

Evolution by Natural Induction

Richard Watson^{1*}, Michael Levin² and Tim Lewens³

¹ Dept. Computer Science/ Institute for Life Sciences, University of Southampton, U.K.. SO17 1BJ
(*corresponding author: R.A.Watson@soton.ac.uk)

² Allen Discovery Centre, Tufts University, Boston, MA, USA. (Michael.Levin@tufts.edu)

³ Dept. of History and Philosophy of Science, Cambridge University, U.K. (tml1000@cam.ac.uk)

It is conventionally assumed that all evolutionary adaptation is produced, and could only possibly be produced, by natural selection. Natural induction is a different mechanism of adaptation. It occurs in dynamical systems described by a network of interactions, where connections give way slightly under stress and the system is subject to occasional perturbations. This differential adjustment of connections causes reorganisation of the system's internal structure in a manner equivalent to associative learning familiar in neural networks. This is sufficient for storage and recall of multiple patterns, learning with generalisation and solving difficult constraint problems (without any natural selection involved). Various biological systems (from gene-regulation networks, to metabolic networks, to ecosystems) meet these basic conditions and therefore have potential to exhibit adaptation by natural induction. Here (and in a follow-on paper) we consider various ways that natural induction and natural selection might interact in biological evolution. For example, in some cases, natural selection may act not as a source of adaptations but as a memory of adaptations discovered by natural induction. We conclude that evolution by natural induction is a viable process that expands our understanding of evolutionary adaptation.

1. Introduction: Possible mechanisms of adaptation.

Living things exhibit adaptive complexity in the sense of an appearance of design. For our purposes, we can take this to mean a system with many interdependent parts in specific arrangement, that provides a solution to some problem or serves some function that is, in some sense, special or difficult to attain (Lewens and Watson 2024, Watson 2024). It is this that Darwin sought to explain with evolution by natural selection.

It is generally taken for granted that natural selection is the sole and complete explanation for the adaptive complexity of living things (Dawkins 1983, Dennett 1996, Stoltzfus 2012). Indeed, this is so taken for granted that biological adaptation is often defined as the product of evolution by natural selection, and anything that is not the product of natural selection simply is not considered to be an adaptation. Furthermore, it is widely believed that natural selection is the only possible mechanism that could in principle cause any system to improve, or increase in adaptive complexity (absent a designer) (Weismann 1893, Dawkins 1983, Stoltzfus 2012, Lewens and Watson 2024).

However, natural selection is not the only natural mechanism by which a system can become adaptively organised (Buckley, Lewens et al. 2024, Lewens and Watson 2024). Learning is an obvious and well-known mechanism of adaptation (and not the same as selection¹) familiar in the context of organisms with brains and in connectionist models of machine learning (i.e. neural networks) (Lewens and Watson 2024). Nevertheless, it is usually assumed that learning requires specific learning machinery, such as neurons and brains, and is thus only relevant in these special contexts. If an ability

¹ In Lewontin's terms, selection is a *variational* process whereas learning is a *transformational* process (Lewontin 1980; 1970).

to learn is rare in the tree of life, and a late arrival in the evolutionary story, then although an ability to learn can be a product of evolution (Mitchell 2023) it cannot be a part of how evolution works in general.

It has become clear that this view of learning, and the natural intelligence it provides, is not correct. In fact, wherever we look in the tree of life, we find that living things exhibit some capacity of 'basal cognition', e.g. an ability to store and recall past states or behaviours, learn from experience, and problem-solve (Baluska and Levin 2016, Manicka and Levin 2019, Manicka and Levin 2019, Lyon, Keijzer et al. 2021, Lyon, Keijzer et al. 2021), in many different 'problem spaces' (e.g. morphological, transcriptional, physiological, behavioural) (Fields and Levin 2022). These capabilities appear in animals without neurons or brains, fungi, plants, single-celled eukaryotes and even bacteria (Gershman, Horvitz and Tenenbaum 2015, Baluska and Levin 2016, Vallverdu, Castro et al. 2018, Levin 2019, Levin and Dennett 2020, Biswas, Manicka et al. 2021, Levin, Keijzer et al. 2021, Lyon, Keijzer et al. 2021, Lyon, Keijzer et al. 2021, Fields and Levin 2022, Gunawardena 2022). If learning is more widespread in the tree of life, then its effects in biological evolution are not necessarily rare or late, and it could be a significant player in the evolutionary process (Levin 2023). Nevertheless, if learning is a derivative product of natural selection (e.g. because it requires special mechanisms, even if they are not always neural), then regardless of how widespread it might be, natural selection remains as the 'prime mover' in the narrative of adaptation.

But it turns out that learning does not necessarily need anything special at all. Many kinds of physical materials exhibit a kind of learning and memory, without any biology involved (Zhong, Schwab and Murugan 2017, Keim, Paulsen et al. 2019, Wong, Cleland et al. 2023). Various physical systems (from macro-molecules (Zhong, Schwab and Murugan 2017), to soap bubbles (Mukherji, Kandula et al. 2019), to paper folds (Dudte, Vouga et al. 2016)) are sufficient to exhibit learning, including in unsupervised learning scenarios (i.e. without a 'teacher' to prescribe desired outputs) (Stern, Arinze et al. 2020, Stern, Pinson and Murugan 2020, Stern, Hexner et al. 2021, Stern and Murugan 2022). Chemical circuits can do probabilistic inference (Katz, Springer and Fontana 2018, Katz and Fontana 2022) and learning is now studied in a range of non-living materials (Cragg and Temperley 1955, Kaygisiz and Ulijn 2024). England (England 2015, Chvykov, Berrueta et al. 2021) describes a general propensity for dynamical systems with many degrees of freedom to be internally (re)organised by a driving force. This causes them to find an arrangement that moves more easefully with the pattern of driving or its past contact with the environment. In so doing it becomes better able to resolve the tensions that are created by that kind of forcing when it experiences it again. Recently, it has been shown that, with only very general assumptions, this kind of non-biological learning has non-trivial cognitive capabilities – forming an associative memory that can store and recall multiple past states, learning with generalisation, and adaptive problem-solving capable of solving difficult combinatorial optimisation problems (Buckley, Lewens et al. 2024). Its behaviour recognizes familiar patterns, anticipates other patterns that are similar, and moves more easily into configurations that mirror the behaviour of its environment. This can be demonstrated in any dynamical system described by a network of partially elastic (viscoelastic) interactions that give way under stress, including something as simple as a network of masses connected randomly by (imperfectly elastic) springs. The learning occurs without selection or design either in the 'set-up' (construction) of the system or at 'run time' (operation) (Buckley, Lewens et al. 2024, Lewens and Watson 2024).

These systems do not reproduce (in part or whole), and so they do not exhibit adaptation understood in any way that defines that notion with respect to reproductive fitness. However, they can find configurations that resolve complex frustrated constraints between system components, increase the goodness of fit between the system and its environment, and find solutions to non-trivial

combinatorial optimisation problems (Buckley, Lewens et al. 2024). The system does not merely find the nearest locally optimal configuration (in the manner of a local gradient descent process), but learns to find better solutions with experience, eventually finding solutions that are of exceptional quality compared to the average of local optima (Buckley, Lewens et al. 2024). Buckley et al (2024) call the mechanism behind this kind of adaptive process “natural induction” (Section 2).

In short, the belief that natural selection is the only possible natural source of adaptive organisation is not correct; Adaptation by natural induction is a different process of adaptation (Lewens and Watson 2024). But is adaptation by natural induction relevant to evolutionary adaptation in biology as we know it? Various biological systems (from gene regulation networks to ecosystems) meet the necessary conditions for natural induction to occur and therefore have potential to exhibit adaptation by this mechanism (Watson, Buckley et al. 2010, Watson, Wagner et al. 2014, Kounios, Clune et al. 2016, Watson and Szathmary 2016, Kouvaris, Clune et al. 2017, Nash 2021, Watson and Levin 2023, Buckley, Lewens et al. 2024). However, since they are biological, these systems also involve natural selection, at some level of organisation or another. So, what is the nature of the possible interaction between natural induction and natural selection in biological evolution?

Lewens and Watson point out that there are common arguments used in evolutionary thinking that depend on the assumption that natural selection holds an exclusive position in producing and explaining biological adaptation. For example, the argument that “system x cannot be adapted because x does not have the necessary and sufficient conditions for natural selection to occur” is frequently applied to ecosystems, loose social groups, and the biosphere as a whole (Dawkins 1982, Wilson 2016). Similarly, the following argument is frequently applied to many mechanisms including adaptive phenotypic plasticity, niche construction, developmental bias, epigenetic inheritance, symbiosis/symbiogenesis and self-organisation: “mechanism y cannot be a new source of adaptive organisation in biological evolution, and therefore cannot be something that changes our understanding of how evolution works in any substantial way, because to the extent that y introduces change that is directed toward adaptive outcomes this can only be because it is itself a product of evolution by natural selection” (Dawkins 1983, Margulis and Fester 1991, Maynard Smith and Szathmary 1997, Laland, Uller et al. 2014, Laland, Uller et al. 2015, Watson and Thies 2019). Even though these arguments rely on the faulty premise that natural selection is the only process that can produce adaptation, it might still be true that, for example, the biosphere is not adapted because neither natural selection nor natural induction (nor any other adaptive mechanism) applies in this particular case. These possibilities therefore need to be assessed on a case-by-case basis (Lewens and Watson 2024).

However, before getting to empirical questions, there are several conceptual issues that need to be addressed. In this paper, we discuss these issues and articulate some of the possibilities for the way that natural selection and natural induction might interact with each other. We limit attention here to cases where evolution is constituted by genetic change in a population, and all the genetic change that occurs arises through random genetic variation and natural selection (other scenarios are possible (Livnat 2013, Melamed, Nov et al. 2022, Livnat and Love 2024)). This might seem to eliminate any controversy, ensuring that natural selection is a complete explanation for any evolutionary adaptation that occurs. Nonetheless, we discuss some scenarios where natural selection would not be the source of the adaptive organisation that evolves even within these limitations. This is because adaptive evolution is possible in these cases only because natural induction is shaping the selective

field (i.e. the distribution of selection coefficients over the genome²) in which natural selection occurs. In other words, natural induction tells natural selection what to do: it leads it to find adaptive organisations it would not otherwise discover. Natural selection thus provides a necessary sub-mechanism for genetic evolution to occur, but in these cases, natural selection does not describe the mechanism of adaptation.

In this paper we discuss the particular scenario of phenotype-first evolution, where natural selection canalises adaptations already found by natural induction within individuals' lifetimes (West-Eberhard 2003, West-Eberhard 2005, Schwander and Leimar 2011). In this scenario it should be clear that although, under different circumstances, natural selection over many generations could provide adaptations at the level of organismal traits, natural selection is not the source of the adaptation in this case. This is because the adaptation arises entirely without natural selection, before any natural selection occurs (i.e. within one lifetime), and because the adaptation found by natural induction can be better than the one found by natural selection.

In the following section, we give some more background on what natural induction is, how it works and how it differs from natural selection. In Section 3, we discuss some potential objections to the idea that anything other than natural selection can be relevant to biological evolution. In Section 4, we describe the phenotype-first evolutionary scenario where natural selection and natural induction could interact together in biological evolution. Section 5 Concludes.

In a companion paper (Watson, Levin and Lewens 2025) we discuss some other possible ways that natural selection and natural induction might interact in biological evolution. These are more subtle because natural selection is integral to the operation of natural induction in these cases and so it cannot be said that the adaptation arises without natural selection. Nonetheless, these adaptations cannot be attributed to natural selection because they occur at organisational scales or timescales where natural selection does not apply. Specifically, we discuss the evolution of ecosystem adaptations (where the ecosystem is not an evolutionary unit) and the evolution of evolvability (where characteristics of evolvability evolve before selection acts on phenotypes enabled by this evolvability). This second paper also provides more in-depth discussion on the similarities and contrasts between natural induction and natural selection.

2. Adaptation by natural Induction

Buckley et al. demonstrate adaptation by natural induction in a random network of viscoelastic connections (e.g. springs that give-way slightly under stress) (Buckley, Lewens et al. 2024). Being viscoelastic means that the connections are partly elastic and partly viscous or plastic. For example, when springs are stretched too far, or for too long, they do not go back all the way to their original length when released – they are slightly deformed by the forcing they have experienced. The basis of the phenomenon is that, as the system interacts with the environment, the springs give-way or deform differentially according to the stress this interaction creates in the network. The structure of the

² We introduce the term 'selective field' to refer to a distribution of selective coefficients over the genome but also to connote a notion that allows for these selective conditions to be constructed or optionally applied to genetic variations in a manner specific to the organism. The conventional notion of a fitness landscape, in contrast, generally refers to the effects of ambient environmental conditions external to the organism. And although the notion of niche construction permits the possibility that the action of an organism might alter selection, this is enacted through the external environment. Our use of selective field has more emphasis on conditions internal to the organism. This allows for the possibility that, without any change to the external environment, the properties of an organism (e.g. its learning capabilities) may 'exert an influence' that alters the distribution of selection coefficients that apply to the genetic variations it is carrying.

network thus holds a residue, an imprint, or memory of past experiences that modifies its future interactions with the external environment. Specifically, these structural changes make the system more likely to revisit configurations it has visited in the past. Moreover, because the changes are 'associative' (i.e. the length of a spring affects *correlations* among mass positions), this memory can generalise. The underlying principle is the same as connectionist models of learning in neural networks – involving changes to connections and linkages that control associations. This causes it to show the formation of an associative memory, learning with generalisation, and adaptive problem-solving behaviour (Buckley, Lewens et al. 2024).

These behaviours can also be shown in a simple kind of artificial neural network, the Hopfield network with Hebbian learning (Hopfield 1982, Hopfield 1984, Hopfield and Tank 1986, Tank and Hopfield 1987, Watson, Buckley and Mills 2011), and for the same reasons. Hebbian learning is often described as 'neurons that fire together wire together', meaning that the more two neurons are excited at the same time, the more the connection between them is changed in the direction that makes their future co-excitation easier or more likely. In physical systems, the more general principle is that 'parts that accommodate together integrate together'; they align their behaviors, join-up, and respond as one – just because this is the easeful way to be (given the pattern of stresses applied to the system). The result is a system that 'mirrors' the structure and behaviour of its environment – a model holding knowledge of its past experiences. Dynamically, the principle is simply positive feedback on correlations, or 'correlation becomes causation'. The consequence of these changes to connections is to enlarge the dynamical basin of attraction for patterns that include this correlation, making it more likely that the system will visit such patterns from a given distribution of initial conditions. And from this, memory, learning and problem-solving behaviours follow.

These adaptive abilities may seem to require learning machinery that has been selected or designed for this purpose. One might wonder, for example, why anything like a Hopfield network should appear spontaneously, and why relations between elements should change in a Hebbian manner rather than anti-Hebbian. However, the network assumptions are very general³ and the mechanistic principle involved is a natural process of relaxation or energy minimisation applicable to any physical system. When this applies to the connections of a network, its effect is to facilitate a 'second-order' energy minimisation (i.e. relaxing the constraints that limit its ability to relax), giving it the ability to relax more easily in future into the same correlations. This provides positive feedback on correlations that is functionally equivalent to Hebbian learning in the Hopfield network (Watson, Buckley and Mills 2011, Watson, Mills and Buckley 2011). For example, in the spring network, the more two masses are forced to be a particular distance apart (e.g. by external factors or other spring forces in that context) the more the spring between them accommodates to that length by material yielding, making it easier for them to have that correlated relation in future. It is important to recognise that this is not a uniform system-wide relaxation of all connections but one that is specific: a relaxation of one spring and not another, in a manner that accommodates to the particular historical experience of the system. In other words, it is a relaxation that creates specific organisation.

Buckley et al (Buckley, Lewens et al. 2024) describe the general conditions for natural induction to occur: a) a dynamical system described by a set of state variables and a network of

³ The textbook Hopfield network is a fully connected set of nodes with symmetric connections (Hopfield 1982; 1984). These conditions can also be relaxed to partially connected and, within limits, connections do not have to be entirely symmetric (Watson et al. 2014). The behaviour of each node is simply that it produces an output if its inputs are sufficiently strong, or more generally, the output is a non-linear weighted sum of the inputs. This describes many kinds of physical interaction networks and chemistries, both biological and non-biological (Ashby 1952; Farmer 1990; Fontana 1990; Varela & Bourgine 1992; Gross & Sayama 2009)

interactions (i.e. A pushes B, or a change in A causes a change or reaction in B, and vice versa), b) the system is subjected to occasional shocks or perturbations, or periodic/intermittent stresses (e.g. that disturb or randomise the values of the state variables), c) the interactions are not perfectly elastic but slightly 'viscoelastic', i.e. they give-way under stress. These conditions are likely met by many components of living materials at multiple scales.

In general terms, the mechanism of adaptation by natural induction involves changes that occur at different timescales, as follows:

- 1) The values of the state variables (e.g. positions of masses) are modified by contact with the environment (e.g. external forcing) and by a network of interactions between themselves (e.g. the springy connections). These dynamics are (at least partly) described by the local minimisation of an energy function, or relaxation of the system. On short timescales, the dynamics reach a local equilibrium. We will call this 'first-order relaxation' of the system.
- 2) On longer timescales, the state variables are repeatedly perturbed out of equilibrium (e.g. by external shocks that randomise the positions of the masses). In between each disturbance, the system relaxes again (to a possibly different equilibrium).
- 3) Meanwhile, on a much slower timescale, the connections in the network give-way slightly (e.g. by material creep in the springs). This occurs slowly over a timescale that incorporates many disturbances and re-equilibrations of the states, such that the accommodation of the connections to the configurations of the state variables occurs over a distribution of many local equilibria. Note that the slow-changing connections do not control individual state variables; they affect the *interactions* between the state variables and hence the *combinations* of states the system will exhibit. We call this 'second-order relaxation' because it is a relaxation of the interaction terms that determine how the first-order variables relax.

For any dynamical system described by a network of viscoelastic interactions, 1 and 3 are natural processes of energy minimisation (first and second order, respectively) that do not require any special contrivance. The application of shocks or perturbations (2) and the timescales of these different processes need to be such that the system spends most of its time at better-than-random configurations or equilibria, and many such configurations are visited over the timescale that the connections relax.

