhibiting the production of new and potentially useful phenotypic variants for future selective environments (Fig 2.2 d, e).

Below, we investigate the conditions under which an evolutionary process can avoid canalising the past and remain appropriately flexible to respond to novel selective environments in the future. To do so, we test whether techniques used to avoid under-fitting and over-fitting that improve generalisation to unseen test sets in learning models will likewise alleviate canalisation to past phenotypic targets and improve fit to novel selective environments in evolutionary systems. For this purpose, we choose the time scale of environmental change to be moderate ($K = 20000$). This constitutes our control experiment in the absence of environmental noise and/or any selective pressure on the cost of connections. In the following evolutionary scenarios, simulations were run for 150 epochs. This demonstrates prediction d,e, and f from Table 2.1.

2.3.1.2 Noisy Environments (Training with Noisy Data)

In this scenario, we investigate the evolution of generalised developmental organisations in noisy environments by adding Gaussian noise, $n_\mu \sim N(0, 1)$ to the respective target phenotype, S, at each generation. The level of noise was scaled by parameter κ. In order to assess the potential of noisy selection to facilitate phenotypic generalisation, we show results for the optimal amount of noise (here $\kappa = 35 \times 10^{-4}$). Later, we will show how performance varies with the amount of noise.

We find that the distribution of potential phenotypic variants induced by the evolved development in noisy environments was still biased in generating past phenotypic patterns (Fig 2.2 f). However, it slightly improved fit to other selective environments in the class compared with Fig 2.2 e. The evolved developmental structure was characterised by more suitable variability, displaying higher propensity, compared to the control, in producing those variants from the class that were not directly selected in the past.

Masking spurious details in the training set by adding noise to the training samples during the training phase is a general method to combat the problem of over-fitting in learning systems. This technique is known as 'training with noise' or 'jittering' (Bishop et al., 2006) and is closely related to the use of intrinsic noise in deep neural networks;

a technique known as 'dropout' (Hinton et al., 2012). The intuition is that when noise is applied during the training phase, it makes it difficult for the optimisation process to fit the data precisely, and thus it inhibits capturing the idiosyncrasies of the training set. Training with noise is mathematically equivalent to a particular way of controlling model complexity known as Tikhonov regularisation (Bishop et al., 2006).

2.3.1.3 Favouring Weak Connectivity (L_2-regularisation)

In this scenario, the developmental structure was evolved under the direct selective pressure for weak connectivity — favouring regulatory interactions of small magnitude, i.e., L_2-regularisation (see Methods). Weak connectivity is achieved by applying a direct pressure on the cost of connections that is proportional to their magnitude. This imposes constraints on the evolution of the model parameters by penalising extreme values.

Under these conditions natural selection discovered more general phenotypic distributions. Specifically, developmental generalisation was enhanced in a similar manner as in the presence of environmental noise, favouring similar weakly generalised phenotypic distributions. The distribution of potential phenotypic variants induced by development displayed higher propensity in producing useful phenotypic variants for potential future selective environments (Fig 2.2 g).

2.3.1.4 Favouring Sparse Connectivity (L_1-regularisation)

In this scenario, the developmental structure was evolved under the direct selective pressure for sparse connectivity — favouring fewer regulatory interactions, i.e., L_1-regularisation. Sparse connectivity is achieved by applying an equal direct pressure on the cost of connections. This imposes constraints on the evolution of the parameters by decreasing all non-zero values equally, and thus favouring models using fewer connections.

We find that under these conditions the evolution of generalised developmental organisations was dramatically enhanced. The evolved phenotypic distribution (Fig 2.2 h) was a perfect representation of the class (Fig 2.2 i). We see that the evolved developmental process under the pressure for sparsity favoured the production of novel phenotypes that were not directly selected in the past. Those novel phenotypes were not arbitrary, but characterised by the time-invariant intra-modular regularities common to past selective environments. Although the developmental system was only exposed to three selective environments, it was able to generalise and produce all of the phenotypes from the class by creating novel combinations of previously-seen modules. More notably, we see that the evolved developmental process also pre-disposed the production of that phenotypic pattern that was missing under the conditions for weak connectivity and environmental noise due to strong developmental constraints.

Moreover, the parsimonious network topologies we find here arise as a consequence of a direct pressure on the cost of connections. The hypothesis that sparse network can arise through a cost minimisation process is also supported by previous theoretical findings advocating the advantages of sparse gene regulation networks (Leclerc, 2008). Accordingly, natural selection favours the emergence of gene-regulatory networks of minimal complexity. In Leclerc (2008), Leclerc argues that sparser GRNs exhibit higher dynamical robustness. Thus, when the cost of complexity is considered, robustness also implies sparsity. In this study, however, we demonstrated that sparsity gives rise to enhanced evolvability. This indicates that parsimony on the connectivity of the GRNs is a property that may facilitate both robustness and evolvability.

Favouring weak or sparse connectivity belongs in a general category of *regularisation* methods that alleviate over-fitting by penalising unnecessary model complexity via the application of a parsimony pressure that favours simple models with fewer assumptions on the data, i.e., imposing a form of Occam's razor on solutions (e.g., the Akaike (Akaike, 1974) and Schwarz et al. (1978) Bayesian information criteria, limiting the number of features in decision trees (Deng and Runger, 2012), or limiting the tree depth in genetic programming (Soule and Foster, 1998)). The key observation is that networks with too few connections will tend to under-fit the data (because they are unable to represent the relevant interactions or correlations in the data); whereas networks with more connections than necessary will tend to over-fit the idiosyncrasies of the training data, because they can memorize those idiosyncrasies instead of being forced to learn the underlying general pattern.

2.3.2 How Generalisation Changes over Evolutionary Time

We next asked why costly interactions and noisy environments facilitate generalised developmental organisations. To understand this, we monitor the match between the phenotypic distribution induced by the evolved developmental process and the ones that describe the past selective environments (training set) and all potential selective environments (test set) respectively over evolutionary time in each evolutionary setting (see Methods). Following conventions in learning theory, we term the first measure 'training error' and the second 'test error'. This demonstrates predictions c, e and f from Table 2.1.

The dependence of the respective errors on evolutionary time are shown in Fig 2.3. For the control scenario (panel A) we observe the following trend. Natural selection initially improved the fit of the phenotypic distributions to both distributions of past and future selective environments. Then, while the fit to past selective environments continued improving over evolutionary time, the fit to potential, but yet-unseen, environments started to deteriorate (see also Fig B in Supporting Figures in S1 Appendix A). The evolving organisms tended to accurately *memorise* the idiosyncrasies of their past environments,

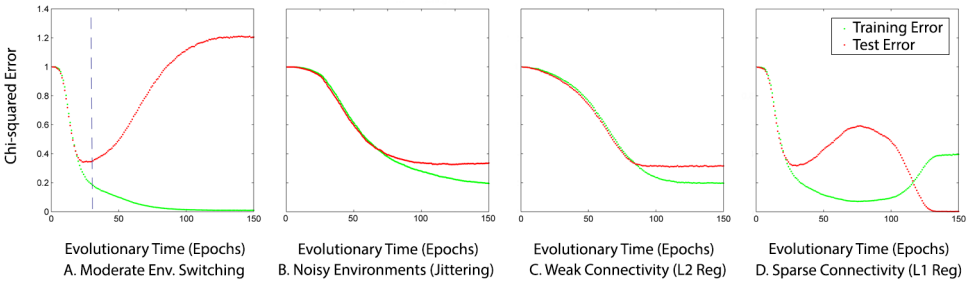

FIGURE 2.3: **How generalisation changes over evolutionary time.** The match between phenotypic distributions generated by evolved GRN and the target phenotypes of selective environments the developmental system has been exposed to (training error) and all selective environments (test error) against evolutionary time for (A) moderate environmental switching, (B) noisy environments, (C) favouring weak connectivity and (D) favouring sparse connectivity. The vertical dashed line denotes when the ad-hoc technique of early stopping would be ideal, i.e. at the moment the problem of over-fitting begins. Favouring weak connectivity and jittering exhibits similar effects on test error as applying early stopping.

at the cost of losing their ability to retain appropriate flexibility for the future, i.e., over-fitting. The dashed-line in Fig 2.3 A indicates when the problem of over-fitting begins, i.e., when the test error first increases. We see that canalisation can be opposed to the evolution of generalised phenotypic distributions in the same way over-fitting is opposed to generalisation. Then, we expect that preventing the canalisation of past targets can enhance the generalisation performance of the evolved developmental structure. Indeed, Fig 2.3 B,C,D confirm this hypothesis (predictions a-c from Table 2.1).

In the presence of environmental noise, the generalisation performance of the developmental structure was improved by discovering a set of regulatory interactions that corresponds to the minimum of the generalisation error curve of 0.34 (Fig 2.3 B). However, natural selection in noisy environments was only able to postpone canalisation of past targets and was unable to avoid it in the long term (see SI). Consequently, stochasticity improved evolvability by decreasing the speed at which over-fitting occurs, allowing for the developmental system to spend more time at a state which was characterised by high generalisation ability (see also Fig A in The Structure of Developmental Organisation in S1 Appendix A). On the other hand, under the parsimony pressure for weak connectivity, the evolving developmental system maintained the same generalisation performance over evolutionary time. The canalisation of the selected phenotypes was thus prevented by preventing further limitation of the system's phenotypic variability. Note that the outcome of these two methods (Fig 2.3 B and C) resembles in many ways the outcome as if we stopped at the moment when the generalisation error was minimum, i.e., early stopping; an ad-hoc solution to preventing over-fitting (Bishop et al., 2006). Accordingly, learning is stopped before the problem of over-fitting begins (see also Fig A in The Structure of Developmental Organisation in S1 Appendix A). Under parsimony pressure for sparse connectivity, we observe that the generalisation

error of the evolving developmental system reached zero (Fig 2.3 D). Accordingly, natural selection successfully exploited the time-invariant regularities of the environment properly representing the entire class (Fig 2.2 h). Additionally, Fig D in Supporting Figures in S1 Appendix A shows that the entropy of the phenotypic distribution reduces as expected over evolutionary time as the developmental process increasingly canalises the training set phenotypes. In the case of perfect generalisation to the class (sparse connectivity), this convergence reduces from 16 bits (the original phenotype space) to four bits, corresponding to four degrees of freedom where each of the four modules vary independently. In the other cases, overfitting is indicated by reducing to less than four bits.

2.3.3 Sensitivity Analysis to Parameters Affecting Phenotypic Generalisation

As seen so far, the generalisation ability of development can be enhanced under the direct selective pressure for both sparse and weak connectivity and the presence of noise in the selective environment, when the strength of parsimony pressure and the level of noise were properly tuned. Different values of λ and κ denote different evolutionary contexts, where λ determines the relative burden placed on the fitness of the developmental system due to reproduction and maintenance of its elements, or other physical constraints and limitations, and κ determines the amount of extrinsic noise found in the selective environments (see Evaluation of Fitness).

In the following, we analyse the impact of the strength of parsimony pressure and the level of environmental noise on the evolution of generalised developmental organisations. Simulations were run for various values of parameters λ and κ. Then, the training and generalisation error were evaluated and recorded (Fig 2.4). This demonstrates prediction (g) from Table 2.1.

We find that in the extremes, low and high levels of parsimony pressures, or noise, gave rise to situations of over-fitting and under-fitting respectively (Fig 2.4). Very small values of λ, or κ, were insufficient at finding good regulatory interactions to facilitate high evolvability to yet-unseen environments, resulting in the canalisation of past targets, i.e., over-fitting. On the other hand, very large values of λ over-constrained the search process hindering the acquisition of any useful information regarding environment's causal structure, i.e., under-fitting. Specifically, with a small amount of L_1-regularisation, the generalisation error is dropped to zero. This outcome holds for a wide spectrum of the regularisation parameter $ln(\lambda) \in [0.15, 0.35]$. However, when λ is very high (here $\lambda = 0.4$), the selective pressure on the cost of connection was too large; this resulted in the training and the generalisation errors corresponds to the original 'no model' situation (Fig 2.4 C). Similarly, with a small amount of L_2-regularisation, the generalisation error quickly drops. In the range $[10, 38]$ the process became less sensitive to changes in

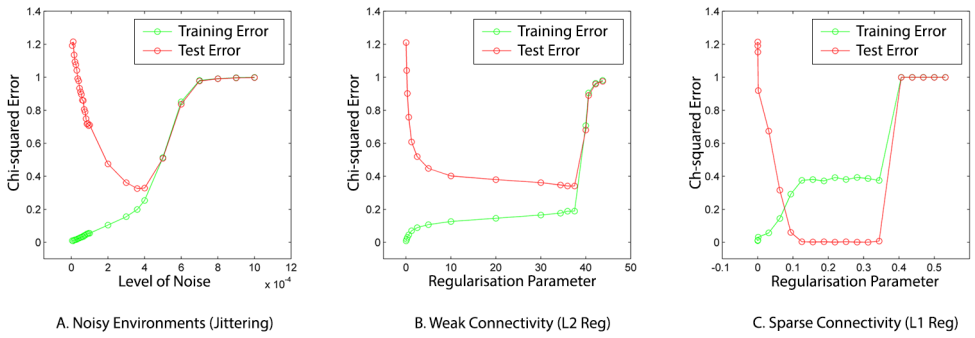

FIGURE 2.4: **Role of the strength of parsimony pressure and the level of environmental noise.** The match between phenotypic distributions and the selective environments the network has been exposed to (training error) and all possible selective environments of the same class (generalisation error) for (A) noisy environments against parameter κ and under the parsimony pressure weak (B) and sparse (C) connectivity against parameter λ.

λ, resulting in one optimum at $\lambda = 38$ (Fig 2.4 B). Similar results were also obtained for jittering (Fig 2.4 A). But the generalisation performance of the developmental process changes 'smoothly' with κ, resulting in one optimum at $\kappa = 35 \times 10^{-4}$ (Fig 2.4 A). Inductive biases need to be appropriate for a given problem, but in many cases a moderate bias favouring simple models is sufficient for non-trivial generalisation.

2.3.4 Generalised Developmental Biases Improve the Rate of Adaptation

Lastly we examine whether generalised phenotypic distributions can actually facilitate evolvability. For this purpose, we consider the rate of adaptation to each of all potential selective environments as the number of generations needed for the evolving entities to reach the respective target phenotype.

To evaluate the propensity of the organisms to reach a target phenotype as a systemic property of its developmental architecture, the regulatory interactions were kept fixed, while the direct effects on the embryonic phenotype were free to evolve for 2500 generations, which was empirically found to be sufficient for the organisms to find a phenotypic target in each selective environment (when that was allowed by the developmental structure). In each run, the initial gene expression levels were uniformly chosen at random. The results here were averaged over 1000 independent runs, for each selective environment and for each of the four different evolutionary scenarios (as described in the previous sections). Then, counts of the average number of generations to reach the target phenotype of the corresponding selective environment were taken. This was evaluated by measuring the first time the developmental system achieved maximum fitness possible. If the target was not reached, the maximum number of generations 2500 was assigned.

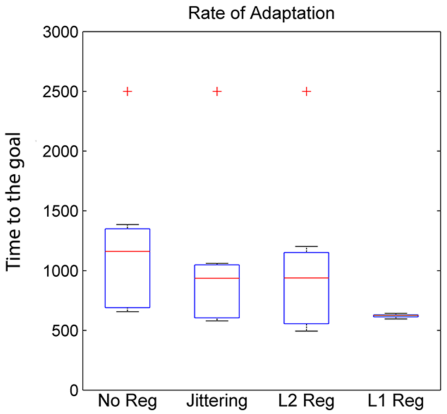

FIGURE 2.5: **Generalised developmental organisations improve the rate of adaptation to novel selective environments.** Boxplot of the generations taken for the evolved developmental systems to reach the target phenotype for all potential selective environments under different evolutionary conditions. The developmental architecture is kept fixed and only the direct effects on the embryonic phenotype are free to evolve. Organisms that facilitate generalised phenotypic distributions, such as the ones evolved in noisy environments or under the direct pressure on the cost connections, adapt faster to novel selective environments exhibiting enhanced evolvability. The outliers indicate the inability of the corresponding evolved developmental structures to reach that selective target due to strong developmental constraints.

We find that organisms with developmental organisations evolved in noisy environments or the parsimony pressure on the cost of connections adapted faster than the ones in the control scenario (Fig 2.5). The outliers in the evolutionary settings of moderate environmental switching, noisy environments and favouring weak connectivity, indicate the inability of the developmental system to express the target phenotypic pattern for that selective environment due to the strong developmental constraints that evolved in those conditions. This corresponds to the missing phenotype from the class we saw above in the evolved phenotypic distributions induced by development (Fig 2.2 e, f, g). In all these three cases development allowed for the production of the same set of phenotypic patterns. Yet, developmental structures evolved in the presence of environmental noise or under the pressure for weak connectivity exhibited higher adaptability due to their higher propensity to produce other phenotypes of the structural family. In particular, we see that for the developmental process evolved under the pressure for sparsity, the rate of adaptation of the organisms was significantly improved. The variability structure evolved under sparsity to perfectly represent the functional dependencies between phenotypic traits. Thus, it provided a selective advantage guiding phenotypic variation in more promising directions.

2.4 Conclusions

The above experiments demonstrated the transfer of predictions from learning models into evolution, by specifically showing that: a) the evolution of generalised phenotypic distributions is dependent on the time-scale of environmental switching, in the same way that generalisation in online learning algorithms is learning-rate dependent, b) the presence of environmental noise can be beneficial for the evolution of generalised phenotypic distributions in the same way training with corrupted data can improve the generalisation performance of learning systems with the same limitations, c) direct selection pressure for weak connectivity can enhance the evolution of generalised phenotypic distributions in the same way L_2-regularisation can improve the generalisation performance in learning systems, d) noisy environments result in similar behaviour as favouring weak connectivity, in the same way that Jittering can have similar effects to L_2-regularisation in learning systems, e) direct selection pressure for sparse connectivity can enhance the evolution of generalised phenotypic distributions in the same way that L_1-regularisation can improve the generalisation performance in learning systems, f) favouring weak connectivity (i.e., L_2-regularisation) results in similar behaviour to early stopping, g) the evolution of generalised phenotypic distributions is dependent on the strength of selection pressure on the cost of connections and the level of environmental noise, in the same way generalisation is dependent on the level of inductive biases and h) in simple modularly varying environments with independent modules, sparse connectivity enhances the generalisation of phenotypic distributions better than weak connectivity, in the same way that in problems with independent features, L_1-regularisation results in better generalisation than L_2-regularisation..

Learning is generally *contextual*; it gradually builds upon what *concepts* are already known. Here these concepts correspond to the repeated modular sub-patterns persisting over all observations in the training set which become encoded in the modular components of the evolved network. The inter-module connections determine which combinations of (sub-)attractors in each module are compatible and which are not. Therefore, the evolved network representation can be seen as dictating a higher-order conceptual (combinatorial) space based on previous experience. This enables the evolved developmental system to explore permitted combinations of features constrained by past selection. Novel phenotypes can thus arise through new combinations of previously selected phenotypic features explicitly embedded in the developmental architecture of the system (Watson et al., 2014). Indeed, under the selective pressure for sparse connectivity, we observe that the phenotypic patterns generated by the evolved developmental process consisted of combinations of features from past selected phenotypic patterns. Thus, we see that the 'developmental memories' are stored and recalled in combinatorial fashion allowing generalisation.

We see that noisy environments and the parsimony pressure on the cost of connections led to more evolvable genotypes by internalising more general models of the environment into their developmental organisation. The evolved developmental systems did not solely capture and represent the specific idiosyncrasies of past selective environments, but internalised the regularities that remained time-invariant in all environments of the given class. This enabled natural selection to 'anticipate' novel situations by accumulating information about and exploiting the tendencies in that class of environments defined by the regularities. Peculiarities of past targets were generally represented by weak correlations between phenotypic characters as these structural regularities were not typically present in all of the previously-seen selective environments. Parsimony pressures and noise then provided the necessary selective pressure to neglect or de-emphasise such spurious correlations and maintain only the strong ones which tended to correspond to the underlying problem structure (in this case, the intra-module correlations only, allowing all combinations of fit modules). More notably, we see that the parsimony pressure for sparsity favoured more evolvable developmental organisations that allowed for the production of a novel and otherwise inaccessible phenotype. Enhancing evolvability by means of inductive biases is not for granted in evolutionary systems any more than such methods have guarantees in learning systems. The quality of the method depends on information about past targets and the strength of the parsimony pressure. Inductive biases can however constrain phenotypic evolution into more promising directions and exploit systematicities in the environment when opportunities arise.

In this study we demonstrated that canalisation can be opposed to evolvability in biological systems the same way under- or over-fitting can be opposed to generalisation in learning systems. We showed that conditions that are known to alleviate over-fitting in learning are directly analogous to the conditions that enhance the evolution of evolvability under natural selection. Specifically, we described how well-known techniques, such as learning with noise and penalising model complexity, that improve the generalisation ability of learning models can help us understand how noisy selective environments and the direct selection pressure on the reproduction cost of the gene regulatory interactions can enhance context-specific evolvability in gene regulation networks. This opens-up a well-established theoretical framework, enabling it to be exploited in evolutionary theory. This equivalence demystifies the basic idea of the evolution of evolvability by equating it with generalisation in learning systems. This framework predicts the conditions that will enhance generalised phenotypic distributions and evolvability in natural systems.

2.5 Methods

We use computer simulations on GRN models developed in previous work (Watson et al., 2014). GRN models are well-understood and extensively studied for understanding the evolution of complex developmental systems (Aldana et al., 2007; Levine and Davidson,

2005; Davidson and Levine, 2008; Longabaugh et al., 2005; Erwin and Davidson, 2009; Weaver et al., 1999). We choose a GRN model since prior work has shown significant insights into the evolution of variational properties of the biological systems that stems from genotype-phenotype map. As a result, studying GRN models is directly relevant to the evolution of evolvability and mutational robustness and the current work on the evolution of developmental flexibility (Wagner and Altenberg, 1996; Aldana et al., 2007; Quayle and Bullock, 2006; Lenski et al., 2006). Moreover, this type of models provide understanding towards empirical networks that are increasingly derived from both genomic and developmental biology (Jeong et al., 2000; Alon, 2003).

In addition, it is appreciated that GRNs function in the same way artificial neural networks do (Wagner, 1996; Vohradský, 2001a,b; Watson et al., 2014; Fierst and Phillips, 2015). In fact, the dynamics of GRNs is formalised in the same way as neural networks, and both systems are governed by the equations of motion (Wessels et al., 2001). Watson et al. demonstrated a further result, more important to our purposes here; that the way regulatory interactions *evolve* under natural selection is mathematically equivalent to the way neural networks *learn* (Watson et al., 2014). GRN models of development are a good candidate for our work here to test our predictions on how the same conditions that tend to improve the generalisation performance of learning systems can also enhance the evolution of generalised phenotypic distributions. In this work, we adopt the same developmental model (same parameters) as in (Watson et al., 2014). For the purpose of this work, one developmental architecture is sufficient to demonstrate the transfer of predictions from learning models into evolution and more precisely that the evolution of evolvability is possible under noisy selective environments and/or the direct selection pressure on the reproduction cost of the gene regulatory interactions. For completeness, we also present a sensitivity analysis over the evolutionary conditions used to enhance the evolution of generalised phenotypic distributions. Different model architectures are beyond the scope of this Chapter. Feed-forward networks are presented and studied in the next two Chapter of the thesis. Future work on the role of model architecture in the evolution of adaptive potential is proposed in the last Chapter of the thesis.

The source code used for the current chapter can be found here:
https://github.com/KostasKouvaris/Evolvability.

2.5.1 Evolution of GRNs

We model the evolution of a population of GRNs under strong selection and weak mutation where each new mutation is either fixed or lost before the next arises. This emphasises that the effects we demonstrate do not require lineage-level selection (Palmer and Feldman, 2012; Masel and Trotter, 2010; Rajon and Masel, 2011) — i.e., they do not require multiple genetic lineages to coexist long enough for their mutational

distributions to be visible to selection. Accordingly a simple hill-climbing model of evolution is sufficient (Watson et al., 2014; Kashtan et al., 2007).

The population is represented by a single genotype $[G, B]$ (the direct effects and the regulatory interactions respectively) corresponding to the average genotype of the population. Similarly, mutations in G and B indicate slight variations in population means. Consider that G' and B' denote the respective mutants. Then the adult mutant phenotype, P'_a, is the result of the developmental process, which is characterised by the interaction B', given the direct effects G'. Subsequently, the fitness of P_a and P'_a are calculated for the current selective environment, S. If $f_S(P'_a) > f_S(P_a)$, the mutation is beneficial and therefore adopted, i.e., $G_{t+1} = G'$ and $B_{t+1} = B'$. On the other hand, when a mutation is deleterious, G and B remain unchanged.

The variation on the direct effects, G, occurs by applying a simple point mutation operator. At each evolutionary time step, t, an amount of μ_1 mutation, drawn from $[-0.1, 0.1]$ is added to a single gene i. Note that we enforce all $g_i \in [-1, 1]$ and hence the direct effects are hard bounded, i.e., $g_i = min\{max\{g_i + \mu_1, -1\}, 1\}$. For a developmental architecture to have a meaningful effect on the phenotypic variation, the developmental constraints should evolve considerably slower than the phenotypic variation they control. We model this by setting the rate of change of B to lower values as that for G. More specifically, at each evolutionary time step, t, mutation occurs on the matrix with probability 1/15. The magnitude μ_2 is drawn from $[-0.1/(15N^2), 0.1/(15N^2)]$ for each element b_{ij} independently, where N corresponds to the number of phenotypic traits.

2.5.2 Evaluation of Fitness

Following the framework used in (Kashtan et al., 2009), we define the fitness of the developmental system as a benefit minus cost function.

The benefit of a given genetic structure, b, is evaluated based on how close the developed adult phenotype is to the target phenotype of a given selective environment. The target phenotype characterises a favourable direction for each phenotypic trait and is described by a binary vector, $S = \langle s_1, \ldots, s_N \rangle$, where $s_i \in \{-1, 1\}, \forall i$. For a certain selective environment, S, the selective benefit of an adult phenotype, P_a, is given by (modified from (Watson et al., 2014)):

$$b = w(P_a, S) = \frac{1}{2}\left(1 + \frac{P_a \cdot S}{N}\right), \tag{2.1}$$

where the term $P_a \cdot S$ indicates the inner product between the two respective vectors. The adult phenotype is normalised in $[-1, 1]$ by $P_a \leftarrow P_a/(\tau_1/\tau_2)$, i.e., $b \in [0, 1]$.

The cost term, c, is related to the values of the regulatory coefficients, $b_{ij} \in B$ (Dekel and Alon, 2005). The cost represents how fitness is reduced as a result of the system's

effort to maintain and reproduce its elements, e.g., in *E. coli* it corresponds to the cost of regulatory protein production. The cost of connection has biological significance (Kashtan et al., 2009; Dekel and Alon, 2005; Clune et al., 2013b; Striedter, 2006; Cherniak et al., 2004), such as being related to the number of different transcription factors or the strength of the regulatory influence. We consider two cost functions proportional to i) the sum of the absolute magnitudes of the interactions, $c = \|B\|_1 = \sum_{i=1}^{N^2} |b_i j|/N^2$, and ii) the sum of the squares of the magnitudes of the interactions, $c = \|B\|_2^2 = \sum_{i=1}^{N^2} b_{ij}^2/N^2$, which put a direct selection pressure on the weights of connections, favouring sparse (L_1-regularisation) and weak connectivity (L_2-regularisation) respectively (Russell et al., 1995).

Then, the overall fitness of P_a for a certain selective environment S is given by:

$$f_S(P_a) = b - \lambda c, \tag{2.2}$$

where parameter λ indicates the relative importance between b and c. Note that the selective advantage of structure B is solely determined by its immediate fitness benefits on the current selective environment.

2.5.3 Chi-squared Error

The χ^2 measure is used to quantify the lack of fit of the evolved phenotypic distribution $\hat{P}_t(s_i)$ against the distribution of the previously experienced target phenotypes $P_t(s_i)$ and/or the one of all potential target phenotypes of the same family $P(s_i)$. Consider two discrete distribution profiles, the observed frequencies $O(s_i)$ and the expected frequencies $E(s_i)$, $s_i \in S, \forall i = 1, \ldots, k$. Then, the chi square error between distribution O and E is given by:

$$\chi^2(O, E) = \sum_i \frac{(O(s_i) - E(s_i))^2}{E(s_i)} \tag{2.3}$$

S corresponds to the training set and the test set when the training and the generalisation error are respectively estimated. Each $s_i \in S$ indicates a phenotypic pattern and $P(s_i)$ denotes the probability of this phenotype pattern to arise.

The samples, over which the distribution profiles are estimated, are uniformly drawn at random (see Estimating the Empirical Distributions). This guarantees that the sample is not biased and the observations under consideration are independent. Although the phenotypic profiles here are continuous variables, they are classified into binned categories (discrete phenotypic patterns). These categories are mutually exclusive and the sum of all individual counts in the empirical distribution is equal to the total number of observations. This indicates that no observation is considered twice, and also that the

categories include all observations in the sample. Lastly, the sample size is large enough to ensure large expected frequencies, given the small number of expected categories.

2.5.4 Estimating the Empirical Distributions

For the estimation of the empirical (sample) probability distribution of the phenotypic variants over the genotypic space, we follow the Classify and Count (CC) approach (Forman, 2008). Accordingly, 5000 embryonic phenotypes, $P(0) = G$, are uniformly generated at random in the hypercube $[-1, 1]^N$. Next, each of these phenotypes is developed into an adult phenotype and the produced phenotypes are categorised by their closeness to target patterns to take counts. Note that the development of each embryonic pattern in the sample is unaffected by development of other embryonic patterns in the sample. Also, the empirical distributions are estimated over all possible combinations of phenotypic traits, and thus each developed phenotype in the sample falls into exactly one of those categories. Finally, low discrepancy quasi-random sequences (Sobol sequences; (Galanti and Jung, 1997)) with Matousek's linear random scramble (Matoušek, 1999) were used to reduce the stochastic effects of the sampling process, by generating more homogeneous fillings over the genotypic space.

Chapter 3

How can adaptive plasticity reliably evolve when plasticity is not selected for?

Abstract

Phenotypic plasticity is an important mechanism of how organisms respond to environmental heterogeneity. The evolution of adaptive plasticity however remains a contentious issue in evolutionary theory. In nature, organisms encounter environmental changes across and within generations. When organisms experience multiple environments over their lifetime, plasticity is expected to be evolutionarily advantageous and thus selected for. Yet, when organisms experience one environment in their lifetime, it is not clear how adaptive plasticity arises. Natural selection concerns immediate fitness benefits and as such selection would always favour the canalisation of the current selected phenotype, namely, non-plastic responses. In this work, we show that adaptive plasticity can reliably increase in cases where non-plastic phenotypes are optimal in the short-term without lineage selection. Environmental sensitivity can systematically arise as a by-product of selection towards phenotypic targets in different selective environments across generations. The plausibility of adaptive plasticity to arise depends on the expected adaptive change accumulated at each given environment, and is thus influenced by the rates of mutation and environmental change. We argue that such sensitivity is analogous to the sensitivity of prediction to learning rates in learning systems. This allows us to characterise general conditions that can facilitate the evolution of adaptive plasticity under natural selection.

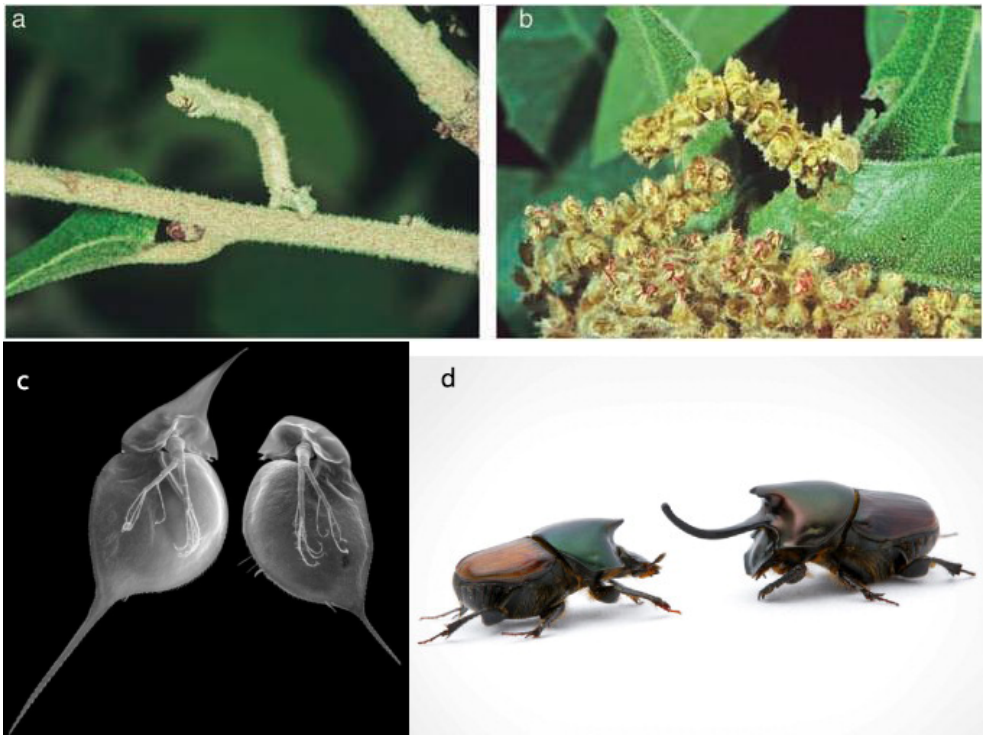

FIGURE 3.1: Plastic responses to environmental cues. (a, b) Seasonal morphological differences in *Nemoria arizonaria* caterpillars. Summer broods and spring broods develop to resemble oak twig and oak catkins respectively. Adapted from (Greene, 1989). (c) *Daphnia lumholtzi* (water fleas) induce tail spines and protective helmet in the presence of predaceous chemicals. Adapted from (Agrawal, 2001). (d) Presence or absence of horns in *Onthophagus nigriventris* (dung beetles). Adapted from (Pfennig et al., 2010).

3.1 Introduction

Most biological organisms have evolved some degree of responsiveness to external environmental stimuli, i.e., phenotypic plasticity (West-Eberhard, 2003; Nijhout, 2015). Phenotypic plasticity can be defined as the capacity of an individual genotype to express different phenotypic outcomes when exposed in variable environmental conditions (Schlichting, 1986; Via, 1993; Scheiner, 1993; Via et al., 1995; Pigliucci, 2001; Nicoglou, 2015). Whether phenotypic plasticity is adaptive or not depends on the environment plasticity is expressed (Ghalambor et al., 2007). Phenotypic plasticity is adaptive, if the plastic response yields fitness benefits for the organism in the given environmental context, and is non-adaptive otherwise. Typical examples of adaptive plasticity include i) differences in the body size of female and male *Onthophagus taurus* (beetles) given different nutritive environments (Moczek, 1998), ii) inducible defenses, e.g., tail spines and protective crests in *Daphnia* against *Notonecta* (bug predator) (Dodson, 1989; Moczek et al., 2011), and iii) morphological differences between winged and wingless aphids or grasshoppers (Fairbairn, 1997) (see also Fig. 3.1).

There are two distinct conceptions of phenotypic plasticity which are not sufficiently distinguished in the literature (Via, 1993; De Jong, 2005). On one hand, phenotypic plasticity is conceived as a quantitative trait (Schlichting, 1986; Scheiner, 1993; DeWitt and Scheiner, 2004). As such, plasticity is a property of the genotype and subject to selection like any other evolvable trait. According to this view, adaptive plasticity evolves as a direct product of selection. Hence, plasticity is of little theoretical significance, since it does not provide any new insights into how evolution works. On the other hand, plasticity is considered to be a process of development (West-Eberhard, 1989, 2003; Schlichting, 2004). This means that plasticity is an inherent property of the developmental system that can facilitate adaptive evolution. The important question here is not about *why* plasticity has evolved. Such capacity that facilitates the expression of alternative phenotypes to potentially better match different environments would be ultimately evolutionarily advantageous, and thus favoured (De Jong, 2005). Yet, not all plasticity is adaptive (Pigliucci et al., 2006).

Despite the plethora of theoretical studies, the role of plasticity in facilitating adaptive evolution remains a contentious issue (Via, 1993; Pigliucci, 2001; De Jong, 2005; Ghalambor et al., 2007). Part of this controversy stems from the fact that most of the variation induced by environmental variability tends to be deleterious (Ghalambor et al., 2007; De Jong, 2005; Van Kleunen and Fischer, 2005; West-Eberhard, 2003). That is, there is a high demand for developmental homeostasis and thus buffering against environmental variability. Since natural selection concerns with immediate fitness benefits, selection would always favour the canalisation of the currently selected phenotypes in order to maintain functional stability and reliability (Waddington, 1942; Schmalhausen, 1949; Pigliucci, 2001). Adaptive plasticity, however, requires flexible developmental structures that track changes in the environment. How can then plasticity be adapted by natural selection when non-plastic phenotypes are optimal in the short-term?

The environments encountered by a population dictate which phenotypes are realised by the individuals along with their fitness consequences. Environments in nature are characterised by high spatial and temporal heterogeneity (Levins, 1968; West-Eberhard, 2003; Baythavong, 2011; DeWitt, 2016; Hamann et al., 2016; Lande and Arnold, 1983). Organisms can therefore encounter environmental changes across and within generations upon which selection can act. Fine-grained environmental variation occurs when individuals experience multiple selective environments during their lifetime. Instead, in coarse-grained environments, individuals experience a single selective environment during their lifetime, and thus only one trait value is exposed and subject to individual selection.

Local adaptation of populations can take place either through plasticity or adaptive genetic differentiation (Hamann et al., 2016; West-Eberhard, 2003). When individuals can express multiple phenotypes in response to intra-generational environmental change (i.e., fine-grained environmental variability), a highly plastic individual that produces

high fit phenotypes across different habitats would be favoured by natural selection. In fine-grained environments, adaptive phenotypic plasticity is directly selected for. Hence, phenotypic differentiation can arise among different environments without the need for genetic differentiation (Van Tienderen, 1991; Sultan and Spencer, 2002; Alpert and Simms, 2002; Kawecki and Ebert, 2004; Baythavong, 2011; Hamann et al., 2016).

On the other hand, in situations where within-generation environmental variation is absent or individuals produce one phenotype over their life-time period (e.g., size of maturity), only one component of plasticity can be expressed. Theoretical studies suggest that the selective pressures that maximise fitness to the local habitat of the population leads to adaptive genetic differentiation in coarse-grained environments (Linhart and Grant, 1996; Sultan and Spencer, 2002; Byars et al., 2007; Lande and Arnold, 1983; Volis et al., 2015). It is thus often expected for natural selection to favour plastic genotypes in fine-grained environmental variability and genotypic differentiation in coarse-grained environmental variability (Hamann et al., 2016).

The quantitative genetic model proposed by Lande (2009) is a good example of how adaptive phenotypic plasticity fails to evolve in coarse-grained environments, when plasticity is not directly selected for. According to the model, a population is exposed to a different environment every generation (i.e., coarse-grained environmental variability). The population thus experiences a selective pressure towards each local phenotypic optimum. As a result, the population adapted to its current local habitat and generated high (cryptic) random genetic variance with respect to plasticity. An increase in phenotypic variation due to plasticity indeed is useful for adaptive evolution. However, the population failed to evolve to appropriate environmental sensitivity such as to adaptively respond to new environmental cues in a systematic manner, that is, the environmental signals were not instructive.

In this paper, we demonstrate that plasticity does not need to be a direct product of selection. Adaptive plasticity can arise as a by-product of selection within an intrinsically varying environment. Specifically, we show that adaptive plasticity can systematically evolve in cases where plasticity is not expected to increase, that is, plasticity is costly and not needed in any individual. To do so, we explore the conditions under which natural selection favours developmental organisations that facilitate adaptive plastic responses in fine- and coarse-grained environments. The results do not depend on the assumption that plasticity is inherently costly, nor the fact that some individuals within the population happen to have a better reaction norm that is selected between all environments (lineage-level selection). Lastly, we argue that the plausibility of adaptive plasticity to arise depends on the expected adaptive change accumulated at each given environment, the same way as prediction in learning systems is sensitive to the learning rates. This analogy allows a characterisation of general conditions that can facilitate the evolution of adaptive plasticity under natural selection.

3.2 Experimental Setup

In the following experiments, we consider the evolution of adaptive phenotypic plasticity through computer simulations. Here we assume a more conservative hypothesis; plasticity is not only unnecessary, but also costly. The experimental setup involves a population that experiences environmental heterogeneity, where each individual receives information from the environment, e, and develops into an adult phenotype, P^a, upon which selection can act (see Developmental Network Model). We evaluate the fitness of each individual, f, based on how close the developed phenotype, P^a, is to the respective target phenotype, ϕ_E, which is determined by the current environment, E (see Evaluation of Fitness).

We model a heterogeneous varying environment such as each selective environmental state is characterised by a single trait optimum, e, associated with a single environmental cue, ϕ. For simplicity, we consider the phenotypic targets to have a linear function with the environments (see Environmental Variability). In addition, we assume that the lifespan of the individuals is fixed and equal among them. This assumption allows us to control the granularity of environmental variability with a single parameter, K. If $K < 1$ the population encounters $1/K$ on average environments per generation indicating fine-grained environmental variability. On the other hand, if $K >= 1$ the population encounters K environments per generation indicating coarse- ($K = 1$) or extra-coarse-grained ($K > 1$) environmental variability.

Our experiments are carried out as follows. The population is evolved by exposure to a limited number of selective environments (here 10, see Evolutionary Process). We then characterise the adaptive potential of plasticity, by estimating the reaction norm for each individual in the population separately and then comparing it to the optimal reaction norm (see Evaluation of Reaction Norms). Reaction norms are graphical representations that describe the association of phenotypic responses to environmental change. Accordingly, reaction norms can be defined as sets of phenotypes that would be expressed if the given individual is exposed to the respective set of environments. Here we consider reaction norms that are linear, that is, the developmental system is described by a single linear unit (see Developmental Network Model). Linear reaction norms among individuals can differ in their slope (i.e., degree of plasticity) and offset (i.e., breeding value).

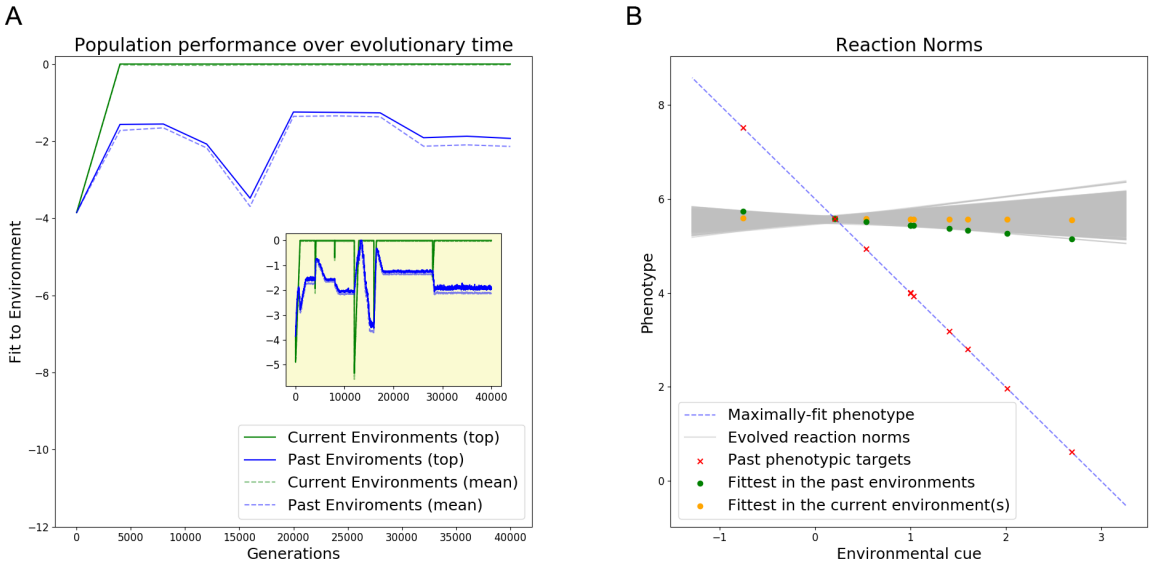

FIGURE 3.2: **Natural selection fails to evolve adaptive phenotypic plasticity in extra-coarse-grained environments.** A. Goodness of fit of the evolving reaction norms to the current environment (green) and all past selective environments (blue) over evolutionary time. The evaluation is performed at the end of each environmental period. Solid lines indicate the adaptive potential of the best individual in the population, whereas dashed lines indicate the mean performance of the population. The inset plot corresponds to the initial 40000 generations illustrating the performance of the evolving reaction norms for the first 10 environmental switches. An increase in performance on the current environment is accompanied with a decrease in the adaptive potential of the evolving plasticity over the past selective environments. B. Individual reaction norms of the evolved population. Grey lines indicate reaction norms of the evolved individuals. The red marks correspond to trait optima of the past experienced selective environments, while the blue dashed line corresponds to the optimal reaction norm. The fittest reaction norm in the evolved population for the current environment (yellow dots) corresponds to a flat horizontal line indicating environmental buffering.

3.3 Results

In the following series of experiments, a population of 1000 individuals is considered. Unless otherwise stated, the mutation rate, σ_μ, and the selection coefficient, ω, are chosen to be 0.01 and 0.2 respectively. In addition, we choose the associated cost of plasticity, λ, to be 0.1, while the environment change from one to another every $K = 4000$ generations. The same qualitative results hold in the long-run, when the population is evolved over 5000000 generations (see S1 Appendix B).

3.3.1 Natural selection fails to evolve adaptive plasticity in extra-coarse-grained environments.

In this scenario, we assess the evolution of adaptive phenotypic plasticity in situations where individuals encounter a single selective environment during their life-time. The

individuals alternate from one environment to another every $K = 4000$. This guarantees that the population experiences extra-coarse-grained environmental variability, i.e., inter-generational environmental change.

Overall, we find that natural selection fails to evolve adaptive reaction norms in such environmental conditions (Figure 3.2). We see that the evolving population quickly adapted to its local habitat through genetic change by the end of each environmental period, namely, in less than 4000 generations. Specifically, the evolving individuals reliably produced phenotypic trait values that perfectly match the optimal trait value of their current environment by the end of each environmental switch. This is indicated by the negligible residuals errors of the evolving reaction norms, which corresponds to the optimal fit for each of the single selective environments (green curve in Figure 3.2, A).

Furthermore, we see that an increase in the goodness of fit to the current selective environment is associated with a decrease in the goodness of fit to the other past selective environments (Figure 3.2, A (inset)). Plasticity is not only unnecessary, but also costly and thus selected against. Note that goodness of fit to past environments keeps decreasing even after the goodness of fit is optimised for the current environment. Selection in a single environment only optimised the trait value for the currently encountered environment, leaving residual genetic variation that may be expressed in different selective environments (Figure 3.2, B). We see that individual reaction norms within the same population vary and show different degrees of plasticity, as long as their realised character is locally adapted. Such potential variation may lead to an evolutionary advantage for those reaction norms that happen to be genetically closer to the optimal reaction norm (green dots in Figure 3.2, B). As a result, they might show some long-term adaptability to environment change, even if they do not entirely match the optimal reaction norm. Nevertheless, natural selection fails to differentiate and select these adaptive reaction norms. This is illustrated by the flat horizontal norm of reaction of the average individual, that is, the population on average is environmentally insensitive (yellow dots in Figure 3.2, B). The results suggest that the evolved genotypes were unable to track any changes in the environment and tend to canalise the trait values that are optimal for their current environment.

3.3.2 Natural selection favours the evolution of adaptive plasticity in fine-grained environments.

In this scenario, we assess the evolution of adaptive phenotypic plasticity in situations where individuals encounter multiple environmental states during their life-time (here 10). We further assume that the phenotypic traits are labile during individuals' lifespan and that only one character is expressed at a time. The population thus experiences fine-grained environmental variability, i.e., intra-generational environmental change.

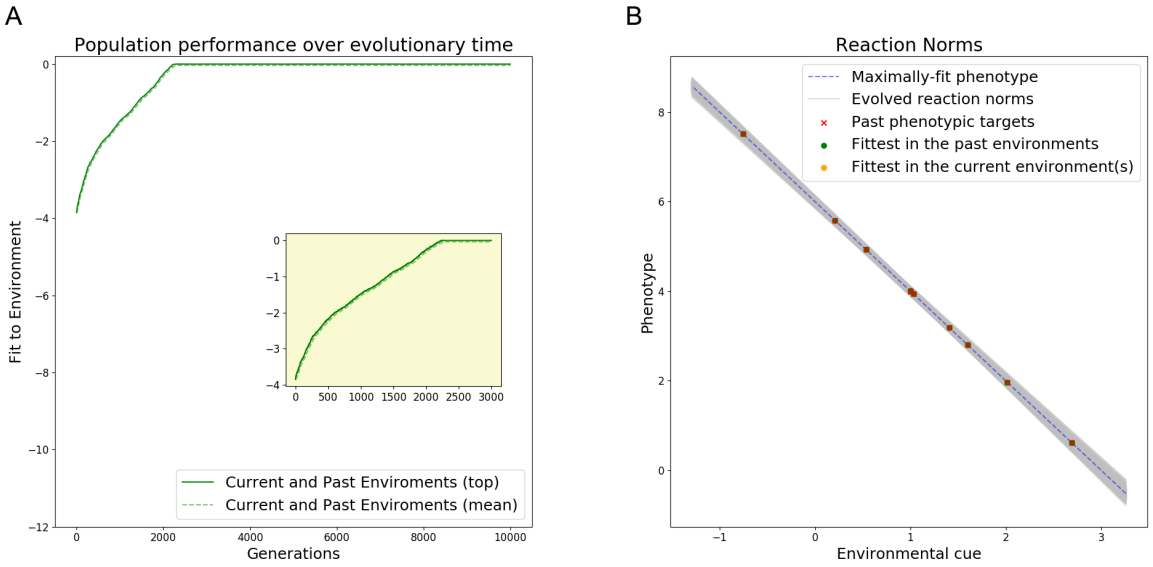

FIGURE 3.3: **Natural selection promotes the evolution of adaptive plasticity in fine-grained environments.** A. Goodness of fit of the evolving reaction norms to the current environment (green) and all past selective environments (blue) over evolutionary time. The evaluation is performed at the end of each environmental period. Solid lines indicate the adaptive potential of the best individual in the population, whereas dashed lines indicate the mean performance of the population. The inset plot corresponds to the initial 3000 generations. The adaptive potential of the evolving plasticity over the past selective environments is increased along with the performance on the current environment. B. Individual reaction norms of the evolved population. Grey lines indicate reaction norms of the evolved individuals. The red marks correspond to trait optima of the past experienced selective environments, while the blue dashed line corresponds to the optimal reaction norm. The yellow dots indicate the fittest reaction norm in the evolved population.

We find that natural selection favours the evolution of adaptive reaction norms in such environmental conditions (Figure 3.3). We see that the population quickly evolved appropriate environmental sensitivity to adapt to its local habitats in less than 3000 generations (Figure 3.3, A (inset)). This means that the evolving individuals produced phenotypic trait values that perfectly match the optimal trait value of all environmental states they encountered during their lifetime.

We observe that the goodness of fit to current and past environments decreased to zero, indicating an optimal fit to an intrinsically varying environment (Figure 3.3, A). In fine-grained environments, selection optimises all realised trait values for all currently encountered environments (i.e., selection minimises the distance between the realised phenotypes and their respective targets). We see that selection simultaneously acts on multiple aspects of the reaction norm, and thus favour reaction norms that are closer to the optimal one. In addition, the residual genetic variation in fine-grained environments was decreased when compared to the case of coarse-grained environmental variability (Figure 3.3, B). This is also indicated by the narrow gap between the top and the mean performance curve in Figure 3.3, A. Note that the reaction of the average (yellow dots)

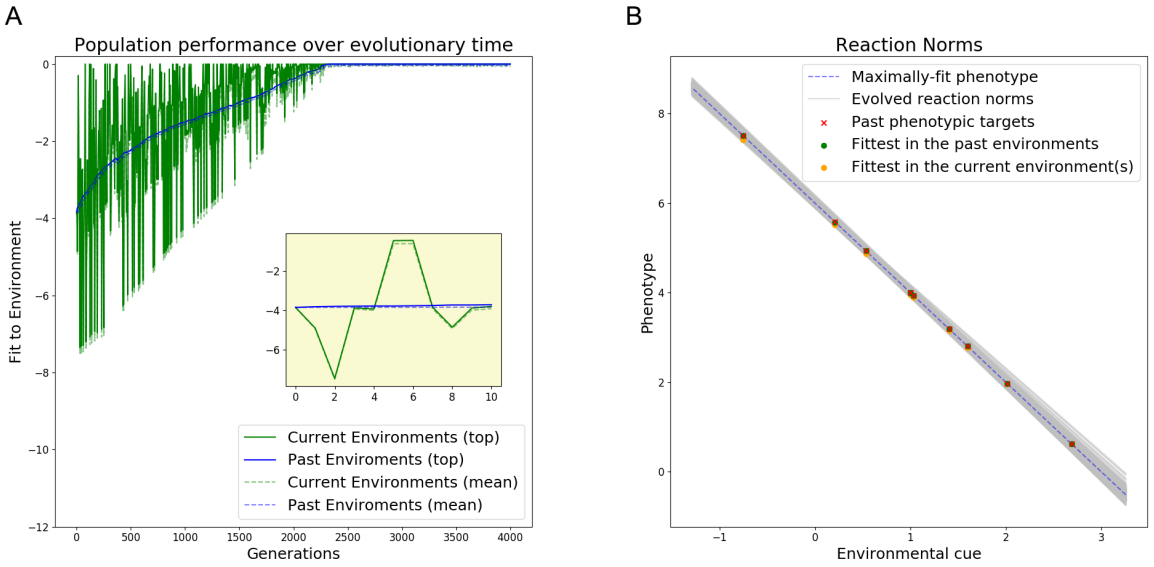

FIGURE 3.4: **Natural selection promotes the evolution of adaptive plasticity in coarse-grained environments for high rates of environmental change.** A. Goodness of fit of the evolving reaction norms to the current environment (green) and all past selective environments (blue) over evolutionary time. The evaluation is performed at the end of each environmental period. Solid lines indicate the adaptive potential of the best individual in the population, whereas dashed lines indicate the mean performance of the population. B. Individual reaction norms of the evolved population. Grey lines indicate reaction norms of the evolved individuals. The red marks correspond to trait optima of the past experienced selective environments, while the blue dashed line corresponds to the optimal reaction norm. The population evolved optimal reaction norms without a direct selection pressure for plasticity. Environmental targets rapidly change from one to another every generation, $K = 1$.

and best individual (green dots) are perfectly aligned and match the optimal reaction norm.

3.3.3 High rate of environmental change can enhance the evolution of adaptive plasticity in coarse-grained environments.

In this scenario, we assess the impact of the rate at which selective environments switch from one to another on the evolution of adaptive reaction norms. To do so, the population is exposed to a different environment selected at random every generation, $K = 1$. Note that the individuals still experience a single environment during their lifetime, i.e., inter-generational environmental variability.

We find that when the environmental targets rapidly alternated from one to another, the population evolved appropriate environmental sensitivity to adapt to its local habitats (Figure 3.4). We observe that the goodness of fit to current and past environments decreased to zero, indicating an optimal fit to an intrinsically varying environment (Figure 3.4, A). In coarse-grained environments, selection optimises only the trait value for the

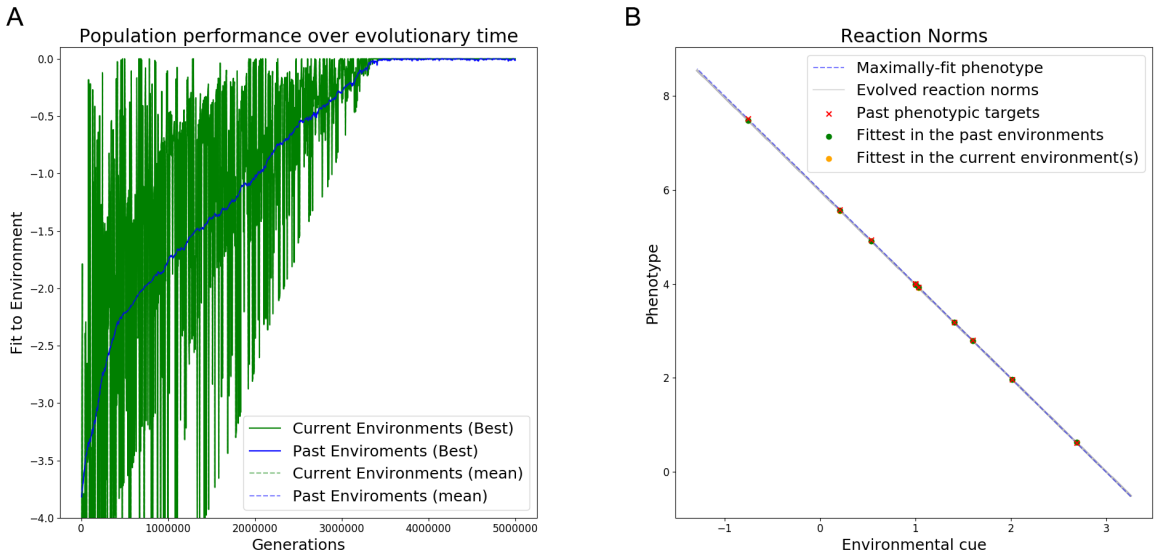

FIGURE 3.5: **Natural selection promotes the evolution of adaptive plasticity in coarse-grained environments for low mutation rates.** A. Goodness of fit of the evolving reaction norms to the current environment (green) and all past selective environments (blue) over evolutionary time. The evaluation is performed at the end of each environmental period. Solid lines indicate the adaptive potential of the best individual in the population, whereas dashed lines indicate the mean performance of the population. B. Individual reaction norms of the evolved population. Grey lines indicate reaction norms of the evolved individuals. The red marks correspond to trait optima of the past experienced selective environments, while the blue dashed line corresponds to the optimal reaction norm. The fittest reaction norm in the evolved population for the current environment (yellow dots) is aligned with direction of environmental change, namely, the organism can track environmental changes and realise optimal phenotypic responses. The environment changes every $K = 4000$ generations, while the mutation rate is chosen at $\sigma_\mu = 0.0001$.

currently encountered environment. We see however that plasticity increases despite the immediate fitness benefits of non-plastic individuals. In addition, the residual genetic variation was decreased when compared to the case of extra-coarse-grained environmental variability (Figure 3.4, B). This is also indicated by the narrow gap between the top and the mean performance curve in Figure 3.4, A. As a result, the ability of natural selection to favour the evolution of adaptive phenotypic plasticity depends on the time-scale of environmental change in coarse-grained scenarios.

3.3.4 Low mutation rates can enhance the evolution of adaptive plasticity in coarse-grained environments.

In this scenario, we assess the impact of mutation rates on the evolution of adaptive reaction norms when individuals are exposed to a single environment during their lifetime, i.e., inter-generational environmental variability. We consider a small mutation rate, $\sigma_\mu = 0.0001$.

We find that when evolution occurs through small genetic changes, the population evolved to plastically adapt to its local habitats (Figure 3.5). We observe that both the goodness of fit to current and past environments decreased to zero, indicating that selection discovered an optimal reaction norm with respect to both current and past environments (Figure 3.5, A). As above, we see that plasticity increases despite the immediate fitness benefits of non-plastic individuals. In addition, the residual genetic variation was significantly decreased when compared to the case of extra-coarse-grained environmental variability (Figure 3.5, B). This is also indicated by the narrow gap between the top and the mean performance curve in Figure 3.5, A.

3.3.5 Lineage selection is not necessary for the evolution of adaptive plasticity.

We next ask whether adaptive plasticity arises as a result of lineage selection. We emphasise that the conditions that favour the evolution of plasticity without direct selection do not rely on lineage-level selection. That is multiple genetic lineages do not need to coexist long enough for their reaction norms to become visible to selection.

To assess whether lineage selection is needed for the evolution of adaptive plasticity, we model the evolution of a population under strong selection and weak mutation. Each new mutation is either fixed or lost before the next arises. A simple hill-climbing model of evolution is thus sufficient (see Hill-climbing Model). The population is evolved for 20000000 generations.

We find that the same qualitative results hold as in all experiments above (Figure 3.6 and 3.7). We expect that selection for increased fitness in the current environment would result in less plasticity. However, assuming strong selection and weak mutation, we observe a systematic increase in plasticity for i) fine-grained environments (Figure 3.6 B), ii) rapid environmental change in extra-coarse-grained environments (Figure 3.7 A) and low mutation rates in extra-coarse-grained environments (Figure 3.7 B). Under these conditions, the evolving system fails to evolutionary track changes in the environments without plasticity. This can result from the environment switching rapidly from one state to another, or from the evolving system not being able to genetically change fast enough relative to the rate of environmental change. Consequently, resolving the tension between selection for immediate fitness benefits and the evolution of adaptive plasticity does not require an explanation of model or lineage selection.