The initial behaviour of the system (1 and 2) effects a simple gradient descent process that finds a distribution of locally optimal solutions to the intrinsic constraint problem of the system, i.e. the system of constraints among the state variables imposed either by the pattern of external forcing or the internal network interactions. This is nothing special: intuitively it is analogous to a ball repeatedly rolling downhill in the (high-dimensional) energy function of the system. However, this behaviour changes over time because the changes to the connections (3) alter the shape of the 'hill'. Specifically, changes to the connections change the energy function of the system causing it to move to different configurations (in first-order relaxations) (Watson 2024). In particular, this causes it to find equilibria it has already visited more frequently, and also other equilibrium states that are (in a sense we will discuss below) generalisations that are similar to equilibrium states it has visited before. The interesting thing is that these other states can be extremely good quality solutions to the original constraint problem of the system and better than any states found in multiple trials of first-order relaxation (Buckley, Lewens et al. 2024).

This kind of adaptation is very different from natural selection. Mechanistically, it does not involve any differential survival or reproduction, no random variation and selective retention, and no heritable variation in fitness. Neither does it involve any differential survival or persistence of entities in a population, e.g. of interactors rather than replicators (Hull 1980) (see ‘E’ in the next subsection). Because natural induction has different mechanisms, the adaptation it produces may therefore arise in circumstances where natural selection does not apply. It also has different adaptive capabilities from natural selection because natural selection (in basic scenarios) finds, at best, only a locally optimal configuration. In fact, in several important senses, natural induction is not just different from natural selection, it is its opposite or complement (Table 1).

Natural selection	Natural induction
Variational process (differential retention and multiplication of variants in a population)	Transformational process (configurational or conformational rearrangements or accommodation of a single system) (Lewontin 1970, Lewontin 1980)
Change in frequencies of entities (non-relational)	Change of relationships between entities (not frequencies)
Separates the things to retain from the things to discard	Differentially eases the frustrations between things
Specific retention (selection) and non-specific variation (e.g. random mutation)	Specific variation (directed transformation) and non-specific retention (just one system, no population)
Amplifying differences/asymmetries	Amplifying complementarity/symmetries
Deductive aspect is foregrounded, inductive aspect is implicit/ in background.	Inductive aspect is foregrounded, deductive aspect is in background.
Invites attention on the differential retention of existing structures, e.g. “What persists exists” (Pross 2011, Pross 2016).	Invites attention on the creation of new structures by transforming the way existing things associate, e.g. “What relates creates” as developed by Watson elsewhere (Watson 2023), and in (Watson, Levin and Lewens 2025).

Table 1: Some complementary aspects of natural selection and natural induction.

The name “natural induction” refers to the type of logical inference that underlies this type of learning. Whereas paradigm cases of deductive inference often involve the drawing of specific conclusions from general premises, induction is the complement; drawing general conclusions from specific examples (see Discussion in (Watson, Levin and Lewens 2025)). Expanding on this a little helps us to see the conceptual complementarity between natural selection and natural induction, and that both have an elegant logical basis. In particular, although any adaptive process must both produce something new (inductively) and ensure this agrees with experience (deductively), the emphasis of evolution by natural selection is firmly deductive (and its inductive aspect is generally implicit) (Popper 1979) whereas the emphasis of evolution by natural induction is explicitly inductive. Specifically, like many learning processes (including all neural network learning), natural induction uses induction to generate novel adaptive patterns, based on their underlying patterns of feature combinations (see Discussion in (Watson, Levin and Lewens 2025)). Whereas brains and machine learning systems are selected or designed for the purpose of doing inductive learning of this kind, natural induction occurs naturally, i.e. without selection or design, in dynamical systems with suitable properties.

The necessary conditions for natural induction are obviously not met by all dynamical systems, and natural induction therefore does not apply universally. However, these conditions are not onerous. Many types of dynamical systems are described by a network of interactions, many

interactions are not perfectly elastic, and the stress applied to a system is usually not constant. Thus, to the extent that these conditions apply, natural induction will occur without any contrivance, design or selection. Finally, whether the natural induction that occurs produces significant adaptation, or good solutions to a given problem, depends on the reliability of the associative induction principle, i.e. that combinations of features that were good together in past experience tend to be good together in the future (see Discussion in (Watson, Levin and Lewens 2025)). This is not always true, but it is often valuable. Both 'naïve proximity' and associative models work from a common principle of 'things that are similar to what worked in the past may work well in the future', but associative generalisation is a level-up from naïve proximity because it allows the proximity in model space to be different from the original feature space (see Discussion in (Watson, Levin and Lewens 2025)).

Biology is full of dynamical systems described by networks of interactions. These include, for example, chemical reaction networks, cytoskeleton networks, enzymatic networks, protein interaction networks, gene-regulation networks, metabolic networks, endocrine networks, morphogen networks, biomechanical networks, bioelectric networks, social interaction networks and ecological networks (not to mention neural networks). In many cases, the dynamics of these systems may be subject to shocks or perturbations, or external stresses that are variable – for example, changes in environment, changes in prey or predator population density, a repeated exposure to a toxin, or periodic exposure to some other kind of stressor or intermittent favourable condition. And in most cases, the interaction coefficients of these networks are not constant but change over time as a result of the system's experience/history (i.e. its behaviour in interaction with an environment). For example, the depletion of a shared resource (Gawne, McKenna and Levin 2020, Smiley and Levin 2022), and switching to alternatives, changes the interaction between ecological competitors; or social ties may weaken or break under stress; or chemical reaction rates may change as enzyme or reagent stocks change, or back-reactions become more significant. If so, then biology is potentially full of networks, at many different levels of organization, that provide adaptation by natural induction. Learning of this kind may occur in such networks without selection being involved, and thus they may act as a source of adaptive organization in their own right. This provides a potential mechanism to explain the widespread basal cognition that is observed in many different substrates independent of natural selection.

3. Some initial considerations

Having established that natural induction is a different mechanism of adaptation—capable of producing the appearance of design, and not the same as natural selection—we now want to discuss whether natural induction is relevant to biological evolution. Some possible objections to this idea are worth dealing with directly.

--

- A) "We already have a complete explanation for the complexity and diversity of life, so we do not need another".

Richard Dawkins once stated that "What Darwin achieved was nothing less than a complete explanation of the complexity and diversity of all life" (Dawkins 2009). If one really believed this, it might be tempting to disregard the possibility that natural selection and natural induction are both or jointly relevant to biological evolution on the grounds of parsimony. Of course, there are many open questions in evolutionary biology: the origin of life (Walker 2017), the pace and tempo of evolution (Gould and Eldredge 1977, Grunert, Holden et al. 2021), evolutionary transitions in individuality (Maynard Smith and Szathmary 1997, Okasha 2006, Szathmary 2015, West, Fisher et al. 2015, Watson,

Levin and Buckley 2022), and many others. One might retort that although such issues are not completely understood, there is nothing that threatens the completeness of natural selection to explain them all *in principle*. This is really nothing more than a hopeful gesture to the possibility of an explanation in terms of selection, and does nothing by itself to show that an alternative scientific explanation might not be required instead. What is more, one of these open questions is the difficulty in identifying a quantity that natural selection systematically increases (Grafen 2009, Birch 2014, Okasha and Paternotte 2014). Since selection sometimes leads to a reduction in complexity, and even a reduction in mean fitness, across a population (Bedau 2009, Grafen 2009, Okasha and Paternotte 2014), some appeal to difference-making factors, distinct from selection itself, is therefore necessary to explain when selection promotes adaptive complexity.

There are further reasons to doubt the all-sufficiency of selection. Natural selection explanations depend on the axiomatic mechanisms of variation, selection and inheritance. However, these mechanisms are not fixed over evolutionary time. Like any theory, attempts to use the theory of natural selection to explain changes in its own axioms run into problems of circularity. This is apparent in discussions around the evolution of mutation rates and sexual recombination (changing genetic variability) and evolvability (by changing the distribution of phenotypic variation), niche construction (modifying selection), and the evolution of extended inheritance mechanisms (e.g. epigenetics, niche inheritance and cultural inheritance) (Laland, Uller et al. 2015, Watson and Thies 2019).

More importantly, a simplistic rivalry between theories is not useful or necessary. The more interesting possibility is that the two mechanisms interact with each other, providing a deeper understanding of how biological evolution works than either mechanism does on its own. Since the conditions for natural induction are different from those for natural selection, induction may apply in cases where selection does not. Since their adaptive capabilities are different, induction may be able to explain adaptations that selection cannot. And since induction is not about differential survival or reproductive fitness, the adaptations it explains may be different in kind.

B) “Any mechanism that is not motivated by survival and reproduction, directly or indirectly, cannot be relevant to biological evolution because differential survival is the only intrinsic source of directionality”.

This sentiment is reinforced by the idea that the natural direction of change in physics is toward disorder (entropy increase) and natural selection is the only mechanism that goes in the other direction. This response simply begs the question against our position: we have made the case that natural induction provides a counter example to the assertion that natural selection is the only possible natural source of adaptive organisation (Dawkins 1983, Lewens and Watson 2024). There is of course no magic, or anything unphysical, in either natural induction or natural selection. Cycles of reproduction in parallel provide the opportunity for differential frequency change in natural selection. And cycles of repeated disturbance and relaxation in series provide the opportunity for transformational change in natural induction.

C) “We can measure the allele frequency change that occurs in natural populations, it agrees with the selection coefficients and variation in the population, so natural selection is a complete explanation for biological evolution. There is no deviation from predictions that needs further explanation and no room for any other mechanism to do anything”.

One way this objection might miss the point is a case where natural induction orchestrates the action of natural selection – such that natural selection is a sub-mechanism of evolution by natural induction.

A metaphor is useful. Suppose we were interested in how my computer does computation and we consider the theory that it works by electricity. We develop a theory of electricity that describes the relationship between voltage (potential difference), resistance and the flow of charge (current). Then we go poking electrodes into my laptop and we find whenever we measure resistance, voltage and current they agree perfectly with our theory. We even discover a source of the electricity and show that when it is removed or isolated, no computation happens, i.e. electricity is essential for computation. Moreover, in the computer there are no other moving parts besides current flow. So, it seems like a theory of electricity describes everything that happens, and in this sense, it is a complete explanation for computation!⁴ One might have to concede that computation is not apparent in every case where electricity flows, e.g. in a short circuit or in the domestic lighting circuits of my home, but nonetheless we might hold onto the idea that computation is still occurring in some primitive sense whenever charge flows. This enables us to maintain the idea that computation is 'just current flow (in an appropriate context)'.

And yet, although a theory of electricity is 'not wrong', it does not explain computation at all. It is the organisation of the causal flows that is important to computation, not the electricity. A theory of computation needs many levels of abstraction from electricity, to logic gates and circuits, to operating systems, to software programs, to algorithms, none of which is accounted for by a theory of electricity. These levels are necessary not merely to make it easier for us to explain or understand the system as post hoc observers, but also to collapse the effective degrees of freedom as necessary for any practical mode of intervention (whether by a 'programmer', variation and selection, machine learning, or by self-modifying code). Electricity is, at best, one level of sub-mechanism within computational systems. In fact, electricity is not even required in principle, since you could make a computer to run the same algorithms out of mechanical parts like cogs and levers, or hydraulics and valves. Indeed, from the point of view of computation, a theory of electricity is barely more than a theory that 'things go where they are pushed' (i.e. charge flows in proportion to the potential difference, allele frequencies change in proportion to selection coefficients). This is general (i.e. equally applicable to mechanical levers or fluid hydraulics) but does essentially no explanatory work in a theory of computation. This is not just because it is an inappropriate level of explanation for our limited cognitive capacities, but because it makes no reference to the causal structures essential to the algorithm. Notice that no amount of evidence of current flow (nor its agreement with our theory of electricity) is sufficient to prove that computation is the flow of charge or caused by the flow of charge.

Could the theory of evolution by natural selection be a sub-mechanism of evolutionary adaptation like the flow of electrical charge is a sub-mechanism of computation? It is true that allele frequencies change as a function of the differential selection coefficients acting on them; however, this does not necessarily explain the organisation of the causal flows that are important. For example, it does not distinguish when the action of natural selection does or does not increase complexity or adaptive organisation (Bedau 2009).

Here is a complementary way to put this point: Suppose we grant that there is nothing in biology that is not, at root, a consequence of interactions between fundamental physical particles. Despite this, evolutionary theory attributes evolutionary change not to physical laws but to the substrate-neutral algorithm of natural selection (without suggesting there is anything wrong with the laws of physics). It is possible that biology, like computation, has many different levels of causal flows (not just the levels of physics and natural selection), and their organised structure matters to the algorithm

⁴ To extend this analogy further, in the style of generalised Darwinism, perhaps we find evidence that this is not a special case but something much more general; We find that the predictions of the theory of electricity also holds for all measurements you can make inside a brain when it is doing cognition.

of biological evolution (like the organised structure of algorithms matters to computation and the direction of current flow) (Watson 2012).

- D) “Even if it were true that adaptative phenotypic changes might be acquired by an organism within-lifetime, if these changes are not inherited to the next generation then this does not matter to evolution. Conversely, if heritable differences do occur and are undirected (not the result of acquired adaptive change), then it is simply evolution by natural selection that determines evolutionary outcomes, as always”.

Both statements are incorrect because adaptative phenotypic changes may modify the direction of selection that applies to random genetic variation. In this case, acquired phenotypic changes matter to evolutionary outcomes, even if they are not (directly) inherited. And although it is true that undirected variation and natural selection occurs, this does not distinguish between one evolutionary outcome and another (short circuit vs computation, changes in gene frequencies vs evolution of adaptive complexity).

- E) “Natural selection, when understood in sufficiently general terms, is sufficient to explain all cases of adaptation, including those special cases that might be natural induction. In particular, natural selection can explain adaptation in systems (such as ecosystems) that do not reproduce. Thus, it may be true that learning networks do not reproduce, but it could still be the case that they adapt via natural selection.”

One might worry that our argument that natural induction is distinct from natural selection is overly reliant on a particularly narrow interpretation of natural selection and, in particular, on the observation that neither learning networks nor their parts reproduce. If one assumes (as some theorists do, e.g. (Godfrey-Smith 2009, Wilson 2016)) that natural selection requires reproduction, it follows that learning networks do not adapt by natural selection. However, this assumption about the conditions for natural selection could be too hasty: several prominent theorists have made a case that natural selection does not require reproduction. The grounds offered for this claim are varied: to give just a few examples, Van Valen (1989), followed by Bouchard (2008, 2011), argue for a form of selection that merely requires relative success of entities in terms of growth, not reproduction; Bourrat (2023) has pointed out that the Price Equation gives an account of selection compatible with relative success in growth of entities with respect to a subsequent time-step, again without requiring that those entities reproduce; Papale and Doolittle (2024), following and elaborating Hull (1980), have argued that natural selection at a level requires ‘differential extinction and proliferation of interactors’, regardless of whether those interactors reproduce. This allows for cases where the organisational relationship between the interactors and the replicators they may contain might not be simple. For example, in the case of a sexually reproducing organism (interactor), there may be many reproducing genes (replicators). Likewise, an ecological community (interactor) may be composed of many reproducing individuals from multiple different species. But this broader understanding of adaptation under natural selection still requires a population of interactors such that something differential may distinguish them – if not differential reproduction, then differential survival or persistence.

Some of us have endorsed similar arguments, especially in the context of cultural evolution (e.g. (Lewens 2023, Lewens 2024)). But we do not think that our reasons for viewing natural induction as a separate process to natural selection fall if one acknowledges that reproduction is not necessary for natural selection. Consider, to begin with, the stance of Papale and Doolittle (2024): because they follow Hull, they require the existence of replicators as a precondition for natural selection to act. If replicators are understood as entities of which distinct token copies are made (see (Godfrey-Smith

2000) for discussion), there are no units within a persisting learning network that are replicators. Perhaps one could recover a replicator interpretation by making the case that token patterns of activity within the network at an instant in time are responsible for the production of resembling patterns of activity at a later time step (Fernando, Goldstein and Szathmary 2010). While the replicator/interactor framework dispenses with the idea that the relevant interactors must reproduce, we think it has not been shown how the replicator/interactor framework could be applied to a learning network.

When theorists drop the requirement of differential reproduction for natural selection (in favour, for example, of differential persistence), the manner in which they invoke natural selection in the explanation of adaptation typically inherits significant features of more standard natural selection explanations. The most obvious of these is that, in the standard case, natural selection explains adaptation in a variational, rather than transformational, manner. To be more specific, natural selection explanations tend to appeal to the proliferation of suitable ‘platforms’ that can act as the bases for further potentially beneficial (but undirected) mutations (Neander 1995, Lewens 2004, Godfrey-Smith 2009). In Doolittle’s case (2014), for example, he stresses that if an entity is able to survive, rather than perishing, the added time in which it persists gives it further opportunities to change. This also means it is more likely to change in ways that display further adaptations, than if it had perished earlier: each time-step that an entity’s persistence affords it can be thought of as a further ‘platform’ in which additional positive change can occur. And one might add that in the case of growth, rather than mere survival, the number of sub-units within the growing entity that might acquire beneficial mutations also increases. In both cases, one retains the basic schema underpinning natural selection explanations of adaptation, showing how it applies even in cases where reproduction is absent.

We do not take issue here with these approaches to natural selection explanations; our reason for insisting on a distinction between natural induction and natural selection is that we want to highlight the differences between how adaptation is explained in the context of learning networks, and how it is explained even in these extended approaches to natural selection. True, a persisting network has more opportunity to change (both for better and worse) than a network that perishes. But in natural induction there is no differential comparison or competition (in persistence or reproduction) between one network and another – there is no population of networks. It is a purely transformational process applicable to a single network (no population of interactors). It is not the case that the network changes in several random ways, with the beneficial ones being retained and the detrimental ones being discarded. Rather, the incremental adjustments that the network undergoes are directed by the application of the stress. Focusing on natural selection in any form also obscures the crucial way in which the organisation of the connections within a learning network are associative. Because the changes to connections affect the interactions between the state variables, this enables the network as a whole to acquire a form of associative memory (capable of storing and recalling multiple patterns), perform associative generalisation (generating novel patterns from the same class), and problem-solve rapidly (improving the quality of solutions generated with experience). This kind of associative learning, and these adaptive capabilities, are not explained by merely increasing the number of platforms for variation.