3.3.6 The role of model complexity in the evolution of plasticity.

Lastly, we investigate the role of complexity of the evolving reaction norms in the evolution of adaptive phenotypic plasticity. For simplicity, we have considered a simple

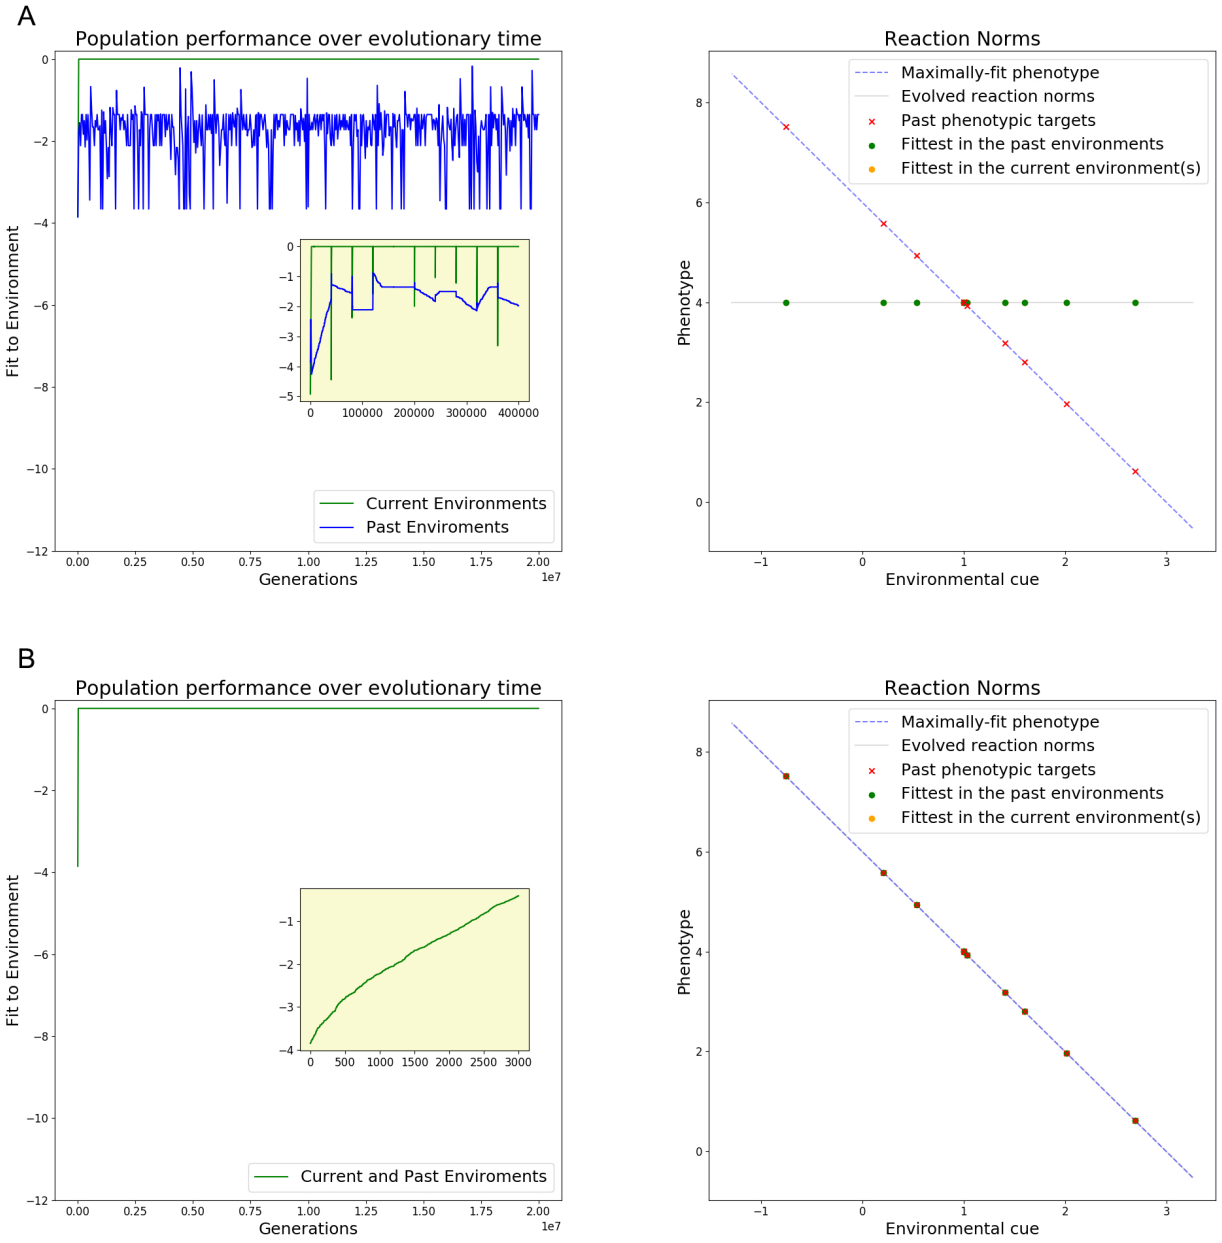

FIGURE 3.6: **The evolution of adaptive plasticity does not rely on lineage selection.** Left. Goodness of fit of the evolving reaction norms to the current environment (green) and all past selective environments (blue) over evolutionary time. Right. Evolved reaction norm (grey line). The red marks correspond to trait optima of the past experienced selective environments, while the blue dashed line corresponds to the optimal reaction norm. A. Coarse-grained environments ($K = 40000$). B. Fine-grained environments.

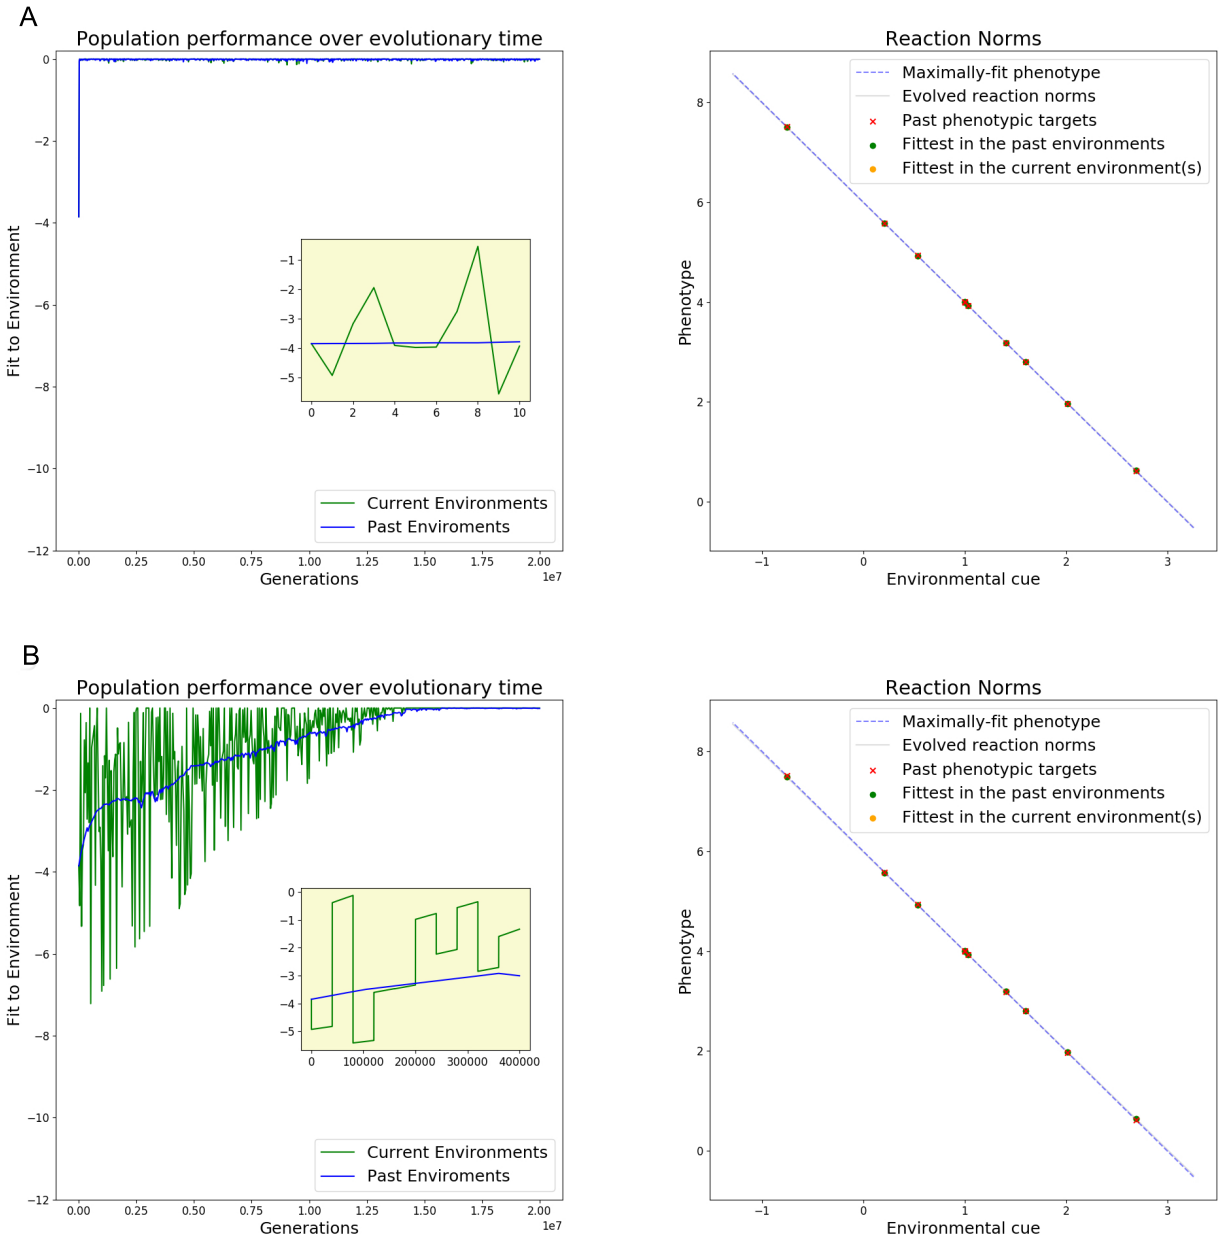

FIGURE 3.7: **The evolution of adaptive plasticity does not rely on lineage selection.** Left. Goodness of fit of the evolving reaction norms to the current environment (green) and all past selective environments (blue) over evolutionary time. Right. Evolved reaction norm (grey line). The red marks correspond to trait optima of the past experienced selective environments, while the blue dashed line corresponds to the optimal reaction norm. A. Fast environmental switching ($K = 1$). B. Low mutation rate ($\sigma_\mu = 10^{-5}$).

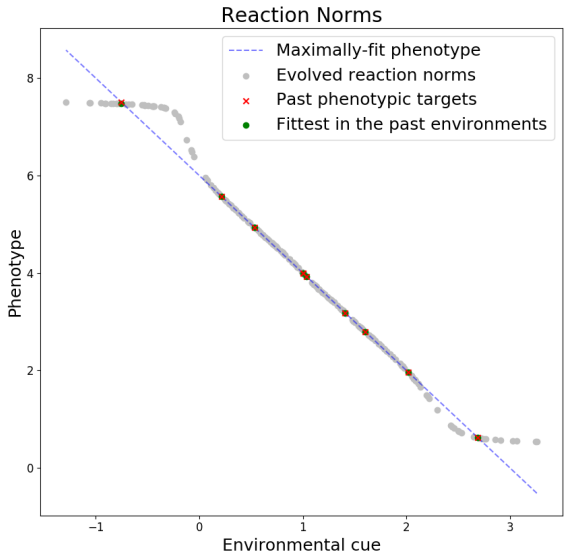

FIGURE 3.8: **Highly complex models tend to fail to capture the underlying regularities of environmental change.** Evolved reaction norm of a highly complex developmental system. Grey dots indicate the evolved reaction norm. The red marks correspond to trait optima of the past experienced selective environments, while the blue dashed line corresponds to the optimal reaction norm. The developmental model is characterised by a network of 4 layers of $10, 15, 2$ and 1 genes respectively. The mutation rate is chosen to be 0.2. The population is evolved for 4000000 generations in fine-grained environments under the assumption of strong selection and weak mutation. Selection optimises the phenotypic traits in all current and past environments, but leaves high phenotypic variance across a wider range of environments.

developmental model, which can be represented by a linear reaction norm. We can write the reaction norm as a linear equation, where the adult developed phenotype for a given environment x is given by $\theta(x) = \theta_1 * x + \theta_0$. The coefficients, θ_1 and θ_0 correspond to the slope and the intercept of the line respectively. The slope indicates sensitivity to environmental cues, and thus plasticity. Furthermore, the environment is also described in a linear fashion, where the phenotypic targets and the cues from the environment are characterised by a linear relationship, $g(x) = g_1 * x + g_0$. A linear developmental model is thus expected to be sufficient for the evolution of plasticity in the given problem.

In this scenario, we consider a developmental model of higher complexity. Specifically, the developmental structure is described by a gene network of 4 layers of $10, 15, 2$ and 1 genes respectively. The population is evolved for 4000000 generations in fine-grained environments under the assumption of strong selection and weak mutation. The mutation rate is chosen to be 0.2.

We find that natural selection discovered a complex developmental structure with the right environmentally-induced potential variability for the set of environments under selection. We see that selection optimises the life-history traits that are realised in all previously-seen environments. But selection leaves a lot of residual phenotypic variation

to be expressed in yet-unseen environments. We also see that highly complex models tend to increase the individuals' adaptability to their frequently encountered environments. However, they tend to lower adaptive potential of the individuals in novel environments that fall outside the range of their recent history. Consequently, we see that high model complexity creates a tension between local adaptation (i.e., what is useful now) and long-term adaptability (i.e., what is useful in the future).

3.4 Discussion and Conclusions

In this study, we investigate how adaptive plasticity can arise without a direct selection pressure. When retaining adaptive phenotypic plasticity in the short-term is unnecessary or even costly, then we would expect natural selection to favour a decrease in plasticity. However, we demonstrate that adaptive plasticity can reliably increase without lineage selection when there is no direct selection *for* plasticity, or plasticity is even selected against, under certain circumstances. Specifically, a systematic selection *of* plasticity in coarse-grained environment occurs when the population experiences rapid environmental fluctuations or is characterised by low mutation rates. In such cases, environmental sensitivity arises as a by-product of selection towards phenotypic targets in different selective environments across generations. Lastly, we show that the complexity of the developmental systems is crucial in understanding the evolution of adaptive plasticity.

3.4.1 Inter- vs intra-generational environmental variability.

The realised phenotypic variation upon which selection can act is determined by the environments the individual is exposed to during its lifetime (West-Eberhard, 2003). Whether organisms encounter environmental changes occurs across and within generations affect what can be seen by natural selection, and thus has important evolutionary consequences for plasticity. In situations where environmental variability within each generation is absent or individuals produce one phenotype over their life-time period (e.g., size of maturity), then only one component of the norm of reaction is expressed. Consequently, only the respective realised trait value (character state) is exposed and subject to individual selection in the given environment. Natural selection thus fails to see and evolve plasticity. Alternatively, when individuals encounter multiple environments during their life-time, multiple aspects of evolving reaction norm are expressed. In such situations, the evolution of plasticity is in a time scale that natural selection can see, it thus is not surprising that natural selection would favour it.

Indeed, we see that the slope of the average individual in a population evolves under extra-coarse-grained environmental variability tends to zero. Individual selection concerns with immediate fitness benefits. Plasticity here is not only unnecessary for the

individual to track environmental change, but is also costly. Hence, non-plastic individuals are always fitter. These results signify that no information is transferred from one generation to another, and thus selection fails to capture relevant information about the environment. On the other hand, under fine-grained environmental variability, the population quickly evolves towards the optimal reaction norm. Individual selection acts on multiple aspects of the reaction norm and thus there is direct selection pressure for plasticity (Via et al., 1995). Note that although plasticity is still inherently costly, in such situations plasticity is necessary for the individual to track environmental change.

Furthermore, selection tends to leave higher residual genetic variation when the environments change across rather than within generations. Such cryptic genetic variation can potentially aid future adaptive change. However, we see that it is undirected and does not entail any information gained from past selection. Thus, it does not offer an explanation on how past selection can reliably enhance the adaptive potential of the population. Similar results are presented by Lande (2009).

3.4.2 Insights from learning theory.

In learning theory, inter- and intra-generational environmental variation can be seen as analogous to online and batch learning respectively in learning theory (Bishop et al., 2006). In the former case, one training sample is presented at a time during the learning phase; while in the latter, all training samples are simultaneously presented. In online learning, the performance of the learner is thus evaluated over one training sample at a time, while, in batch learning, the performance of the learner is evaluated over all samples in the training set. This leads to qualitative differences in the optimisation process (learning) as the derived objective functions are different. As such, online and batch learning algorithms can yield different models with different predictive performance. This analogy between learning and evolution provides useful insights and helps to characterise general conditions that favour the evolution of adaptive plasticity under natural selection when plasticity is not directly selected for.

From learning theory perspective, it is not surprising that selection favours adaptive plasticity in fine-grained environments (i.e., batch learning). When the population was exposed to many environments within each generation, natural selection could act on multiple components of the norm of reaction the same way the optimisation process in batch learning is informed by error signals over all training samples. Similarly, we know from learning theory that capturing the regularities in a set of observations from presenting a single observation at a time (i.e., online learning) is possible. As such, the evolution of adaptive reaction norms in coarse- and extra-coarse-grained environments is also possible as a by-product of past selection.

Yet, we know from learning theory that prediction in online learning is sensitivity to the learning rates. When the learning rates are high, the learning system tends to focus on the current training sample and forgets about the previously-seen samples losing information about its past experience. Similarly, whether adaptive plasticity can evolve in coarse-grained environments should be sensitive to the amount of genetic change the population accumulates in each selective environment (i.e., learning rate). As a result, we expect that parameters that control the amount of adaptation in each selective environment, such as the rate of mutation and environmental change, to also affect the evolution of adaptive plasticity in coarse-grained environments.

The longer a population is exposed to a given selective environment, the higher the expected adaptation accumulated to that environment would be. The rate of genetic change in each environment (i.e., learning rate) can thus be controlled by the rate of environmental change (considering a fixed mutation rate). Accordingly, slow and fast environmental switches correspond to fast and slow learning rates respectively. When the environment changes very slowly (i.e., very high learning rates), natural selection evolves genetic representations that match the physiology of the organism to its environment. When the environment rapidly changes from one state to another (i.e., low learning rates), exposure to each environment state is too short to accumulate sufficient information about each of them. Thus, adaptation is also too slow in relation to the rate of environmental change. In such situations, we see that plasticity evolves as a suitable strategy to compensate for the inability of the organisms to cope with rapid environmental fluctuations through adaptive genetic differentiation.

Similarly, the higher the mutation rate is, the higher the genetic change the population accumulates to a given environment (i.e., learning rate), considering a fixed timescale of environmental change. When the mutation rate is very high (i.e., very high learning rates), individuals can evolutionarily track environmental change across generations. When the mutation rate is low (i.e., low learning rates), individuals fail to track environmental fluctuations through adaptive genetic differentiation, and thus plasticity arises.

Organisms can track environmental change either through genetic change, or through environmental induction. We see that high expected adaptation accumulated in each environment tends to lead to local adaptation through adaptive genetic differentiation. On the other hand, when adaptation is too slow, adaptive plasticity arises as a mechanism to cope with environmental fluctuations. This suggests that parameters that affect the speed of adaptation, such the population size and the strength of selection, can also affect the evolution of adaptive plasticity. Increasing the population size can increase the discovery of beneficial mutations (Goodell et al., 1997; Mustajärvi et al., 2001), while increasing the strength of selection can increase the spread of the beneficial mutations in the population. Both cases thus tend to lead to local adaptation through genetic change. Therefore, we expect that plasticity is more likely to arise in small populations

and/or under weak selection (Marshall and Jain, 1968; Froesch et al., 1985; Kery et al., 2000; Paschke et al., 2003).

3.4.3 The role of model complexity.

The complexity of the developmental systems is crucial in understanding the evolution of adaptive plasticity. Even in situations where the evolution of phenotypic plasticity is possible, whether the evolved plasticity is adaptive in novel environmental situations is also a matter of whether the model complexity of the developmental systems matched the complexity of environmental change. On one hand, a developmental model should be complex enough to capture the underlying regularities of the environment. On the other hand, highly complex models can lead to situations where plasticity is adaptive with respect to past experienced environments, but not in new environment situations. This holds even under conditions that enhance the evolution of adaptive phenotypic plasticity, such as fine-grained environmental variability and low mutation rates.

The high model complexity of development allows for high potential phenotypic variability which is not explained by the environment. This leads to situations where optimising performance to current and past environments does not necessarily optimise the behaviour of the system to new environments. Such behaviour is not obvious when linear models are considered in an environment that also varies in a linear fashion. In addition, we know that a linear reaction norm would only require selection over two environments to minimise its potential residual genetic variance (De Jong, 2005). Consequently, optimal reaction norms can arise as a result of direct selection in at least two environments. Such (under-)representations of development diminish the role of complexity in the evolution of plasticity.

3.5 Methods

In this work, we choose a set of environments that changes in a linear way. As a result, the optimal reaction norm is also a linear function. The choice of studying a linear reaction norm is two-fold. Firstly, linear reaction norms are widely-used in experimental biology, and thus the findings in the current Chapter can be directly applicable to experimental biological research (Lande and Arnold, 1983; Martin et al., 2011; Westneat et al., 2011; Sultan and Bazzaz, 1993; Heino et al., 2002). Secondly, linear reactions norms serve as theoretical models of plasticity that are well-understood in the biological literature (Chevin et al., 2010; De Jong and Bijma, 2002; Via, 1993; Via et al., 1995; Pigliucci, 2001). The aim of the current work is to demonstrate the possibility of the evolution of adaptive plasticity when plasticity is selected against. So, a simple and well-established model is used to question the assumptions made on the evolution of adaptive plasticity,

while providing conceptual clarity. Models of plasticity of higher complexity, such as higher degree polynomials, are beyond the scope of the current Chapter. A theoretical treatment of such models along with their potential implications in understanding the plasticty first hypothesis is presented in the following Chapter of the thesis.

The phenotypic targets are directly proportional to the respective cue. For the slope and the intercept of the optimal reaction norm, we choose a negative (-2) and a positive (6) value respectively. This ensures that the relationship between selective environment and cues remains constant across environmental states. Considering that the initial value for both the slope and the intercept for each organism is 0, a positive slope is initially favoured to better match the exact phenotypic trait values. We choose a negative slope to be opposed to that initial tendency of selection and ensure that when the slope of the reaction norm evolves to match the optimal reaction norms is an artifact of our model parameter choice. For the rest of the parameter values, such as the amount of mutation and the rate of environmental change, we show how they affect the evolution of adaptive plasticity in different evolutionary settings (i.e., different experiments).

The source code used for this chapter can be found here:
https://github.com/KostasKouvaris/Plasticity.

3.5.1 Developmental Network Model

We consider genotype-phenotype models, where the information is transferred in one direction, from the environment, e, to the adult developed phenotype, P, through gene cascades. Specifically, we describe each developmental system by a continuous, non-linear and feed-forward gene network, that is, it does not allow for feed-back connections). This model is inspired by transcriptional factors networks, where the expression level for each gene depends on 1) the activity of the genes that is connected to and 2) its pattern of connections, namely, how strongly the respective gene is connected to its neighbouring nodes. Each connection here is defined by a real-valued weight, $w \in \mathbb{R}$. Positive and negative values denote excitatory and inhibitory connections respectively. A value of zero denotes the absence of the respective connection between the two genes.

The network is structured in layers of genes. The gene expression level from one layer can be transferred to another without any lateral connectivity. The set of environmental cues denotes the input layer, l_0, the set of genes that receives the cues denotes the first layer, l_1, and so on. The last layer, l_N, denotes a set of phenotypic traits that corresponds to the adult developed phenotype, P, of the respective individual, upon which selection can act. Each layer, i, is characterised by K_i nodes, $p_k^i \in \mathbb{R}, \forall k \in \{1, .., K_i\}$. K_0 indicates the number of environmental cues received by the system, $e = P^0 =< p_1^0, e_2^0, \ldots, e_{K_0}^0 >$, while K_N indicates the number of phenotypic traits that comprise the adult phenotype, $P^a = P^N =< p_1^N, p_2^N, \ldots, p_{K_N}^N >$. The regulatory interactions are represented by a

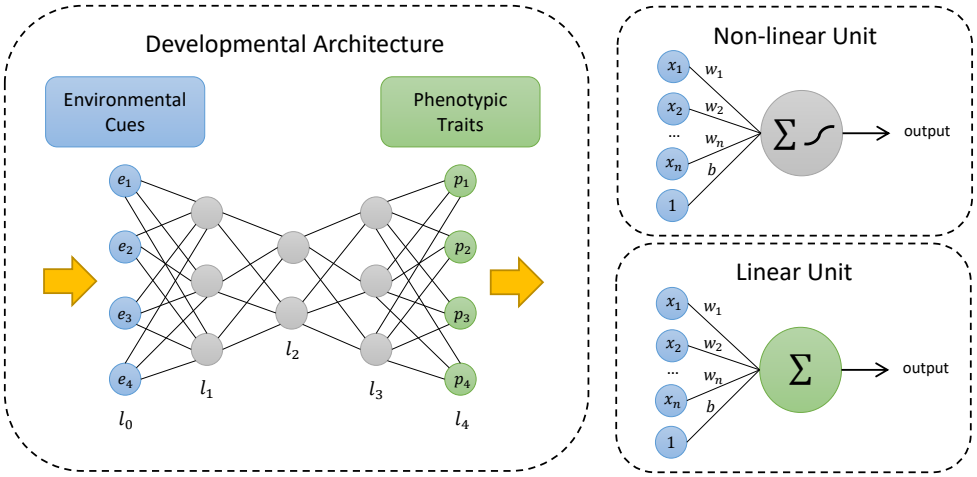

FIGURE 3.9: **Graphical representation of the developmental architecture.** Hypothetical developmental network of 4 layers. The network is structure in layers of genes, l_0 to l_4. The information is transferred in one direction from the environmental cues (input), l_0, to the adult phenotype (output), l_4. The blue circles indicate the cues received by the developmental system from the environment, while the green circles indicate the adult developed phenotypic traits upon which selection can act. The activity expression levels of the genes (grey circles) is described by a non-linear bounded function, while the realised phenotypic traits are described by a linear function. Each unit (excluding the input units) is characterised by a bias nodes, whose activity is always on. The bias affects how sensitive the respective gene to the net activity of the genes that is connected to.

set of matrices $W^i \in \mathbb{R}^{K_{i-1} \times K_i}$, along with an associated set of biases vectors, $B^i \in \mathbb{R}^{1 \times K_i}, \forall i \in \{1, .., N\}$.

In the first time step, a set of genes receives information from the environment. Thereafter, at every developmental step the expression level of the genes at each layer, l_i is estimated solely based on the expression of the genes at the previous layer l_{i-1}. The developmental cascades are estimated by the following set of equations, $\forall i \in \{1, .., N-1\}$:

$$p_k^i = \sigma(\sum_j w_{jk}^{i-1} p_j^{i-1} - b_k^{i-1}), \tag{3.1}$$

The right-hand side of equation (A.1) corresponds to the interaction term, the activity of which is limited by a non-linear, monotonic and bounded (sigmoid) activation function, $\sigma(x) = tanh(x)$. Finally, the output layer is consisted of linear units, where the phenotypic trait values are given by:

$$p_k^a = \sum_j w_{jk}^{N-1} p_j^{N-1} - b_k^{N-1}. \tag{3.2}$$

3.5.2 Evolutionary Process

We model the evolution of a population of developmental organisms as follows. Firstly, a pair of parents is selected with probability proportional to their current fitness (Hancock, 1994; Lipson and Pollack, 2000). The first parent is selected at random with probability $f/\bar{f}$, where $\bar{f}$ corresponds to mean fitness in the current population. Then, a second parent is selected in the same way. The parents should not be identical. For each pair of parents, a single offspring is generated by recombining the genotypes of the parents, followed by genetic mutation. Recombination is performed by selecting each connection from parent A or B at random with equal probability, that is, uniform crossover. Then, for each connection w_{ij}, an amount of mutation, μ, is added, where $\mu \sim \mathcal{N}(0, \sigma_\mu)$ and $\sigma_\mu = 0.01$. This process is repeated until a new population of the same size is constructed. All matrices W^i and B^i are initialised at zero.

3.5.3 Evaluation of Fitness

We assume a more conservative hypothesis where plasticity is not only unnecessary due to short-term fitness benefits of non-plastic individuals, but also inherently costly. Following previous work De Jong (2005); Kashtan et al. (2009); Draghi and Whitlock (2012), we define the overall fitness of a developmental system, f, in terms of a benefit-minus-cost function, d, as follows:

$$f = exp(\frac{d}{2\omega}), \tag{3.3}$$

where d corresponds to the net effect over all selective environments encountered by the individual during its life-time. ω is a scaling factor on the relation between f and d. We choose $\omega = 0.2$, which corresponds to a scenario of strong selection (see (Draghi and Whitlock, 2012)).

We model the fitness benefits a developmental system as follows. We consider that the individuals experience a distribution of selective environments during their lifetime with occurring probabilities, $q^{E_1}, q^{E_2}, .., q^{E_N}$. Each environment contributes to the selection process in proportion to its occurrence (De Jong and Bijma, 2002). Then, the overall fitness benefits of an individual over all experienced environments in its lifetime is determined by the arithmetic mean of the fitness benefits in each environment, b^{E_i}, weighted by the occurrence, q^{E_i}, of each environment:

$$b = \sum_i q^{E_i} b^{E_i}. \tag{3.4}$$

In cases of coarse-grained environmental variability, where each individual encounters a single environment in its lifespan, $q^{E_i} = 1$ for the respective environment, $i = j$, and

$q^{E_i} = 0$ for $i \neq j$. On the other hand, in cases of fine-grained environmental variability, we assume a uniform distribution of environments experienced during individual's lifespan, that is, $q^{E_i} = 1/K$.

The benefit of a given genetic structure, b^{E_i}, for each environment, E_i, is determined based on how close the developed adult phenotype is to the target phenotype, ϕ^{E_i}, of the given selective environment, E_i. The target phenotype characterises a collection of favourable trait values and is described by a vector, $\phi^{E_i} = \langle \phi_1, \ldots, \phi_N \rangle^{E_i}$, where $\phi_j \in \mathbb{R}, \forall j$. For a given selective environment, E_i, the selective benefit of an adult phenotype, P^a, is thus given by:

$$b^{E_i} = w(P^a, \phi^{E_i}) = -\|P^a - \phi^{E_i}\|_2^2, \tag{3.5}$$

where $\| * \|_2^2$ corresponds to the euclidean distance between the two respective vectors. Note that the selective advantage of respective developmental structure is solely determined by its immediate fitness benefits on the currently encountered selective environment(s).