Finally, even in the case where changes to interactions in the network are, as a matter of fact, provided by natural selection, this does not alter the fact that the adaptation at the network level occurs via this process of associative learning (i.e. by natural induction) and not by amplifying platforms of variation. The companion paper (Watson, Levin and Lewens 2025) discusses such cases in more detail.

In this pair of papers, we consider three different scenarios in which genetic evolution occurs by random genetic variation and natural selection. But in each case, we consider how natural induction might shape the selective field in which variants are selected and thereby lead genetic evolution to find adaptive organisations it would not otherwise discover. Natural selection thus provides a necessary sub-mechanism for genetic evolution to occur (like the electricity metaphor above), but in these cases, natural selection is not the mechanism of adaptation, i.e. it does not distinguish between an evolutionary outcome that is adaptive and one that is not. In the scenario discussed in this paper, phenotype-first evolution, this is because natural selection merely canalises adaptations already found by within lifetime adaptation. The companion paper (Watson, Levin and Lewens 2025) discusses two additional scenarios where this is because natural selection does not apply at that organisational scale (ecosystem adaptation) or timescale (evolution of evolvability), respectively.

4. Induction-first evolution

Here we consider one way in which natural induction and natural selection might interact; where, natural induction provides adaptive phenotypes within lifetimes and genetic evolution canalises them over subsequent generations.

Phenotype-first evolution describes a process of evolution where adaptive phenotypes are first discovered by some mechanism of adaptive phenotypic plasticity and then subsequent genetic evolution canalises (or assimilates/accommodates) this phenotype into heritable genetic changes (Waddington 1953, Waddington 1957, West-Eberhard 2003, West-Eberhard 2005, Lande 2009, Schwander and Leimar 2011). It is not doubted that many phenotypic adaptations occur within a lifetime, without genetic change, and that subsequent genetic evolution may canalise them. However, the suggestion that this constitutes a substantive change to evolutionary thinking is often dismissed with the assumption that, if such a mechanism of phenotypic plasticity is adaptive, this can only be because it is itself an adaptation provided by prior natural selection. But since natural induction is a different natural mechanism of adaptation, mechanisms of within-lifetime adaptation are not necessarily products of prior natural selection, and adaptive plasticity can be a genuinely different source of adaptive phenotypes.

A different issue is that adaptive phenotypic plasticity on its own does not constitute evolutionary change (because acquired phenotypes are not directly heritable, and no genetic change to the population has occurred), and if subsequent genetic change to a population depends on natural selection, based on random variation and selective retention, one might argue that all the evolutionary change, all the genetic change in the population, was actually provided by natural selection (objection D above). Whilst it might be correct that all the genetic change in the population was provided by natural selection, this does not mean that natural selection was the process that discovered the phenotypic adaptation that evolves (neither does it mean that the result of genetic evolution without prior phenotypic evolution would be to, sooner or later, find the same solution).

To understand the difference in the work that natural selection is doing with and without adaptive phenotypic plasticity, consider a thought experiment. Suppose you are learning a new yoga pose blindfolded. In one scenario (selection-first), you try adjusting your position by trial and error, with only verbal 'hotter/colder' feedback from the teacher. In a second scenario (phenotype-first), your teacher physically adjusts you into the correct position. In both cases you need an ability to

remember how to produce positions. In the selection-first case you need to do this differentially, i.e. remembering changes that afford improvement but not others (over many trials). In the phenotype-first case, trial and error is not required – you can simply remember the pose you are placed into. Clearly, it is much easier to remember the pose when you are put into it.

If an organism has no phenotypic plasticity, then it may adapt to a new environment only by genetic trial and error, and ‘hotter/colder’ feedback from selection (selection-first scenario). If, in contrast, an organism has adaptive phenotypic plasticity, then it may adapt to a new environment, in the first instance, without genetic change (phenotype-first scenario). Moreover, in this context, genetic change that canalises this new phenotype (without changing the phenotype) both constitutes genetic evolution and secures this adaptive phenotypic outcome, even though it did not discover this adaptive phenotype. In both cases, genetic variation and selection are required for genetic evolution to ‘remember and reproduce’ the phenotype (and it is assumed variation is random). In the second case, it is sufficient for genetic evolution to simply canalise the adaptive phenotype that has already been found by the adaptive plasticity. In the first case, on the other hand, canalisation of the current phenotype would not result in an adaptive phenotypic outcome – thus, in the absence of adaptive plasticity, differential retention/amplification of genetic changes that confer directly beneficial phenotypic change (from many samples) is required.

Clearly the yoga analogy presupposes a benevolent teacher, and the phenotype-first example presupposes a mechanism of adaptive phenotypic plasticity (independent of selection). Induction-first evolution suggests that natural induction provides this adaptive directionality for phenotypic plasticity (not prior selection). One potential concern, however, is that if natural induction is not a product of natural selection, how do we know that its adaptive direction is aligned with organismic fitness? This thinking is the wrong way around. Instead of imagining that natural induction must somehow anticipate how to lead evolutionary trajectories to confer a fitness advantage, we must instead accept that the idea of phenotype-first evolution really does mean that the phenotype is leading. That is, if phenotypic changes come first then, whatever they are, they dictate the direction of subsequent natural selection to a large extent. For example, if phenotypic changes cause gene-expression to take a particular novel pattern that is locally optimal in some respect, then genetic changes that disrupt this pattern are selected against and genetic changes that are compatible with it, or canalise it, are not (there may also be genetic changes that change the space of phenotypic possibilities, but this might be relatively rare). This does not mean that phenotypic changes anticipated the genetic changes that were necessary to produce an adaptive outcome. On the contrary, phenotypic changes determine which genetic changes confer adaptive outcomes. Consider, for example, that when human tracheal cells are placed in an unfamiliar context, more than 9000 genes change expression (Gumuskeya, Davey et al. 2024). The genetic mutations that are beneficial in this new context do not necessarily have any meaningful correspondence to those that were beneficial in the previous context. This helps us understand that phenotype-first evolution is not necessarily finding the same adaptations that would have eventually been found by selection-first genetic evolution in the absence of that phenotypic plasticity (nor adaptations that have already been explored in the previous history of genetic evolution). Indeed, if it did, it would be considerably less interesting inasmuch as it would make no difference to evolutionary outcomes in the long term (Watson, Levin and Buckley 2022, Watson and Levin 2023, Watson 2024).

To take a very simple example, suppose that the relationship between gene-expression patterns and organismic stress (e.g. metabolic inefficiency) in a given environment is described by a rugged potential surface with many local optima. Individual changes to gene expression, as produced by selection-first genetic evolution acting on individual mutations, will tend to find the nearest local

optimum in this surface. In contrast, if a within-lifetime adaptive process finds a different local optimum, now the selection acting on individual genetic mutations will follow a direct gradient toward this configuration. Note that in both cases, genetic change is following local gradients to reduce organismic stress in that environment, but the results are not the same because mutations are being selected in the context of different gene-expression backgrounds.

So, if the outcomes of genetic evolution with and without phenotypic plasticity are not necessarily the same, then in what direction is natural induction leading evolution? The direction of natural induction is to find configurations that better resolve the constraints, tensions or frustrations between a system's component parts that are created by its behaviour in interaction with its environment (Watson 2024, Watson, Levin and Lewens 2025). Over a distribution of past experiences (which may occur within lifetime), this is sufficient to provide significant problem-solving capability, i.e. finding configurations that resolve more constraints than those found by a naïve hill-climbing or a simple gradient ascent process (Buckley, Lewens et al. 2024). In other words, natural induction may be a leader of genetic evolution, and moreover, it is not unreasonable to imagine that it leads to better-adapted phenotypes than genetic evolution without natural induction.

The possibility that within-lifetime adaptation can guide genetic evolution to superior adaptations it would not otherwise find is also involved in the Baldwin effect (Baldwin 1896, Hinton and Nowlan 1987), first described by Simpson (Weber and Depew 2003). Both phenotype-first evolution and the Baldwin effect involve genetic assimilation of a phenotype or behaviour that is first found within lifetime. But one potential limitation of the Baldwin effect is that within-lifetime adaptation is provided by a learning mechanism which is assumed to involve neural cognition. This is a potential problem because genetic mutations can only canalise an acquired adaptation if they can incrementally reduce the amount of learning required, and this is not for granted when the learning space and the genetic space are very different (Mayley 1996). That is, it is not clear that genetic mutations having the effect of canalising learned solutions will exist if the learned solutions involve neural cognition. The emphasis of the induction-first scenario described above is slightly different. The idea is not that you can learn an adaptation cognitively and then reduce the amount of learning required through changes to some unrelated heritable substrate. The idea is that the within-lifetime adaptation 'shows natural selection how to do it' in quite a direct way. That is, the within-lifetime adaptation and the evolutionary adaptation happen in almost the same substrate. For example, if a mechanism of within-lifetime adaptation involved direct modifications to the gene-expression profile of an organism, then canalising genetic evolution can trivially retain the adaptive phenotype. In other words, in a classic scenario of the Baldwin effect, i.e. involving neural cognition, the assimilation process is highly inefficient. Since the change in behaviour at the organismic level is so abstract compared to the genetic changes available, genetic evolution needs to in effect 'rediscover' the learned solution in this new substrate using only the warmer/colder selective signal. In contrast, when the learned solution is effected via a learning gene-regulation network, for example, the change in selective feedback on genes is high-dimensional, 'showing' genetic evolution exactly what to do to recreate this same solution, making it much more efficient.

The timing of events is worth some discussion. If we assume that mutations occur, followed by within-lifetime phenotypic plasticity, and then selection applies, then plasticity appears to be unambiguously 'downstream' of genetic change – such that genetic evolution could be the cause of the adaptive plasticity but adaptive plasticity could not be the cause of genetic evolution. Put differently, by the time the plasticity has altered the phenotype, it seems too late to influence the mutations that might be inherited. Thus, absent a mechanism for direct inheritance of the acquired phenotype, the adapted phenotype is necessarily forgotten and only the original genotype (which it

did not influence) remains. However, although the acquired phenotype cannot alter the mutations a genotype carries, it can alter the selection that applies to them. In a population of plastic individuals, each of which acquires an adaptive phenotype, the mutations that survive to the next generation are those that are compatible with this adaptive phenotype and not those that disrupt it. Thus, even though the mutations an individual carries arose before that particular phenotype was acquired in that individual, this does not mean that the mutations caused the adaptive phenotype. In fact, when the adaptive phenotype arises prior to genetic change in the population (West-Eberhard 2003), the opposite is possible, i.e. that the adaptive phenotype caused the mutations *that persist in the population*, inasmuch as it applied the selection that determined which mutations were retained.

This scenario does have the limitation, however, that only aspects of the adaptive configuration that are repeatedly rediscovered in many individuals over many generations can be genetically canalised. Another possibility that alleviates this limitation is that an adaptive change in acquired gene expression results in epigenetic marks on the DNA that allow for the inheritance of this pattern over multiple generations (Jablonka and Lamb 1995, Jablonka and Lamb 1998, Jablonka and Lamb 2010, Jablonka 2013) and for it to thereby spread through the population. Then this epigenetic change may be canalised over further generations into more permanent changes in the DNA by altering selection, as before. This possibility, with the inclusion of epigenetic inheritance, has the advantage that it does not require an adaptive phenotype to be repeatedly rediscovered in multiple individuals or multiple generations.

Note that, in general, variations can be either deleterious, neutral or beneficial, and in a conventional account of genetic evolution, evolutionary adaptation requires that at least some mutations are beneficial. In contrast, in a phenotype-first evolution scenario, it is possible that adaptive genetic evolution occurs without beneficial mutations. That is, if neutral mutations canalise the current phenotype⁵ and fix occasionally, and deleterious mutations are discarded by selection (or fix less often), this is sufficient for adaptive evolution to occur if (and only if) the phenotype is already adapted by some other means. This might have some significance for 'constructive neutral evolution' (Stoltzfus 2012), neutral theory in ecology (Hubbell 2005) and neutral evolution more generally (Kimura 1979, Duret 2008).

Discussion and Conclusions

In this pair of papers, we limit our attention to cases where evolution is constituted by genetic change in a population, and all genetic change occurs through random genetic variation and natural selection. Nonetheless, we have discussed a scenario where natural selection is not the source of the adaptive organisation that evolves. In this case, natural induction is shaping the selective field in which natural selection occurs. In other words, natural induction tells natural selection what to do and leads it to find adaptive organisations it would not otherwise discover. We noted that electrical current flow provides a necessary sub-mechanism for computation to occur yet does not describe the mechanism of computation. Likewise, in these cases, natural selection provides a necessary sub-mechanism for genetic evolution to occur, yet it does not describe the mechanism of adaptation.

Because natural induction is a different mechanism of adaptation, independent of natural selection, it potentially provides a way for phenotype-first evolution to properly be a leader of evolutionary change, hence induction-first evolution. In this case, natural selection may act not as a

⁵ This assumes that neutral mutations (i.e. that do not disrupt the current phenotype) are also likely to increase the robustness of a phenotype to subsequent genetic mutations. This does not seem unreasonable given that the effects of environmental perturbations and genetic perturbations are likely to be correlated by virtue of a shared developmental architecture (Brun-Usan et al. 2020).

source of adaptations but as a memory of adaptations discovered by natural induction. If so, we must also consider the possibility that phenotype-first plasticity does not result in the same outcomes as natural selection alone (without phenotypic plasticity). One reaction to this is to assert that phenotypic plasticity can only be a leader of evolutionary adaptation when its direction of change is properly aligned with Darwinian fitness, and that this can only occur when natural selection has set the reward criteria of the learning process – hence putting natural selection back in the driving seat. We do not think this is correct. First, what we want to understand is how evolution occurs and its trends. If within-lifetime plasticity takes phenotypes in a particular direction, and this shapes the direction of subsequent natural selection, then that is the answer. Whilst a certain basic compatibility with Darwinian fitness is expected (e.g. not decreasing it), this still allows for phenotypic plasticity to change the availability of phenotypes, modify fitness payoffs, and introduce path dependencies of various kinds. Second, whether this is a satisfying answer depends on how well it explains the complexity and diversity of life that we observe. Arguably, since natural selection naturally finds local adaptive peaks (at best) and natural induction can find combinations of traits that better resolve constraints between component parts, the latter may explain complex adaptations more easily. Third, the differences between the adaptive behaviour of induction and selection are conceptually interesting. It is sometimes the case that natural selection increases Darwinian fitness by reducing organismal complexity and the conditions that result in complexity increase are poorly understood (Bedau 2009). In the second paper we discuss further what it is that natural induction is adapting *for* (Watson, Levin and Lewens 2025).

In the second paper we present two other ways in which natural induction and natural selection may interact. These are more subtle because natural selection is integral to the operation of natural induction in these cases and so it cannot be said that the adaptation arises without natural selection. Nonetheless, these adaptations cannot be attributed to natural selection because they occur at organisation scales or timescales where natural selection does not apply. Specifically, we discuss the evolution of ecosystem adaptations (where the ecosystem is not an evolutionary unit) and the evolution of evolvability (where characteristics of evolvability evolve before selection acts on phenotypes enabled by this evolvability). This second paper also provides more in depth discussion on a) how the scenarios of both papers expands our understanding of adaptive evolution and impacts a number of important issues in evolutionary biology (such as the origin of life and evolutionary transitions in individuality), b) some further discussion on what it is that natural induction is doing if it is not directly aligned with Darwinian fitness, and c) the underlying similarity and contrast between induction and selection and their relationship to different types of inference (induction and deduction, respectively).

New conceptual spaces and theory also imply new modes of intervention – new ways to communicate with living systems at different levels of organisation, e.g. for bioengineering and biomedicine. Instead of our only tools being bottom-up molecular engineering or top-down artificial selection, we might approach the management, prediction and control of biological systems on a spectrum of ‘persuasion’ (Levin 2021). In bio-inspired computation, the genetic algorithm, based on neo-Darwinian principles, has limited adaptive capabilities and is surpassed by many kinds of machine learning and artificial intelligence methods. Natural induction connects the algorithmic principles of evolution with machine learning, inspiring new computational methods (Caldwell, Watson et al. 2018, Caldwell, Knowles et al. 2021, Caldwell, Knowles et al. 2022).

If natural induction is operating in biological evolution, there are broad implications that warrant further investigation. Whilst it may remain true that ‘nothing in biology makes sense except in the light of evolution’, if natural selection is not the sole mechanism of evolutionary adaptation, a

lot of evolutionary biology may need a rethink (Laland, Uller et al. 2014). First, natural induction gives us new conceptual and theoretical foundations with which we might be able to make sense of many aspects of biology that are currently pre-theoretic. For example, we are able to import existing theory from learning systems into 'design principles' for biological organisation that go further than 'whatever survives and reproduces best' (Power, Watson et al. 2015, Kouvaris, Clune et al. 2017). Also, the partitioning of evolutionary change into that which is caused by selection and 'everything else' might be expanded to provide a theoretical basis for the non-uniform generation of variants. This potentially pulls our focus away from the 'blind' and 'dumb' mechanics of selection, and instead attends to how evolutionary processes can be informed by past experience. This offers a different understanding of creativity beyond the differential retention of 'accidents'. These are small steps in the right direction to rehabilitate the phenotype in evolutionary thinking (Marín and Wade 2025) and address the active role of phenotypes in guiding evolutionary processes (Laland, Uller et al. 2014, Laland, Uller et al. 2015, Uller, Moczek et al. 2018, Feiner, Feldman et al. 2024).

More conceptually, these shifts enable us to recognise that what is interesting about organisms (the features that Darwin wanted to explain, and Paley before him) is not only their survival or fecundity. It is also their design-like properties, the mutual information between the structure of the organism and environment (Kirchhoff, Parr et al. 2018, Parr, Pezzulo and Friston 2022) that draws our attention. Evolution by natural induction links conventional selection-based principles with notions of design and mutual information (Vanchurin, Wolf et al. 2021).