The cost represents how fitness is reduced as a result of the system's effort to maintain and reproduce its ability to be sensitive to environmental changes. The plasticity cost term, c, is related to the values of the regulatory coefficients, $w_{ij} \in W^1$. The cost function here is proportional to the sum of the squares of the magnitudes of the interactions, $c = \|W^1\|_2^2 = \sum_i w_{ij}^2$.

Then, the overall performance of a developmental structure over a distribution of selective environments is given by:

$$d = b - \lambda c, \tag{3.6}$$

where parameter λ indicates the relative importance between b and c.

3.5.4 Environmental Variability

The population experience environmental heterogeneity where individuals receive information from the environment about its state. The population is exposed to a different selective every every K generations. K is chosen to be considerably smaller than the overall number of generations simulated. Here we assume that the lifespan of the individuals is fixed and equal among them. As a result, environmental granularity is solely determined by the parameter K. $K < 1$ indicates fine-grained environmental variability where the population encounters $1/K$ on average environments per generation. On the other hand, $K >= 1$ indicates coarse- ($K = 1$) or extra-coarse-grained ($K > 1$) environmental variability where the population encounters K environments per generation.

The environment here has two distinct roles West-Eberhard (2003). It can perform as an ''agent of development" affecting what phenotype is produced from a given genotype, or as an ''agent of selection" evaluating how well- or poorly-adapted the produced phenotypes are. We model a heterogeneous varying environment such as each selective environmental state is characterised by a single trait optimum, represented by a real number ϕ. Each target optimum is associated with an environmental cue, represented by a real number e that vary between 0 and 1. Environmental fluctuation is normally distributed, $e \sim \mathcal{N}(0, 1)$. For simplicity, we consider the phenotypic targets to have a linear relationship with the environment, $\phi = g(e) = g_1 * e + g_0$. Hence, the targets are directly proportional to the respective cue. We choose $g_1 = -2$ and $g_0 = 6$. This ensures that the environment changes in a systematic manner (i.e., share common regularities), and thus can provide informative cues.

3.5.5 Evaluation of Reaction Norms

Norms of reaction, or reaction norms, are used to visualise phenotypic plasticity (Sarkar, 2004). Reaction norms are graphical representations that describe the association of phenotypic responses to environmental change by plotting phenotypic values across the range of possible selective environments. Accordingly, reaction norms can be defined as sets of phenotypes that would be expressed if the given individual is exposed to the respective set of environments. Consequently, reaction norms can be used to characterise the potential of the environmentally-induced phenotypic variation of a single genotype. Note that, even if two individuals are characterised by the same reaction norm, their realised phenotype is determined by the environment they actually experience.

We evaluate the adaptive potential of the population due to plasticity by estimating how close the reaction norm of each individual in the population is to the (theoretical) optimal reaction norm. The optimal reaction norm here corresponds the function that describes how the phenotypic trait optimum changes along with the environment (Environmental Variability). So, we evaluate how adaptive reaction norms are based on how well they fit the optimal reaction norm. The goodness of fit, $Perf_D$ of a given reaction norm, D, is estimated as a function of the phenotypic trait values in each of the past selective environments (here 10), E_i, that quadratically decreases with the distance from each phenotypic optimum, ϕ^{E_i}:

$$Perf_D = \sum_{E_i} \|D(e^{E_i}) - \phi^{E_i}\|_2^2, \tag{3.7}$$

where $\| * \|_2^2$ corresponds to the euclidean distance between the two respective vectors. The evaluation of goodness of fit is performed for each individual at the end of each environmental period. The average and best performance in the population are reported.

3.5.6 Hill-climbing Model

A hill-climbing evolutionary model simulates a scenario of strong selection and weak mutation. Mutations in the weights W and B indicate slight variations in population means with W' and B' denoting the respective mutants. Then the adult mutant phenotype, $P^{a\prime}$, is the result of the developmental process, which is characterised by the parameters W' and B', given a cue from the environment, C. Subsequently, the fitness values, f and f' of P^a and $P^{a\prime}$ respective are calculated the distribution of the current selective environments. If $f' > f$, the mutation is beneficial and therefore adopted, i.e., $W_{t+1} = W'$ and $B_{t+1} = B'$. Otherwise, the mutation is deleterious and W and B remain unchanged.

Chapter 4

Extending plasticity-first evolution with learning theory

Abstract

The plasticity-first hypothesis suggests that plasticity may precede and even enhance the process of adaptive evolution. Although the plasticity-first hypothesis has a long history in evolutionary biology, it still remains a contentious issue. Part of the controversy stems from the fact that natural selection cannot favour plastic responses that are adaptive to yet-unseen environmental situations. Such capacity would imply a form foresight from natural selection, which is inconsistent with the existing evolutionary theory. As such, we argue that the plausibility of the plasticity-first hypothesis relies on the same conceptual issues as the evolution of evolvability. Here we argue that innovation in evolving systems is actually possible by recognising that this is equivalent to generalisation in learning systems. Learning systems do not really 'see into the future', but can nonetheless make useful predictions by generalising past experience. Such predictive performance follows simply from the ability to represent structural regularities in previously seen observations that are also true in the yet-unseen ones. We demonstrate that understanding the evolution of plasticity as analogous to task learning can provide useful insights about the evolutionary conditions under which plasticity can shape and guide adaptive evolution. The proposed conceptual framework can provide testable hypotheses for both theoretical and experimental work.

4.1 Plasticity-first hypothesis: Points of controversy

Evolution by natural selection is a process by which new variants spread in a population because of the benefits they confer to survival or reproductive success (Lewontin, 1970).

A prerequisite for adaptive evolution by means of natural selection is the availability of *recurrent* and *heritable* variation that gives rise to fitness differences in a population (West-Eberhard, 2003; Rachootin and Thomson, 1981; Rossiter, 1996). Explaining, however, how novel complex traits arise remains one of the most enduring problems in evolutionary biology (Gould, 2002; Mayr, 1960; Wagner and Lynch, 2010; Moczek et al., 2011). New characters or combination of characters can originate through genetic mutation, recombination of the genetic material, or environmental induction. All three are potent sources of evolutionarily significant variation as long as the novel genetic or environmental inputs are recurrent (West-Eberhard, 2003; Moczek et al., 2011).

Most researchers would argue that novelty solely arises through genetic change for the simple reason that evolution is defined as change in the allele frequencies (Carroll, 2008; Ghalambor et al., 2007; Pigliucci and Murren, 2003). This traditional view regards mutation as the prime mover and the only legitimate source of evolutionary novelty. As a result, evolutionary forces that act upon non-heritable variation had been historically considered as 'unimportant' (Endler and McLellan, 1988; Ghalambor et al., 2007; Wright, 1931; Simpson, 1953). Indeed, the process of adaptive evolution through mutation is only conceptually clear when a new mutation triggers reliable and advantageous phenotypic effects across genetic and environmental contexts. Mutation-driven evolution however is slow since mutations occur in single individuals and beneficial variants are frequently lost due to drift (De Beer, 1958; Lande, 1978; West-Eberhard, 2003; Moczek et al., 2011).

Yet, genetic change is not the only potent source of innovation. Many researchers state that environmental induction precedes and even facilitates the formation of evolutionary novelties, i.e., plasticity-first hypothesis (West-Eberhard, 2003; Pigliucci et al., 2006; Moczek et al., 2011; Schwander and Leimar, 2011; Laland et al., 2015; Gilbert et al., 2015; Sultan, 1995). In contrast to the gene-centric view of evolution, the evolution through environmental induction relies on recurrence of the environment and heritable variation in response to environmental change. Organisms can accommodate environmental cues in a functional manner which would lead to fitness differences (West-Eberhard, 2003). Natural selection can then differentiate among developmental variants and favour those genotypes which reliably and robustly produce functional responses, enhancing fitness in stressful environments (West-Eberhard, 2003; Badyaev, 2005; Ghalambor et al., 2015). Environmental induction can affect many individuals within the same generation, thereby making evolution proceed faster. Although the plasticity-first hypothesis (i.e., 'genes as followers') has a long history, it still is controversial for both theoretical and epistemological reasons (Pigliucci and Murren, 2003; De Jong, 2005; Pigliucci et al., 2006; Ghalambor et al., 2007; Moczek et al., 2011; Schlichting and Wund, 2014; Levis and Pfennig, 2016; Schwander and Leimar, 2011; Laland et al., 2015; Gilbert et al., 2015).

A common deflationary response to the theoretical significance of plasticity-first hypothesis is that plasticity merely serves as an intermediate step to a new genetically modified

and stable state that is phenotypically invariant (e.g., De Jong (2005); Ghalambor et al. (2007); Lande (2009)). Environmentally induced traits might indeed show some level of phenotypic adaptation in a different environment, even if this environment has not been encountered before. Such plastic traits lead to enhanced fitness for the respective individuals and thus genotypic differentiation appears at the population level. Subsequently, these traits will be selected and fixed in the population by allelic substitutions, i.e., genetic assimilation. As a result, the population becomes adapted to its new environment. In this model, such adaptive plastic responses can appear in the same random way as random mutations. The fact that individuals initially responded in a plastic manner is thus of little theoretical importance. Consequently, evolution still remains a matter of selection acting upon standing genetic variation with plasticity not providing any new insights into how evolution works.

The plasticity-first hypothesis would however be of theoretical importance if plasticity could guide future phenotypic exploration in promising directions. Plastic individuals should be capable of responding *directionally* to novel environments and show a predisposition towards producing phenotypically plastic variants that are better than random. Environmental induction should hence be able produce functional solutions to environmental problems that have not yet been seen by natural selection. Such capacity must however depend upon previous selection. In other words, natural selection has favoured this kind of plasticity in the past. True adaptive innovations then appear to be impossible, because it would imply that natural selection has a form of foresight, inconsistent with the existing theory.

Understanding the plausibility of the plasticity-first hypothesis requires an explanation about the predisposition of individuals to readily produce adaptive plastic phenotypes in new environments, which fall outside the range of their past history. However, most of the variation induced by the environment tends to be deleterious (Ghalambor et al., 2007; De Jong, 2005; Van Kleunen and Fischer, 2005). That is, there is a high demand for developmental homeostasis and thus buffering against environmental variability. Since natural selection concerns about immediate fitness benefits, selection would always favour the canalisation of past selected phenotypes in order to maintain functional stability and reliability (Waddington, 1942; Schmalhausen, 1949; Pigliucci, 2001). As a result, plasticity has been thought to slow down and even constrain adaptive evolution by shielding genetic variation from selection (Grant et al., 1977; Falconer et al., 1981; Sultan, 1995; Levin, 1988).

Next, we argue that such controversies stems from practical and theoretical considerations of how plasticity is modeled and represented in the literature.

4.2 Competing models of plasticity and their impact on the plasticity-first hypothesis

Advocates of plasticity-first have yet to produce discriminatory evidence in favour of the hypothesis (Via et al., 1995; De Jong, 2005; Futuyma, 2013; Levis and Pfennig, 2016). Although there is some evidence from lab studies (Waddington, 1953; Suzuki and Nijhout, 2006), evidence for an active role of plasticity is indirect (Ghalambor et al., 2007; Schlichting and Wund, 2014). These candidate case studies can equally well be explained by scenarios that rely on mutation or standing genetic variation without plasticity. We need a theory that can make clear, empirically and rigorously testable predictions of the plasticity-first hypothesis (Levis and Pfennig, 2016).

The seminal work of West-Eberhard (2003) on the plasticity-first hypothesis emphasises both the development and ecological processes in evolution. Development is a constructive process which makes extensive use of exploratory and selective processes. Since phenotypes develop, the responses of developing systems to new inputs become underdetermined by the genes they carry. On the ecological side, natural selection can work more efficiently on environmentally, as opposed to mutationally, induced variants (West-Eberhard, 2003). Despite that, theoretical and empirical studies have proceeded largely without considering plasticity as a developmental process (Bradshaw, 1965; Pigliucci and Murren, 2003; De Jong, 2005; Pigliucci et al., 2006; Ghalambor et al., 2007; Lande, 2009). A mechanistic model dealing with shaping of phenotypic outcomes by environmental cues is required to explain how plasticity molds the amount and direction of phenotypic variation realised in previously-unseen selective environments (Draghi and Whitlock, 2012).

Most of the research that addresses the role of plasticity in adaptive evolution deploys statistical tools used in quantitative genetics but not mechanistic models (Pigliucci, 2005; Pigliucci et al., 2006). Deducing possible causal mechanisms, historical paths or predicting potential future outcomes cannot be adequately studied using quantitative genetics, since many causal paths may lead to very similar outcomes. Similarly, empirical studies are typically concerned with demonstrating that ancestral plasticity mirrors derived phenotypes in ancestor-descendant comparisons (Doughty and Reznick, 2004; Gotthard and Nylin, 1995; Haugen and Vøllestad, 2000; Levis and Pfennig, 2016). For example, the quantitative genetic model by Lande (2009) reduces developmental processes to linear genotypic reaction norms. Arguably, such representations make plasticity a minor add-on to evolutionary theory, consistent with the general verbal criticism. Despite the manifestation of phenotypic innovation in extreme environments through plasticity, selection acts upon the genetic determinants of the slope of the reaction norm, i.e., standing genetic variation. These results are consistent with both plasticity-first and gene-first scenarios, and thus do not provide evidence in favour of one hypothesis over the other.

Furthermore, the lack of an adequate theory for the role of development in adaptive evolution has led to metaphors (e.g., in Ghalambor et al. (2007)). Reaction norms had been considered to be pieces of strings that selection keeps tight in a shape that is fit over the current and previously experienced environments. The string is loose and moves freely outside the history of recent environments, representing the additional cryptic genetic variation generated when the organism is environmentally stressed. Adaptive change outside the range occurs when the tension of the string extends to include a new environment, once the environment is encountered and thus selection can act upon. This agrees well with Lande (2009) model, and hence faces the same conceptual issues when explaining plasticity-first scenarios. Such (under-)representations of development diminish the constructive role of development. The idea that the string moves freely outside the range of recent history neglects the constraints and biases imposed upon potential phenotypic variation by the developmental process. Greater phenotypic variation may indeed arise in stressed environmental conditions, but it does not need to be entirely random (or undirected).

On the other hand, mechanistic models are powerful enough to capture the underlying process and the causal mechanisms of the modeled system (Bolker, 2008; Liberles et al., 2013). This improves our understanding of the systemic behaviour, something that could also be described by a phenomenological statistical model, but we also gain insight about the underlying mechanisms that give rise to such behaviour. Therefore, mechanistic models can offer an explanation in addition to a description of the relationship between variables of interest. They are also expected to work better when deployed for predictions through extrapolation beyond the observed conditions. Mechanistic models can describe the effects of systemic properties of development on the production of potentially useful environmentally-induced variation and how those properties change as a result of selection. As such, they can provide sufficient clarity on the causes and the consequences of plasticity in initiating and shaping adaptation. Over the past two decades, mechanistic models of development have been investigating the role of developmental plasticity in adaptation (Draghi and Whitlock, 2012; Ancel et al., 2000; Espinosa-Soto et al., 2011; Fierst, 2011). Yet, these studies do not explain adequately the nature of potential variation as mechanistic models do not typically provide general solutions.

Recent theory suggests a deep analogy between learning and evolution (Watson and Szathmáry, 2015b). Evolution by natural selection to algorithmic learning. Learning models can provide mechanistic models for evolutionary processes for which we can exploit learning theory to characterise their general properties. This theory has been successfully applied before in studies of the evolution of evolvability (Watson and Szathmáry, 2015b; Watson et al., 2016; Kounios et al., 2016; Kouvaris et al., 2017). Here we argue that the scientific problem domain for plasticity-first is the evolution of evolvability and hence it faces the same conceptual roadblocks. Both scenarios concern with future fitness benefits, especially in novel environmental situations. Selection though

cannot favour potential phenotypic variants for yet-unrealised fitness benefits. Such capacity appears to be impossible because it implies foresight. Furthermore, the need for a high level of developmental homeostasis makes it less obvious how selection can facilitate adaptive responses to environmental variation. This creates a tension between canalisation of past selected phenotypes and developmental flexibility for future environments. The same considerations have been made about tension between robustness and evolvability in the evolution of evolvability.

Learning theory can help to overcome the aforementioned conceptual issues. Our goal here is to provide conceptual clarity about how certain assumptions can affect the interpretation of evolution and obscure the important role of plasticity in the adaptive evolution. In this paper, we use learning theory to demonstrate the possibility of generalisation and innovation in evolving systems. Learning models can provide a better understanding on how the underlying structure of a developmental system affects: 1) the extend of diversity induced or revealed when the system experiences environments outside of its recent history; 2) how environmental cues shape phenotypic responses; and 3) changes in developmental organisation facilitating potential phenotypic traits. This improves our understanding of the evolution of adaptive phenotypic plasticity and its role in shaping and potentiating adaptation to novel environments. We specifically address the issue of the tension between canalisation of phenotypes selected in past environments and anticipation of phenotypes that will be fit in future environments by recognising that this is equivalent to prediction in learning systems. Such predictive ability follows simply from the ability to represent structural regularities in previously seen observations that are also true in the yet-unseen ones. Understanding the evolution of reaction norms as a form of task learning can provide useful insights about the conditions that enhance the evolution of adaptive phenotypic plasticity under natural selection. We recognise that learning systems do not really 'see into the future' but can nonetheless make useful predictions by generalising past experience. We thereby demystify the notion that short-sighted natural selection can produce novel phenotypic responses adaptive for previously-unseen selective environments. Moreover, we are now also able to characterise general conditions where this is possible.

4.3 The evolution of environmental induction as task learning

Here we argue that the way organisms evolve to systematically incorporate environmental influence can be seen and better understood from learning theory perspective as a task learning process.

4.3.1 The role of development

A crucial part in understanding how environmentally induced variants can be systematically adaptive in new environments is to account for development as systemic property of the organism. Plasticity is an inherent property of the developmental system that can facilitate evolution (West-Eberhard, 1989, 2003; Schlichting, 2004).

Environmental cues influence which phenotypic variants are expressed. Development (i.e., a single genotype) can thus accommodate a set of (conditionally-expressed) alternative phenotypes across a range of possible environments. Uncorrelated responses to environmental stimuli, i.e., random variation do not necessarily indicate a plasticity-first scenario. The system might discover a fit solution by accident. Instead, plasticity-first scenarios are evident when the organisms are capable of producing environmentally-induced phenotypes that are systematically better than random. To understand plasticity-first evolution, we need to examine not only the amount of potential environmentally-induced phenotypic variation, but also its direction.

Such a bias towards functional phenotypes is a strong indication that the developmental system holds useful information about how the environment changes. Organisms can acquire knowledge about their environments by internalising it into their ontogenetic structures through genetic evolution. The structure of development biases both (potential) genetically and environmentally induced phenotypic distributions. Indeed, internal and regulatory factors in development dictate the functionality of development and as such act like hidden parameters of a model that determine the overall behaviour of the system. Those parameters are necessary and sufficient to store information about the environment.

4.3.2 Learning as a Search Problem

So what is learning? A simple way of defining learning is as any change in the learning system that lead to higher performance the second time the system experience the same task (Simon, 1983). More formally, learning can be defined as inferring a target function, or concept, from a limited number of instances, or examples, drawn from that function, or concept (Valiant, 1984). The instances are presented to the learner in an input-output form. Let X be an instance space; a set of all possible input-output pairs over which a target function can be defined. The set of these exemplary observations is known as *training set*. Then, the task of the learner is to find an approximation of the given target function.

Learning can be viewed as a search problem in a multidimensional space for a state that optimises a pre-determined *objective function*, J, known as, error, cost or criterion function, with respect to the underlying problem function (Hassoun, 1995). In polynomial

curve fitting, for example, the search space corresponds to all possible weight vectors of the respective coefficients. Hence, we can see learning as searching for a vector of coefficients that optimises an associated objective function. Consider $\theta \in \Theta$ represents the vector of the parameter values of the learning system (e.g., the set of coefficients of the polynomial), where Θ corresponds to the set of all possible parameter states. In learning theory, this space is known *hypothesis space* and refers to a family of functions, or concepts, $H = \{h_\theta : \theta \in \Theta\}$. We can then define the objective function as $J = J(\theta, x, z)$, where x indicates an input pattern and z a teacher or reward signal. The input signals are assumed to be generated from the environment following an unknown distribution D over X with a probability density $p(x, z)$. The learning signals depend on both the actual output of the learning system and the desired output, or reward signal respectively. After experiencing all of the training examples, the learning system attempts to find a hypothesis h that best approximate the target function. The learner then has to guess the target values based on the hypothesis h. The general idea is that the closer the answers are to the target values for a given hypothesis h, the closer that hypothesis is to the target concept. Thus, we see that learning refers to any *change* in the information processing of the system that is *adaptive* with respect to a given task.

4.3.3 Evolution of reaction norms: Learning

Evidence that plasticity has evolved can be seen as a change in the reaction norm (Doughty and Reznick, 2004). The evolution of reaction norms can be seen as a form of learning, if we view the part of the genome that describes the process of regulation (i.e., the program) as the corresponding hypothesis in computational learning (Fig. 4.1). Consider environmental cues E that are associated with, but different from the target phenotypes S characterising the selective environments. This enables us to investigate the evolution of reaction norms and more notably, the conditions that can improve the shape of reaction norms to better match the structure of the environment. We can thus view the evolution of plasticity of a given genotype as inferring input-output associations between environmental goals and their respective environmental cues. Development acts as a discriminative model. This type of problems falls in the category of supervised learning problems. The training data consist of a set of input vectors (i.e., environmental cues) and their respective target vectors (i.e., phenotypic optima). Consequently, the evolution of developmental organisations in a fitness maximisation manner by natural selection is directly analogous to a learning process by error minimisation, namely, searching the hypothesis space for better hypotheses that generate outputs closer to the desired ones.

Evolution can accumulate information regarding its past behaviour and respond to selection differently in the future (Frank, 2009; Shalizi et al., 2009). Given the fact that organisms experience a limited number of selective environments, different past selected

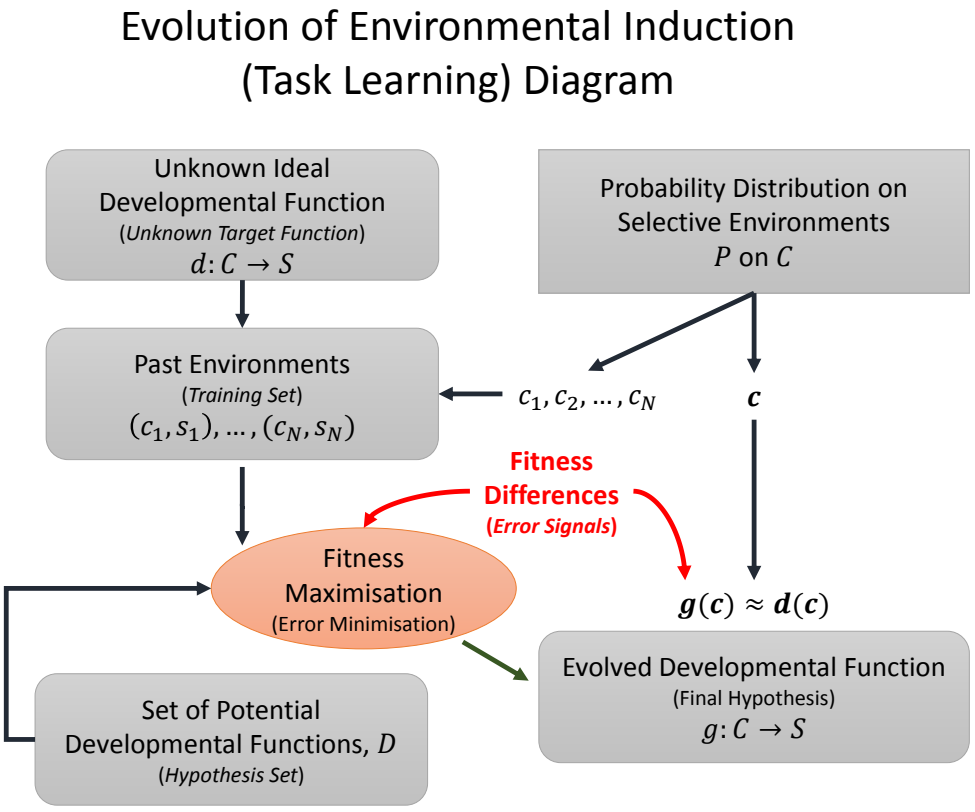

FIGURE 4.1: **Evolution of reaction norms diagram.** The diagram illustrates the evolution of adaptive plasticity as task learning process. Selective environments are assumed to follow a distribution D. Each environment generates a vector cues which is associated with the respective target vector, i.e., phenotypic optimum. The genetic determinants of the reaction norms (e.g., the set of coefficients of a polynomial) are described by the hypothesis space with genomes corresponding to different hypotheses. Evolving plasticity in a fitness maximisation manner by natural selection is directly analogous to a learning process by error minimisation, which involves searching the hypothesis space for better hypotheses that generate outcomes closer to the desired ones. This type of problems falls in the category of supervised learning problems.

targets provide different representativeness of the underlying environmental structure. Evolution of plasticity thus exhibits a path-dependent dynamical behaviour. Future evolutionary trajectories and hence adaptive potentials are dependent on the whole history of all previous states (Riedl and Jefferies, 1978).

4.4 Phenotypic Innovation: The problem of Induction

Here we show how selection to past selective environments affects the ability of development to predispose the production of adaptive phenotypes in new environments, outside the range of their past history.

There are countless distinct possibilities that can occur in real-world environments. Generalisation from past experience is therefore a crucial part of coping with previously-experienced and new situations (Luger, 2005; Valiant, 2013; Bishop et al., 2006; Abu-Mostafa et al., 2012). The system should ideally improve its performance on the next repetition of an instance experienced before, but also on similar instances, or tasks, in the same domain. Yet interesting domains tend to be enormous, but the system is provided with a fraction of all possible instances or the class. The system then has to generalise from that limited experience to new unseen instances drawn from the same class. This is known as the problem of induction. Yet, given the limited number of training examples, there would potentially be multiple hypotheses that explain the given observations equally well, i.e., under-determination. It is quite difficult in practice to get the hypothesis with the minimum possible true error. In addition, different training sets provide different representativeness for the target concept. Since it is unrealistic to practically provide the learner with the perfect set of training examples, the expectation of learning the true class is relaxed. The system must generalise by selecting those aspects and features of its past experience that can be directly relevant to and useful for the future. The corresponding selection criteria are generally known in learning theory as *inductive biases*.

Depending on the fitness benefits phenotypic plasticity provides, natural selection can either favour the buffering or canalisation of the respective phenotype against environmental variation, or modify the respective plastic responses in a certain range of phenotypic variation adaptive to particular environments. However, there could be multiple developmental structures (e.g., polynomial coefficients) that give exactly the same or closely similar behaviour. Natural selection is oblivious to such kind of structural differences, since they do not provide fitness differences – or at least significant statistical difference to distinguish between different reaction norms. Favouring developmental structures that give adaptive plastic responses to new environments seems impossible without a form of 'foresight'. It is not clear how natural selection would promote future adaptive exploration, favouring fitness benefits that have not been realised before.

4.4.1 Past and Future Environments Should Share Common Regularities

From learning theory perspective, the ability of learning systems to generalise is neither mysterious, nor taken for granted (Luger, 2005; Valiant, 2013; Bishop et al., 2006; Abu-Mostafa et al., 2012). Generalisation is possible when the context of the new environmental instances is not fundamentally different from the ones that it was originally made. This does not imply that the environment remains the same, but that there exist certain underlying regularities that remain time-invariant. Those regularities in fact dictate how the environment tend to change. There must be functional relationships

that characterise new possible instances arise and remain relatively the same over time. Learning systems do not really 'see into the future' but can nonetheless make useful predictions by generalising past experience. Similarly, plastic traits can be optimised to easily adapt to novel environmental situations considering that the phenotypic optima over past and future environments are described by similar functions, that is, past and future environments share common regularities. The potential of an organism to generate a range of fit alternative phenotypes in previously-unseen selective environments is possible despite the short-sighted nature of natural selection. Adaptive phenotypic plasticity can thus arise as an epiphenomenon of selection towards optimal phenotypic values in different selective environments as long as the regularities that are capture by the developmental organisms in the past are still relevant in the future.

The regularities that describe the environment should be detectable. Yet, the existence of such regularities is not sufficient to explain how it is feasible for an organism to acquire a detection function in the first place. If we want to analyse the capacity of a developmental organisation, and its evolution, in producing new and useful potential phenotypes in novel environmental contexts, a separation between past and future potential selective environments is essential. The past and current selective environments experienced by a population play the role of exemplary data that the system learns from, while previously-unseen environments play the role of new test data over which the ability of the system to generalise is evaluated. In learning theory, the former set is known as the training set and the latter as the test set.

4.4.2 Complexity of Developmental Model

The complexity of the developmental model indicates how flexible the model is. Model complexity is associated with the internal parameters of developmental systems, i.e., degrees of freedom. Such parameters include the number or the magnitude of the regulatory interactions, the number of interacting genes, topological features, such as feed-forward and recurrent (feed-back) structures or non-linearity. The complexity of the developmental system is a crucial determinant for the capacity of the evolving organisms to produce adaptive responses to novel situations. Considering different models of reaction norms (i.e., hypothesis space) can give rise to different quantitative and qualitative evolutionary outcomes.

Previous studies tend to encompass mathematical descriptions of plasticity that perfectly match the underlying relationship between environmental cues and targets (Lande, 2009; De Jong, 2005, 1995; Via et al., 1995). For example, the proposed model by Lande (2009) incorporates linear reaction norms in a linearly varying environment, while De Jong (2005) chooses the degree of the polynomial based on the number of the selective environments, e.g., a linear reaction norm in at least 2 environments or a quadratic reaction

norm in at least 3 environments. The complexity of the developmental model is appropriately tuned based on the number of points we want to fit. Such assumptions weaken the importance of the complexity of development and its vital role in induction to new environments. Different developmental functions (i.e., hypothesis spaces) can capture different aspects of the problem.

Firstly, models should be complex enough to capture the underlying regularities of the environment, which means that the regularities must be detectable by the evolving system. This condition is necessary for any form of information acquisition (learning) about the environment. Simple developmental models are characterised by a few degrees of freedom and tend to produce robust behaviour by remaining appropriately inflexible to environmental variation. If however the developmental model is too simple then it might be incapable of capturing noteworthy information about environmental change. This can leads to situations of under-fitting, where the system fails to fit the past observations (Luger, 2005; Valiant, 2013; Bishop et al., 2006; Abu-Mostafa et al., 2012). In the polynomial approach for example, we can think of it as fitting a horizontal (flat) line to a number of observations. The best case scenario would be for the intercept to adapt in such way that it provides the best fitness on average over the set of past selection environments. It means that the distance between the single phenotype expressed and the desired phenotypic optima in each environment is minimised (e.g., Fig. 4.2 A). Contrary to prior beliefs (De Jong, 2005), reaction norms do not have to perfectly fit past observations either. A linear model can still be tuned so that its slope and intercept fit past selected targets to a certain extend and still provide a behaviour that is better than random (e.g., Fig. 4.2 B).

The evolution of environmental induction is not solely affected by the complexity of the development, but also by the number of past selected targets. Another good example of under-fitting is the genetic model in Lande (2009). We know that a line can be fully define by two points. As a result, selection in two environments is sufficient to give rise to an optimal linear reaction norm. However, in Lande (2009) the population was exposed and subject to individual selection in each given environment separately. This allows for cryptic genetic variation that affects the slope of the reaction norm to build up, since just the threshold (i.e., parameter) is sufficient to fully represent the phenotypic trait value for a given environment. As a result, the evolving population failed to capture any useful information about how the environment tends to change even for those two environments, i.e., the slope of environmental change. The developmental system failed to evolve environmental sensitivity, implying that the potential environmentally-induced phenotypic variability was unaffected by past selection.

Nevertheless, highly complex models are not sufficient for good generalisation performance in novel environmental instances. If a model is characterised by high number of degrees of freedom the model tend to learn the idiosyncrasies of the past, resulting in over-fitting (Abu-Mostafa et al., 2012). Very flexible developmental models would tend

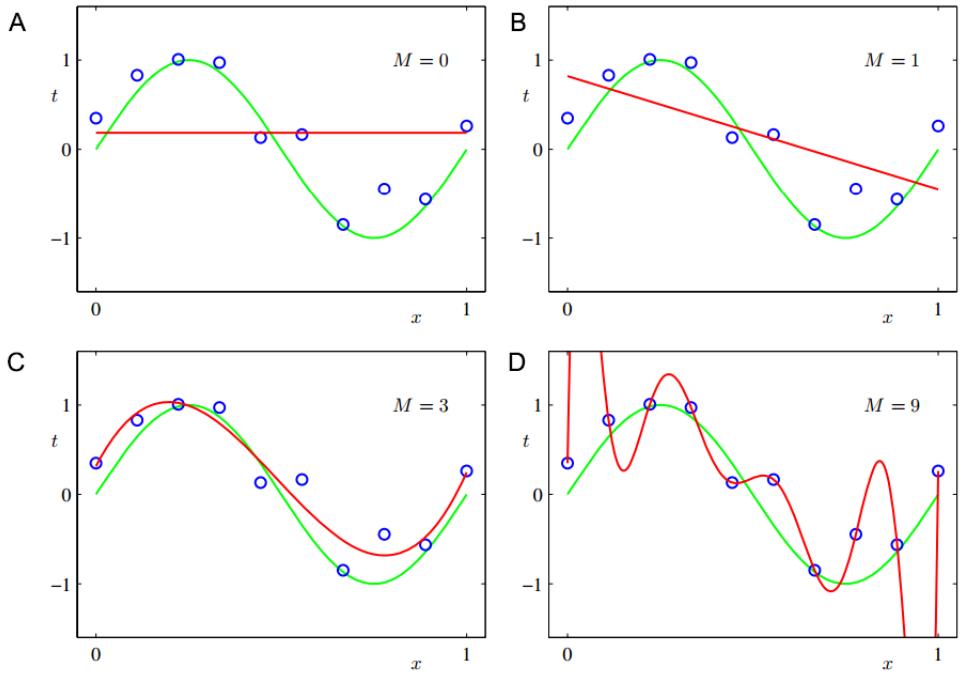

FIGURE 4.2: **Fitting polynomial reaction norms of various degrees, M.** The red curve corresponds to the fitted polynomial, while the green curve indicates the underlying regularity of the environment. The blue circles correspond to past phenotypic targets. Low and high degree polynomials tend to under- and over-fit respectively past phenotypic targets, respectively. Adapted from (Bishop et al., 2006) (Appendix C).

to describe the stochastic effects from the past selective environments, rather than the underlying regularities of the environment that remain time-invariant (e.g., Fig. 4.2 D). Consequently, situations of over-fitting oppose to the capacity of the evolving systems to generalise well by fitting irrelevant information. The problem of over-fitting can be considered as a problem of *model selection*. It begins when the test error starts increasing, while the training error decreases (Fig. 4.4 A). The test error represents the error that incurred from new observations. As such, the test error indicates how well the evolved developmental system can do in future situations, not yet experienced situations. Often dramatically, the training error tends to under-estimate the test error. This happens when the model is characterised by high complexity, and thus over-fitting. However, for simple models the training error can be a good estimation of the test error, i.e., under-fitting.

Developmental models of appropriate complexity are less sensitive to the sequence of the environments they experience and tend to avoid over-representing aspects of the past environment that are not relevant to future potential environments. Model-free approaches require an impractically high number of past samples to prevent high variance. Alternatively, selection in a few past environments may allow for hidden variation in the norm of reactions, regardless of the complexity of the model. If such organisms were exposed to a stressed environment, additional residual genetic variation would be

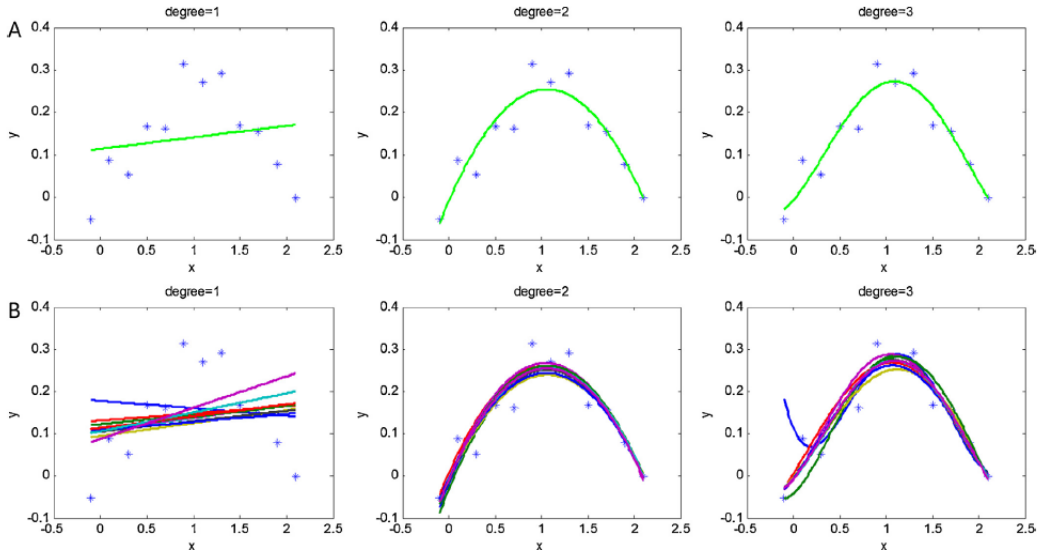

FIGURE 4.3: **Fitting polynomials of degree 1,2 and 3 respectively.** (A) A polynomial function of degree 1 (a line) under-fits the data, while a polynomial function of degree 3 over-fits the data. (B) The models are obtained using 'leave-one-out' method. Polynomial models of degree 2 tend to produce more stable solution, showing less variability when compared to polynomial models of degree 1 or 3. Adapted from (Deng et al., 2015).

expressed (e.g., Clausen (1940); De Jong (2005)). Moreover, character-state models of development where discrete separate traits are used for each environment, do not allow for extrapolation to new environments. Such developmental systems are characterised by 'infinite' degrees of freedom and can only perform rote learning, which results in inconsistent reaction norms. The variation being expressed outside the range of the past environments is undirected by past selection, and thus is completely random. Selection in a large number of environments and simpler developmental models can eliminate residual genetic variation leading to more promising evolutionary avenues.

4.4.3 Bias-variance trade-off

The problem of under- and over-fitting are two sides of the same coin. To better understand them, we have to introduce theory of the trade-off between bias and variance in model learning.

In principle, the generalisation error has two sources of error that prevent learning from generalising well to new observations outside the set of observations used in the learning phase: i) the bias and the ii) the variance. The bias reflects the ability of the model to accurately fit the data, whereas the variance reflects the stability of the model. Increasing the bias will decrease the variance, and vice versa. Furthermore, increasing the complexity of the model decreases the bias and increases the variance (Fig. 4.4 B). The more complex the model is, the more fitting ability it has, and the more unstable it tends to be.

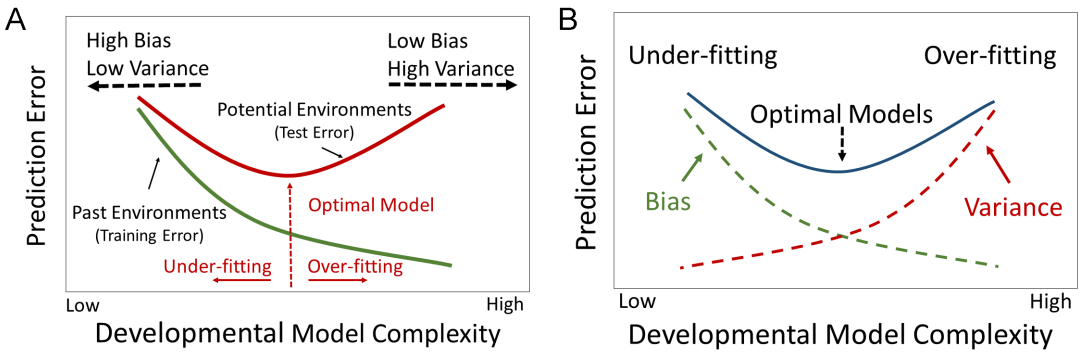

FIGURE 4.4: **Prediction error versus developmental model complexity.** (A) Behaviour of training and test error as model complexity varies. The prediction error has two sources of errors: the bias and the variance. The error sources are affected by the complexity of the model. A small difference between the training and the test error indicates under-fitting, while a large difference in errors indicates over-fitting. The prediction error can be decreased with model complexity to a certain extend. Over-fitting occurs when the training error continuously decreases, while the test error increases. This means that the model learns to represent the idiosyncrasies of the training data at the expense of generalisation performance. (B) The bias-variance dilemma in learning. Bias and variance decreased and increased, respectively, with increasing the complexity of the model. High bias indicates under-fitting, while high-variance indicates over-fitting. Hence, models of 'moderate' complexity should be favoured. Modified from (Deng et al., 2015).

In this context, simple developmental models are characterised by limited flexibility and unable to capture the true underlying environmental structure from past selective environments, leading to under-fitting. In contrast, highly complex developmental models are sensitive to the particularities of the past selective environments and capture spurious properties of their past targets, resulting in over-fitting. A useful indicator of the tension between bias and variance is the gap between the training and test error curves (Fig. 4.4 A). A small gap indicates under-fit reaction norms. This means that the evolving reaction norms tend to be consistent and insensitive to the set of past selective environments (low variance), but inaccurate on average (high bias). On the contrary, a big gap indicates over-fit reaction norms. This suggests that the evolving reaction norms are very accurate in fitting past targets (low bias), but inconsistent and very sensitive on average to the idiosyncrasies of previously encountered environments (high variance).

4.5 Conditions that facilitate plasticity-first scenarios

Conditions that affect the complexity of developmental system also affect the evolution of plasticty, and its ability to produce adaptive responses to new environments. According to learning theory, imposing inductive biases that penalise the complexity of models can improve the generalisation performance of the derived learning models on specific tasks. Adaptive phenotypic plasticity can thus be enhanced in the same fashion by evolving

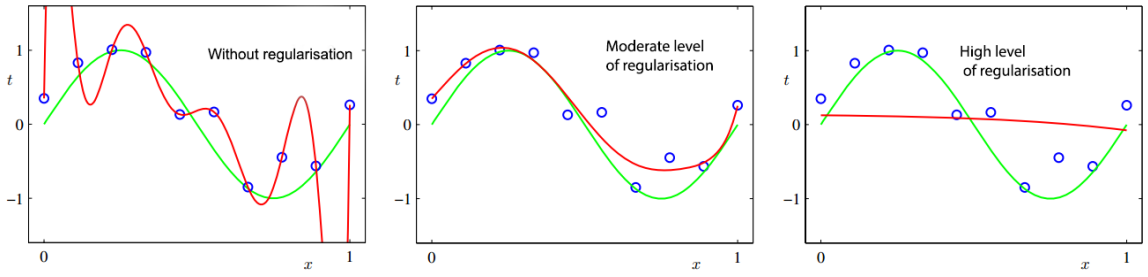

FIGURE 4.5: **Penalising model complexity improves generalisation.** Fitting a polynomial of degree 9, in situations of no, moderate and high level of regularisation. High parsimony pressure tends to lead to situations of under-fitting. Modified from (Bishop et al., 2006) (see Appendix C).

more parsimonious genetic structures which give rise to useful phenotypic variants in the context of potential future selective environments. Here we review inductive biases with biological analogues that favour parsimonious models, i.e., Occam's razor.

4.5.1 The role of cost and constraints in developmental evolution

Organisms which are characterised by a high degree of plasticity pay a fitness decrement compared to less plastic organisms (Murren et al., 2015; Van Tienderen, 1991; Moran, 1992; DeWitt et al., 1998). Consequently, the evolution and the nature of plasticity may be biased by their associate costs (Pigliucci, 2001). Constraints and limits on plasticity have many important ecological and evolutionary consequences (DeWitt et al., 1998; Lande, 2014; Ernande and Dieckmann, 2004). For example, the cost of plasticity may affect the level of plasticity (León, 1993; Auld et al., 2010; Schleicherová et al., 2014), or the ability to maintain genetic diversity for plasticity (Tauber and Tauber, 1992; Auld et al., 2010). The majority of the costs tend to be environment-specific (Auld et al., 2010). Varying the relative amounts of costs associated with the maintenance and production of plasticity can lead to different evolutionary consequences resulting in a compromise between a flat and an optimal reaction norm (Ernande and Dieckmann, 2004). For instance, maintenance costs led to populations that are closer to the adaptive peak of the most frequently experienced selective environment. As long as the inherent cost of fitness for maintaining and producing plasticity is sufficiently small (Sultan, 1995; Sultan and Spencer, 2002), plasticity is shaped by past selection and maintained, as suggested in the standing theory.

According to learning theory, associated costs with internal parameters of the developmental systems responsible for plasticity can be viewed as inductive biases that provide a parsimony pressure towards simpler, more general models (Bishop et al., 2006; O'Reilly and Munakata, 2000). The complexity of the evolved developmental model decreases as the parsimony pressure increases. When the learning system is highly complex, it tends to capture the idiosyncrasies of the exemplary data used in the training phase and not the

underlying structure of the process that generated those data. This prevents the learned model from generalising well to new situations. A parsimony pressure that favours simpler models would then enhance generalisation by preventing over-fitting (Fig. 4.5 B). A direct selection pressure on the cost can have an effect of feature selection, or imposing restrictions for smoothness, thereby preventing the system from fitting the noise in the training data. Accordingly, an evolving developmental system that is characterised by high complexity would tend to capture the irrelevant information (noise) in the past selective environments and not the underlying regularities (relevant information) in the environment, producing inappropriate phenotypic responses to new previously-unseen environments (Fig. 4.5 A). As a result, a fitness cost of plasticity can have a positive impact on the evolution of adaptive phenotypic plasticity enhancing the ability of the system to produce useful phenotypic responses to new environmental situations. It has been suggested before that selection can act on multiple character states when reaction norms differ on their reproductive and maintenance cost (Van Tienderen, 1991). Hene, selection would favour reaction norms that are less costly to maintain, acting indirectly on components of the reaction norm that are not currently expressed.

4.5.2 The role of intrinsic and extrinsic noise in developmental evolution

Stochasticity refers to the property of biological systems to behave in a probabilistic manner and is a fundamental aspect of how biological systems function and evolve(Losick and Desplan, 2008). Recent studies suggest that apart from the role of noise in physiology and differentation, noise can also play an essential role at the evolutionary level and can even facilitate adaptive developmental evolution by increasing potential phenotypic variability from single genotypes (Eldar and Elowitz, 2010; Losick and Desplan, 2008; Raj and van Oudenaarden, 2008; Thattai and Van Oudenaarden, 2004; Pélabon et al., 2004; Paenke et al., 2007). The development of an embryonic phenotype into an adult one is often influenced by many factors which can affect organism's fitness in current and future environments (Eldar and Elowitz, 2010). Such factors include genetic or transgenerational epigenetic effects, developmental noise and environmental stochastic fluctuations, and are investigated in the same framework (McNamara et al., 2016). Thus, noise naturally arises during the evolution of plasticity from either intrinsically developmental processes or extrinsically by being present in the environment as a source of unobserved variability (Swain et al., 2002; Forde, 2009). For instance, the inherent stochasticity of gene expression levels due to transcription and translation introduces a level of randomisation in development, defined as intrinsic noise (Eldar and Elowitz, 2010). At the same time, noisy environmental cues can introduce unreliable and not perfectly informative cues into development producing extrinsic noise (McNamara et al., 2016; Kight and Swaddle, 2011).

In learning theory, noise can enhance phenotypic innovation in two ways. Firstly, developmental noise can make it difficult or even impossible for natural selection to precisely match the phenotypic optima of the respective selective environments (Yoshimura and Shields, 1995). Indeed, it is known that training with noise (jittering), can improve the generalisation performance of the learning system (Bishop et al., 2006). Intuitively, noise smears out the training data and thus it hinders the model from fitting the training data precisely. It therefore prevents the irrelevant information from being learnt, and thus avoid over-fitting. Secondly, developmental noise helps to remove unnecessary inter-dependencies between various components of development, favouring flexible developmental structures. Techniques, such as DropOut (Hinton et al., 2012), tend to alleviate the problem of overfitting by preventing the emergence of complex co-adaptations among feature-detectors, thereby favouring simpler models. Recent theoretical work demonstrate that DropOut belongs into the same class of regularisation as jittering (Bishop, 2008). This suggests that noisy environments will have a similar effect on costly regulatory connections, expanding the conditions supporting the evolution of potentially adaptive phenotypes.

4.5.3 Inductive biases are always good in moderation

Like generalisation behaviour in learning, the potential of an organism to generate a range of fit alternative phenotypes in different environments is sensitive to the level of parsimony pressure. Inductive biases determine which models are favourable. In the extremes, low and high levels of parsimony pressure can give rise to diametrically opposed effects. On one hand, very strong inductive biases can impose strong constraints preventing the evolution of development to even internalise useful information about the current selective environments (i.e., training set) (Fig. 4.5 C). On the extreme, the evolved developmental system is incapable of environmental sensitivity, producing a single phenotype regardless of any environmental change. The evolving developmental system thus fails to locally adapt to the environment resulting situations of under-fitting. On the other hand, very low inductive biases permits very flexible development models to capture aspects of the current environment which are irrelevant with respect to future potential environments (i.e., test set) (Fig. 4.5 A). The evolving developmental system thus reduces it adaptability over evolutionary time. Finally, inductive biases can be seen as indirectly controlling the complexity of the developmental model, and thus the bias-variance tradeoff. Likewise, high bias (i.e., robust models) tends to lead to situations of underfitting, while high variance (i.e., flexible models) tends to lead to situations of overfitting (Bishop, 2008; O'Reilly and Munakata, 2000).

4.6 Combinatorial generalisation: Beyond extrapolation

Deploying linear statistical models permits simple extrapolations, but fails to capture cases of non-trivial generalisation. While linear models can allow for a certain level generalisation, they cannot capture any non-linear relations between environmental cues and phenotypic optima. For instance, linear correlation based models are not sufficient for learning a task which is characterised by overlapping patterns. This, for example, can be seen in the XOR problem which belongs in general family of non-linearly separable problems (O'Reilly and Munakata, 2000). Here, linear developmental systems fail to adapt to situations with a strong overlapping between different patterns of environmental cues. Similarly, dynamical systems characterised by linear interactions fail to give rise to distributions of multiple phenotypic attractors and generate complete and distinct phenotypic patterns (Watson et al., 2014).

However, there are more sophisticated and powerful forms of generalisation than simply extrapolation of previous learned behaviour outside the past 'known domain. The system can create new knowledge by discovering and subsequently systematising emerging patterns from past experience. It can separate that knowledge from the originally observed context to be re-deployed in new situations. Generally, learning is contextual and hence knowledge is built upon previously acquired concepts. In an evolutionary setting, concepts may correspond to repeated phenotypic sub-patterns or associations between environmental cues and phenotypic traits which persist over environmental change and become encoded in the structural components of the evolved developmental system. These concepts can arise through internal selection and provide re-usable building blocks which can be re-deployed in new contexts, i.e., evolution as 'tinkering' (Jacob, 1977). The evolved system's representation can be seen as dictating a higher-order conceptual (combinatorial) space based on previous experience and facilitating an abstraction of previously acquired knowledge. This enables natural selection to explore permitted combinations of features constrained by past selection. Hence, novel phenotypic responses can arise through new combinations of previously selected concepts explicitly embedded in the developmental architecture of the system. We see that generalisation can occur by storing and recalling previously-acquired concepts in combinatorial fashion (Watson et al., 2014; Watson and Szathmáry, 2016; Kouvaris et al., 2017; Kounios et al., 2016).

Modularity plays an important role in combinatorial generalisation as it facilitates variability that respects the natural decomposable structure of the selective environment. This means that modularity keeps those things together that must be kept together, and separates those things that are independent. Canalisation occurs on the correlations/associations between intra-modular phenotypic traits. This gives rise to modular network representations characterised by phenotypic sub-attractors. As a result, the system remains appropriately flexible to explore new combinations of modular patterns. Contrary

to previous beliefs based on the concept of extrapolation (De Jong, 2005; Ghalambor et al., 2007), this form of induction can allow predictions for values in environments that are not necessarily neighbouring. It is even possible, especially when multivariate models are considered, that certain phenotypes, perhaps distant in the space of phenotypic traits, to become accessible with small environmental variation.

4.7 Conclusions

Traditional views suggest that developmental plasticity buffers against environmentally induced variation and thus hinders adaptive evolution. Even if novel traits start as conditionally expressed alternative phenotypes, plasticity is of little theoretical importance, since the realised phenotypic variation tends to be undirected and does not provide any new insights into how evolution works. Instead, explaining whether plasticity can precede or even facilitate adaptive evolution requires an understanding of how development can reliably predispose the production of adaptive phenotypic variants in new environments. As such, the plausibility of the plasticity-first hypothesis relies on the same conceptual arguments as the evolution of evolvability. The discovery of new and useful phenotypes depends on the amount, but also the direction of the potential environmentally-induced phenotypic variation. Past selected developmental biases can constrain future evolutionary exploration into more promising phenotypic regions. Learning theory can explain how the developmental system can systematically evolve to facilitate new potential, not yet-realised phenotypic traits that are adaptive to new environmental conditions as a by-product of past selection towards past phenotypic targets. Consequently, a direct selection on the model is not necessary. Optimal models can be indirectly favoured when certain inductive biases, such as the cost of connections or developmental noise, are provided.

Chapter 5

Conclusions and Future Work

5.1 Understanding the evolution of adaptive potential by means of inductive biases

The capacity of the organisms to produce adaptive phenotypic variation in new environmental situations is crucial for the survival of a lineage and its future evolution. In general, phenotypic variability is dictated by the underlying structure of development. Development is the constructive ontogenetic process that translates a genotype into a phenotype. Ontogeny, however, is genetically and environmentally determined. Hence, genetic mutation and environmental induction can be both potent sources of evolutionarily significant variation. The current dogma in evolutionary theory emphasises on mutation as the main driver of evolutionary change. Adaptive variants arise through random genetic variation and then spread in the population by natural selection. Such adaptive traits might appear, but not necessarily so. Mutations occur in single individuals and beneficial variants are often lost due to drift. To explain how adaptive phenotypes can reliably arise, we need to understand how developmental biases and plasticity can guide phenotypic exploration, by channeling undirected genetic variation into directed adaptive phenotypic variation.

The evolution of adaptive potential has been previously studied from the perspectives of the evolution of evolvability (Draghi et al., 2010; Wagner, 2013; Watson et al., 2014; Altenberg, 1994; Kashtan et al., 2007; Parter et al., 2008; Pigliucci, 2008; Hendrikse et al., 2007; Kirschner and Gerhart, 1998; Aldana et al., 2007; Houle, 1992; Wagner and Altenberg, 1996) and the evolution of adaptive plasticity (West-Eberhard, 2003; Moczek et al., 2011; Schwander and Leimar, 2011; Laland et al., 2015; Gilbert et al., 2015; Sultan, 1995; Pigliucci and Murren, 2003; De Jong, 2005; Pigliucci et al., 2006; Ghalambor et al., 2007; Schlichting and Wund, 2014; Levis and Pfennig, 2016). In this thesis, I argue that these two concepts have theoretical unifying themes and face the

same conceptual issues, and thus they can be integrated and better understood under a unified framework.

Understanding the plausibility of the plasticity-first hypothesis and the evolution of evolvability requires an explanation about the predisposition of individuals to readily produce adaptive plastic phenotypes in new environments, which fall outside the range of their past history. Yet, the evolution of such adaptive potential is un-Darwinian. Selection cannot favour traits for fitness benefits that have not yet been realised. It is thus unclear why selection would favour biological structures in the present environments that promote good phenotypic variants in future previously unseen environments. Such capacity would imply that natural selection has a form of foresight, inconsistent with the existing theory. Furthermore, selection would always favour for more stable and reliable developmental systems. We therefore expect selection to make past selected phenotypes more resistant against genetic and environmental variation. Such developmental canalisation seems to be intrinsically opposed to an increase in phenotypic variability. This creates an evolutionary tension between robust and flexible developmental structures – a prominent issue in both the evolution of evolvability and the evolution of adaptive plasticity. Then, how is it possible for natural selection to evolve developmental organisations that facilitate higher potential for adaptive phenotypic variation in the future?

A resolution is provided by learning theory. Evolving systems and learning systems share common underlying, conceptual and algorithmic, principles (Watson and Szathmáry, 2016). In this work, I show that this analogy between learning and evolution can help us make predictions about the evolutionary conditions relevant to the evolution of adaptive potential. Specifically, I demonstrate that the evolution of adaptive potential in evolving systems is actually possible the same way, and with the same limitations, as generalisation is possible in learning systems. More importantly, the tension between canalisation of past environments and anticipation of future environments is resolved, by recognising that induction in learning systems merely requires the ability to represent structural regularities in previously seen observations that are also true in the yet-unseen ones. We see that the ability of short-term selection to discover and exploit such regularities is analogous to the ability of learning systems to generalise from limited past experience. Generalisation in learning systems is commonplace and not considered mysterious. It is not really about the past or the future, but about generalising from the data you have seen to the test cases you have not. But it is also understood that successful generalisation in learning systems is not for granted and requires certain well-understood conditions. Understanding the evolution of developmental organisations that can facilitate adaptive plasticity and evolvability by means of inductive biases can provide useful insights and testable hypotheses about the conditions that enhance the evolution of adaptive potential under natural selection.

In Chapter 1, I show how generalised developmental organisations can arise through considerations of parsimony. Firstly, I demonstrate that the failure of natural selection to enhance evolvability is directly analogous to the learning problems of under- and over-fitting, and the subsequent failure to generalise. Evolvability can thus be enhanced the same way generalisation in learning systems is improved. I show that the same conditions that alleviate the problems of under- and over-fitting in learning systems successfully predict which biological conditions that enhance evolvability. Specifically, I demonstrate that adaptive developmental biases can evolve under i) moderate rates of environmental switching, ii) the direct selective pressure on the cost of the regulatory interactions and iii) the presence of extrinsic environmental noise. Under these conditions, the variational structure of the organism systematically evolves internal representations of the statistical correlation structure of past selective environments that remain time-invariant in the future. The dimensionality of the phenotypic space is decreased by facilitating the production of certain, potentially useful combinations of phenotypic traits, without a need for reduction in individual traits' variability. Natural selection can therefore create reusable parts that facilitate faster exploration in new directions. Modularity becomes an important characteristic of generalised developmental structures. Modular developmental structures exhibit high robustness *within* functional modules, by confining epistatic dependencies or pleiotropic effects caused by genetic perturbations within the modules, and at the same time, high flexibility *between* modules, by allowing for different clusters of characters to evolve independently of others.

In Chapter 2 and 3, I show that the same general conditions that enhance the evolution of evolvability also enhance the evolution of adaptive phenotypic plasticity. Unlike the evolution of adaptive developmental biases, which is formally analogous to model learning, we now think of the evolution of adaptive plasticity as a form of task learning. Hence, evolving appropriate environmental sensitivity for past and previously-unseen selective environments is characterised by the same limitations as a learning system faces when learning a given task. In Chapter 2, I demonstrate that plasticity reliably increases when environmental change occurs over generations and when plasticity is inherently costly. I conclude that the plausibility of adaptive plasticity to arise relies on the amount of adaptive change accumulated by the population at each given environment. Adaptive plasticity can reliably evolve as a by-product of selection towards phenotypic targets in different selective environments across generations under i) low rates of mutation, ii) high rates of environmental change and iii) complex enough developmental structures. This is analogous to prediction in learning systems being sensitive to the learning rate and the underlying complexity. Although here the evolved plasticity is adaptive in both past and future environments, this capacity is a trivial form of generalisation and is limited to our underlying assumption of linearity in both the given problem and developmental model. When highly complex models are considered, the evolving system tends to canalise past selected targets. In Chapter 3, I analyse the importance of this distinction between model and problem space, and how it affects our interpretation about the role

of plasticity in adaptive evolution. I show that development can reliably predispose the production of adaptive phenotypic variants in new environments, allowing plasticity to precede and even facilitate adaptive evolution. I conclude that the plausibility of the plasticity-first hypothesis relies on the same conceptual arguments as the evolution of evolvability, and thus the scientific problem domain for plasticity-first is the evolution of evolvability.

Overall, the findings presented in this thesis suggest that evolving systems and learning systems are different instantiations of the same algorithmic principles as existing results from the learning domain can be successfully transferred to the evolution domain. The principles of induction from learning theory can provide a better understanding of the predisposition of developmental systems to new phenotypic responses, which are genetically or environmentally-induced and are adaptive to novel environmental situations. Equating adapative potential with generalisation in learning systems demystifies the ability of short-sighted natural selection to produce novel phenotypes fit for previously-unseen selective environments. Although generalisation in learning system is not guaranteed, it does not require a mystical ability to 'see into the future'. This predictive ability simply stems from the ability to represent deep structural regularities in previously seen observations that remain time-invariant, and thus are also true in the yet-unseen ones. Understanding that the evolution of adaptive developmental biases and the evolution of adaptive plasticity are formally analogous to model and task learning respectively provide valuable insights on the mechanistic nature of evolution of evolvability and adaptive phenotypic plasticity as well as their interplay. The link between learning and evolution appears to be more than an analogy and opens-up a well-developed theoretical framework ready to be exploited in evolutionary theory.