We conclude that evolution by natural induction is not an adaptive process derived from natural selection, nor a redescription of what we already knew about evolution by natural selection; it is a different algorithm with different adaptive capabilities. And by accounting for the interaction of natural selection and natural induction together, operating on different timescales, we can explain more than we could with natural selection alone. For example, instead of insisting that every kind of morphological and behavioural adaptation, including every mechanism of developmental plasticity is an adaptation honed by many previous generations of genetic variation and selection, we can see that the learning capabilities of life may solve these problems and then show genetic evolution the way.

Acknowledgements

We thank Chrisantha Fernando for the original insight that Hebbian principles can apply in non-neural substrates, Denis Noble, Jim Shapiro and Eva Jablonka for insights into 'top-down' pathways, phenotypic adaptation and non-genetic inheritance, and Kevin Lala, Tobias Uller and Adi Livnat for discussion that invites expansion of evolutionary thinking. Thank you to Sylvain Charlat for many long discussions on the conceptual and philosophical distinctions involved in this work. Also, Frederick Nash, Kostas Kouvaris, Louis Kounios, David Prosser, Daniel Power, Miguel Brun-Usan, Jamie Caldwell, Chris Buckley, Adam Davies, Rob Mills, Christoph Thies, Taztio Tissot and Nam Le for the many computational models and conversations that co-constructed this larger conceptual vision before 'evolution by natural induction' had a name.

References

- Ashby, W. R. (1952). Design for a brain: The origin of adaptive behaviour, Springer Science & Business Media.
- Baldwin, J. M. (1896). "A new factor in evolution." Adaptive individuals in evolving populations: Models and algorithms: 59-80.
- Baluska, F. and M. Levin (2016). "On Having No Head: Cognition throughout Biological Systems." Front Psychol 7: 902.

- Bedau, M. A. (2009). The evolution of complexity. Mapping the future of biology: Evolving concepts and theories, Springer: 111-130.
- Birch, J. (2014). "Has Grafen formalized Darwin? Commentary on Grafen's 'The Formal Darwinism project in outline'." Biology & Philosophy **29**: 175-180.
- Biswas, S., S. Manicka, E. Hoel and M. Levin (2021). "Gene regulatory networks exhibit several kinds of memory: quantification of memory in biological and random transcriptional networks." iScience **24**(3): 102131.
- Bouchard, F. (2008). "Causal processes, fitness, and the differential persistence of lineages." Philosophy of science **75**(5): 560-570.
- Bouchard, F. (2011). "Darwinism without populations: a more inclusive understanding of the "Survival of the Fittest"." Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences **42**(1): 106-114.
- Bourrat, P. (2023). "A Pricean formalization of Gaia." Philosophy of Science **90**(3): 704-720.
- Brun-Usan, M., A. Rago, C. Thies, T. Uller and R. A. Watson (2020). "Developmental models reveal the role of phenotypic plasticity in explaining genetic evolvability." bioRxiv.
- Buckley, C. L., T. Lewens, M. Levin, B. Millidge, A. Tschantz and R. A. Watson (2024). "Natural Induction: Spontaneous adaptive organisation without natural selection." Entropy **26**(9).
- Caldwell, J., J. Knowles, C. Thies, F. Kubacki and R. Watson (2021). Deep Optimisation: Multi-scale Evolution by Inducing and Searching in Deep Representations. International Conference on the Applications of Evolutionary Computation (Part of EvoStar), Springer, Cham.
- Caldwell, J., J. Knowles, C. Thies, F. Kubacki and R. Watson (2022). "Deep Optimisation: Transitioning the Scale of Evolutionary Search by Inducing and Searching in Deep Representations." SN Computer Science **3**(3): 253.
- Caldwell, J., R. A. Watson, C. Thies and J. D. Knowles (2018). "Deep optimisation: Solving combinatorial optimisation problems using deep neural networks." arXiv preprint arXiv:1811.00784.
- Chvykov, P., T. A. Berrueta, A. Vardhan, W. Savoie, A. Samland, T. D. Murphey, K. Wiesenfeld, D. I. Goldman and J. L. England (2021). "Low rattling: A predictive principle for self-organization in active collectives." Science **371**(6524): 90-95.
- Cragg, B. G. and H. N. Temperley (1955). "Memory: the analogy with ferromagnetic hysteresis." Brain **78**(2): 304-316.
- Dawkins, R. (1982). The extended phenotype: The long reach of the gene, Oxford University Press.
- Dawkins, R. (1983). "Universal darwinism." Evolution from molecules to men: 403-425.
- Dawkins, R. (2009). The Genius of Charles Darwin, Channel 4: 60 mins.
- Dennett, D. C. (1996). "Darwin's dangerous idea." Recherche **27**(293): 100.
- Doolittle, F. W. (2014). "Natural selection through survival alone, and the possibility of Gaia." Biology & Philosophy **29**(3): 415-423.
- Dudte, L. H., E. Vouga, T. Tachi and L. Mahadevan (2016). "Programming curvature using origami tessellations." Nature materials **15**(5): 583-588.
- Duret, L. (2008). "Neutral theory: the null hypothesis of molecular evolution." Nature education **1**(1): 803-806.
- England, J. L. (2015). "Dissipative adaptation in driven self-assembly." Nature nanotechnology **10**(11): 919-923.

- Farmer, J. D. (1990). "A Rosetta stone for connectionism." Physica D: Nonlinear Phenomena **42**(1-3): 153-187.
- Feiner, N., M. Feldman, S. F. Gilbert, K. N. Lala and T. Uller (2024). Evolution evolving: The developmental origins of adaptation and biodiversity, Princeton University Press.
- Fernando, C., R. Goldstein and E. Szathmary (2010). "The neuronal replicator hypothesis." Neural computation **22**(11): 2809-2857.
- Fields, C. and M. Levin (2022). "Competency in Navigating Arbitrary Spaces: Intelligence as an Invariant for Analyzing Cognition in Diverse Embodiments."
- Fontana, W. (1990). Algorithmic chemistry, Los Alamos National Lab., NM (USA).
- Fontana, W. (2018). Functional self-organization in complex systems. Pattern Formation In The Physical And Biological Sciences, CRC Press: 43-63.
- Gawne, R., K. Z. McKenna and M. Levin (2020). "Competitive and Coordinative Interactions between Body Parts Produce Adaptive Developmental Outcomes." Bioessays **42**(8): e1900245.
- Gershman, S. J., E. J. Horvitz and J. B. Tenenbaum (2015). "Computational rationality: A converging paradigm for intelligence in brains, minds, and machines." Science **349**(6245): 273-278.
- Godfrey-Smith, P. (2000). "The replicator in retrospect." Biology and Philosophy **15**(3): 403-423.
- Godfrey-Smith, P. (2009). Darwinian populations and natural selection, Oxford University Press.
- Gould, S. J. and N. Eldredge (1977). "Punctuated equilibria: the tempo and mode of evolution reconsidered." Paleobiology **3**(2): 115-151.
- Grafen, A. (2009). "Formalizing Darwinism and inclusive fitness theory." Philosophical Transactions of the Royal Society B: Biological Sciences **364**(1533): 3135-3141.
- Gross, T. and H. Sayama (2009). Adaptive Networks: Theory, Models and Applications, Springer Publishing Company, Incorporated.
- Grunert, K., H. Holden, E. R. Jakobsen and N. C. Stenseth (2021). "Evolutionarily stable strategies in stable and periodically fluctuating populations: The Rosenzweig–MacArthur predator–prey model." Proceedings of the National Academy of Sciences **118**(4): e2017463118.
- Gumuskaya, G., N. Davey, P. Srivastava, A. Bender, L. Pio-Lopez, D. Hazel and M. Levin (2024). "The morphological, behavioral, and transcriptomic life cycle of Anthrobots." Advanced Science: 2409330.
- Gunawardena, J. (2022). "Learning outside the brain: Integrating cognitive science and systems biology." Proceedings of the IEEE **110**(5): 590-612.
- Hinton, G. E. and S. J. Nowlan (1987). "How learning can guide evolution." Complex systems **1**(3): 495-502.
- Hopfield, J. J. (1982). "Neural networks and physical systems with emergent collective computational abilities." Proceedings of the national academy of sciences **79**(8): 2554-2558.
- Hopfield, J. J. (1984). "Neurons with graded response have collective computational properties like those of two-state neurons." Proceedings of the national academy of sciences **81**(10): 3088-3092.
- Hopfield, J. J. and D. W. Tank (1986). "Computing with neural circuits: a model." Science **233**(4764): 625-633.
- Hubbell, S. P. (2005). "Neutral theory in community ecology and the hypothesis of functional equivalence." Functional ecology **19**(1): 166-172.

- Hull, D. L. (1980). "Individuality and selection." Annual review of ecology and systematics **11**: 311-332.
- Jablonka, E. (2013). "Epigenetic inheritance and plasticity: the responsive germline." Progress in biophysics and molecular biology **111**(2-3): 99-107.
- Jablonka, E. and M. J. Lamb (1995). Epigenetic inheritance and evolution: the Lamarckian dimension, Oxford University Press, USA.
- Jablonka, E. and M. J. Lamb (1998). "Epigenetic inheritance in evolution." Journal of evolutionary biology **11**(2): 159-183.
- Jablonka, E. and M. J. Lamb (2010). "Transgenerational epigenetic inheritance."
- Katz, Y. and W. Fontana (2022). "Probabilistic Inference with Polymerizing Biochemical Circuits." Entropy (Basel) **24**(5).
- Katz, Y., M. Springer and W. Fontana (2018) "Embodying probabilistic inference in biochemical circuits." arXiv:1806.10161.
- Kaygisiz, K. and R. V. Ulijn (2024). "Can Molecular Systems Learn?" ChemSystemsChem **n/a**(n/a): e202400075.
- Keim, N. C., J. D. Paulsen, Z. Zeravcic, S. Sastry and S. R. Nagel (2019). "Memory formation in matter." Reviews of Modern Physics **91**(3): 035002.
- Kimura, M. (1979). "The neutral theory of molecular evolution." Scientific American **241**(5): 98-129.
- Kirchhoff, M., T. Parr, E. Palacios, K. Friston and J. Kiverstein (2018). "The Markov blankets of life: autonomy, active inference and the free energy principle." Journal of The royal society interface **15**(138): 20170792.
- Kounios, L., J. Clune, K. Kouvaris, G. P. Wagner, M. Pavlicev, D. M. Weinreich and R. A. Watson (2016). "Resolving the paradox of evolvability with learning theory: How evolution learns to improve evolvability on rugged fitness landscapes." arXiv preprint arXiv:1612.05955.
- Kouvaris, K., J. Clune, L. Kounios, M. Brede and R. A. Watson (2017). "How evolution learns to generalise: Using the principles of learning theory to understand the evolution of developmental organisation." PLoS computational biology **13**(4): e1005358.
- Laland, K., T. Uller, M. Feldman, K. Sterelny, G. B. Müller, A. Moczek, E. Jablonka, J. Odling-Smee, G. A. Wray and H. E. Hoekstra (2014). "Does evolutionary theory need a rethink?" Nature **514**(7521): 161-164.
- Laland, K. N., T. Uller, M. W. Feldman, K. Sterelny, G. B. Müller, A. Moczek, E. Jablonka and J. Odling-Smee (2015). "The extended evolutionary synthesis: its structure, assumptions and predictions." Proceedings of the Royal Society B: Biological Sciences **282**(1813): 20151019.
- Lande, R. (2009). "Adaptation to an extraordinary environment by evolution of phenotypic plasticity and genetic assimilation." Journal of evolutionary biology **22**(7): 1435-1446.
- Levin, M. (2019). "The Computational Boundary of a "Self": Developmental Bioelectricity Drives Multicellularity and Scale-Free Cognition." Front Psychol **10**: 2688.
- Levin, M. (2021). "Technological Approach to Mind Everywhere (TAME): an experimentally-grounded framework for understanding diverse bodies and minds."
- Levin, M. (2023). "Darwin's agential materials: evolutionary implications of multiscale competency in developmental biology." Cellular and Molecular Life Sciences **80**(6): 142.
- Levin, M. and D. C. Dennett (2020). "Cognition all the way down." Aeon Essays. Retrieved.
- Levin, M., F. Keijzer, P. Lyon and D. Arendt (2021). "Uncovering cognitive similarities and differences, conservation and innovation." Philos Trans R Soc Lond B Biol Sci **376**(1821): 20200458.

- Lewens, T. (2004). "Organisms and artifacts: Design in nature and elsewhere."
- Lewens, T. (2023). Equations at an exhibition: on the cultural Price equation. Evolutionary thinking across disciplines: problems and perspectives in generalized Darwinism, Springer: 353-373.
- Lewens, T. (2024). Cultural Selection, Cambridge University Press.
- Lewens, T. and R. A. Watson (2024). Universal Darwinism Revisited. (under review).
- Lewontin, R. (1980). Models of natural selection. Vito Volterra Symposium on Mathematical Models in Biology: Proceedings of a Conference Held at the Centro Linceo Interdisciplinare, Accademia Nazionale dei Lincei, Rome December 17–21, 1979, Springer.
- Lewontin, R. C. (1970). "The units of selection." Annual review of ecology and systematics **1**(1): 1-18.
- Livnat, A. (2013). "Interaction-based evolution: how natural selection and nonrandom mutation work together." Biology direct **8**: 1-53.
- Livnat, A. and A. C. Love (2024). "Mutation and evolution: Conceptual possibilities." BioEssays **46**(2): 2300025.
- Lyon, P., F. Keijzer, D. Arendt and M. Levin (2021). "Basal cognition: multicellularity, neurons and the cognitive lens." Philosophical Transactions of the Royal Society B: Biological Sciences **376**(1821).
- Lyon, P., F. Keijzer, D. Arendt and M. Levin (2021). "Reframing cognition: getting down to biological basics." Philos Trans R Soc Lond B Biol Sci **376**(1820): 20190750.
- Manicka, S. and M. Levin (2019). "The Cognitive Lens: a primer on conceptual tools for analysing information processing in developmental and regenerative morphogenesis." Philos Trans R Soc Lond B Biol Sci **374**(1774): 20180369.
- Manicka, S. and M. Levin (2019). "Modeling somatic computation with non-neural bioelectric networks." Sci Rep **9**(1): 18612.
- Margulis, L. and R. Fester (1991). Symbiosis as a source of evolutionary innovation: speciation and morphogenesis, Mit Press.
- Marín, C. and M. J. Wade (2025). "Bring back the phenotype." New Phytologist.
- Mayley, G. (1996). "Landscapes, learning costs, and genetic assimilation." Evolutionary Computation **4**(3): 213-234.
- Maynard Smith, J. and E. Szathmary (1997). The major transitions in evolution, Oxford University Press.
- Melamed, D., Y. Nov, A. Malik, M. B. Yakass, E. Bolotin, R. Shemer, E. K. Hiadzi, K. L. Skorecki and A. Livnat (2022). "De novo mutation rates at the single-mutation resolution in a human HBB gene region associated with adaptation and genetic disease." Genome Research **32**(3): 488-498.
- Mitchell, K. J. (2023). "Free agents: how evolution gave us free will."
- Mukherji, S., N. Kandula, A. Sood and R. Ganapathy (2019). "Strength of mechanical memories is maximal at the yield point of a soft glass." Physical review letters **122**(15): 158001.
- Nash, F. K., L.; Thies, C.; Kouvaris, K.; Tarapore, D.; Watson, R. (2021). "Scaling-up variation and evolvability: The causes and consequences of developmental hierarchy." BioArxiv.
- Neander, K. (1995). "Pruning the tree of life." The British Journal for the Philosophy of Science **46**(1): 59-80.
- Okasha, S. (2006). Evolution and the levels of selection, Oxford University Press.
- Okasha, S. and C. Paternotte (2014). "Adaptation, fitness and the selection-optimality links." Biology & Philosophy **29**: 225-232.

- Papale, F. and W. F. Doolittle (2024). "Towards a more general theory of evolution by natural selection: a manifesto." Philosophy, Theory, and Practice in Biology **16**(1).
- Parr, T., G. Pezzulo and K. J. Friston (2022). Active inference: the free energy principle in mind, brain, and behavior, MIT Press.
- Popper, K. R. (1979). Objective knowledge: An evolutionary approach, Clarendon press Oxford.
- Power, D. A., R. A. Watson, E. Szathmary, R. Mills, S. T. Powers, C. P. Doncaster and B. Czapp (2015). "What can ecosystems learn? Expanding evolutionary ecology with learning theory." Biol Direct **10**: 69.
- Pross, A. (2011). "Toward a general theory of evolution: Extending Darwinian theory to inanimate matter." Journal of Systems Chemistry **2**: 1-14.
- Pross, A. (2016). What is life?: How chemistry becomes biology, Oxford University Press.
- Schwander, T. and O. Leimar (2011). "Genes as leaders and followers in evolution." Trends in Ecology & Evolution **26**(3): 143-151.
- Smiley, P. and M. Levin (2022). "Competition for finite resources as coordination mechanism for morphogenesis: An evolutionary algorithm study of digital embryogeny." Biosystems **221**: 104762.
- Stern, M., C. Arinze, L. Perez, S. E. Palmer and A. Murugan (2020). "Supervised learning through physical changes in a mechanical system." Proceedings of the National Academy of Sciences **117**(26): 14843-14850.
- Stern, M., D. Hexner, J. W. Rocks and A. J. Liu (2021). "Supervised learning in physical networks: From machine learning to learning machines." Physical Review X **11**(2): 021045.
- Stern, M. and A. Murugan (2022). "Learning without neurons in physical systems." arXiv preprint arXiv:2206.05831.
- Stern, M., M. B. Pinson and A. Murugan (2020). "Continual learning of multiple memories in mechanical networks." Physical Review X **10**(3): 031044.
- Stoltzfus, A. (2012). "Constructive neutral evolution: exploring evolutionary theory's curious disconnect." Biol Direct **7**: 35.
- Szathmary, E. (2015). "Toward major evolutionary transitions theory 2.0." Proc Natl Acad Sci U S A **112**(33): 10104-10111.
- Tank, D. W. and J. J. Hopfield (1987). "Collective computation in neuronlike circuits." Sci Am **257**(6): 104-114.
- Uller, T., A. P. Moczek, R. A. Watson, P. M. Brakefield and K. N. Laland (2018). "Developmental bias and evolution: A regulatory network perspective." Genetics **209**(4): 949-966.
- Vallverdu, J., O. Castro, R. Mayne, M. Talanov, M. Levin, F. Baluska, Y. Gunji, A. Dussutour, H. Zenil and A. Adamatzky (2018). "Slime mould: The fundamental mechanisms of biological cognition." Biosystems **165**: 57-70.
- Van Valen, L. M. (1989). "Three paradigms of evolution." Evolutionary theory **9**(2): 1-17.
- Vanchurin, V., Y. I. Wolf, M. I. Katsnelson and E. V. Koonin (2021). "Towards a Theory of Evolution as Multilevel Learning." arXiv preprint arXiv:2110.14602.
- Varela, F. J. and P. Bourgine (1992). Introduction: Toward a practice of autonomous systems: Proceedings of the First European Conference on Artificial Life, MIT press.
- Waddington, C. (1957). The Strategy of the Genes, George Allen & Unwin.

- Waddington, C. H. (1953). "Genetic assimilation of an acquired character." Evolution: 118-126.
- Walker, S. I. (2017). "Origins of life: a problem for physics, a key issues review." Reports on Progress in Physics **80**(9): 092601.
- Watson, R. (2024). "Agency, Goal-Directed Behavior, and Part-Whole Relationships in Biological Systems." Biological Theory **19**(1): 22-36.
- Watson, R., C. L. Buckley, R. Mills and A. Davies (2010). "Associative memory in gene regulation networks."
- Watson, R. and M. Levin (2023). "The collective intelligence of evolution and development." Collective Intelligence **2**(2): 26339137231168355.
- Watson, R. A. (2012). Is evolution by natural selection the algorithm of biological evolution? ALIFE 2012: The Thirteenth International Conference on the Synthesis and Simulation of Living Systems, MIT Press.
- Watson, R. A. (2023). Songs of Life and Mind. YouTube: <https://www.youtube.com/playlist?list=PLVmJximp0I4OJdT9bsFlebu0HjPAjtIEN>.
- Watson, R. A., C. Buckley and R. Mills (2011). "Optimization in "self-modeling" complex adaptive systems." Complexity **16**(5): 17-26.
- Watson, R. A., M. Levin and C. L. Buckley (2022). "Design for an individual: connectionist approaches to the evolutionary transitions in individuality." Frontiers in Ecology and Evolution: 64.
- Watson, R. A., M. Levin and T. Lewens (2025). "Evolution by Natural Induction II: Further interactions with natural selection." under submission.
- Watson, R. A., R. Mills and C. L. Buckley (2011). "Global adaptation in networks of selfish components: emergent associative memory at the system scale." Artif Life **17**(3): 147-166.
- Watson, R. A. and E. Szathmary (2016). "How can evolution learn?" Trends in Ecology and Evolution(31): 147-157.
- Watson, R. A. and C. Thies (2019). "Are Developmental Plasticity, Niche Construction, and Extended Inheritance Necessary for Evolution by Natural Selection? The Role of Active Phenotypes in the Minimal Criteria for Darwinian Individuality." Evolutionary Causation: Biological and Philosophical Reflections: 197-226.
- Watson, R. A., G. P. Wagner, M. Pavlicev, D. M. Weinreich and R. Mills (2014). "THE EVOLUTION OF PHENOTYPIC CORRELATIONS AND 'DEVELOPMENTAL MEMORY'." Evolution **68**(4): 1124-1138.
- Weber, B. H. and D. J. Depew (2003). Evolution and learning: The Baldwin effect reconsidered, Mit Press.
- Weismann, A. (1893). "THE ALL-SUFFICIENCY OF NATURAL SELECTION." The Contemporary review, 1866-1900 **64**: 596-610.
- West-Eberhard, M. J. (2003). Developmental plasticity and evolution, Oxford University Press.
- West-Eberhard, M. J. (2005). "Developmental plasticity and the origin of species differences." Proceedings of the National Academy of Sciences **102**(suppl_1): 6543-6549.
- West, S. A., R. M. Fisher, A. Gardner and E. T. Kiers (2015). "Major evolutionary transitions in individuality." Proc Natl Acad Sci U S A **112**(33): 10112-10119.
- Wilson, D. S. (2016). Two Meanings of Complex Adaptive Systems. Complexity and Evolution: Toward a New Synthesis for Economics. D. S. Wilson and A. Kirman, The MIT Press: 0.

Wong, M. L., C. E. Cleland, D. Arend Jr, S. Bartlett, H. J. Cleaves, H. Demarest, A. Prabhu, J. I. Lunine and R. M. Hazen (2023). "On the roles of function and selection in evolving systems." Proceedings of the National Academy of Sciences **120**(43): e2310223120.

Zhong, W., D. J. Schwab and A. Murugan (2017). "Associative pattern recognition through macromolecular self-assembly." Journal of Statistical Physics **167**: 806-826.

Evolution by Natural Induction II: Further interactions with natural selection

Richard Watson^{1*}, Michael Levin² and Tim Lewens³

¹ Dept. Computer Science/ Institute for Life Sciences, University of Southampton, U.K.. SO17 1BJ
(*corresponding author: R.A.Watson@soton.ac.uk)

² Allen Discovery Centre, Tufts University, Boston, MA, USA. (Michael.Levin@tufts.edu)

³ Dept. of History and Philosophy of Science, Cambridge University, U.K. (tml1000@cam.ac.uk)

Natural induction is a different adaptive mechanism from natural selection. It occurs in dynamical systems described by a network of interactions, where connections give-way slightly under stress and the system is subject to occasional perturbations. Various biological systems (from gene-regulation networks, to metabolic networks, to ecosystems) meet these basic conditions and therefore have potential to exhibit adaptation by natural induction. In this pair of papers, we consider various ways that natural induction and natural selection might interact. In a scenario of induction-first evolution, discussed in the first paper, a phenotypic adaptation is clearly not a product of natural selection because it appears before any genetic evolution has taken place. This paper discusses more subtle cases, where an interaction of both natural selection and natural induction is involved in producing adaptations. Nonetheless, natural selection alone does not explain these adaptations because they occur at scales where natural selection does not apply. Specifically, in the case of ecosystem adaptation, natural selection does not apply at that organisation scale. And in the evolution of evolvability, natural selection does not apply at the relevant temporal scale. We conclude that evolution by natural induction is a viable process that expands our understanding of evolutionary adaptation.

Introduction

This paper is the second of a pair discussing the interaction of natural induction (Buckley, Lewens et al. 2024) and natural selection in biological evolution. The first paper (Watson, Levin and Lewens 2025) describes the mechanism and principle of natural induction and how it differs from natural selection. It then discusses how natural induction and natural selection may interact together to provide a scenario of induction-first evolution. Specifically, natural induction provides the mechanism of rapid within-lifetime phenotypic adaptation, and then relatively slow genetic evolution over many generations may canalise or assimilate this phenotype into genotypic changes (West-Eberhard 2003, Schwander and Leimar 2011). In this kind of interaction, it is clear that natural selection cannot be responsible for the adaptive phenotype if it arises before any genetic change or selection has taken place.

In this second paper, we briefly recap some of the general principles and broader conditions of natural induction. Then we discuss some other ways in which natural induction and natural selection might interact in evolutionary processes. These interactions are more conceptually tricky to categorise than the induction-first evolution scenario because the adaptations of interest here do not occur without the involvement of natural selection. This makes it tempting to retreat to a conventional assumption that natural selection is, after all, the adaptive mechanism that drives the process. However, we do not think that this interpretation is correct. This can be seen because the relevant adaptations occur at scales where

natural selection does not apply. We discuss two specific scenarios. The first concerns the interaction of natural selection and natural induction in ecosystem adaptation. Here the interaction produces adaptation at the ecosystem level even though natural selection does not apply at this organisational scale. The second concerns the interaction of natural selection and natural induction in the evolution of developmental architecture. Although the relevant adaptation in this case belongs to an organism, which is a bone fide evolutionary unit, natural selection does not apply at the relevant timescale for the evolution of evolvability observed.

In this paper we also provide a more in-depth discussion of some of the implications. In particular, we discuss a) how the scenarios of both papers expand our understanding of adaptive evolution and impact a number of important issues in evolutionary biology (such as the origin of life and evolutionary transitions in individuality), b) some further discussion on what it is that natural induction is doing if it is not directly aligned with Darwinian fitness, and c) the underlying similarity and contrast between natural induction and natural selection and their relationship to different types of inference (induction and deduction, respectively).

1. Natural induction and its conditions

As described in the first paper of this pair (Watson, Levin and Lewens 2025), adaptation by natural induction is a mechanism of adaptation that does not depend on natural selection. Buckley et al. (2024) provide a demonstration of natural induction in a network of randomly connected springs. These springs are not perfectly elastic but slightly *viscoelastic*, meaning that they give way slightly under stress. The system behaviour, in interaction with an environment (or a different set of springs representing a set of problem constraints), naturally finds local equilibria that correspond to locally optimal solutions to the problem variables. If the state of the system is occasionally perturbed it visits a distribution of such solutions over time. In general, however, these configurations are not particularly special in their ability to resolve the constraints between problem variables – they are merely local optima. However, as the lengths of the springs give way slightly over time, in response to the frustrations that the system behaviour cannot naturally resolve, this changes the organisation of the springs (lengthening/shortening some and leaving others unchanged). The new internal organisation changes the future dynamics of the system and the way it behaves in interaction with its environment (or the problem constraints). In particular, it causes the system to develop an associative memory of configurations it has previously visited, it exhibits learning with generalisation, and over time, it improves in its ability to find configurations that resolve the original problem constraints exceptionally well (Buckley, Lewens et al. 2024).

Since this process demonstrably improves a system's ability to problem-solve, and natural selection is not involved (the network as a whole is not reproducing, nor are any of its component parts, nor is there any differential survival of interactors), we conclude that natural selection is not the only possible source of adaptive organisation that can arise naturally, i.e. without design (Lewens and Watson 2024).

In the first paper, we explain how to separate a notion of adaptation from a notion of adaptation by natural selection, and the various ways that this separation might be objected to. In particular, we discuss the objections that 'adaptation is defined as the product of natural selection', 'if natural selection is required for adaptive change then clearly it is responsible for the adaptation' and 'change in survivability or fecundity is the only sense of directionality that is biologically meaningful'. We explain that, although a

narrow sense of adaptation can be useful in some contexts, the explanandum of biology (e.g. the observation of exquisite design) is conceptually separable from any putative explanans (e.g. the mechanism of differential survival or reproduction). This separation is essential if we wish to assess the supposed monopoly of natural selection as an explanation for biological adaptation. Second, we explain that a process can contain an essential sub-process (e.g. cognition in the brain and computation in a computer can both require the flow of electric charge) without this meaning that the process is nothing more than the sub-process (e.g. that cognition is nothing more than the flow of electric charge). Likewise, even if natural selection is essential for biological evolution, it does not follow that biological evolution is nothing more than natural selection. Third, we propose that an organization that resolves, to some exceptional extent, a difficult set of simultaneous constraints (i.e. optimization of a state configuration with respect to a constraint problem) provides a meaningful and quantifiable sense of adaptation that may stand in for the appearance of design. Whilst such organizations may be useful for survival or reproduction, they are not, in themselves, a statement about survival or reproduction. This provides a directionality for biological adaptation that is separate from changes in reproductive fitness. In particular, suppose a set of constraints is highly frustrated (meaning that it is difficult to resolve one constraint without violating another), and whose resolution constitutes a difficult optimization problem. We can then take the process of falling to a local attractor (resolving only some constraints) to be the definition of an unexceptional resolution. Instead, we require optimization capability superior to this as a higher standard for bone fide adaptation.

The conditions for natural induction to apply are not onerous: A dynamical system described by a network of viscoelastic interactions, and occasional perturbations. These conditions are clearly different from the conditions standardly required for natural selection to occur (e.g. heritable variation in reproductive success). This means that natural induction may apply in cases where natural selection does not (and vice versa). Biological systems of various kinds are dynamical systems that can be described by networks of interactions, these interactions may not be perfectly elastic but likely to give-way under stress, and the stress or forcing applied to the system is often varying. Accordingly, biological systems at multiple levels of organisation have a potential to exhibit adaptation by natural induction. The question addressed in this pair of papers is: How does this interact with evolution by natural selection?

Because they are different mechanisms, natural induction and natural selection have different adaptive capabilities. In particular, whereas natural selection acting on independently varying features will naturally arrive at the nearest local optimum in an adaptive landscape, natural induction has the potential to learn to do better than this with experience – finding configurations that are exceptionally good at resolving problem constraints (Buckley, Lewens et al. 2024).

The general underlying principles of natural induction include a separation of timescales, whereby state variables (e.g. positions of masses) change quickly, and interaction variables (e.g. spring lengths) change relatively slowly. The system's behaviour can therefore be described as two kinds of local energy minimisation – fast 'first-order' energy minimisation that relaxes the position of state variables, and slow 'second-order' energy minimisation that relaxes the connections or interactions that are the effective parameters of the state variable dynamics. The feedback between the two creates a system that learns to relax better with experience, causing it to improve its ability to resolve problem constraints over time. This is related to the ability of complex dynamical systems to self-organise in a manner that dissipates work (from a driving source) more easily (England 2015, Chvykov, Berrueta et al. 2021). A significant difference is that natural induction exhibits a capacity for adaptation, not just organisation, inasmuch as

it improves its ability to find solutions to a difficult constraint optimisation problem with experience (Buckley, Lewens et al. 2024).

In the scenario of induction-first evolution discussed in the first paper, we noted that natural induction does not necessarily produce the same adaptive phenotypes as natural selection. Natural induction really is the leader of the process, shaping the natural selection that follows to whatever it is that natural induction has learned. This implies a separation between the optimisation of Darwinian fitness and whatever natural induction does; a separation we explore further in the following examples.

A useful way to access the mindset of these examples is to recognise that natural induction is a generalised algorithm that is essentially substrate independent (like Darwinism can be generalised to other substrates (Aldrich, Hodgson et al. 2008)). It does not matter whether the dynamical system involves masses and springs, or gene expression levels and gene-regulatory connections, or species densities and interspecific fitness interactions – the algorithm is the same, and the source of the adaptive impetus is the same. So long as the fast variables are controlled by relatively slow changing interaction variables, the interaction variables give way or relax under stress, and the fast variables are occasionally perturbed – adaptation by natural induction will result. It also does not matter what the specific mechanism is that changes the slow variables – whether they are metals that yield under forcing (to reduce tension in the springs) (Buckley, Lewens et al. 2024), gene-interactions that evolve under natural selection (to enable expression levels to respond more fully to directional selection) (Watson, Wagner et al. 2014), ecological interactions that evolve under natural selection (to reduce competitive interactions and thereby increase individual growth rate) (Power, Watson et al. 2015), or for that matter, friendships in a social network that are modified to increase an individual utility function (Davies, Watson et al. 2011). However, as discussed below, notice that in some cases this means that natural selection and natural induction may be both involved in the same system at the same time.

2. The inductive evolution of ecosystem-level adaptations

The possibility of ecosystem adaptation, over and above the sum of adaptations occurring at the level of the individual populations contained therein, is an old debate (Clements 1916, Gleason 1917, Lovelock and Margulis 1974, Dawkins 1982, Lovelock and Margulis 2007, Betts and Lenton 2008, Doolittle 2014). It is clear that the individual populations within an ecosystem can evolve and be adapted to one another. It is also clear that this sometimes results in somewhat predictable system-level dynamics, such as succession dynamics (Clements 1916) that have been likened to the development of organismic phenotypes. But while there may be special cases where the ecosystem is a proper Darwinian unit (e.g. vertical inheritance of gut microbiota (Rohner and Moczek 2024)), in general, ecosystems are not subject to adaptation by natural selection as reproductive units in themselves. This underpins long-standing scepticism of the Gaia hypothesis, whose opponents have argued that the biosphere as a whole cannot be adapted unless there is a population of biospheres undergoing differential survival and reproduction (Doolittle 1981, Dawkins 1982). A different possibility, utilising a more general interpretation of natural selection, is that ecosystems might play the role of ‘interactors’ (Doolittle 2014, Papale and Doolittle 2024) where “the differential extinction and proliferation of interactors cause[s] the differential perpetuation of the replicators that produced them” (Hull 1980). This does not require the ecosystem to be a reproducer in the conventional sense and thus potentially broadens the scope of systems for which natural selection may contribute to adaptive explanation. This does still depend on a population of ecosystems with, if not differential reproduction, then at least differential survival, growth or persistence.

However, adaptive organisation at the ecosystem level could be provided by natural induction even when there is a 'population of one' and therefore no differential survival, growth or reproduction of ecosystems. This might help us to better-understand what is essential, and what is not, for an ecosystem to exhibit adaptive organisation. Power et al investigated this possibility modelling a single ecosystem as a network of interspecific competitive interactions with a generalised Lotka-Volterra model (Power, Watson et al. 2015). They allowed random variation and selection on individuals within each species to modify the interaction coefficients. Crucially, there is no system-level selection, only individual-level selection within each component species population. This acts, as always, to increase relative growth rate of individuals with respect to other individuals in the same species. But what is the effect of the coevolution between one species and another? A natural expectation is that an individual has an advantage over other individuals in its species if it can find a way to reduce the competitive influence it experiences from other species (e.g. through character displacement that decreases niche overlap). If we assume that it is not possible to be free from all competition, then reduced competition with one species causes increased competition with others, so selection at the individual level finds a trade-off. This is sensitive to combinations of species densities; The more two species are in high density at the same time, the greater the rate of evolutionary decrease in their competitive interaction (Power, Watson et al. 2015). This differential easing of frustrated relationships provides the Hebbian dynamics or associative learning that effects natural induction. The effect is that the ecosystem forms an associative memory of multiple past ecological states (Power, Watson et al. 2015).