5.2 Future research directions

The current work establishes a conceptual framework for the evolution of adaptive potential using principles of induction from learning theory. Here I suggest a series of new directions of research that can expand our understanding on how evolution can learn from past selective environments and generalise to new future environments.

5.2.1 The role of developmental architecture in the evolution of adaptive potential

In Chapter 1, I considered a single layer recurrent network model to investigate how past developmental biases can affect the amount and the nature of phenotypic variation in future selective environments. In Chapter 2, I considered feed-forward network models

of one or more layers to analyse the impact of environmental sensitivity on the long-term adaptability of the organisms when environmental change occurs within and across generations. Such assumptions about the architecture of the developmental networks however form different hypothesis spaces, and thus hold different potentials for the evolvability of the organism. The functionality of the developmental systems is heavily determined by their underlying structure, and therefore their network architecture. Since there is no scientific consensus on the preference of developmental models (Aldana et al., 2007), future studies should test the effectiveness of the (exogenous) inductive biases and conditions that enhance evolvability (e.g., environmental noise and the direct pressure on the cost of connection) in different models of development. How is the evolution of adaptive potential affected by the topological properties of development? Note that the structure of network itself imposes a form of (endogenous) inductive biases to the learning process.

According to learning theory, fully recurrent networks are characterised by a simple feedback mechanism, where the output of the network recurrently serves as a new input for the network until the network configuration is converged to a stable state (Hassoun, 1995). The system performs an error correction process that induces a dynamical behaviour at the system level with information being retrieved as an evolution of the state in high-dimensional space. Recurrent networks therefore have an inherent tendency to form attractors by mapping neighbourhoods of vectors (basins of attractions) to single vectors. Such systems are also known as associative memories (Hassoun, 1995), whose stored stable states corresponds to individual memories. Through their internal dynamics they can perform pattern completion by retrieving the full stored pattern from partial or noisy information. In contrast, feed-forward networks map points to points do not allow for loops in their architecture. Information in feed-forward networks is unidirectional and the network's activity (state) is transformed from layer to layer. These differences between feed-forward and recurrent network suggest a number of predictions.

For example, we expect that that recurrent developmental networks to be more robust than feed-forward networks as they form basins of attractions and thus are naturally more noise resistant, hence providing a buffer against genetic and environmental perturbations. Directional selection also favours flexibility (Houle, 1992), while stabilising selection favours robustness (e.g., Meiklejohn and Hartl (2002)). Future studies investigating the evolutionary origins of such architectures in evolving developmental systems would therefore be beneficial. On the other hand, recurrent networks keep an internal memory which can be used to process sequences of environmental events. From the perspective of associative sequence learning, we can then study how organisms can evolve internal representations in their developmental architecture that reflect the temporal correlational structure of the environment. Moreover, we know that certain techniques that improve the generalisation performance of the learning system, such as DropOut (Hinton et al., 2012), are more effective in multi-layered networks (Bayer et al., 2013;

Pham et al., 2013). Intrinsic noise has the effect of defocusing the co-dependencies between features when acting between layers, while intrinsic noise tends to destabilise the activity within a recurrent layer. In general, we expect that inductive biases for sparsity and modularity, such as L_1-regularisation, to work better in recurrent networks for small training sets, while inductive biases for smoothness, such as jittering, DropOut and L_2-regularisation, to work better in feed-forward networks for large training sets (Lai et al., 2015).

5.2.2 Adaptive developmental accommodation: How to generalise across tasks

In this work, I assume that each phenotype consists of a collection of one or more phenotypic traits. Phenotypic traits may include morphological or physiological characteristics of the organisms, such as height, eye colour and size of a bird's beak. In Chapter 1, I consider a class of phenotypic patterns that describe modularly varying environmental goals. In Chapter 2 and 3, I consider tasks as mappings between environmental cues and optimal phenotypic patterns for each environment. The fitness of each individual is determined on the basis of the proximity of the developed adult phenotype to the target phenotypic pattern of a given selective environment. Each selective environment represents an instance of a single task, or a target function. Here, I successfully demonstrated that evolvable developmental representations can arise when future selective environments share common regularities as the past ones. Indeed, we know from learning theory that in such situations, the learner can generalise to new instances drawn from the same task domain as the examples used for training, i.e., within-task generalisation.

However, phenotypes can also represent functions. Phenotypic functions may include behavioural or physiological mechanisms, such as predator detection, mimicry, cognition and learning. Future studies should thus investigate to what extent the same considerations presented in this thesis can be applied in the case of generalisation across different tasks, i.e., not just within-task generalisation. This is also known as inductive transfer learning (Pan and Yang, 2010; Torrey and Shavlik, 2009). Across-task generalisation is important because it signifies flexibility of the knowledge acquired from previous environments. Given a limited past experience, a developmental system that effectively performs novel tasks has higher potential coping with new environmental challenges, i.e., evolvability.

Recent studies on the evolution of evolvability consider environmental goals as specific tasks (i.e., input-output mappings) that the organisms have to accomplish (Kashtan and Alon, 2005; Kashtan et al., 2007; Clune et al., 2013b; Parter et al., 2008). These tasks vary in modular fashion to guarantee that the environment changes in a systematic

manner. In such situations, it was demonstrated that evolution is capable of discovering biological designs that enhance adaptive behaviour to new environmental settings (Parter et al., 2008; Clune et al., 2013b). Yet, the extent to which that past selection can capture the general structure of the environment is unclear. From machine learning perspective, when the organisms face a new environmental challenge they have to transfer developmental systematicities acquired from previous adaptation in past environments. The idea is that prior knowledge acquired by solving one task can be beneficial in learning how to solve another task when the source and target tasks are related. In the aforementioned studies, the fitness of the organisms is evaluated over all possible instances of the respective task in each selective environment during its lifetime, which is biologically implausible. Thus, a potential starting point for future investigations could be to relax this assumption and investigate whether the same results stand.

5.2.3 Can evolution perform inductive inference?

In Chapter 3, I argue that adaptive developmental organisations can facilitate the expression of environmentally-induced phenotypes which are useful for future adaptation. Such plastic responses can prevent extinction when the population experiences strong directional selection moving into new environmental situations. Adaptive plasticity can reliably move the population closer to the new adaptive peak and thus decreases the pressure of directional selection. This first step in adaptation is crucial for the survivability and the subsequent evolution of the population. It however is equally important to understand whether and how fast the population can accumulate genetic mutations in the direction of selection and ultimately adapt to the new environment and reach the new adaptive peak (Quayle and Bullock, 2006; Kashtan et al., 2007; Ghalambor et al., 2007; Pigliucci, 2008; Parter et al., 2008; Kounios et al., 2016). Chapter 1 has addressed this, but here the focus is on the role of plasticity in speeding up future evolution.

Generally, the rate of adaptation depends on how close the realised phenotype is to the target optimum in the new environment (Price et al., 2003). Yet, Draghi and Whitlock (2012) show that plasticity does not solely affect the initial response of the population when encountering a new environment, but has more enduring effects on the adaptive potential of the population. Plasticity also modifies the genetic and mutation patterns of the population facilitating evolvability in the direction of past environmental variation. This theory enables prompts new intriguing questions regarding the evolution of adaptive potential. For example, whether and how can natural selection shape evolvability in new directions of selection that have not been directly selected in the past?

Draghi and Whitlock (2012) consider a developmental model of two phenotypic traits in an environmental setting where the two trait optima are strongly, positively correlated. This ensures that the environment changes systematically. Evolvability in their study is measured by how fast the population responds to directional selection. However, when

two traits are considered, there is only one direction that can be tested for evolvability, that is, the alignment of the line of least resistance is to the past direction of environmental change. Future studies should focus on understanding how these past selected correlations can increase the rate of adaptation to new directions of selection through genetic change.

A first step would be to extend the aforementioned model to include at least three phenotypic traits (i.e., $p =< p_A, p_B, p_C >$). This allows us to investigate scenarios about the adaptive potential of the population to rapidly respond in directions of selection that have not been experienced before. Consider the following scenario. A population is exposed to two selective environments. The first environment is characterised by a strong positive correlation between the first two traits (i.e., AB+), while the second environment is characterised by strong negative correlation (i.e., BC-). We expect that the population will increase its evolvability along with the environmental change of its previous environments. It is unknown whether the population will increase its rate of adaptation in a situation where the population experiences a strong directional selection towards a negative correlation between the first and the third trait (i.e., AC-). This setting ensures consistency with past environmental regularities and defines the problem.

Can natural selection extract such information from its past selective environments? Such capacity implies that natural selection can perform a form of inductive inference. Indeed, we know from learning theory that this is the simplest form of inference: If A then B, if B then C, therefore, if A then C. This different type of generalisation has not been tested in the current work. This suggests that the system can potentially use a form of 'logic' to cope with novel environmental situations. If this hypothesis is true, then the evolving system can appropriate respond to partial information and produce a consistent whole. For example, even under a directional selection for p_A, the developmental system could be able to generate appropriate responses for both p_B and p_C. Such scenarios could be also extended and tested for the immediate plastic response of the evolving developmental system to environmental change. Note that this type of induction does not indicate that the system performs formal mathematical logic, but rather a form of probabilistic inference.

5.2.4 How individual developmental biases can boost evolvability at the population level: Ensemble learning

Here I study the capacity of the organisms to produce useful potential phenotypic variation as a systemic property of development focussing on the individual organism. Future studies should explore how such capacity at the individual level can improve evolvability at the population level. How can locally adapted (imperfect) developmental representations of the past (e.g., specialists (DeWitt et al., 1998)) increase the adaptive potential of the population for future evolution?

We know from learning theory that many weak learners (over subsets of the training set) can collectively form a strong learner at the group level. This is known as ensemble learning or boosting (Breiman, 1996b,a; Bishop et al., 2006). Boosting techniques, such as bagging, are primary used to reduce the bias (Breiman, 1996a), but also to reduce the variance through model averaging (Breiman, 1996b). Such techniques can thus alleviate the problem of over-fitting. In general, weak learners are characterised by low bias, and hence high variance. Yet, ensemble learning theory suggests that the individuals do not necessarily have to capture the general regularities of their environment, as they can collectively form a strong group and increase their performance at a given task.

Similarly, not all of the individual organisms encounter the same environmental situations within their lifetime. The evolved developmental models are thus characterised by low bias and high variance. Individual selection in spatially and temporarily heterogeneous environments causes individual organisms to locally adapt by capturing the idiosyncrasies of their local habitats, resulting in over-fitting. Despite that, we expect that the average response of the population as a whole to new environments to be enhanced by averaging the potential of each individual through a fitness maximisation process based on the accuracy of each weak learner. We expect that environmental heterogeneity will improve the adaptive potential of the population, when compared to population where all individuals experience the same set of environments in their lifetime.

Appendix A

Appendix: Supporting Information for Chapter 1

A.1 Supporting Figures

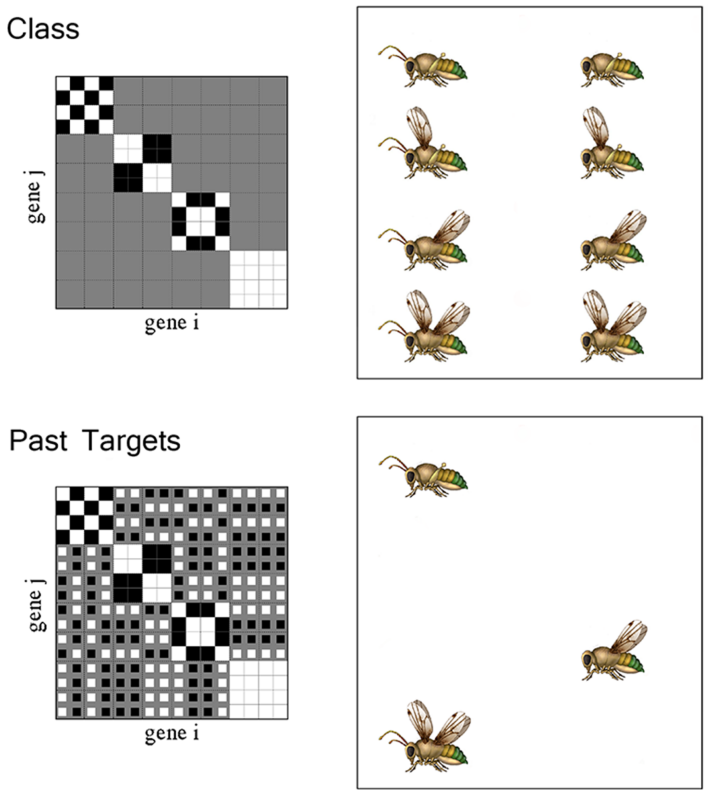

FIGURE A.1: **The underlying correlational structure of the class and the training set.** (Top) Hinton diagram of the variance-covariance matrix and phenotypic distribution of all potential future phenotypic targets. The true underlying structure of the given problem set which is comprised of all 8 possible phenotypic targets is described by the block diagonal interaction matrix. Accordingly, the traits within each module that encode for each functional part of the organism (e.g., front wings) are strongly correlated with each other (positively or negatively depending on the combination of signs in the particular phenotypic pattern used), and no correlations between one module and another (e.g., the production of halteres is functionally independent from the production of front wings). (Bottom) Hinton diagram of the variance-covariance matrix and phenotypic distribution of past phenotypic targets. The structure of the training set which is comprised of 3 phenotypic targets is described by an interaction matrix with non-zero off-diagonal elements. Those elements correspond to spurious correlations that describe functional phenotypic dependencies between modules that are present in the past selected phenotypic targets (e.g., the production of front wings is positively correlated with the production of antennas). Such developmental structures will appropriately represent the 3 past selected targets, but fail to generate all 8 phenotypes from the class. The colour and the size of the squares in Hinton's representation indicate the sign and the magnitude of the respective correlations.

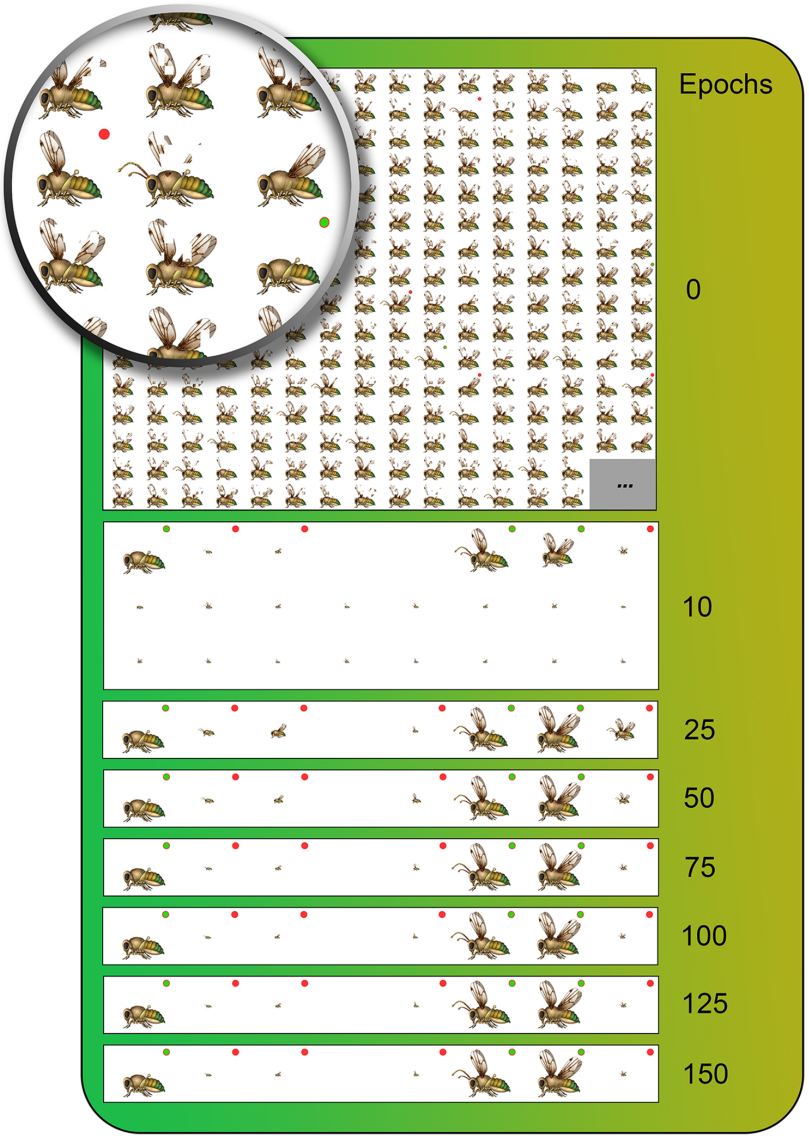

FIGURE A.2: **The evolution of phenotypic distribution for moderate environmental switching.** Pictorial representation of the phenotypic distributions induced by the evolving developmental process over evolutionary time for moderate environmental switching. Green circles indicate past selected targets, while red circles indicate previously-unseen phenotypes from the same phenotype family as the past ones. Phenotypes outside of the class are represented by distorted mosaic images. The size of the insect-like creatures indicates the propensity of development to express the respective phenotype. At the beginning (epoch 0), development equally predisposes the production of all possible phenotypic patterns (here 2^{12}), i.e., no developmental biases. The evolving developmental structure initially starts canalising only phenotypes from the class. After epoch 25 however it further canalises the production of past selected phenotypes, by reducing the propensity of producing those phenotypes from the class that were not selected in the past, i.e., over-fitting.

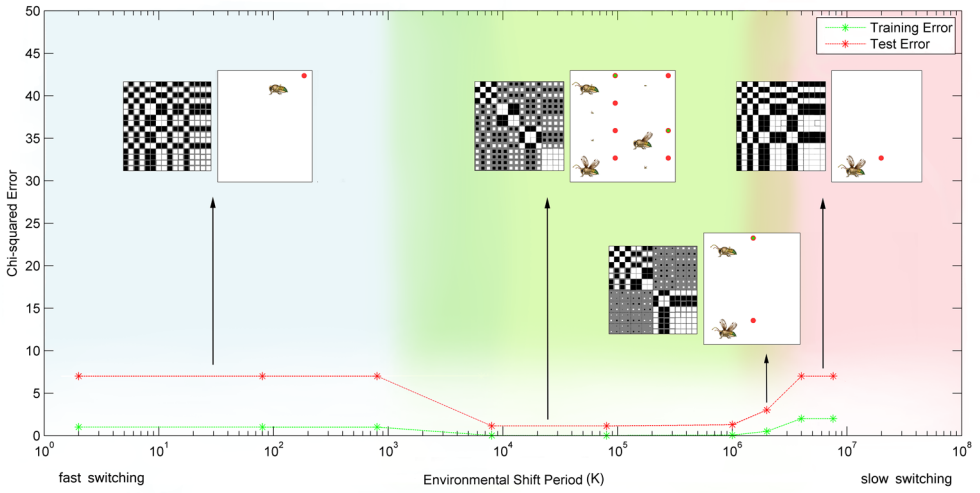

FIGURE A.3: **Fast and Slow Environmental Switching Fail to Evolve Developmental Memory.** The match between phenotypic distributions and the selective environments the network has been exposed to (training error) and all selective environments (generalisation error) against different environmental switching intervals (K). The insets illustrate the Hinton diagram of the evolved interaction matrix for each regime (indicated by different background colour) and the respective phenotypic distribution induced by the evolved developmental process.

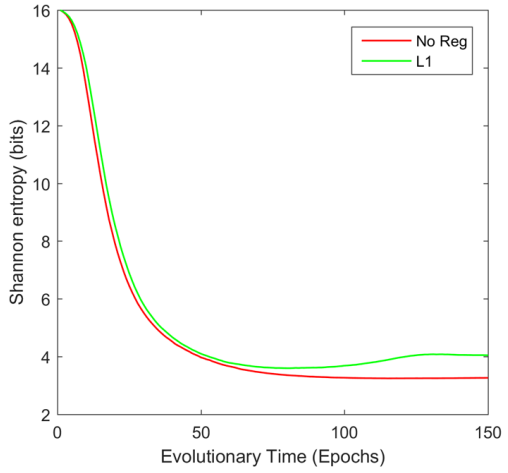

FIGURE A.4: **Entropy of the phenotypic distribution reduces over evolutionary time.** Shannon entropy Shannon (2001) of the phenotypic distribution induced by the evolving developmental process for moderate environmental switching and sparse connectivity. Overfitting is indicated by reducing to less than four bits. For the case of sparse connectivity entropy converges to four bits indicating that each of the four modules vary independently. The sample size was 5×10^5.

A.2 Developmental Model

Following previous work Watson et al. (2014), we describe the development of the embryonic phenotype to an adult phenotype by a continuous, non-linear and recurrent (i.e., it allows for feed-back connections) model of gene-regulatory networks Vohradský (2001a,b).

At each developmental time step, t, the phenotype of an individual organism is characterised by a collection of phenotypic traits, $P_t = \langle p_{t,1}, \ldots, p_{t,N} \rangle$, where $p_{t,i} \in \mathbb{R}, \forall i$. The genotype is comprised of two parts: the direct effects on the embryonic phenotypic traits, $G_t = \langle g_{t,1}, \ldots, g_{t,N} \rangle$, where $g_{t,i} \in \{-1, 1\}, \forall i$ and the regulatory interactions between the genes, b_{ij}, that determine the dynamical developmental process Wagner (1989); Lipson et al. (2002); Kashtan et al. (2009). The regulatory interactions are represented by the matrix B.

The dynamics of the expression level for each gene depend on 1) the gene expression levels of the genes that is connected to and 2) the its pattern of connections, i.e., how strongly the respective gene is connected to its neighbouring genes. In the first time step, the embryonic phenotype is solely characterised by the direct effects of G ($P_0 = G$). Thereafter, at every developmental step the phenotypic traits are developed under the following set of difference equations Wessels et al. (2001); Watson et al. (2014):

$$p_{t+1,i} = p_{t,i} + \tau_1 \sigma(\sum_j b_{ij} p_{t,j}) - \tau_2 p_{t,i}, \tag{A.1}$$

where $\tau_1 = 1$ and $\tau_2 = 0.2$ indicate the maximal expression rate and the constant rate of degradation of the given gene product respectively. The second term in the right-hand side of equation (A.1) corresponds to the interaction term, the activity of which is limited by a non-linear, monotonic and bounded (sigmoid) activation function, $\sigma(x) = tanh(\alpha x)$, where $\alpha = 0.5$. Then, over a fixed number of developmental time steps, T (here $T = 10$), the embryonic phenotype is transformed into an adult phenotype, $P_a = P_T$, upon which selection can act. Both G and B are initialised at zero.

A.3 Varying Selective Environments

In this work, a set of related phenotypic targets is considered from the same family (as in Parter et al. (2008); Watson et al. (2014)). This guarantees that the environment changes in a systematic manner (i.e., shares common regularities invariant over time) — something which is ubiquitous in natural environments.

Since we are interested in modelling phenotypic variability, traits that are under constant selection are omitted from our model. We choose a simple family of modularly-varying

targets. Modularity is widespread in the natural world and provides a simple way to test for generalised developmental oragnisations that are biological relevant Clune et al. (2013b); Callebaut and Rasskin-Gutman (2005); Carroll (2001); Alon (2006); Wagner et al. (2007); Lipson et al. (2002); Kashtan et al. (2009). For simplicity, to model selection that varies in a modular manner, we assume an extreme form of modularity, namely separable modules Watson (2006). Accordingly, selection on any trait is strongly interdependent with selection on other traits in the same module, but independent of selection on traits in other modules. Specifically, when a change in the environment occurs, if the direction of selection on a given trait changes, the direction of selection on all other traits in the same module also changes (this defines the modules). Selection thus favours two complementary states for each module that confer high fitness in different environments. Since the selection on each module is independent of selection on other modules, this means that there are 2^k possible high-fitness phenotypes, where k corresponds to the number of modules.

Here we assume a class of phenotypes consisted of equal sized modules (4 modules of 4 phenotypic traits each). The particular patterns chosen are irrelevant. So we pick one phenotype of 16 traits arbitrarily, here $(-+-+--++-+++----)$, and divide it into 4 equal modules (i.e., $(-+-+)$, $(--++)$, $(-++-)$ and $(----)$). Accordingly, for the phenotypic patterns that belong in the class, each module (block) can have 2 states: A or B; denoting a particular phenotypic sub-pattern or sub-goal (e.g., here the sub-goal for the first module can be either $(-+-+)$ (A) or $(+-+-)$ (B)). The class is thus comprised of 16 different modular patterns; all possible combinations of the sub-patterns (blocks) (see A.5 in S3 Appendix).

The time-invariant regularities here are the correlations between traits within any one module. The actual underlying structure of the given problem can thus be described by the block diagonal interaction matrix (see B.1 in S1 Appendix). The colour and the size of the squares in Hinton's representation indicate the sign and the magnitude of each correlation respectively. This clearly shows that selection on the traits within each module are strongly correlated with each other (positively or negatively depending on the combination of signs in the particular phenotypic pattern used), and no correlations between one module and another.

Complementary patterns here are also stable states of the evolved dynamical system as a result of Equation A.1. The map described in Equation A.1 is an odd function (i.e., symmetric with respect to the origin) since $f(-x) = -f(x)$. Accordingly, if R is a stable state of the system, i.e., $R = f(R)$, then $-R$ is also stable since $-R = -f(R) = f(-R)$. In order to focus on the more interesting (non-trivial) attractors that may arise, we limit the phenotypic space so as to ignore complementary targets (i.e., thus removing 8 of the patterns). Specifically, without loss of generality, we consider the phenotypic targets in which the sub-pattern in the last slot (trait positions:$13-16$) corresponds to state A: $\{-,-,-,-\}$, i.e., we focus on the top-half of the class as arranged in the lower

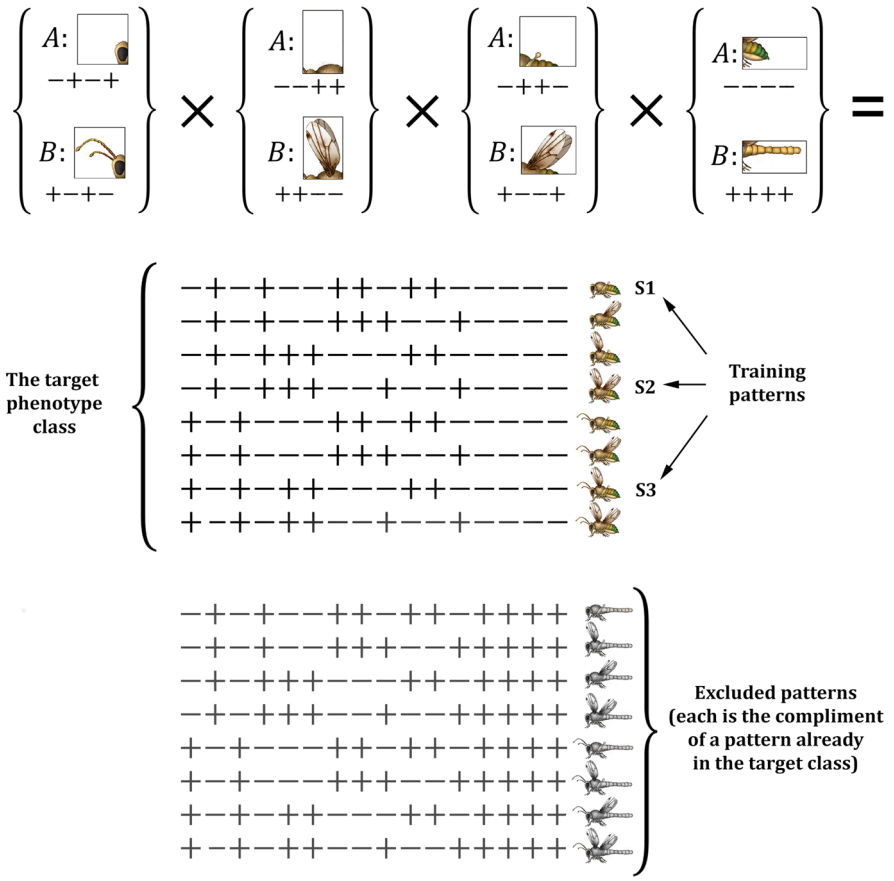

FIGURE A.5: **Modularly-varying environment.** Target phenotypes varying from one another in a modular fashion. Each target phenotype consists of 4 modules of 4 phenotypic traits (i.e., 16 phenotypic traits in total). Each module can take two (complementary) states: A or B; denoting particular sub-patterns favoured by selection in different selective environments. The complete set of phenotypes is thus comprised of $2^4 = 16$ phenotypes, differing from one another in a modular fashion. The signs of phenotypic traits correspond to the direction favoured by selection in a given environment. Eight of the 16 possible phenotypes are designated as the target class (the other eight are merely the complement of a pattern already in the target class). For the main experiments, three patterns from the target class are used as 'training' patterns, i.e., selected for.

part of Fig 1. Accordingly, each member of the other half of the class is the bit-wise complement of a member in the top half.

In this work, we want to examine the ability of the developmental system to 'learn' from past selective environments and generalise to new environments by producing novel phenotypes within the same class. Accordingly, to assay generalisation and the conditions that promote it, the population is evolved by exposure to a limited number of selective environments (< 8, i.e., a strict sub-set of the class). Otherwise, generalisation would not be relevant, since the population would have been exposed to all possible selective environments (i.e., all phenotypes in the class are presented). For this paper, we use the

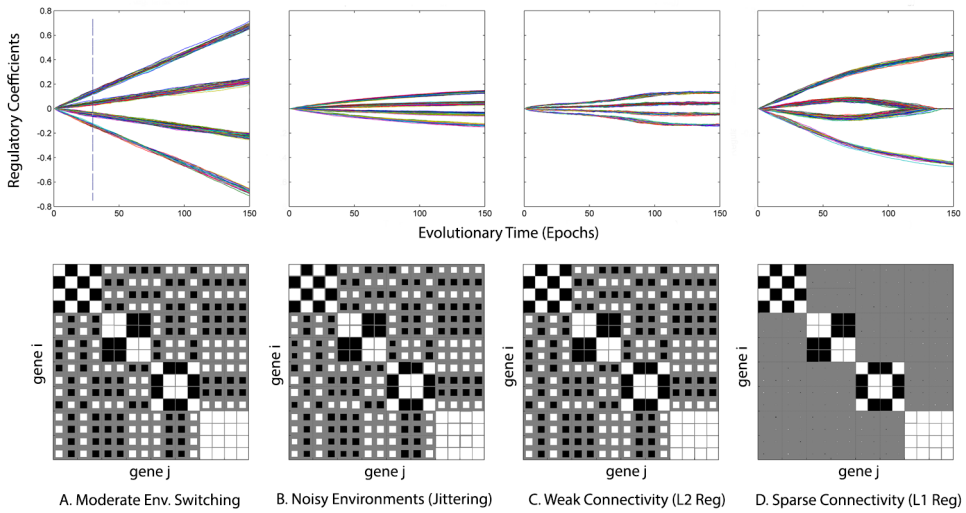

FIGURE A.6: **Evolution of regulatory coefficients in noisy environments and under parsimony pressure.** The evolution of regulatory coefficients over evolutionary time and the Hinton diagram of the evolved regulatory coefficients (after epoch 150) for (A) moderate environmental switching, (B) noisy environments, (C) favouring weak connectivity and (D) favouring sparse connectivity. The vertical dashed line denotes when the ad-hoc technique of early stopping is used, i.e., the moment the problem of over-fitting begins. Favouring sparsity ignores the weak spurious correlations of the finite sampling noise and maintains the time-invariant ones.

following example from this problem domain as a training set:

$$\begin{aligned} S_1 &= \{-,+,-,+\},\{-,-,+,+\},\{-,+,+,-\},\{-,-,-,-\}. \\ S_2 &= \{-,+,-,+\},\{+,+,-,-\},\{+,-,-,+\},\{-,-,-,-\}. \\ S_3 &= \{+,-,+,-\},\{+,+,-,-\},\{-,+,+,-\},\{-,-,-,-\}. \end{aligned} \tag{A.2}$$

In Favouring Sparse Connectivity in Different Training Sets, we explore sensitivity to this particular choice by examining generalisation from training on all possible proper subsets of the class.