To show that this can produce adaptation in the sense of the appearance of design, Power showed that the ecosystem can solve difficult combinatorial optimisation problems. To provide an intuitive feel for how difficult it is to find solutions to these problems, he created ecological constraint problems that represented Sudoku puzzles that humans find difficult (Power 2019). Each species represents a particular number in a particular cell of the puzzle. The set-up is that the initial interspecific interaction coefficients reflect the constraints of the puzzle – the rules of the game. For example, a "4" in one cell has a strongly competitive interaction with all the other numbers in that cell and with "4"s in other cells on the same row, column or box. The fixed numbers in the puzzle are represented by species that are clamped to high density. Given competitive interactions with this structure, the natural Lotka-Volterra ecological dynamics causes species that are in conflict to change density until an equilibrium is reached. This is better than a random configuration of species but not a good solution to the puzzle (i.e. many pairs of strongly competitive species remain in high density at the same time, meaning that many constraints of the puzzle remain unresolved). If the species densities are repeatedly perturbed to different initial conditions, the ecosystem finds a distribution of equilibria, none of which are particularly good solutions. However, as individual-level natural selection differentially relaxes the interspecific interaction coefficients over a distribution of these equilibrium states, the ecosystem learns dynamics that discover better solutions to the puzzle. Eventually a network of interactions is found that leads the ecosystem dynamics to find species configurations that solve the puzzle from any initial conditions, i.e. quickly re-finding a solution after each perturbation. This discovered perfect solutions for seven out of ten of the most difficult puzzles tested (known to be difficult for human experts), and it found good (but not perfect) solutions for the other three (Power 2019).

Notice that the behaviour of a system undergoing natural induction is not merely to return to an attractor after perturbation (ecological 'resilience'), nor merely to move out of one attractor into another arbitrary one (perhaps, 'fragility'). This would not constitute adaptation. In fact, the system exhibits a

systematic ability to learn from experience – undergoing transformational change in its organization that improves its capacity to find good solutions over time. This is not accounted for by resilience (or increasing resilience) or by fragility (or increasing fragility), with or without perturbations. Rather, it requires that the changing organization of the relationships among component parts takes on a specific structure that ‘mirrors’ (learns) the structure of the problem. Because this involves changes to connections, the resulting learning includes associative generalization. Accordingly, the system does not simply revisit the configurations it has visited in the past (memory), it also increases the likelihood of visiting other specific configurations – those with similar underlying regularities (generalisation). This provides an ability for the network to generate novel but non-arbitrary configurations, capturing patterns that are common to good solutions. In this way, the dynamics of the system does not merely find a better solution over time (once, as it were), but improves its ability to discover better solutions over time, such that, in the end, it can re-find exceptionally good solutions from any starting position quickly and reliably. It has not simply arrived at a good solution, but rather, knowledge of how to find a good solution is now incorporated into its succession dynamics.

In the sense that the species configurations found by the ecosystem constitute systems of many interdependent parts, which resolve many difficult-to-solve simultaneous constraints, they might reasonably constitute the appearance of design. And in the sense that the system improves its ability to find these configurations with experience, with solutions much better than the initial dynamics, it might reasonably be said to have adapted. If so, what is responsible for this adaptation? Is it the natural selection or the natural induction? Certainly, the ability of the ecosystem to find interactions that solve the puzzle is not possible if the interactions do not change, and without the action of natural selection the interactions do not change. Thus, selection is essential for the adaptation of the ecosystem. However, all this really means is that natural selection here contributes an element of the explanation of adaptation, not that natural selection is the general mechanism of adaptation.

Consider what happens when the network of springs solves an optimisation problem in the absence of natural selection. Material deformation of the springs is required to provide the second order relaxation of the system, and natural induction will not occur without it, but a theory of material deformation does not describe the adaptive algorithm. Likewise, in the case of the ecosystem, the work natural selection is doing is nothing more than the equivalent of relaxing the springs (i.e. the interspecific interaction coefficients). A theory of evolutionary change under natural selection (e.g. rate of allele frequency change is proportional to selection coefficients) does not describe the adaptive mechanism that is at work any more than material deformation explains the adaptation in the spring network, or a theory of electricity explains computation.

Our point here is not that Gaia is like a Darwinian organism: it does not reproduce, and it cannot be assigned a reproductive fitness. However, since natural induction provides a different mechanism of adaptation, and a different sense of adaptation, the possibility of ecological adaptation can be re-examined. Since this can occur in a population of one, without system-level reproduction, it potentially applies to the biosphere as a whole (Lovelock and Margulis 1974, Lenton and van Oijen 2002, Lenton, Daines et al. 2018, Lenton and Latour 2018). The biosphere might not be pondering Sudoku puzzles, but the notion that it can be organised for self-regulation (Lovelock and Margulis 1974, Lenton 2002), and perform resilient ecosystem functions (Oliver, Heard et al. 2015), for example, is consistent with the expectation that it will canalise non-arbitrary equilibria, that improve the ability to resolve conflicting constraints among its components.

3. The evolutionary induction of developmental structure and evolvability

In the ecosystem example, adaptation at this level of organisation is surprising because the ecosystem is not a Darwinian unit. Next, we consider the possibility that natural induction might provide adaptation at the same organisational scale as natural selection, but at a temporal scale where natural selection does not apply.

A set of models have been developed involving the evolution of gene-regulation networks (GRNs) (Watson, Buckley et al. 2010, Watson, Wagner et al. 2014, Kounios, Clune et al. 2016, Watson, Mills et al. 2016, Kouvaris, Clune et al. 2017, Rago, Kouvaris et al. 2019, Brun-Usan, Rago et al. 2020, Brun-Usan, Thies and Watson 2020, Brun Usan, Rago et al. 2021, Nash 2021). Unlike the inductive evolution of ecological adaptations discussed previously, here there is selection at the network level. An individual in the population is a network (not just a node in the network, as in the ecosystem model) and selection is applied to the phenotypes of the network as a whole, i.e. the gene-expression pattern produced by one network is compared to that of another. The patterns produced by the network are selected for their closeness to a particular target phenotype (Watson, Wagner et al. 2014) or for their fitness in a complex epistatic function (Kounios, Clune et al. 2016).

This work shows that there is a deep functional equivalence between the action of natural selection, when acting on the connections of the GRN, and Hebbian learning acting on neural connections (Watson, Wagner et al. 2014, Watson, Mills et al. 2016, Watson and Szathmary 2016). A single evolved GRN can hold a memory of several phenotypes selected in the past, recognise and repair corrupted phenotypes, and generalise to create new combinations of features (Watson, Wagner et al. 2014, Kouvaris, Clune et al. 2017). The work also shows that the conditions that increase an ability to generalise generated patterns in learning systems (e.g. costly connections) also translate into conditions that increase an ability to generalise adaptive phenotypes in evolving networks (Kouvaris, Clune et al. 2017). The ability to form an associative memory of fit phenotypes can also enhance evolvability in rugged fitness landscapes, making it possible to evolve high-fitness phenotypes that are otherwise unevolvable (Kounios, Clune et al. 2016). An important control experiment is provided by evolving a direct encoding of phenotypes (i.e. where an individual produces a vector of phenotypic traits each of which is controlled by a single, non-pleiotropic gene). Whereas the gene-regulation network has potential to evolve interactions that represent which traits 'go together' (see Discussion), the direct encoding cannot. In rugged fitness landscapes, the direct encoding merely finds the nearest locally optimal phenotype. In contrast, the GRN, evolved over many such locally optimal phenotypes, slowly improves its ability to find high-fitness phenotypes – more reliably and more quickly, and to find phenotypes that are higher in fitness than any that are found with the direct encoding (Kounios, Clune et al. 2016). This illustrates how an ability to learn a generalised model in a learning system corresponds to evolvability in the evolutionary model (Watson and Szathmary 2016).

Since, in this case, selection is applied at the level of the whole network, perhaps we are not surprised that natural selection finds good networks, capable of adapting quickly and reliably. But this does this mean it is correct to attribute all the adaptation that occurs to natural selection. The evolution of evolvability is a notoriously tricky topic in evolutionary theory (Wagner and Altenberg 1996, Gerhart and Kirschner 2007, Pigliucci 2008, Watson 2021). The crux of the problem is that natural selection is short-sighted – it has no ability to favour characteristics because they might facilitate *future* innovations. It is the differential fitness of phenotypes already produced that determines changes in allele frequencies,

not the (unrealised) potential to produce fit phenotypes in the future. More exactly, the timescale relevant to the operation of natural selection is one where competing characteristics coexist in the population long enough to deliver on their adaptive potential by producing phenotypes with different actualised fitnesses. In these models, for conceptual clarity, evolution is modelled under *strong selection weak mutation* assumptions where such coexistence is not possible (Kounios, Clune et al. 2016). So, how does the evolution of evolvability occur? One way of understanding it is that it is some sort of fortuitous side-effect of selection for something else. But if that were the case, it would be selection *of* characteristics that facilitated evolvability but not selection *for* characteristics that facilitated evolvability – in which case, any increase in evolvability would be merely fortuitous happenstance (Watson 2021). Accordingly, selection for evolvability (in the sense of a potential for future innovation) is conceptually problematic in a conventional natural selection framework, and in these models it is not the reason for an increase in evolvability.

However, adaptation that facilitates future innovation is not conceptually problematic in the context of learning. That is, in machine learning, it is very clear that genuine learning requires more than performing well against examples that have been seen during training – this only requires memorization (a.k.a. rote learning). Genuine learning requires generalization – an ability to use experience from the past and apply it to something new in the future. To an evolutionary thinker this seems like magical thinking. How can a system possibly be adapted for an environment that it has not previously encountered? In fact, in simple cases, this is obviously not a problem. A learner can be well-adapted to an environment that it has not encountered previously simply because it is well adapted to a past environment that is similar. Likewise, natural selection may bring improved fitness to a new environment without natural selection happening in the new environment yet, simply because the new environment is similar to past environments where selection has previously applied. This is a rather trivial sense of generalisation, however, depending on ‘naïve proximity’ discussed below (i.e. distance or differences in individual features).

In learning terms, however, “similarity” has other more sophisticated possibilities than naïve proximity, and generalization can be much more interesting. In particular, by learning underlying regularities such as associative relationships (not just individual features), a system can learn to respond appropriately to conditions that are novel. That is, it can be adapted to specific cases that are not nearby to anything in previous experience. When this happens, a system can be adapted for a future situation (known as the test set) that is non-trivially different from those it has previously encountered (known as the training set). Note that the reason the learning system changes in this way cannot be because it has foreknowledge of the test set – it does not. Rather generalization depends on an ability to identify underlying regularities from those patterns it has previously encountered and in common with those it has not (see Discussion for how this relates to inductive bias). In the context of biological adaptation, this means that transformational change in the organization of a system under natural induction can result in high fitness in a future environment, without assuming that the new environment is similar to old environments in a naïve sense of proximity. Of course, the reason the biological system changes in this way cannot be because it has already been exposed to the future environments – it has not been. Rather, generalization depends on an ability to identify underlying regularities from those environments it has previously encountered and in common with those it has not.

Thus, inductive generalisation (see Discussion) means that a system does well against cases it has not already had experience of. In effect, in the same sense that learning systems can generalise from the training set to the test set, natural induction is also not short-sighted (Kounios, Clune et al. 2016, Kouvaris, Clune et al. 2017). This is precisely what is enabled in the evolving GRN. Because it learns what goes together, instead of learning individual trait contributions, this enables it to learn how to evolve better (to enhance its future evolvability). Hence, natural induction is important because it brings improved fitness to the new environment without natural selection happening in the new environment yet. Many previous papers have also modelled the evolution of evolvability by evolving gene-regulation networks in environments that alternate conditions (Wagner and Altenberg 1996, Clune, Misevic et al. 2008, Parter, Kashtan and Alon 2008, Pigliucci 2008, Clune, Mouret and Lipson 2013). The concept of natural induction invites further investigation to establish which of these are dependent on naïve proximity and which are utilising associative generalisation. More importantly, natural induction gives us a conceptual framework that frees us from the assumption that long-term evolvability must involve selection that acts on the same timescale, or otherwise depend on fortuitous side-effects.

Again, as in the ecosystem examples, it is true that no evolution would occur without changes to the connections in the network, and without the action of natural selection, the interactions could not change. In this sense, selection is essential for the adaptation of the GRN. However, again, this only means that natural selection contributes to adaptation, not that natural selection is the general mechanism of adaptation. Again, one might counter that natural selection *in the proper context*, i.e. when acting on the pleiotropic interactions of a genotype phenotype map, does explain everything that evolves here. But again, this adaptive process requires the multiple timescales of interacting relaxation described in Section 2, acting on state variables and their interaction coefficients (first- and second-order variables), and that is the adaptive process of natural induction. This is true regardless of whether the sub-mechanism that changes the interaction coefficients is material yielding or natural selection. It is this that helps us understand how natural induction provides adaptations that facilitate the long-term evolution of evolvability, acting beyond a timescale where natural selection applies. Whereas adaptation from selection is post-hoc, not predictive (i.e. individuals are fit because they have already passed through the sieve of selection), learning uses structural regularities of past experience to anticipate adaptive responses to future environments. Thereby, when biological evolution instantiates the principles of natural induction, it demystifies the evolution of evolvability.

These models help us understand the role of development in evolution more generally. It is clear that the structure and dependencies of developmental architecture matters to evolvability and the possibility of phenotypic innovation (Wagner and Altenberg 1996, Gerhart and Kirschner 2007, Moczek 2008, Moczek, Sultan et al. 2011). To the extent that natural selection operates in basic configurations like a direct encoding (on a vector of genes or a vector of non-pleiotropic phenotypic traits) it provides a process functionally equivalent to first-order relaxation of a dynamical system against an energy function or objective function. Even in the case where there is pleiotropy and fitness epistasis, we can still characterise natural selection as a response to the average excess of individual alleles (given all this context) (Fisher 1941), but the result of this is limiting when framed this way. All that can be explained this way is (at best) the ascent of a local fitness gradient to a local adaptive peak.

In contrast, the framework of evolution by natural induction involves a separation between the (genetic) space where variation and selection is applied and the (feature) space where adaptation occurs (Hull 1980, Pigliucci 2010, Watson and Szathmary 2016, Buckley, Lewens et al. 2024). For example, gene

expression profiles are controlled indirectly by evolved gene-regulatory interactions (or ecological species profiles are controlled indirectly by evolved interspecific interaction coefficients). The separation of genotypic replicators from phenotypic interactors is already appreciated (Hull 1980). But it is additionally important to recognise that in this model of natural induction, the replicator/genes control the interaction terms of the developmental process that produces the interactor/phenotype. This means that the slow evolution of the genotype-phenotype map is functionally equivalent to associative learning and thus capable of non-trivial generalisation. It is this second-order process that permits the evolution of evolvability (and the evolution of ecosystem-level adaptations). Developmental structure is where evolution stores past experience that informs future evolution. This helps us to understand the essential role of development in evolutionary innovation, not just as a contextual contingency or a complicating factor, but an integral part of the processes that explain evolutionary innovation.

4. Discussion

In this pair of papers, we have discussed three different ways in which natural selection and natural induction might interact in biological evolution. In these examples, natural selection is necessary for genetic evolution to occur, but it is not the source of adaptive organisation. In induction-first evolution, a refinement of phenotype-first evolution (West-Eberhard 2003, West-Eberhard 2005, Watson, Levin and Lewens 2025), natural induction provides a mechanism of adaptation that operates within lifetimes, such as a mechanism of basal cognition, independent of natural selection. Then natural selection may canalise the adaptive solutions natural induction has already found. In the case of ecosystem adaptation, natural selection does not provide adaptations at this organisational scale because the ecosystem is not a Darwinian unit. And in the case of the evolution of evolvability, natural selection does provide adaptations at this timescale because natural selection does not apply at the relevant timescale, i.e. alternative evolvability characteristics do not coexist long enough for selection to see the innovative phenotypes they will produce. Both the latter cases can be understood as examples where natural selection does not apply to the type of adaptation observed, either because the adaptation is not at the right organisational scale or the right timescale for natural selection to apply, respectively. These are both different from the induction-first evolution scenario because, in that case, an adaptive configuration was found before any natural selection took place, whereas in the ecosystem adaptation model and developmental evolvability models, adaptation by natural induction cannot occur without natural selection.

In these examples, natural selection and natural induction should not be viewed as rival theories. In the same way that a theory of electricity is not a rival for a theory of computation (Watson, Levin and Lewens 2025), natural selection is not a rival for a theory of evolution by natural induction. Likewise, in the same way that current flow is necessary for computation or cognition to occur (in computers and brains as we know them), it is not sufficient to explain the algorithm that is operating. In fact, like a theory of current flow, a theory of natural selection might be nothing more than a theory that says allele frequencies go where they are pushed (i.e. by selection coefficients). What really matters to the algorithm of biological evolution is the organisation of those flows, the organisation of the selective field acting on variations, that results in adaptation.

The origin of life and evolutionary transitions in individuality

In pre-biotic systems where there are no self-reproducing structures, evolution by natural selection does not apply, but there may be many active components and reactions between them. Before the onset of

natural selection, it is generally supposed that these networks of reactions cannot have any adaptive organisation – it is merely a collection of independent parts ‘doing their own thing’ (albeit in the context of one another). After self-reproducing systems are established, natural selection could explain adaptation at the network level, organising the structure of reactions in this network to facilitate its survival and reproduction. Subsequently, we may view these components as parts within a larger system whose activity is orchestrated for the purposes of supporting Darwinian fitness. Of course, this does not violate whatever rules of interaction describe the physical and chemical processes that occur within this system – these still follow the same rules as before. But their organisation is different. They are still ‘doing their own thing’, but what they are doing in this arrangement serves the interests of the new Darwinian unit.

A potential limitation of this account for the origin of evolution is that all the molecular machinery necessary to establish the first self-reproducing system has to get organised ‘all at once’ – as if it is one big coincidence (Damer and Deamer 2020). A stronger account might be provided if there was some way for the network architecture to acquire a non-arbitrary adaptive organisation before it was a Darwinian unit (England 2015). One interesting theory is that this might involve repeated cycles of wet-dry cycling (Damer and Deamer 2020, Song, Simonis et al. 2024). These external perturbations cause the system to repeatedly relax into crystallised, low-energy states, from many different initial conditions (solute states). We speculate that, if this causes inelastic changes in molecular arrangements that change slowly over multiple cycles, this could facilitate a process of adaptation by natural induction, providing adaptive organisation before a Darwinian unit is established.

The major transitions in evolution (Maynard Smith and Szathmary 1997), or evolutionary transitions in individuality (ETIs) (Okasha 2006, Godfrey-Smith 2009), such as the transition from individual self-replicating molecules to chromosomes, and from unicellular life to multicellular organisms, have the characteristic that “entities that were capable of independent replication before the transition can replicate only as part of a larger whole after the transition” (Maynard Smith and Szathmary 1997). It might be assumed that the new level of evolutionary process created in an ETIs is not as difficult to explain as the origin of life because at least we begin the transition with components that are already capable of adaptation, i.e. already units of natural selection at their own level of organisation. This is not really sufficient to explain adaptation at the new higher level of organisation, however, because natural selection is not operating at this new level until after the transition. There may be many evolving components in networks of ecological relationships, but before the onset of natural selection at the new level, it is generally supposed that these ecological networks cannot have any adaptive organisation (Wilson 2016). It is merely a collection of independent units ‘doing their own thing’ (albeit in a context of one another). After natural selection is established at the collective level, this could organise the structure of ecological interactions in this network to facilitate the survival and reproduction of the collective. Subsequently, it is common to imagine these components as parts within a larger system whose activity is orchestrated for the purposes of supporting Darwinian fitness at the collective level (e.g. now a multicellular organism). Of course, this does not violate whatever interactions or constraints describe the (previously ecological, now developmental) processes that occur within this system – these still follow the same rules as before, e.g. maximising inclusive fitness (West and Gardner 2013). But their organisation is different. They are still ‘doing their own thing’, but in this new context (e.g. with functional integration, homogenous genetics and a population bottleneck (Godfrey-Smith 2009)) what they are doing serves the interests of the new Darwinian unit (even if it over-rules their direct fitness interests, e.g. in the case of

somatic cells). The correspondence with the origin of life should be clear; both the origin of life problem and the ETI problem are the problem of explaining an organisation that instantiates a unit of selection, before selection applies at that level (Watson, Levin and Buckley 2022).

If there were external perturbations that caused an ecosystem to repeatedly relax into climax communities, from many different initial conditions, and natural selection on ecological interactions caused (inelastic) changes to interspecific interactions that change slowly over multiple cycles, this could facilitate a process of adaptation by natural induction, as described in Section 2. Unlike a scenario based on natural selection only, this has potential to explain higher-level units that do not merely canalise the nearest local equilibrium (or evolutionarily stable state), but instead find configurations that are better than locally optimal, including adaptations that serve the interests of the whole whilst opposing the interests of the parts (e.g. germ-soma separation) (Watson, Levin and Buckley 2022, Watson 2024).

When collective phenotypes (Queller 2014, Corning and Szathmary 2015) have non-decomposable features that feed back onto the reproductive structures that control the evolution of their parts (Watson, Levin and Buckley 2022), the new Darwinian unit can be established. An abstract model of this illustrated how ecological dynamics (with perturbations) shaped the selection of new symbiotic relationships following Hebbian principles, enabling the symbiogenesis of new evolutionary units that were especially effective at solving ecological constraints (Watson, Palmius et al. 2011). This suggests the idea that natural selection and natural induction might enable a ‘multi-layer sandwich’ – where natural induction repeatedly shapes the ecological interactions among existing Darwinian units, this then enables adaptive organisations that facilitate evolutionary units at a new level of organisation, which then participate in a network of relationships adapted by natural induction at the next level, and so on (Watson, Mills et al. 2016, Vanchurin, Wolf et al. 2021). In this way, natural induction and natural selection may work together, balancing and potentiating each other, to ratchet-up increasing levels of biological complexity. In a machine learning context, it can be shown that a similar principle provides a very effective algorithm of ‘deep’ constraint optimisation, capable of solving problems that cannot be solved by single-level evolution (Caldwell, Knowles et al. 2022).

Adaptations for what?

We have mentioned that the trajectories of change under evolution by natural induction can be different from those under selection-first evolution by natural selection (Watson, Levin and Lewens 2025). We also mentioned that ecosystems have Darwinian fitness only in rare circumstances (e.g. vertical inheritance of holobiont communities), so natural induction cannot generally be aligned or misaligned with it. Also, there is a rather subtle sense in which adaptation for evolvability is potentially not fully aligned with conventional fitness too. For example, a phenotype that generalises well (in a way that will confer future innovation but has not yet done so) is not solely explained by direct alignment with maximising immediate fitness (Watson 2021). So, if the thing that natural induction increases is not the same thing as Darwinian fitness, what is it and what is their relationship?

One sense of directionality that might be useful in understanding the action of natural induction is that it causes a system to find configurations that allow its component parts to do what they were already doing more easefully (Watson 2024). This may also be related to dissipation driven adaptive organisation which turns work (in far from equilibrium systems) into organisations that dissipate energy more readily (England 2015, Chvykov, Berrueta et al. 2021). A dynamical system described by a number of active components in interaction with each other will often not be able to find globally minimal energy

configurations due to constrained or frustrated interactions (Buckley, Lewens et al. 2024, Watson 2024). But the modified dynamics of the system, after a natural induction process has adjusted the connections, causes the system to get better at resolving those frustrations. In other words, the interactions in the system become less frustrated or more harmonious. This is sufficient to solve complex combinatorial optimisation problems in a quantifiably objective manner (Buckley, Lewens et al. 2024, Watson 2024), and has some intuitive relevance to an appearance of design.

This kind of directionality is perhaps closer to ‘becoming a good model of the world’ or ‘minimising surprise’ (Schwartenbeck, FitzGerald et al. 2013, Friston, Levin et al. 2015, Friston, FitzGerald et al. 2016, Kirchhoff, Parr et al. 2018, Kuchling, Friston et al. 2020, Parr, Pezzulo and Friston 2022), rather than maximising survival or fecundity. Of course, an accurate model of the world might be useful for Darwinian fitness (Mitchell 2023), but this does not mean that Darwinian fitness was the ‘prime mover’ in the process. The “survival of the fittest” starting point is such a deep assumption in evolutionary thinking it is difficult to imagine anything else. But survival and reproduction are the agenda only because selection is the leader of the story. If learning is the leader, then self-preservation is not the basic organising principle, and although existing is a pre-requisite for learning, bear in mind that the ecosystem example demonstrates learning without being a Darwinian individual. A tendency for dynamical systems to come into a state of complementary organisation or “mutual knowing” with their environment (Froese and Di Paolo 2010, Friston, Levin et al. 2015, Fields and Levin 2020, Friston, Parr et al. 2024, Fernando, Banarse and Osindero 2025) might provide a more appropriate basis.

Deduction and induction in evolutionary concepts

Unlike deduction, inductive inferences draw conclusions that are not logically entailed by the observations (Table 2) (Fisher 1935, Popper 1979). However, induction is necessary for learning with generalisation. Whereas ‘rote learning’ (without generalisation) is just a kind of memorisation of what has already passed, generalisation means that you can apply what you have learned in novel situations that have not been previously encountered (Mitchell 1997). This therefore necessitates inductively established conclusions that go beyond past observation. Unlike evolutionary theory, learning theory recognises this explicitly. When the parameters of a model are modified to agree better with the training set¹, we can be certain that they agree better with the training set. However, this is the uninteresting aspect of the learning process. What is interesting is the ability to infer a productive generalisation that will work well on the test set, and this cannot be directly deduced from the training set performance. For this reason, in learning theory, inductive bias is understood to be integral and not seen as something that should ideally be eliminated. Appropriate inductive bias arises from the pre-existing structural regularities in common to the system doing the learning and the system that is being learned about (their underlying symmetries or complementarity). This can be simple and natural. For example, the most common and effective inductive bias is a preference for simple models (a.k.a. parsimony, or Occam’s razor). That is, out of the space of all possible models or hypotheses that agree with the observations, the ones that are simple, or utilise less model parameters, tend to generalise better.

¹ Mechanistically, this can be implemented by a variation and selection process (Such et al. 2017), although most learning systems utilise more efficient gradient methods.

Deduction	Induction
Paradigm cases use generalisations to infer specifics. E.g. 'all swans are white', 'this bird is a swan', therefore <deduction> 'this bird is white'	Paradigm cases use specifics to infer generalisations. E.g. 'this swan is white', 'that swan is white' therefore <induction> 'all swans are white'
Filter/eliminate false hypotheses	Create/generate new hypotheses
Logically valid, but insufficient for learning with generalisation	Capable of learning with generalisation, but not logically valid
Depends on nothing but evidence (selection). 'Bias'=unwanted preference.	Depends on inductive bias coming from the properties of the model type/material substrate (not just evidence)

Table 2: Some complementary aspects of logical inference types.

For some, the appealing logic of natural selection is deeply related to the logical validity of deduction (Popper 1979). That is, there is no doubt that an individual that has survived the sieve of selection is indeed fit with respect to that selective environment. The selection part of the evolutionary process is formally equivalent to Bayesian updating (Harper 2009, Shalizi 2009, Valiant 2013, Watson and Szathmary 2016). However, in order for evolution by natural selection to produce adaptations that are new (not part of the original standing variation) it must also generate new variants. In conventional evolutionary thinking, this part of the evolutionary process is often deemphasised. Variation is assumed to be undirected, such that all directionality comes from selection.

However, a deductive interpretation of natural selection is not as logically elegant as it seems. First, in a theoretical sense, it is not easy to define what it might mean to generate hypotheses from the set of 'all possibilities'. Second, making a specific generalisation *requires* bias (i.e. the set of all possibilities consistent with the data says nothing specific about cases that go beyond the data by definition). Third, biologically, it is increasingly recognised that generative bias is not simply a pariah (Wagner 1988, Wagner 1996, Wagner and Laubichler 2004, Gerhart and Kirschner 2007, Uller, Moczek et al. 2018) and that the notion of an unbiased distribution of variants might be conceptually wrong-headed (Salazar-Ciudad 2021).

Fourth, there is in any case a way that inductive bias enters (implicitly) into evolutionary thinking. This is via the assumption that new variations arise in the form of small modifications to existing variants. Thus, inasmuch as evolution by natural selection also involves the generation of new variants in a particular way, it does not consider *all* possibilities with equal probability. This means that evolutionary trajectories and outcomes are not determined solely by selection acting on individuals, but also by their underlying representation space (e.g. the neighbourhood of 'small variations'). This is evident when the adaptive abilities of evolution by natural selection are explained as local gradient ascent or 'hill-climbing' (i.e. as climbing local gradients in a fitness landscape) (Wright 1931, Wright 1988, Provine 1989). This conception of evolution is not purely deductive; it exploits a common but very weak inductive bias, i.e. the idea that good things tend to be near to other good things we have already seen or tend to have similar features. We will call this a naïve proximity bias.

Learning theory, in contrast, explicitly recognises that adaptation always involves inductive bias, and this kind of naïve proximity-based inductive bias is just a starting point. More sophisticated types of optimisation can build a more expressive model of the solution space based on past samples and use this to guide future exploration more intelligently. For example, the associative learning demonstrated in learning networks (whether composed of springs or of neurons or of GRN connections, etc.) exploits the

inductive generalisation that good things tend to have similar combinations of features to other good things we have already seen (even if they are not close in terms of naïve proximity) (Watson, Buckley and Mills 2011, Watson, Mills and Buckley 2011). This subtle change from patterns with similar features to patterns with similar *combinations* of features, or learning ‘what traits go together’ (a.k.a. associative learning, or Hebbian learning), gives natural induction adaptive capabilities that are superior to simple hill-climbing. This is still a kind of similarity principle, but it is a distance measure in the similarity of pairwise associations among features, rather than individual features. This change is subtle but important because it allows a model to change the proximity of patterns by ‘folding’ feature space, such that one pattern can be treated as close to another without also including all the patterns in between (the linear interpolation between them) (Watson, Wagner et al. 2014, Watson, Levin and Buckley 2022, Watson and Levin 2023).

This kind of associative induction is fundamental in models of learning (including all neural network learning) and sufficient for many kinds of non-trivial learning behaviours including the formation of an associative model that can store and recall multiple memories, generalising in a way that creates new combinations of features, and learning to solve optimisation problems that cannot be solved with hill-climbing (Hopfield 1982, Hopfield 1984, Hopfield and Tank 1986, Watson, Buckley and Mills 2011). In neural networks, this kind of inductive generalisation is provided by changes to the weighted connections between neurons. In physical systems, this kind of inductive generalisation is provided by a system that accommodates its internal structure (or hidden variables), i.e. the connections between components or their arrangement, to past experience (Zhong, Schwab and Murugan 2017, Stern, Pinson and Murugan 2020, Chvykov, Berrueta et al. 2021, Zhong, Gold et al. 2021, Stern and Murugan 2022). In contrast, a system that merely takes a superficial imprint of past experience in a direct or point-for-point manner (by accommodating only the surface structure or observable variables), like a print in a bed of clay, for example, does not afford this kind of associative learning (Buckley, Lewens et al. 2024). This is equivalent to evolution by natural selection with a direct encoding, i.e. phenotype=genotype (Kounios, Clune et al. 2016).

With this perspective we can see that natural selection and natural induction are not unrelated processes – they both involve refinement of a model with experience (Watson and Szathmary 2016) – but in natural induction the attention is on the (associative) inductive bias, whereas selection-based thinking (based on naïve proximity) tends to overlook the implicit inductive bias that is required for it to explain new adaptations. It is for this reason that explanations based on natural selection tend to presuppose that the iterative application of the generate and test process (involving a population), rather than simply following a gradient (given by an equation of motion), is an essential explainer of evolutionary adaptation. In contrast, when we recognise that inductive bias is always involved in adaptive change, the population becomes less important. Instead, the real adaptive work is provided by the inductive bias, inherent in the way that the problem variables relate to the solution space (e.g. the genotype-phenotype map). When this bias is appropriate, e.g. when combinations of features that were good together in some contexts are good together in other contexts, this allows for a gradient method (without a generate and test process or a population) to provide non-trivial adaptation. Hence, the significant algorithmic contrast between adaptation by natural induction and adaptation under natural selection is not that the former is transformational and the latter uses a population-based variational process – the important difference is that the former uses a slightly deeper inductive model than the latter (pairwise associations instead of naïve proximity).

Lastly, it is clear that the adaptive capabilities of natural selection and natural induction are both limited, but they are not limited in the same way. They apply in different conditions and produce adaptations that are different in kind. This changes where we might be willing to look for adaptive organisation in empirical investigations (e.g. ecosystems and loose social groups) whereas previously this was theoretically prohibited. And in the same way that natural selection reaches across domains to generalised Darwinism, natural induction may also legitimately apply outside of biology but in different ways from natural selection (Alexander, Cunningham et al. 2021, Lent 2021).

Reflecting on induction and selection and their interacting roles in biological evolution, we could accept a simple conclusion that they are two different sources of adaptation that, when they are both present, sometimes interact in complicated ways. But perhaps their interaction is more principled than that. Perhaps, they are more like a pair of processes (deduction and induction) that, in some sense, provide complementary aspects of an adaptive narrative that needs both. For example, natural selection is primarily a contractive process (removing variants that are incompatible with the environmental conditions) but it says little about the generation of new variants or why the selective conditions had that structure and not some other. Conversely, natural induction is an expansive process (creating novel generalisations by transforming the relationships among existing entities), but it says little about where the entities come from that participate in these relationships. Learning requires memory, something that persists, and natural selection provides continuity driven by self-preservation. But a creative process also requires things to change, to transform, to reorganise their components and relations over time. Natural selection on its own is not a theory of biological organisation – it doesn't say anything about the kind of organisations we should expect evolution to produce beyond 'whatever survives and reproduces'. Natural induction, in contrast, is a specific theory of how the organisation of relationships changes over time via principles of connectionist learning. And yet, the direction of this change can be understood in terms of how it better-satisfies the frustrations among the component parts involved in those relationships – e.g. their self-preservation. Biological evolution is a process that negotiates change and non-change at different levels of organisation. Perhaps this suggests that, in the bigger view of life and its evolution, natural selection and natural induction potentiate each other, and balance processes of self-preservation and transformation at multiple different organizational and temporal scales.

Perspectives

In this pair of papers, we first needed to make some conceptual space to even discuss the possibility that natural selection is not necessarily the only explanatory tool needed to make sense of biological evolution. This should not be controversial, there is a lot we do not understand about biological evolution, but evolutionary thinking is often tightly locked-down as if natural selection provided a complete explanation and no further discussion were required. This makes it difficult to have any productive discussion about how other mechanisms might change how we think about biological evolution in any significant way. For example, the influence of phenotype-first evolution is constricted by the presupposition that any form of adaptive plasticity must itself be a product of prior natural selection. We suggest that natural induction releases phenotype-first evolution from this way of thinking and provides a way that it can be a genuine leader of evolutionary change with natural selection acting as a memory of adaptations it discovers (Watson, Levin and Lewens 2025). In this second paper, we have discussed some other more subtle forms of potential interaction between natural induction and natural selection, involving processes that depend on natural selection but producing adaptation that cannot be attributed to natural selection (like computation depends on electricity, in digital computers as we know them, but cannot be attributed to

electricity). Specifically, this is because natural induction is producing adaptations at organisational scales or temporal scales where natural selection does not apply. We do not present natural induction as a competitor for or replacement for evolution by natural selection. In all the scenarios discussed, evolution by natural selection is involved, and to the extent that it is, it works the way it always does. However, the presumption that natural selection necessarily provides a complete and exclusive mechanism of all adaptation in biological evolution is an over-statement. We hope that evolution by natural induction helps to expand the space of explanatory concepts and mechanistic possibilities.