A.4 The Structure of Developmental Organisation

Here we show how costly interactions and noisy environments facilitate the emergence of more general and parsimonious developmental models. For this purpose, we monitor the evolution of regulatory interactions over evolutionary time in each evolutionary setting. The regulatory coefficients here correspond to the free parameters of the developmental model that determine the functional organisation of development.

We first analyse the evolution of regulatory coefficients in the control scenario, i.e., moderate rate of environmental change. A.6 A in S4 Appendix shows that the ontogenetic

interactions evolved under natural selection to reflect the correlations in the previously-experienced selective environments. As seen, the Hinton diagram of the evolved regulatory matrix appropriately matched the variance-covariance matrix of the past phenotypic targets (A.6 in S4 Appendix). The colour and the size of the squares in Hinton's representation indicate the sign and the magnitude of the respective correlations.

Yet natural selection did not directly select either *for* correlations, or *for* matching the exploration distribution to the fitness distribution of the phenotypic variants (i.e., training error minimisation). Natural selection selected *for* immediate fitness differences depending on how well adapted the organism was to its current selective environment; i.e., how close the produced adult phenotype was to the respective target phenotype. The evaluation of the developmental process performed here against the training and the test set was a post hoc analysis, and hence not part of the actual evolutionary dynamics.

In the same fashion as the nervous system Anderson (1983), evolution does not try to analyse anything. It just tries to generate appropriate behaviour. The observed (correlation) learning behaviour of evolution can be seen as a by-product of developmental systems' effort to produce high-fitness phenotypic variants in varied selective environments — optimise the actual functionality of the system. The system does not explicitly aim at inferring the target function, namely, the ideal G-P map that gives rise to proper system functionality in long-term (over certain genetic and environmental conditions). Nevertheless, we see that under certain conditions the system may discover a hypothesis (i.e., set of regulatory coefficients) closer to the target function, by producing phenotypic variants that are fitter in short term.

A.6 B in S4 Appendix shows that under the presence of environmental noise, the regulatory interactions evolved towards smaller in magnitude weights. In particular, we observe that the rate of evolutionary change was decreased with evolutionary time giving rise to a plateau in the test error in Fig 3 B. The set of evolved regulatory coefficients here corresponds to the one we get if we stopped evolution the moment over-fitting begins, i.e., at the vertical dashed line in Fig 3 A. From Hinton diagram we can see that the relative importance between strong and weak correlations remained the same as in the case of the control run, i.e., only the magnitudes changed. Therefore, noise had a beneficial role on the evolution of genetic structures by making it difficult for natural selection to find configurations that over-fit past phenotypic targets.

We observe similar results for the evolution of regulatory interactions under the pressure for weak connectivity (A.6 C in S4 Appendix). In contrast to environmental stochasticity, however, favouring weak connectivity imposes strict constraints on the evolution of regulatory coefficients that prohibit them from growing bigger, i.e., providing a hard bound determined by the strength of parsimony pressure (see below). Accordingly, the regulatory coefficients initially increased until they reached a level that the further increase in the reproduction and maintenance cost of interactions was greater than the

benefit of the developmental structure. Moreover, when properly tuned favouring weak connectivity exhibits the same behaviour as stopping early. Favouring weak connectivity (L_2-regularisation) can be understood as imposing inductive biases (i.e., additional constraints) in the evolution of regulatory interactions, punishing interactions (parameters) with extreme (high) magnitudes by applying a penalty proportional to their current magnitudes (as in weight-decay).

Lastly, A.6 D in S4 Appendix illustrates how favouring sparse connectivity can exhibit a form of feature selection emphasising the relative importance of the strong correlations against the weak correlations. Specifically, we see that only the strongest (time-invariant) correlations persisted, while the weak (spurious) correlations, which arose as a result of the sampling process, were eliminated over evolutionary time. The strong correlations here (i.e., the block diagonal of the interaction matrix) correspond to the actual underlying modular structure of the environmental variation that remain invariant over time. Consequently, if the strength of parsimony pressure is large enough to ignore the spurious correlations, the evolved associations are (almost) identical to the variance-covariance matrix that describes the phenotypes family (see A.6 in S4 Appendix). Favouring sparse connectivity (L_1-regularisation) can be understood as punishing interactions by equally applying a fixed penalty to all of the weights of the network. The amount of reduction is controlled by the hyper-parameter λ (see below); the higher its value, the higher the penalty applied, and hence the higher the level of sparsity. When properly tuned, favouring sparse connectivity leads to many zero weights, and thus the complexity of the model is reduced by removing degrees of freedom.

A.5 Favouring Sparse Connectivity in Different Training Sets

Experiments were also carried out for every possible training set as a strict sub-set of the test set. Firstly, all possible combinations, $\sum_{0 \leq k \leq N} \binom{N}{k} = 2^N$, were explicitly enumerated, where N indicates the number of patterns in the test set. Then, the respective developmental systems were determined following Hebb's rule with and without the selective pressure on the cost of connections (for optimal λ values). Hebbian learning was used here for computational tractability (65536 possible combinations), since it has been shown before that the interaction matrix evolves under natural selection in a Hebbian manner Watson et al. (2014). According to Hebb's rule, the pair-wise interactions are increased (or decreased) if the phenotypic traits are aligned (or not). The Hebbian matrix can be computed by computing the outer-product over the training inputs, i.e., the auto-correlation matrix. For the sake of comparison, the respective coefficient matrices were also tuned to be of the same average magnitude level as in the experiments above. These simulations allow us to draw some more general conclusions.

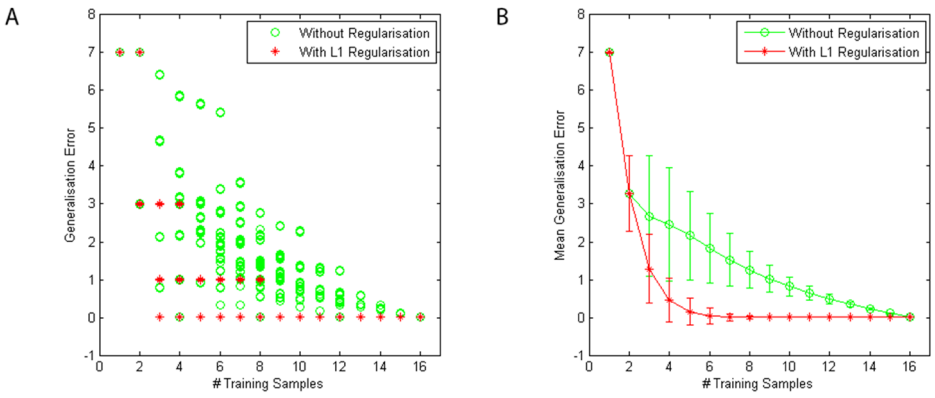

FIGURE A.7: Favouring sparse connectivity enhances phenotypic generalisation. (A) Phenotypic generalisation with and without the parsimony pressure for sparsity (L_1-regularisation) against all possible evolutionary scenarios (training sets), i.e., all possible combinations of distinct past selective environments drawn from the class. (B) Means and error bars of the generalisation performance of the evolved networks with and without the parsimony pressure for sparsity against different numbers of previously experienced selective environments. The cost of connection significantly enhanced evolvability in the majority of the cases. The interaction matrices here were determined using Hebb's rule.

Overall, we find that the cost of connection significantly enhanced evolvability in the majority of the cases (A.7 in S5 Appendix). As the number of observations is increased we observe an increase on average in evolvability, reaching zero generalisation error when $k = N$, even without incorporating the cost of connection. Interestingly, this was also true for some cases of 4, 8 and 12 patterns. We therefore see that different training sets entailed different information about the class, some of which were better representatives than others. For training sets consisted of more than half of the patterns in the class, we also observe that (optimally tuned) parsimony pressure for sparsity certainly resulted in perfect generalisation. On the other hand, in situations like the ones of 1 or 2 patterns the parsimony pressure had no effect on the generalisation performance of the network, and in some situations between 3 to 8 patterns it had little effect.

Appendix B

Appendix: Supporting Information for Chapter 2

B.1 Supporting Figures

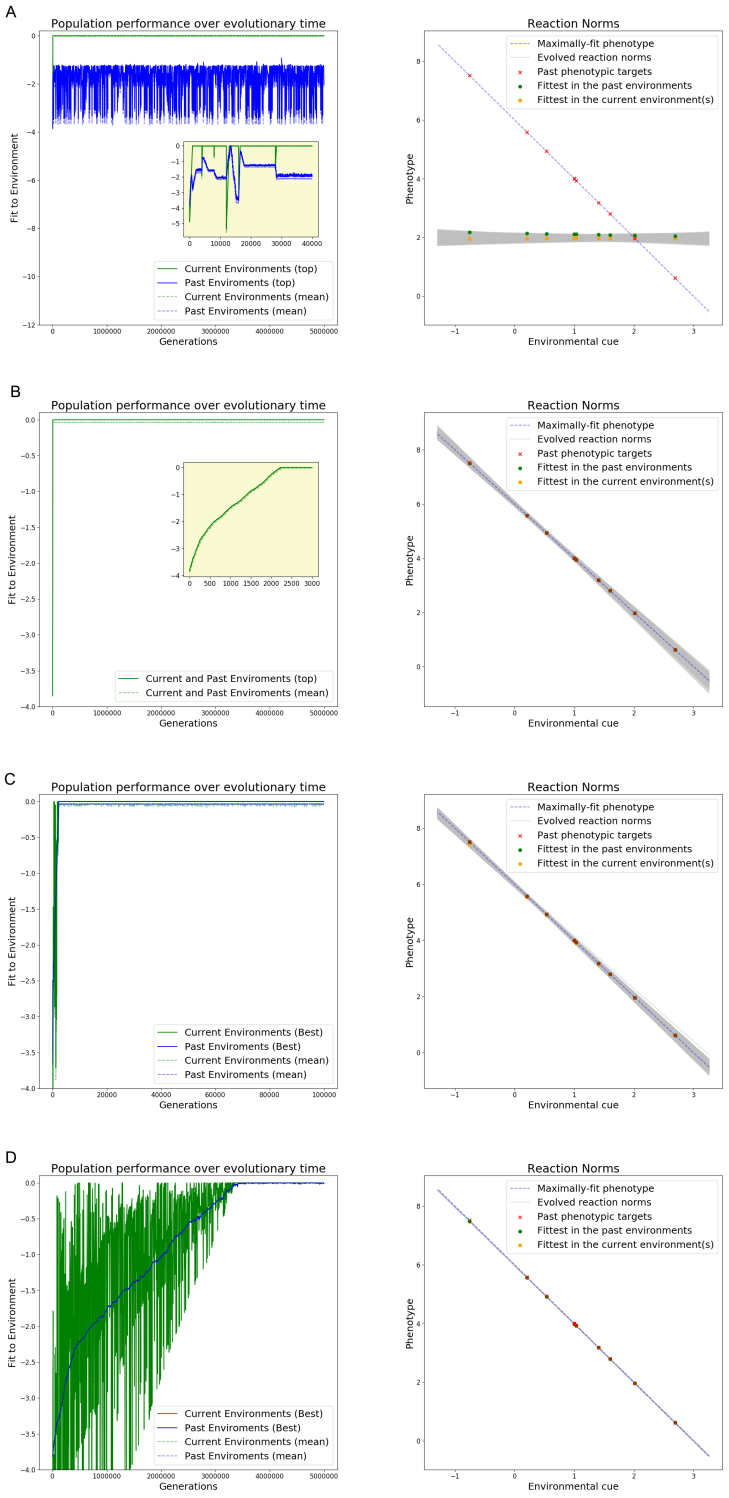

FIGURE B.1: **The evolution of plasticity in the long run.** The population is evolved for 5000000 generations. Left. Goodness of fit of the evolving reaction norms to the current environment (green) and all past selective environments (blue) over evolutionary time. Right. Evolved reaction norm (gray line). The red marks correspond to trait optima of the past experienced selective environments, while the blue dashed line corresponds to the optimal reaction norm. A. Coarse-grained environments ($K = 40000$). B. Fine-grained environments. C. Fast environmental switching ($K = 1$). D. Low mutation rate ($\sigma_\mu = 10^{-5}$).

Appendix C

Appendix: Supporting Information for Chapter 3

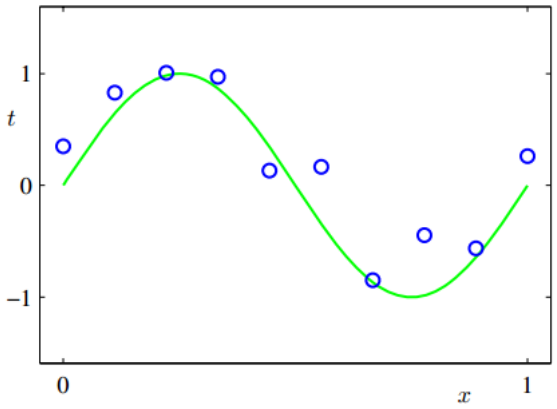

FIGURE C.1: **Training set.** Training set of $N = 10$ data points illustrated by blue circles. The solid green curve represents the underlying structure of the data, function $\sin(2\pi x)$. The aim is to make predictions values of d when new values x are given, without any explicit knowledge of the process that generated the data. Adopted from (Bishop et al., 2006).

C.1 An illustrative example

To illustrate and motivate key concepts from learning theory that are relevant to the evolution of adaptive plasticity, we consider a simple worked example below (Bishop et al., 2006). Consider a developmental system that receives instructive real-values signals, such as temperature and pH, and a selection pressure towards real-values target traits, such as body size or height, in each selective environment. Suppose now that the underlying relationship between environmental cues and target phenotypes is described by a function. For the purpose of this example, we consider the function $d = \sin(2\pi x)$ with the addition of random noise, where x corresponds to the environmental cue and d corresponds to the desired trait value for the respective environment. Consider that the population experiences N environments. Then, we can define a set of all pairs (x_i, d_i) $\forall i \in \{1, \ldots, N\}$, where each i denotes a different selective environment. In learning theory, this set is known as the *training set* (Fig. C.1).

This level of noise can arise from intrinsically developmental processes or extrinsic noise present in the environment as a source of unobserved variability. This captures situations where there is an underlying regularity but the observations are corrupted by noise. Interest here lies in how evolving systems can exploit information in the past selective environments and generate good responses, d, to new environments, when new cues x are given. This problem implicitly is equivalent to inferring the underlying structure of the problem, i.e., discovering the function $\sin(2\pi x)$. This is a generally difficult task since the system has to generalise from a few observations, especially when the observations are also noisy.

Consider that the reaction norm can be mathematically described by a polynomial function. Note that we explicitly choose a hypothesis space which is different from the

function that describes how the environment changes. The complexity of the model is not appropriately tuned based on the number of points we want to fit, or their underlying regularities. Selection can act on the polynomial coefficients such as to maximise the fitness of the individual, producing plastic responses, y that match as much as possible the optimal values d. Fitness function can simply defined here as: $F(\theta) = F_{max} - \sum_1^N (y(x_i, \theta) - d_i)^2$, where θ represents the coefficients of the polynomial. Assume that evolved developmental organisation is described by $y(x_i, \theta^*)$, where θ^* denotes the evolved polynomial coefficients. We investigate now the role of the degree polynomial in fitting the data in the training set. Fig. 4.2 shows different polynomials of various degrees fitting the given data. We see that polynomials of order 0 (i.e., constant value) or 1 performs poorly at fitting the observed data, and as such they are not provide good representations of the underlying process, $\sin(2\pi x)$. On the other hand, when the degree of the polynomial was high ($M = 9$), the resulting polynomial fitted the data very well, but it was not a good representation of the underlying process as well. However, polynomials of degree 3 seem to represent best the underlying regularity in the data. Consequently, we see how the complexity of the model can affect the generalisation ability of the system.

To gain better insight on how model complexity can affect the generalisation ability of the learned model by means of bias-variance tradeoff, consider a uni-variate polynomial function of degree 2 with the addition of random noise. The model complexity here is indicated by the degree of the polynomial function, i.e., the number of polynomial coefficients (degrees of freedom). Note that complexity can be defined in different ways. For instance, here model complexity could refer to the variability or the total magnitude of the polynomial coefficients. Fig. 4.3 (A) illustrates 3 models of different complexity, i.e., polynomial functions of degree 1,2 and 3 respectively. We see that the fitting ability of the model is affected by its complexity. Polynomial functions of degree 1 do not sufficiently fit the data. On the other hand, polynomial functions of degree 3 fit the data very well, but also tend to fit the noise in the data. Polynomial functions of degree 1 and 3 under-fit and over-fit the data respectively. Fig. 4.3 (B) shows results of 12 models using the leave-one-out method, i.e., training the model on 11 points, leaving one out at a time. We see that the polynomials of degree 1 and 3 show high variability, while polynomials of degree 2 tend to be more stable. This stability indicates that when the models are appropriately complex then they are less sensitive to the choice of the training points. They thus tend to avoid fitting the idiosyncrasies of the training set (i.e., noise) and provide more general models of the data.

Lastly, we show how inductive biases can promote more general models. To do so, we consider a polynomial of degree 9. Fig. 4.2 shows that such flexibility results in capturing the idiosyncrasies of the training examples and not the underlying structure of the problem, i.e., over-fitting. Although the model perfectly matches the training samples, it gives rise to large residuals errors outside the range of previous seen observations. This

is because highly complex models tend to capture the irrelevant information (noise) in the data during the training phase. Thus, we expect that a parsimony pressure that favours simpler models to enhance generalisation by preventing such situations of overfitting. Indeed, fig. 4.5 shows that for a moderate level of parsimony pressure can generate models that better represent the underlying relationship over the data. If however the parsimony pressure is excessively high, it prevents the model from learning anything useful, due to high bias levels.

Bibliography

Yaser S Abu-Mostafa, Malik Magdon-Ismail, and Hsuan-Tien Lin. *Learning from data*. AMLBook, 2012.

Christoph Adami, Charles Ofria, and Travis C Collier. Evolution of biological complexity. *Proceedings of the National Academy of Sciences*, 97(9):4463–4468, 2000.

Anurag A Agrawal. Phenotypic plasticity in the interactions and evolution of species. *Science*, 294(5541):321–326, 2001.

Hirotugu Akaike. A new look at the statistical model identification. *IEEE transactions on automatic control*, 19(6):716–723, 1974.

Maximino Aldana, Enrique Balleza, Stuart Kauffman, and Osbaldo Resendiz. Robustness and evolvability in genetic regulatory networks. *Journal of theoretical biology*, 245(3):433–448, 2007.

Uri Alon. Biological networks: the tinkerer as an engineer. *Science*, 301(5641):1866–1867, 2003.

Uri Alon. *An introduction to systems biology: design principles of biological circuits*. CRC press, 2006.

Peter Alpert and Ellen L Simms. The relative advantages of plasticity and fixity in different environments: when is it good for a plant to adjust? *Evolutionary Ecology*, 16(3):285–297, 2002.

Lee Altenberg. The schema theorem and price's theorem. In *FOGA*, volume 3, pages 23–49, 1994.

Lee Altenberg. Genome growth and the evolution of the genotype-phenotype map. In *Evolution and biocomputation*, pages 205–259. Springer, 1995.

Lauren W Ancel, Walter Fontana, et al. Plasticity, evolvability, and modularity in rna. *Journal of Experimental Zoology*, 288(3):242–283, 2000.

James A Anderson. Cognitive and psychological computation with neural models. *Systems, Man and Cybernetics, IEEE Transactions on*, (5):799–815, 1983.

Wallace Arthur. Evolutionary developmental biology: developmental bias and constraint. *eLS*, 2006.

Josh R Auld, Anurag A Agrawal, and Rick A Relyea. Re-evaluating the costs and limits of adaptive phenotypic plasticity. *Proceedings of the Royal Society of London B: Biological Sciences*, 277(1681):503–511, 2010.

Alexander V Badyaev. Stress-induced variation in evolution: from behavioural plasticity to genetic assimilation. *Proceedings of the Royal Society of London B: Biological Sciences*, 272(1566):877–886, 2005.

J Mark Baldwin. A new factor in evolution (continued). *American naturalist*, pages 536–553, 1896.

Justin Bayer, Christian Osendorfer, Nutan Chen, Sebastian Urban, and Patrick van der Smagt. On fast dropout and its applicability to recurrent networks. *arXiv preprint arXiv:1311.0701*, 2013.

Brooke S Baythavong. Linking the spatial scale of environmental variation and the evolution of phenotypic plasticity: selection favors adaptive plasticity in fine-grained environments. *The American Naturalist*, 178(1):75–87, 2011.

Mark A Bedau. *The evolution of complexity*. Springer, 2009.

Mark A Bedau, John S McCaskill, Norman H Packard, Steen Rasmussen, Chris Adami, David G Green, Takashi Ikegami, Kunihiko Kaneko, and Thomas S Ray. Open problems in artificial life. *Artificial life*, 6(4):363–376, 2000.

Chris M Bishop. Training with noise is equivalent to tikhonov regularization. *Training*, 7(1), 2008.

Christopher M Bishop et al. *Pattern recognition and machine learning*, volume 1. springer New York, 2006.

Benjamin M Bolker. *Ecological models and data in R*. Princeton University Press, 2008.

Anthony D Bradshaw. Evolutionary significance of phenotypic plasticity in plants. *Advances in genetics*, 13:115–155, 1965.

Christian Braendle, Charles F Baer, and Marie-Anne Félix. Bias and evolution of the mutationally accessible phenotypic space in a developmental system. *PLoS genetics*, 6(3):e1000877, 2010.

Paul M Brakefield. Evo-devo and constraints on selection. *Trends in Ecology & Evolution*, 21(7):362–368, 2006.

Leo Breiman. Bagging predictors. *Machine learning*, 24(2):123–140, 1996a.

Leo Breiman. Bias, variance, and arcing classifiers. 1996b.

Ingo Brigandt. Typology now: homology and developmental constraints explain evolvability. *Biology & Philosophy*, 22(5):709–725, 2007.

Ingo Brigandt and Alan C Love. Conceptualizing evolutionary novelty: moving beyond definitional debates. *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution*, 318(6):417–427, 2012.

Sean G Byars, Warwick Papst, and Ary A Hoffmann. Local adaptation and cogradient selection in the alpine plant, poa hiemata, along a narrow altitudinal gradient. *Evolution*, 61(12):2925–2941, 2007.

Werner Callebaut and Diego Rasskin-Gutman. *Modularity: understanding the development and evolution of natural complex systems*. MIT press, 2005.

Robert L Carroll. Towards a new evolutionary synthesis. *Trends in ecology & evolution*, 15(1):27–32, 2000.

Sean B Carroll. Chance and necessity: the evolution of morphological complexity and diversity. *Nature*, 409(6823):1102–1109, 2001.

Sean B Carroll. Evo-devo and an expanding evolutionary synthesis: a genetic theory of morphological evolution. *Cell*, 134(1):25–36, 2008.

Christopher Cherniak, Zekeria Mokhtarzada, Raul Rodriguez-Esteban, and Kelly Changizi. Global optimization of cerebral cortex layout. *Proceedings of the National Academy of Sciences of the United States of America*, 101(4):1081–1086, 2004.

Luis-Miguel Chevin, Russell Lande, and Georgina M Mace. Adaptation, plasticity, and extinction in a changing environment: towards a predictive theory. *PLoS biology*, 8 (4):e1000357, 2010.

Curtis P Clausen. *Entomophagous insects*. McGraw Hill Book: London, 1940.

Jeff Clune, Dusan Misevic, Charles Ofria, Richard E Lenski, Santiago F Elena, Rafael Sanjuán, et al. Natural selection fails to optimize mutation rates for long-term adaptation on rugged fitness landscapes. In *GECCO (Companion)*, pages 25–26, 2013a.

Jeff Clune, Jean-Baptiste Mouret, and Hod Lipson. The evolutionary origins of modularity. *Proceedings of the Royal Society b: Biological sciences*, 280(1755):20122863, 2013b.

Michael Conrad. The importance of molecular hierarchy in information processing. *Towards a theoretical biology*, 4:222–228, 1972.

Michael Conrad. Bootstrapping on the adaptive landscape. *BioSystems*, 11(2):167–182, 1979.

Michael Conrad. The geometry of evolution. *BioSystems*, 24(1):61–81, 1990.

Michael Conrad. Towards high evolvability dynamics introduction. In *Evolutionary systems*, pages 33–43. Springer, 1998.

Eric H Davidson and Michael S Levine. Properties of developmental gene regulatory networks. *Proceedings of the National Academy of Sciences*, 105(51):20063–20066, 2008.

Richard Dawkins. The evolution of evolvability. *On growth, form and computers*, pages 239–255, 2003.

Gavin De Beer. Darwin's views on the relations between embryology and evolution. *Botanical Journal of the Linnean Society*, 56(365):15–23, 1958.

G De Jong. Phenotypic plasticity as a product of selection in a variable environment. *American Naturalist*, pages 493–512, 1995.

G De Jong and P Bijma. Selection and phenotypic plasticity in evolutionary biology and animal breeding. *Livestock production science*, 78(3):195–214, 2002.

Gerdien De Jong. Evolution of phenotypic plasticity: patterns of plasticity and the emergence of ecotypes. *New Phytologist*, 166(1):101–118, 2005.

Erez Dekel and Uri Alon. Optimality and evolutionary tuning of the expression level of a protein. *Nature*, 436(7050):588–592, 2005.

Bai-Chuan Deng, Yong-Huan Yun, Yi-Zeng Liang, Dong-Sheng Cao, Qing-Song Xu, Lun-Zhao Yi, and Xin Huang. A new strategy to prevent over-fitting in partial least squares models based on model population analysis. *Analytica chimica acta*, 880: 32–41, 2015.

Houtao Deng and George Runger. Feature selection via regularized trees. In *The 2012 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2012.

Thomas J DeWitt. Expanding the phenotypic plasticity paradigm to broader views of trait space and ecological function. *Current Zoology*, 62(5):463–473, 2016.

Thomas J DeWitt and Samuel M Scheiner. *Phenotypic plasticity: functional and conceptual approaches*. Oxford University Press, 2004.

Thomas J DeWitt, Andrew Sih, and David Sloan Wilson. Costs and limits of phenotypic plasticity. *Trends in ecology & evolution*, 13(2):77–81, 1998.

Marie-Laure Dichtel-Danjoy and Marie-Anne Félix. Phenotypic neighborhood and micro-evolvability. *Trends in Genetics*, 20(5):268–276, 2004.

Stanley I Dodson. The ecological role of chemical stimuli for the zooplankton: predator-induced morphology in daphnia. *Oecologia*, 78(3):361–367, 1989.

PAUL Doughty and DAVID N Reznick. Patterns and analysis of adaptive phenotypic plasticity in animals. *Phenotypic plasticity: Functional and conceptual approaches*, pages 126–150, 2004.

Jeremy A Draghi, Todd L Parsons, Günter P Wagner, and Joshua B Plotkin. Mutational robustness can facilitate adaptation. *Nature*, 463(7279):353–355, 2010.

Jeremy A Draghi and Michael C Whitlock. Phenotypic plasticity facilitates mutational variance, genetic variance, and evolvability along the major axis of environmental variation. *Evolution*, 66(9):2891–2902, 2012.

Avigdor Eldar and Michael B Elowitz. Functional roles for noise in genetic circuits. *Nature*, 467(7312):167, 2010.

John A Endler and Tracy McLellan. The processes of evolution: toward a newer synthesis. *Annual Review of Ecology and Systematics*, 19(1):395–421, 1988.

B Ernande and U Dieckmann. The evolution of phenotypic plasticity in spatially structured environments: implications of intraspecific competition, plasticity costs and environmental characteristics. *Journal of evolutionary biology*, 17(3):613–628, 2004.

Douglas H Erwin and Eric H Davidson. The evolution of hierarchical gene regulatory networks. *Nature Reviews Genetics*, 10(2):141, 2009.

Carlos Espinosa-Soto, Olivier C Martin, and Andreas Wagner. Phenotypic robustness can increase phenotypic variability after nongenetic perturbations in gene regulatory circuits. *Journal of evolutionary biology*, 24(6):1284–1297, 2011.

Daphne J Fairbairn. Allometry for sexual size dimorphism: pattern and process in the coevolution of body size in males and females. *Annual review of ecology and systematics*, 28(1):659–687, 1997.

IR Falconer, RB Jackson, BJ Langley, and MT Runnegar. Liver pathology in mice in poisoning by the blue-green alga microcystis aeruginosa. *Australian Journal of Biological Sciences*, 34(2):179–188, 1981.

Janna L Fierst and Patrick C Phillips. Modeling the evolution of complex genetic systems: The gene network family tree. *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution*, 324(1):1–12, 2015.

JL Fierst. A history of phenotypic plasticity accelerates adaptation to a new environment. *Journal of evolutionary biology*, 24(9):1992–2001, 2011.

Thomas Flatt. The evolutionary genetics of canalization. *The quarterly review of biology*, 80(3):287–316, 2005.

Brian G Forde. Is it good noise? the role of developmental instability in the shaping of a root system. *Journal of experimental botany*, 60(14):3989–4002, 2009.

George Forman. Quantifying counts and costs via classification. *Data Mining and Knowledge Discovery*, 17(2):164–206, 2008.

Steven A Frank. Natural selection maximizes fisher information. *Journal of Evolutionary Biology*, 22(2):231–244, 2009.

Tamar Friedlander, Avraham E Mayo, Tsvi Tlusty, and Uri Alon. Mutation rules and the evolution of sparseness and modularity in biological systems. *PloS one*, 8(8): e70444, 2013.

ER Froesch, Chr Schmid, J t Schwander, and J Zapf. Actions of insulin-like growth factors. *Annual Review of Physiology*, 47(1):443–467, 1985.

DJ Futuyma. Evolution, 3rd edn, 2013.

Silvio Galanti and Alan Jung. Low-discrepancy sequences: Monte carlo simulation of option prices. *The Journal of Derivatives*, 5(1):63–83, 1997.

Alexis Gallagher. *Evolvability: a formal approach.* PhD thesis, University of Oxford, 2009.

Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural computation*, 4(1):1–58, 1992.

John Gerhart and Marc Kirschner. The theory of facilitated variation. *Proceedings of the National Academy of Sciences*, 104(suppl 1):8582–8589, 2007.

Cameron K Ghalambor, Kim L Hoke, Emily W Ruell, Eva K Fischer, David N Reznick, and Kimberly A Hughes. Non-adaptive plasticity potentiates rapid adaptive evolution of gene expression in nature. *Nature*, 525(7569):372–375, 2015.

Cameron K Ghalambor, John K McKay, Scott P Carroll, and David N Reznick. Adaptive versus non-adaptive phenotypic plasticity and the potential for contemporary adaptation in new environments. *Functional ecology*, 21(3):394–407, 2007.

Scott F Gilbert, Thomas CG Bosch, and Cristina Ledón-Rettig. Eco-evo-devo: developmental symbiosis and developmental plasticity as evolutionary agents. *Nature Reviews. Genetics*, 16(10):611, 2015.

Karen Goodell, D Elam, J Nason, and N Ellstrand. Gene flow among small populations of a self-incompatible plant: an interaction between demography and genetics. *American Journal of Botany*, 84(10):1362–1362, 1997.

Karl Gotthard and Sören Nylin. Adaptive plasticity and plasticity as an adaptation: a selective review of plasticity in animal morphology and life history. *Oikos*, pages 3–17, 1995.

Stephen Jay Gould. *The structure of evolutionary theory.* Harvard University Press, 2002.

JN Grant, C Halls, W Avila, and G Avila. Igneous geology and the evolution of hydrothermal systems in some sub-volcanic tin deposits of bolivia. *Geological Society, London, Special Publications*, 7(1):117–126, 1977.

Erick Greene. A diet-induced developmental polymorphism in a caterpillar. *Science*, 243(4891):643–646, 1989.

Cortland K Griswold. Pleiotropic mutation, modularity and evolvability. *Evolution & development*, 8(1):81–93, 2006.

Xun Gu, Zhongqi Zhang, and Wei Huang. Rapid evolution of expression and regulatory divergences after yeast gene duplication. *Proceedings of the National Academy of Sciences of the United States of America*, 102(3):707–712, 2005.

Brian K Hall. Unlocking the black box between genotype and phenotype: cell condensations as morphogenetic (modular) units. *Biology and Philosophy*, 18(2):219–247, 2003.

Brian K Hall and Ryan Kerney. Levels of biological organization and the origin of novelty. *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution*, 318 (6):428–437, 2012.

Benedikt Hallgrímsson and Brian Keith Hall. Variation and variability: central concepts in biology. *Variation: A central concept in biology B*, pages 1–7, 2005.

Elena Hamann, Halil Kesselring, Georg FJ Armbruster, Johannes F Scheepens, and Jürg Stöcklin. Evidence of local adaptation to fine-and coarse-grained environmental variability in poa alpina in the swiss alps. *Journal of Ecology*, 104(6):1627–1637, 2016.

Peter JB Hancock. An empirical comparison of selection methods in evolutionary algorithms. In *AISB Workshop on Evolutionary Computing*, pages 80–94. Springer, 1994.

Thomas F Hansen. Is modularity necessary for evolvability?: Remarks on the relationship between pleiotropy and evolvability. *Biosystems*, 69(2):83–94, 2003.

Mohamad H Hassoun. *Fundamentals of artificial neural networks*. MIT press, 1995.

TO Haugen and LA Vøllestad. Population differences in early life-history traits in grayling. *Journal of evolutionary biology*, 13(6):897–905, 2000.

Mikko Heino, Ulf Dieckmann, and Olav Rune Godø. Measuring probabilistic reaction norms for age and size at maturation. *Evolution*, 56(4):669–678, 2002.

Jesse Love Hendrikse, Trish Elizabeth Parsons, and Benedikt Hallgrímsson. Evolvability as the proper focus of evolutionary developmental biology. *Evolution & development*, 9(4):393–401, 2007.

Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012.

David Houle. Comparing evolvability and variability of quantitative traits. *Genetics*, 130(1):195–204, 1992.

François Jacob. Evolution and tinkering. *Science*, 1977.

Hawoong Jeong, Bálint Tombor, Réka Albert, Zoltan N Oltvai, and A-L Barabási. The large-scale organization of metabolic networks. *Nature*, 407(6804):651, 2000.

Nadav Kashtan and Uri Alon. Spontaneous evolution of modularity and network motifs. *Proceedings of the National Academy of Sciences of the United States of America*, 102 (39):13773–13778, 2005.

Nadav Kashtan, Avi E Mayo, Tomer Kalisky, and Uri Alon. An analytically solvable model for rapid evolution of modular structure. *PLoS computational biology*, 5(4): e1000355, 2009.

Nadav Kashtan, Elad Noor, and Uri Alon. Varying environments can speed up evolution. *Proceedings of the National Academy of Sciences*, 104(34):13711–13716, 2007.

Stuart A. Kauffman. *The origins of order: Self-organization and selection in evolution*. Oxford university press, 1993.

Tadeusz J Kawecki and Dieter Ebert. Conceptual issues in local adaptation. *Ecology letters*, 7(12):1225–1241, 2004.

Marc Kery, Diethart Matthies, and Hans-Heinrich Spillmann. Reduced fecundity and offspring performance in small populations of the declining grassland plants primula veris and gentiana lutea. *Journal of Ecology*, 88(1):17–30, 2000.

Caitlin R Kight and John P Swaddle. How and why environmental noise impacts animals: an integrative, mechanistic review. *Ecology letters*, 14(10):1052–1061, 2011.

Marc W Kirschner and John C Gerhart. Evolvability. *Proceedings of the National Academy of Sciences*, 95(15):8420–8427, 1998.

Marc W Kirschner and John C Gerhart. *The plausibility of life: Resolving Darwin's dilemma*. Yale University Press, 2006.

Ron Kohavi et al. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *IJCAI*, volume 14, pages 1137–1145, 1995.

Loizos Kounios, Jeff Clune, Kostas Kouvaris, Günter P Wagner, Mihaela Pavlicev, Daniel M Weinreich, and Richard A Watson. Resolving the paradox of evolvability with learning theory: How evolution learns to improve evolvability on rugged fitness landscapes. *arXiv preprint arXiv:1612.05955*, 2016.

Kostas Kouvaris, Jeff Clune, Loizos Kounios, Markus Brede, and Richard A Watson. How evolution learns to generalise: Using the principles of learning theory to understand the evolution of developmental organisation. *PLoS computational biology*, 13 (4):e1005358, 2017.

Siwei Lai, Liheng Xu, Kang Liu, and Jun Zhao. Recurrent convolutional neural networks for text classification. In *AAAI*, volume 333, pages 2267–2273, 2015.

Kevin N Laland, Tobias Uller, Marcus W Feldman, Kim Sterelny, Gerd B Müller, Armin Moczek, Eva Jablonka, and John Odling-Smee. The extended evolutionary synthesis: its structure, assumptions and predictions. In *Proc. R. Soc. B*, volume 282, page 20151019. The Royal Society, 2015.

R Lande. Evolution of phenotypic plasticity and environmental tolerance of a labile quantitative character in a fluctuating environment. *Journal of evolutionary biology*, 27(5):866–875, 2014.

Russell Lande. Evolutionary mechanisms of limb loss in tetrapods. *Evolution*, 32(1): 73–92, 1978.

Russell Lande. Adaptation to an extraordinary environment by evolution of phenotypic plasticity and genetic assimilation. *Journal of evolutionary biology*, 22(7):1435–1446, 2009.

Russell Lande and Stevan J Arnold. The measurement of selection on correlated characters. *Evolution*, 37(6):1210–1226, 1983.

Pat Langley. *Elements of machine learning*. Morgan Kaufmann, 1996.

Robert D Leclerc. Survival of the sparsest: robust gene networks are parsimonious. *Molecular systems biology*, 4(1), 2008.

Richard E Lenski, Jeffrey E Barrick, and Charles Ofria. Balancing robustness and evolvability. *PLoS biology*, 4(12):e428, 2006.

Richard E Lenski, Charles Ofria, Robert T Pennock, and Christoph Adami. The evolutionary origin of complex features. *Nature*, 423(6936):139–144, 2003.

Jesús Alberto León. Plasticity in fluctuating environments. *Lecture notes in biomathematics*, pages 105–105, 1993.

BR Levin. The evolution of sex in bacteria. *The evolution of sex*, pages 194–211, 1988.

Michael Levine and Eric H Davidson. Gene regulatory networks for development. *Proceedings of the National Academy of Sciences of the United States of America*, 102 (14):4936–4942, 2005.

Richard Levins. *Evolution in changing environments: some theoretical explorations*. Number 2. Princeton University Press, 1968.

Nicholas A Levis and David W Pfennig. Evaluating 'plasticity-first'evolution in nature: key criteria and empirical approaches. *Trends in ecology & evolution*, 31(7):563–574, 2016.

Richard C Lewontin. The units of selection. *Annual Review of Ecology and Systematics*, pages 1–18, 1970.

David A Liberles, Ashley I Teufel, Liang Liu, and Tanja Stadler. On the need for mechanistic models in computational genomics and metagenomics. *Genome biology and evolution*, 5(10):2008–2018, 2013.

Yan B Linhart and Michael C Grant. Evolutionary significance of local genetic differentiation in plants. *Annual review of ecology and systematics*, 27(1):237–277, 1996.

Hod Lipson and Jordan B Pollack. Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799):974–978, 2000.

Hod Lipson, Jordan B Pollack, and Nam P Suh. On the origin of modular variation. *Evolution*, 56(8):1549–1556, 2002.

Adi Livnat. Interaction-based evolution: how natural selection and nonrandom mutation worktogether. *Biology direct*, 8(1):1, 2013.

Adi Livnat. Simplification, innateness, and the absorption of meaning from context: how novelty arises from gradual network evolution. *Evolutionary Biology*, 44(2):145–189, 2017.

Adi Livnat, Christos Papadimitriou, Jonathan Dushoff, and Marcus W Feldman. A mixability theory for the role of sex in evolution. *Proceedings of the National Academy of Sciences*, 105(50):19803–19808, 2008.

William JR Longabaugh, Eric H Davidson, and Hamid Bolouri. Computational representation of developmental genetic regulatory networks. *Developmental biology*, 283 (1):1–16, 2005.

Richard Losick and Claude Desplan. Stochasticity and cell fate. *science*, 320(5872): 65–68, 2008.

George F Luger. *Artificial intelligence: structures and strategies for complex problem solving*. Pearson education, 2005.

Michael Lynch. The evolution of genetic networks by non-adaptive processes. *Nature Reviews Genetics*, 8(10):803–813, 2007.

Lesley T MacNeil and Albertha JM Walhout. Gene regulatory networks and the role of robustness and stochasticity in the control of gene expression. *Genome research*, 21 (5):645–657, 2011.

David Marr. Vision: A computational investigation into the human representation and processing of visual information, henry holt and co. *Inc., New York, NY*, pages 2–46, 1982.

DR Marshall and SK Jain. Phenotypic placticity of avena fatua and a. barbata. *The American Naturalist*, 102(927):457–467, 1968.

Julien GA Martin, Daniel H Nussey, Alastair J Wilson, and Denis Reale. Measuring individual differences in reaction norms in field and experimental studies: a power analysis of random regression models. *Methods in Ecology and Evolution*, 2(4):362–374, 2011.

Joanna Masel and Meredith V Trotter. Robustness and evolvability. *Trends in Genetics*, 26(9):406–414, 2010.

Jiří Matoušek. *Geometric discrepancy: An illustrated guide*. Springer, 1999.

Ernst Mayr. The emergence of evolutionary novelties. *Evolution after Darwin*, 1:349–380, 1960.

Ernst Mayr and William B Provine. *The evolutionary synthesis: perspectives on the unification of biology*. Harvard University Press, 1998.

John M McNamara, Sasha RX Dall, Peter Hammerstein, and Olof Leimar. Detection vs. selection: integration of genetic, epigenetic and environmental cues in fluctuating environments. *Ecology letters*, 19(10):1267–1276, 2016.

Colin D Meiklejohn and Daniel L Hartl. A single mode of canalization. *Trends in Ecology & Evolution*, 17(10):468–473, 2002.

Henok Mengistu, Joost Huizinga, Jean-Baptiste Mouret, and Jeff Clune. The evolutionary origins of hierarchy. *PLOS Comput Biol*, 12(6):e1004829, 2016.

JAJ Metz, RM Nisbet, and SAH Geritz. How should we define 'fitness' for general ecological scenarios? *Trends in Ecology & Evolution*, 7(6):198–202, 1992.

Armin P Moczek, Sonia Sultan, Susan Foster, Cris Ledón-Rettig, Ian Dworkin, H Fred Nijhout, Ehab Abouheif, and David W Pfennig. The role of developmental plasticity in evolutionary innovation. *Proceedings of the Royal Society of London B: Biological Sciences*, 278(1719):2705–2713, 2011.

Armin Philipp Moczek. Horn polyphenism in the beetle onthophagus taurus: larval diet quality and plasticity in parental investment determine adult body size and male horn morphology. *Behavioral Ecology*, 9(6):636–641, 1998.

Nancy A Moran. The evolutionary maintenance of alternative phenotypes. *The American Naturalist*, 139(5):971–989, 1992.

Gerd B Müller. Evo–devo: extending the evolutionary synthesis. *Nature Reviews Genetics*, 8(12):943–949, 2007.

Gerd B Muller and Günter P Wagner. Novelty in evolution: restructuring the concept. *Annual Review of Ecology and Systematics*, 22(1):229–256, 1991.

Courtney J Murren, Josh R Auld, H Callahan, Cameron K Ghalambor, Corey A Handelsman, Mary A Heskel, JG Kingsolver, Heidi J Maclean, Joanna Masel, Heather Maughan, et al. Constraints on the evolution of phenotypic plasticity: limits and costs of phenotype and plasticity. *Heredity*, 115(4):293, 2015.

Kaisa Mustajärvi, Pirkko Siikamäki, Saara Rytkönen, and Antti Lammi. Consequences of plant population size and density for plant–pollinator interactions and plant performance. *Journal of Ecology*, 89(1):80–87, 2001.

Antonine Nicoglou. The evolution of phenotypic plasticity: Genealogy of a debate in genetics. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 50:67–76, 2015.

H Frederik Nijhout. A developmental-physiological perspective on the development and evolution of phenotypic plasticity. In *Conceptual Change in Biology*, pages 147–173. Springer, 2015.

HF Nijhout and DJ Emlen. Competition among body parts in the development and evolution of insect morphology. *Proceedings of the National Academy of Sciences*, 95 (7):3685–3689, 1998.

Samir Okasha. Multilevel selection and the major transitions in evolution. *Philosophy of science*, 72(5):1013–1025, 2005.

Randall C O'Reilly. *The Leabra model of neural interactions and learning in the neocortex*. PhD thesis, Carnegie Mellon University, 1996.

Randall C O'Reilly and Yuko Munakata. *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. MIT press, 2000.

Ingo Paenke, Bernhard Sendhoff, and Tadeusz J Kawecki. Influence of plasticity and learning on evolution under directional selection. *The American Naturalist*, 170(2): E47–E58, 2007.

Michael E Palmer and Marcus W Feldman. Survivability is more fundamental than evolvability. *PloS one*, 7(6):e38025, 2012.

Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2010.

Merav Parter, Nadav Kashtan, and Uri Alon. Facilitated variation: how evolution learns from past environments to generalize to new environments. *PLoS Computational Biology*, 4(11):e1000206, 2008.

Melanie Paschke, Giorgina Bernasconi, and Bernhard Schmid. Population size and identity influence the reaction norm of the rare, endemic plant cochlearia bavarica across a gradient of environmental stress. *Evolution*, 57(3):496–508, 2003.

Mihaela Pavličev and James M Cheverud. Constraints evolve: Context-dependency of gene effects allows evolution of pleiotropy. *Annual Review of Ecology, Evolution, and Systematics*, 46(1), 2015.

Mihaela Pavlicev, James M Cheverud, and Günter P Wagner. Evolution of adaptive phenotypic variation patterns by direct selection for evolvability. *Proceedings of the Royal Society B: Biological Sciences*, page rspb20102113, 2010.

Mihaela Pavlicev and Thomas F Hansen. Genotype-phenotype maps maximizing evolvability: Modularity revisited. *Evolutionary Biology*, 38(4):371–389, 2011.

Christophe Pélabon, Thomas F Hansen, Matthew L Carlson, and W Scott Armbruster. Variational and genetic properties of developmental stability in dalechampia scandens. *Evolution*, 58(3):504–514, 2004.

David W Pfennig, Matthew A Wund, Emilie C Snell-Rood, Tami Cruickshank, Carl D Schlichting, and Armin P Moczek. Phenotypic plasticity's impacts on diversification and speciation. *Trends in ecology & evolution*, 25(8):459–467, 2010.

Vu Pham, Christopher Kermorvant, and Jérôme Louradour. Dropout improves recurrent neural networks for handwriting recognition. *arXiv preprint arXiv:1312.4569*, 2013.

Massimo Pigliucci. *Phenotypic plasticity: beyond nature and nurture*. JHU Press, 2001.

Massimo Pigliucci. Evolution of phenotypic plasticity: where are we going now? *Trends in Ecology & Evolution*, 20(9):481–486, 2005.

Massimo Pigliucci. Do we need an extended evolutionary synthesis? *Evolution*, 61(12): 2743–2749, 2007.

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

Massimo Pigliucci and Courtney J Murren. Perspective: Genetic assimilation and a possible evolutionary paradox: can macroevolution sometimes be so fast as to pass us by? *Evolution*, 57(7):1455–1464, 2003.

Massimo Pigliucci, Courtney J Murren, and Carl D Schlichting. Phenotypic plasticity and evolution by genetic assimilation. *Journal of Experimental Biology*, 209(12):2362–2367, 2006.

Trevor D Price, Anna Qvarnström, and Darren E Irwin. The role of phenotypic plasticity in driving genetic evolution. *Proceedings of the Royal Society of London B: Biological Sciences*, 270(1523):1433–1440, 2003.

Adrian P Quayle and Seth Bullock. Modelling the evolution of genetic regulatory networks. *Journal of Theoretical Biology*, 238(4):737–753, 2006.

Stan P Rachootin and Keith Stewart Thomson. Epigenetics, paleontology, and evolution. *Evolution today*, 2:181–193, 1981.

Arjun Raj and Alexander van Oudenaarden. Nature, nurture, or chance: stochastic gene expression and its consequences. *Cell*, 135(2):216–226, 2008.

Etienne Rajon and Joanna Masel. Evolution of molecular error rates and the consequences for evolvability. *Proceedings of the National Academy of Sciences*, 108(3): 1082–1087, 2011.

Rupert Riedl and Richard Peter Spencer Jefferies. *Order in living organisms: a systems analysis of evolution*. Wiley New York, 1978.

Frank Rosenblatt. *The perceptron, a perceiving and recognizing automaton Project Para*. Cornell Aeronautical Laboratory, 1957.

MaryCarol Rossiter. Incidence and consequences of inherited environmental effects. *Annual review of ecology and systematics*, 27(1):451–476, 1996.

Stuart Russell, Peter Norvig, and Artificial Intelligence. A modern approach. *Artificial Intelligence. Prentice-Hall, Egnlewood Cliffs*, 25:27, 1995.

SAHOTRA Sarkar. From the reaktionsnorm to the evolution of adaptive plasticity. *Phenotypic plasticity. Functional and conceptual approaches. Editores: TJ DeWitt y SM Scheiner. Oxford University Press, New York. págs*, pages 10–30, 2004.

Samuel M Scheiner. Genetics and evolution of phenotypic plasticity. *Annual review of ecology and systematics*, pages 35–68, 1993.

D Schleicherová, G Sella, S Meconcelli, R Simonini, MP Martino, P Cervella, and MC Lorenzi. Does the cost of a function affect its degree of plasticity? a test on plastic sex allocation in three closely related species of hermaphrodites. *Journal of Experimental Marine Biology and Ecology*, 453:148–153, 2014.

Carl D Schlichting. The evolution of phenotypic plasticity in plants. *Annual review of ecology and systematics*, pages 667–693, 1986.

Carl D Schlichting and Courtney J Murren. Evolvability and the raw materials for adaptation. *Plant Adaptation: Molecular genetics and ecology. NRC research Press, Ottawa*, pages 18–29, 2004.

Carl D Schlichting, Massimo Pigliucci, et al. *Phenotypic evolution: a reaction norm perspective.* Sinauer Associates Incorporated, 1998.

Carl D Schlichting and Matthew A Wund. Phenotypic plasticity and epigenetic marking: an assessment of evidence for genetic accommodation. *Evolution*, 68(3):656–672, 2014.

CD Schlichting. The role of phenotypic plasticity in diversification. *Phenotypic Plasticity: functional and conceptual approaches*, pages 191–200, 2004.

Ivan Ivanovich Schmalhausen. Factors of evolution: the theory of stabilizing selection. 1949.

Tanja Schwander and Olof Leimar. Genes as leaders and followers in evolution. *Trends in Ecology & Evolution*, 26(3):143–151, 2011.

Gideon Schwarz et al. Estimating the dimension of a model. *The annals of statistics*, 6 (2):461–464, 1978.

Cosma Rohilla Shalizi et al. Dynamics of bayesian updating with dependent data and misspecified models. *Electronic Journal of Statistics*, 3:1039–1074, 2009.

Claude Elwood Shannon. A mathematical theory of communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 5(1):3–55, 2001.

James Alan Shapiro. *Evolution: a view from the 21st century.* Pearson Education, 2011.

Herbert A Simon. Why should machines learn? In *Machine learning*, pages 25–37. Springer, 1983.

George Gaylord Simpson. The baldwin effect. *Evolution*, 7(2):110–117, 1953.

J Maynard Smith, Richard Burian, Stuart Kauffman, Pere Alberch, John Campbell, Brian Goodwin, Russell Lande, David Raup, and Lewis Wolpert. Developmental constraints and evolution: a perspective from the mountain lake conference on development and evolution. *Quarterly Review of Biology*, pages 265–287, 1985.

Tom Smith, Phil Husbands, and Michael O'Shea. Evolvability and analysis of robot control networks. In *From Animals to Animats 7: Proceedings of the Seventh International Conference on Simulation of Adaptive Behavior*, volume 7, page 349. MIT Press, 2002.

Terence Soule and James A Foster. Effects of code growth and parsimony pressure on populations in genetic programming. *Evolutionary Computation*, 6(4):293–309, 1998.

Georg F Striedter. Précis of principles of brain evolution. *Behavioral and Brain Sciences*, 29(01):1–12, 2006.

SE Sultan. Phenotypic plasticity and plant adaptation. *Acta Botanica Neerlandica*, 44 (4):363–383, 1995.

SE Sultan and FA Bazzaz. Phenotypic plasticity in polygonum persicaria. i. diversity and uniformity in genotypic norms of reaction to light. *Evolution*, 47(4):1009–1031, 1993.

Sonia E Sultan and Hamish G Spencer. Metapopulation structure favors plasticity over local adaptation. *The American Naturalist*, 160(2):271–283, 2002.

Yuichiro Suzuki and H Frederik Nijhout. Evolution of a polyphenism by genetic accommodation. *Science*, 311(5761):650–652, 2006.

Peter S Swain, Michael B Elowitz, and Eric D Siggia. Intrinsic and extrinsic contributions to stochasticity in gene expression. *Proceedings of the National Academy of Sciences*, 99(20):12795–12800, 2002.

Eors Szathmary and J Maynard Smith. *The major transitions in evolution*. Oxford University Press, 2004.

Catherine A Tauber and Maurice J Tauber. Phenotypic plasticity in chrysoperla: genetic variation in the sensory mechanism and in correlated reproductive traits. *Evolution*, 46(6):1754–1773, 1992.

Mukund Thattai and Alexander Van Oudenaarden. Stochastic gene expression in fluctuating environments. *Genetics*, 167(1):523–530, 2004.

Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

Lisa Torrey and Jude Shavlik. Transfer learning. *Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques*, 1:242, 2009.

Marc Toussaint. On the evolution of phenotypic exploration distributions. In *FOGA*, pages 169–182. Citeseer, 2002.

Marc Toussaint and Werner von Seelen. Complex adaptation and system structure. *BioSystems*, 90(3):769–782, 2007.

Peter Turney. Increasing evolvability considered as a large-scale trend in evolution. 1999.

Leslie Valiant. *Probably approximately correct: nature's algorithms for learning and prospering in a complex world*. Basic Books, 2013.

Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11): 1134–1142, 1984.

Leslie G Valiant. Evolvability. *Journal of the ACM (JACM)*, 56(1):3, 2009.

Mark Van Kleunen and Markus Fischer. Constraints on the evolution of adaptive phenotypic plasticity in plants. *New Phytologist*, 166(1):49–60, 2005.

Peter H Van Tienderen. Evolution of generalists and specialist in spatially heterogeneous environments. *Evolution*, pages 1317–1331, 1991.

Sara Via. Adaptive phenotypic plasticity: target or by-product of selection in a variable environment? *The American Naturalist*, 142(2):352–365, 1993.

Sara Via, Richard Gomulkiewicz, Gerdien De Jong, Samuel M Scheiner, Carl D Schlichting, and Peter H Van Tienderen. Adaptive phenotypic plasticity: consensus and controversy. *Trends in Ecology & Evolution*, 10(5):212–217, 1995.

Jiří Vohradský. Neural model of the genetic network. *Journal of Biological Chemistry*, 276(39):36168–36173, 2001a.

Jiří Vohradský. Neural network model of gene expression. *The FASEB Journal*, 15(3): 846–854, 2001b.

Sergei Volis, Danara Ormanbekova, Kanat Yermekbayev, Minshu Song, and Irina Shulgina. Multi-approaches analysis reveals local adaptation in the emmer wheat (triticum dicoccoides) at macro-but not micro-geographical scale. *PloS one*, 10(3): e0121153, 2015.

Conrad H Waddington. Canalization of development and the inheritance of acquired characters. *Nature*, 150(3811):563–565, 1942.

Conrad H Waddington. Genetic assimilation of an acquired character. *Evolution*, 7(2): 118–126, 1953.

Conrad H Waddington. Canalization of development and genetic assimilation of acquired characters. *Nature*, 183(4676):1654–1655, 1959.

Conrad Hal Waddington. *The strategy of the genes*. Routledge, 2014.

Andreas Wagner. Does evolutionary plasticity evolve? *Evolution*, pages 1008–1023, 1996.

Andreas Wagner. Robustness and evolvability: a paradox resolved. *Proceedings of the Royal Society B: Biological Sciences*, 275(1630):91–100, 2008.

Andreas Wagner. *Robustness and evolvability in living systems*. Princeton University Press, 2013.

Günter P Wagner. The biological homology concept. *Annual Review of Ecology and Systematics*, pages 51–69, 1989.

Günter P Wagner. *Homology, genes, and evolutionary innovation*. princeton university press, 2014.

Gunter P Wagner and Lee Altenberg. Perspective: Complex adaptations and the evolution of evolvability. *Evolution*, pages 967–976, 1996.

Günter P Wagner and Vincent J Lynch. Evolutionary novelties. *Current Biology*, 20(2): R48–R52, 2010.

Günter P Wagner, Mihaela Pavlicev, and James M Cheverud. The road to modularity. *Nature Reviews Genetics*, 8(12):921–931, 2007.

Richard A Watson. *Compositional evolution: the impact of sex, symbiosis and modularity on the gradualist framework of evolution*. Mit Press, 2006.

Richard A Watson, Rob Mills, CL Buckley, Kostas Kouvaris, Adam Jackson, Simon T Powers, Chris Cox, Simon Tudge, Adam Davies, Loizos Kounios, et al. Evolutionary connectionism: algorithmic principles underlying the evolution of biological organisation in evo-devo, evo-eco and evolutionary transitions. *Evolutionary Biology*, pages 1–29, 2015.

Richard A Watson, Rob Mills, CL Buckley, Kostas Kouvaris, Adam Jackson, Simon T Powers, Chris Cox, Simon Tudge, Adam Davies, Loizos Kounios, et al. Evolutionary connectionism: algorithmic principles underlying the evolution of biological organisation in evo-devo, evo-eco and evolutionary transitions. *Evolutionary biology*, 43(4): 553–581, 2016.

Richard A. Watson and Eörs Szathmáry. How can evolution learn? *Trends in Ecology and Evolution*, 2015a.

Richard A Watson and Eörs Szathmáry. How can evolution learn? *interactions*, 32(33): 82, 2015b.

Richard A Watson and Eörs Szathmáry. How can evolution learn? *Trends in ecology & evolution*, 31(2):147–157, 2016.

Richard A Watson, Günter P Wagner, Mihaela Pavlicev, Daniel M Weinreich, and Rob Mills. The evolution of phenotypic correlations and "developmental memory". *Evolution*, 68(4):1124–1138, 2014.

Daniel C Weaver, Christopher T Workman, and Gary D Stormo. Modeling regulatory networks with weight matrices. In *Biocomputing'99*, pages 112–123. World Scientific, 1999.

Lodewyk FA Wessels, Eugene P van Someren, Marcel JT Reinders, et al. A comparison of genetic network models. In *pacific Symposium on Biocomputing*, volume 6, pages 508–519. World Scientific, 2001.

Mary Jane West-Eberhard. Phenotypic plasticity and the origins of diversity. *Annual review of Ecology and Systematics*, pages 249–278, 1989.

Mary Jane West-Eberhard. Evolution in the light of developmental and cell biology, and vice versa. *Proceedings of the National Academy of Sciences*, 95(15):8417–8419, 1998.

Mary Jane West-Eberhard. *Developmental plasticity and evolution*. Oxford University Press, 2003.

David F Westneat, Margret I Hatch, Daniel P Wetzel, and Amanda L Ensminger. Individual variation in parental care reaction norms: integration of personality and plasticity. *The American Naturalist*, 178(5):652–667, 2011.

David H Wolpert and William G Macready. No free lunch theorems for search. Technical report, Technical Report SFI-TR-95-02-010, Santa Fe Institute, 1995.

Sewall Wright. Evolution in mendelian populations. *Genetics*, 16(2):97–159, 1931.

Lev Y Yampolsky and Arlin Stoltzfus. Bias in the introduction of variation as an orienting factor in evolution. *Evolution & development*, 3(2):73–83, 2001.

Jin Yoshimura and William M Shields. Probabilistic optimization of body size: a discrepancy between genetic and phenotypic optima. *Evolution*, 49(2):375–378, 1995.

Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2): 301–320, 2005.