Acknowledgements

We thank Chrisantha Fernando for the original insight that Hebbian principles can apply in non-neural substrates, Bruce Damer for discussions about the relevance to the origin of life, Charlie Munford, Perry Marshall and Josh Knowles for discussions about the nature of induction and its relation to selection, Denis Noble, Jim Shapiro and Eva Jablonka for insights into 'top-down' pathways, phenotypic adaptation and non-genetic inheritance, and Kevin Lala, Tobias Uller and Adi Livnat for discussion that invites expansion of evolutionary thinking. Thank you to Sylvain Charlat for many long discussions on the conceptual and philosophical distinctions involved in this work. Also, Frederick Nash, Kostas Kouvaris, Louis Kounios, David Prosser, Daniel Power, Miguel Brun-Usan, Jamie Caldwell, Chris Buckley, Adam Davies, Rob Mills, Christoph Thies, Taztio Tissot and Nam Le for the many computational models and conversations that co-constructed this larger conceptual vision before 'evolution by natural induction' had a name.

References

- Aldrich, H. E., G. M. Hodgson, D. L. Hull, T. Knudsen, J. Moky and V. J. Vanberg (2008). "In defence of generalized Darwinism." *Journal of evolutionary economics* **18**(5): 577-596.
- Alexander, S., W. J. Cunningham, J. Lanier, L. Smolin, S. Stanojevic, M. W. Toomey and D. Wecker (2021). "The autodidactic universe." *arXiv preprint arXiv:2104.03902*.
- Betts, R. A. and T. M. Lenton (2008). *Second chances for lucky gaia: a hypothesis of sequential selection*, Met Office.
- Brun-Usan, M., A. Rago, C. Thies, T. Uller and R. A. Watson (2020). "Developmental models reveal the role of phenotypic plasticity in explaining genetic evolvability." *bioRxiv*.
- Brun-Usan, M., C. Thies and R. A. Watson (2020). "How to fit in: The learning principles of cell differentiation." *PLoS computational biology* **16**(4): e1006811.
- Brun Usan, M., A. Rago, C. Thies, T. Uller and R. A. Watson (2021). "Development and Selective Grain Make Plasticity 'Take the Lead' in Adaptive Evolution."
- Buckley, C. L., T. Lewens, M. Levin, B. Millidge, A. Tschantz and R. A. Watson (2024). "Natural Induction: Spontaneous adaptive organisation without natural selection." *Entropy* **26**(9).
- Caldwell, J., J. Knowles, C. Thies, F. Kubacki and R. Watson (2022). "Deep Optimisation: Transitioning the Scale of Evolutionary Search by Inducing and Searching in Deep Representations." *SN Computer Science* **3**(3): 253.
- Chvykov, P., T. A. Berrueta, A. Vardhan, W. Savoie, A. Samland, T. D. Murphey, K. Wiesenfeld, D. I. Goldman and J. L. England (2021). "Low rattling: A predictive principle for self-organization in active collectives." *Science* **371**(6524): 90-95.
- Clements, F. E. (1916). *Plant succession: an analysis of the development of vegetation*, Carnegie institution of Washington.
- Clune, J., D. Misevic, C. Ofria, R. E. Lenski, S. F. Elena and R. Sanjuán (2008). "Natural selection fails to optimize mutation rates for long-term adaptation on rugged fitness landscapes." *PLoS computational biology* **4**(9): e1000187.
- Clune, J., J. B. Mouret and H. Lipson (2013). "The evolutionary origins of modularity." *Proc Biol Sci* **280**(1755): 20122863.
- Corning, P. A. and E. Szathmary (2015). "'Synergistic selection': a Darwinian frame for the evolution of complexity." *J Theor Biol* **371**: 45-58.
- Damer, B. and D. Deamer (2020). "The hot spring hypothesis for an origin of life." *Astrobiology* **20**(4): 429-452.
- Davies, A. P., R. A. Watson, R. Mills, C. Buckley and J. Noble (2011). "'If You Can't Be With the One You Love, Love the One You're With': How Individual Habituation of Agent Interactions Improves Global Utility." *Artificial Life* **17**(3): 167-181.
- Dawkins, R. (1982). *The extended phenotype: The long reach of the gene*, Oxford University Press.
- Doolittle, F. W. (2014). "Natural selection through survival alone, and the possibility of Gaia." *Biology & Philosophy* **29**(3): 415-423.
- Doolittle, W. F. (1981). "Is nature really motherly?" *CoEvolution Quat* **29**: 58-63.

- England, J. L. (2015). "Dissipative adaptation in driven self-assembly." Nature nanotechnology **10**(11): 919-923.
- Fernando, C., D. Banarse and S. Osindero (2025). "Wanting to be Understood." arXiv preprint arXiv:2504.06611.
- Fields, C. and M. Levin (2020). "How do living systems create meaning?" Philosophies **5**(4): 36.
- Fisher, R. A. (1935). "The logic of inductive inference." Journal of the royal statistical society **98**(1): 39-82.
- Fisher, R. A. (1941). "Average excess and average effect of a gene substitution."
- Friston, K., T. FitzGerald, F. Rigoli, P. Schwartenbeck and G. Pezzulo (2016). "Active inference and learning." Neuroscience & Biobehavioral Reviews **68**: 862-879.
- Friston, K., M. Levin, B. Sengupta and G. Pezzulo (2015). "Knowing one's place: a free-energy approach to pattern regulation." J R Soc Interface **12**(105).
- Friston, K. J., T. Parr, C. Heins, A. Constant, D. Friedman, T. Isomura, C. Fields, T. Verbelen, M. Ramstead and J. Clippinger (2024). "Federated inference and belief sharing." Neuroscience & Biobehavioral Reviews **156**: 105500.
- Froese, T. and E. A. Di Paolo (2010). "Modelling social interaction as perceptual crossing: an investigation into the dynamics of the interaction process." Connection Science **22**(1): 43-68.
- Gerhart, J. and M. Kirschner (2007). "The theory of facilitated variation." Proc Natl Acad Sci U S A **104 Suppl 1**(Suppl 1): 8582-8589.
- Gleason, H. A. (1917). "The structure and development of the plant association." Bulletin of the Torrey Botanical Club **44**(10): 463-481.
- Godfrey-Smith, P. (2009). Darwinian populations and natural selection, Oxford University Press.
- Harper, M. (2009). "The replicator equation as an inference dynamic." arXiv preprint arXiv:0911.1763.
- Hopfield, J. J. (1982). "Neural networks and physical systems with emergent collective computational abilities." Proceedings of the national academy of sciences **79**(8): 2554-2558.
- Hopfield, J. J. (1984). "Neurons with graded response have collective computational properties like those of two-state neurons." Proceedings of the national academy of sciences **81**(10): 3088-3092.
- Hopfield, J. J. and D. W. Tank (1986). "Computing with neural circuits: a model." Science **233**(4764): 625-633.
- Hull, D. L. (1980). "Individuality and selection." Annual review of ecology and systematics **11**: 311-332.
- Kirchhoff, M., T. Parr, E. Palacios, K. Friston and J. Kiverstein (2018). "The Markov blankets of life: autonomy, active inference and the free energy principle." Journal of The royal society interface **15**(138): 20170792.
- Kounios, L., J. Clune, K. Kouvaris, G. P. Wagner, M. Pavlicev, D. M. Weinreich and R. A. Watson (2016). "Resolving the paradox of evolvability with learning theory: How evolution learns to improve evolvability on rugged fitness landscapes." arXiv preprint arXiv:1612.05955.

- Kouvaris, K., J. Clune, L. Kounios, M. Brede and R. A. Watson (2017). "How evolution learns to generalise: Using the principles of learning theory to understand the evolution of developmental organisation." *PLoS computational biology* **13**(4): e1005358.
- Kuchling, F., K. Friston, G. Georgiev and M. Levin (2020). "Morphogenesis as Bayesian inference: A variational approach to pattern formation and control in complex biological systems." *Phys Life Rev* **33**: 88-108.
- Lent, J. (2021). *The web of meaning: Integrating science and traditional wisdom to find our place in the universe*, New Society Publishers.
- Lenton, T. M. (2002). "Testing Gaia: the effect of life on Earth's habitability and regulation." *Climatic Change* **52**(4): 409-422.
- Lenton, T. M., S. J. Daines, J. G. Dyke, A. E. Nicholson, D. M. Wilkinson and H. T. Williams (2018). "Selection for Gaia across multiple scales." *Trends in Ecology & Evolution* **33**(8): 633-645.
- Lenton, T. M. and B. Latour (2018). "Gaia 2.0." *Science* **361**(6407): 1066-1068.
- Lenton, T. M. and M. van Oijen (2002). "Gaia as a complex adaptive system." *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* **357**(1421): 683-695.
- Lewens, T. and R. A. Watson (2024). *Universal Darwinism Revisited*. (under review).
- Lovelock, J. and L. Margulis (2007). "The gaia hypothesis." *New York*.
- Lovelock, J. E. and L. Margulis (1974). "Atmospheric homeostasis by and for the biosphere: the Gaia hypothesis." *Tellus* **26**(1-2): 2-10.
- Maynard Smith, J. and E. Szathmary (1997). *The major transitions in evolution*, Oxford University Press.
- Mitchell, K. J. (2023). "Free agents: how evolution gave us free will."
- Mitchell, T. M. (1997). *Machine learning*, McGraw-hill New York.
- Moczek, A. P. (2008). "On the origins of novelty in development and evolution." *BioEssays* **30**(5): 432-447.
- Moczek, A. P., S. Sultan, S. Foster, C. Ledón-Rettig, I. Dworkin, H. F. Nijhout, E. Abouheif and D. W. Pfennig (2011). "The role of developmental plasticity in evolutionary innovation." *Proceedings of the Royal Society B: Biological Sciences* **278**(1719): 2705-2713.
- Nash, F. K., L.; Thies, C.; Kouvaris, K.; Tarapore, D.; Watson, R. (2021). "Scaling-up variation and evolvability: The causes and consequences of developmental hierarchy." *BioArxiv*.
- Okasha, S. (2006). *Evolution and the levels of selection*, Oxford University Press.
- Oliver, T. H., M. S. Heard, N. J. Isaac, D. B. Roy, D. Procter, F. Eigenbrod, R. Freckleton, A. Hector, C. D. L. Orme and O. L. Petchey (2015). "Biodiversity and resilience of ecosystem functions." *Trends in ecology & evolution* **30**(11): 673-684.
- Papale, F. and W. F. Doolittle (2024). "Towards a more general theory of evolution by natural selection: a manifesto." *Philosophy, Theory, and Practice in Biology* **16**(1).
- Parr, T., G. Pezzulo and K. J. Friston (2022). *Active inference: the free energy principle in mind, brain, and behavior*, MIT Press.

- Parter, M., N. Kashtan and U. Alon (2008). "Facilitated variation: how evolution learns from past environments to generalize to new environments." *PLoS Comput Biol* **4**(11): e1000206.
- Pigliucci, M. (2008). "Is evolvability evolvable?" *Nature Reviews Genetics* **9**(1): 75-82.
- Pigliucci, M. (2010). "Genotype–phenotype mapping and the end of the ‘genes as blueprint’ metaphor." *Philosophical Transactions of the Royal Society B: Biological Sciences* **365**(1540): 557-566.
- Popper, K. R. (1979). *Objective knowledge: An evolutionary approach*, Clarendon press Oxford.
- Power, D. (2019). "Distributed associative learning in ecological community networks." *University of Southampton PhD Dissertation*.
- Power, D. A., R. A. Watson, E. Szathmary, R. Mills, S. T. Powers, C. P. Doncaster and B. Czapp (2015). "What can ecosystems learn? Expanding evolutionary ecology with learning theory." *Biol Direct* **10**: 69.
- Provine, W. B. (1989). *Sewall Wright and evolutionary biology*, University of Chicago press.
- Queller, D. C. (2014). "Joint phenotypes, evolutionary conflict and the fundamental theorem of natural selection." *Philosophical Transactions of the Royal Society B: Biological Sciences* **369**(1642): 20130423.
- Rago, A., K. Kouvaris, T. Uller and R. Watson (2019). "How adaptive plasticity evolves when selected against." *PLoS computational biology* **15**(3): e1006260.
- Rohner, P. T. and A. P. Moczek (2024). "Vertically inherited microbiota and environment modifying behaviours conceal genetic variation in dung beetle life history." *Proceedings of the Royal Society B* **291**(2021): 20240122.
- Salazar-Ciudad, I. (2021). "Why call it developmental bias when it is just development?" *Biology Direct* **16**: 1-13.
- Schwander, T. and O. Leimar (2011). "Genes as leaders and followers in evolution." *Trends in Ecology & Evolution* **26**(3): 143-151.
- Schwartenbeck, P., T. FitzGerald, R. Dolan and K. Friston (2013). "Exploration, novelty, surprise, and free energy minimization." *Frontiers in psychology*: 710.
- Shalizi, C. R. (2009). "Dynamics of Bayesian updating with dependent data and misspecified models." *Electronic Journal of Statistics* **3**: 1039-1074.
- Song, X., P. Simonis, D. Deamer and R. N. Zare (2024). "Wet–dry cycles cause nucleic acid monomers to polymerize into long chains." *Proceedings of the National Academy of Sciences* **121**(49): e2412784121.
- Stern, M. and A. Murugan (2022). "Learning without neurons in physical systems." *arXiv preprint arXiv:2206.05831*.
- Stern, M., M. B. Pinson and A. Murugan (2020). "Continual learning of multiple memories in mechanical networks." *Physical Review X* **10**(3): 031044.
- Such, F. P., V. Madhavan, E. Conti, J. Lehman, K. O. Stanley and J. Clune (2017). "Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning." *arXiv preprint arXiv:1712.06567*.

- Uller, T., A. P. Moczek, R. A. Watson, P. M. Brakefield and K. N. Laland (2018). "Developmental bias and evolution: A regulatory network perspective." *Genetics* **209**(4): 949-966.
- Valiant, L. (2013). "Probably Approximately Correct: Nature's Algorithms for Learning and Prospering in a Complex World." ISBN-13: 978-0465060726.
- Vanchurin, V., Y. I. Wolf, M. I. Katsnelson and E. V. Koonin (2021). "Towards a Theory of Evolution as Multilevel Learning." *arXiv preprint arXiv:2110.14602*.
- Wagner, G. P. (1988). "The influence of variation and of developmental constraints on the rate of multivariate phenotypic evolution." *Journal of evolutionary biology* **1**(1): 45-66.
- Wagner, G. P. (1996). "Homologues, natural kinds and the evolution of modularity." *American Zoologist* **36**(1): 36-43.
- Wagner, G. P. and L. Altenberg (1996). "Perspective: complex adaptations and the evolution of evolvability." *Evolution* **50**(3): 967-976.
- Wagner, G. P. and M. D. Laubichler (2004). "Rupert Riedl and the re-synthesis of evolutionary and developmental biology: body plans and evolvability." *J Exp Zool B Mol Dev Evol* **302**(1): 92-102.
- Watson, R. (2024). "Agency, Goal-Directed Behavior, and Part-Whole Relationships in Biological Systems." *Biological Theory* **19**(1): 22-36.
- Watson, R., C. L. Buckley, R. Mills and A. Davies (2010). "Associative memory in gene regulation networks."
- Watson, R. and M. Levin (2023). "The collective intelligence of evolution and development." *Collective Intelligence* **2**(2): 26339137231168355.
- Watson, R., N. Palmius, R. Mills, S. Powers and A. Penn (2011). "Can selfish symbioses effect higher-level selection?" *Advances in Artificial Life. Darwin Meets von Neumann*: 27-36.
- Watson, R. A. (2021). "Evolvability." *Evolutionary Developmental Biology: A Reference Guide*: 133-148.
- Watson, R. A., C. Buckley and R. Mills (2011). "Optimization in "self-modeling" complex adaptive systems." *Complexity* **16**(5): 17-26.
- Watson, R. A., M. Levin and C. L. Buckley (2022). "Design for an individual: connectionist approaches to the evolutionary transitions in individuality." *Frontiers in Ecology and Evolution*: 64.
- Watson, R. A., M. Levin and T. Lewens (2025). *Evolution by Natural Induction*. this volume.
- Watson, R. A., R. Mills, C. Buckley, K. Kouvaris, A. Jackson, S. T. Powers, C. Cox, S. Tudge, A. Davies and L. Kounios (2016). "Evolutionary connectionism: algorithmic principles underlying the evolution of biological organisation in evo-devo, evo-eco and evolutionary transitions." *Evolutionary biology* **43**(4): 553-581.
- Watson, R. A., R. Mills and C. L. Buckley (2011). "Global adaptation in networks of selfish components: emergent associative memory at the system scale." *Artif Life* **17**(3): 147-166.
- Watson, R. A. and E. Szathmary (2016). "How can evolution learn?" *Trends in Ecology and Evolution*(31): 147-157.

- Watson, R. A., G. P. Wagner, M. Pavlicev, D. M. Weinreich and R. Mills (2014). "THE EVOLUTION OF PHENOTYPIC CORRELATIONS AND 'DEVELOPMENTAL MEMORY'." Evolution **68**(4): 1124-1138.
- West-Eberhard, M. J. (2003). Developmental plasticity and evolution, Oxford University Press.
- West-Eberhard, M. J. (2005). "Developmental plasticity and the origin of species differences." Proceedings of the National Academy of Sciences **102**(suppl_1): 6543-6549.
- West, S. A. and A. Gardner (2013). "Adaptation and inclusive fitness." Current Biology **23**(13): R577-R584.
- Wilson, D. S. (2016). Two Meanings of Complex Adaptive Systems. Complexity and Evolution: Toward a New Synthesis for Economics. D. S. Wilson and A. Kirman, The MIT Press: 0.
- Wright, S. (1931). "Evolution in Mendelian populations." Genetics **16**(2): 97.
- Wright, S. (1988). "Surfaces of selective value revisited." The American Naturalist **131**(1): 115-123.
- Zhong, W., J. M. Gold, S. Marzen, J. L. England and N. Yunger Halpern (2021). "Machine learning outperforms thermodynamics in measuring how well a many-body system learns a drive." Scientific Reports **11**(1): 9333.
- Zhong, W., D. J. Schwab and A. Murugan (2017). "Associative pattern recognition through macro-molecular self-assembly." Journal of Statistical Physics **167**: 806-826.