entropy

MDPI

Article

Natural Induction: Spontaneous Adaptive Organisation without Natural Selection

Christopher L. Buckley [1], Tim Lewens [2], Michael Levin [3], Beren Millidge [1], Alexander Tschantz [1] and Richard A. Watson [4,*]

[1] Department of Informatics, University of Sussex, Brighton BN1 9RH, UK; c.l.buckley@sussex.ac.uk (C.L.B.); beren@millidge.name (B.M.); alec.tschantz@verses.ai (A.T.)
[2] History and Philosophy of Science, Cambridge University, Cambridge CB2 1TN, UK; tml1000@cam.ac.uk
[3] Department of Biology, Tufts University, Medford, MA 02155, USA; michael.levin@tufts.edu
[4] Electronics and Computer Science/Institute for Life Sciences, University of Southampton, Southampton SO17 1BJ, UK
* Correspondence: r.a.watson@soton.ac.uk

Abstract: Evolution by natural selection is believed to be the only possible source of spontaneous adaptive organisation in the natural world. This places strict limits on the kinds of systems that can exhibit adaptation spontaneously, i.e., without design. Physical systems can show some properties relevant to adaptation without natural selection or design. (1) The relaxation, or local energy minimisation, of a physical system constitutes a natural form of optimisation insomuch as it finds locally optimal solutions to the frustrated forces acting on it or between its components. (2) When internal structure 'gives way' or accommodates a pattern of forcing on a system, this constitutes learning insomuch, as it can store, recall, and generalise past configurations. Both these effects are quite natural and general, but in themselves insufficient to constitute non-trivial adaptation. However, here we show that the recurrent interaction of physical optimisation and physical learning together results in significant spontaneous adaptive organisation. We call this adaptation by natural induction. The effect occurs in dynamical systems described by a network of viscoelastic connections subject to occasional disturbances. When the internal structure of such a system accommodates slowly across many disturbances and relaxations, it spontaneously learns to preferentially visit solutions of increasingly greater quality (exceptionally low energy). We show that adaptation by natural induction thus produces network organisations that improve problem-solving competency with experience (without supervised training or system-level reward). We note that the conditions for adaptation by natural induction, and its adaptive competency, are different from those of natural selection. We therefore suggest that natural selection is not the only possible source of spontaneous adaptive organisation in the natural world.

Keywords: learning; optimisation; adaptation; self-organisation; evolution

check for updates

Citation: Buckley, C.L.; Lewens, T.; Levin, M.; Millidge, B.; Tschantz, A.; Watson, R.A. Natural Induction: Spontaneous Adaptive Organisation without Natural Selection. *Entropy* **2024**, *26*, 765. https://doi.org/10.3390/e26090765

Academic Editor: Antonio Maria Scarfone

Received: 31 July 2024
Revised: 19 August 2024
Accepted: 27 August 2024
Published: 6 September 2024

Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Darwin explicitly focused his work on the Origin of Species on answering the question "How have all those exquisite adaptations of one part of the organic organisation to another part, and to the conditions of life, and of one distinct being to another being, been perfected?" [1]. The answer he gives, of course, is natural selection. Natural selection is a process of adaptation characterised by the differential survival and reproduction of randomly varying types [2,3]. By describing the mechanism of natural selection, Darwin showed that the number of natural processes that could produce adaptive organisation spontaneously (i.e., without a designer) was at least one. Since he showed that it is possible for natural processes to produce adaptive complexity spontaneously, are we sure that the number of such processes is exactly one? Are there others? Or do all putative alternatives

turn out to be either some form of natural selection in another guise or not a genuine source of adaptation at all [4]?

1.1. Natural Sources of Adaptation

If natural selection is understood as a specific type of process (i.e., if it is not defined as any and all processes that produce adaptation), and likewise if adaptation is not defined as the product of natural selection, then it is a valid question to ask whether other processes can produce adaptation. Whilst no well-informed thinker doubts that evolution by natural selection plays a central role in explaining the adaptive organisation of biological organisms, and the role of natural selection in biological evolution need not be threatened by asking this question, the potential ramifications of such a possibility could be important to biological understanding in many areas. These include understanding how natural selection interacts with other biological and physical processes (e.g., development, niche construction, ecological dynamics, extended inheritance mechanisms, learning, organismic agency) [5,6]; understanding how natural selection got started [7,8]; informing open questions such as how natural selection rescales from one level of organisation to another, i.e., evolutionary transitions in individuality [9]; and exploring the possibility of adaptation in biological systems that are not evolutionary units [10–13]. Thus, to expand further the systematic investigation of the question that Darwin first posed, here we investigate the kinds of natural processes that can produce adaptive organisation spontaneously.

It is clear that other mechanisms of adaptation exist in biology, such as adaptive plasticity and animal learning [14,15], and also in engineered systems, such as machine learning mechanisms [16]. These do not necessarily depend on variation and selection for their operation. For example, adaptation like that which occurs via the organisation and reorganisation of synaptic connections in the brain need not involve a selection process (Whilst there is a long history exploring the idea that cognitive processes can indeed be characterised as variation and selection processes [17–22], artificial neural networks show that this type of learning can be implemented using gradient methods that do not involve variation and selection at any level of organisation, including the level of synapses or neurons, nor at the level of the network as a whole). Of course, these examples involve specific adaptive mechanisms that are themselves selected or designed for the purpose of producing adaptive outcomes. Thus, the kind of adaptation that brains (and machine learning systems) demonstrate does show that natural selection is not the only possible adaptive mechanism, but they do not appear to offer a solution to both the chicken problem and the egg problem, namely, to present an adaptive mechanism that is not natural selection and also does not need natural selection to explain how it arose. What adaptive mechanism could be so simple that it does not require selection or design either for its operation (at 'run time') nor to explain its necessary machinery (i.e., for construction or 'set-up')? For lack of any such known alternative in the natural world, natural selection is understood to be the source of all adaptive organisation [4]. Moreover, it has been argued, quite convincingly, that natural selection is the only possible mechanism of adaptation that could occur naturally (even on another planet where we could imagine biology that worked differently), hence "universal Darwinism" [4]. Is it really impossible for there to be other sources of spontaneous adaptation besides natural selection?

If we wish to explore the possibility of a different source of adaptation, i.e., not requiring natural selection either to set up the necessary machinery or in its operation thereafter, we had better be able to show how it works in a system that is not associated with natural selection at all, i.e., that does not depend on any specific characteristics of biological systems or materials. Even if our real interest in adaptation was exclusively biological, it would need to be possible to show adaptation in a non-biological, physical system. This is a stricter condition than Darwin applied—evolution by natural selection is a mechanism of adaptation that presupposes properties of biological systems. It depends on self-reproducing entities that exhibit heritable variation in reproductive success. But given these other (natural and artificial) examples of adaptation in learning processes, one

way to potentially expand the scope of adaptive processes is to ask: what other kinds of systems can learn, and in particular can physical systems exhibit learning spontaneously, i.e., without being selected or designed for that purpose?

1.2. Physical Learning and Physical Optimisation

The principles of learning systems can be very simple, and spontaneous "physical learning" can be readily demonstrated in various kinds of physical systems [23,24]. These include mechanical, material, molecular, and chemical systems, as well as electrical [23,25–31]. Although the details vary, the underlying principle is the differential accommodation or acclimation of system structures to forces acting on them. This might involve natural properties such as differential deformation (yielding, creeping, ageing, dilation, or atrophy) of connections, bonds or flows, or re-arrangements of internal structure in response to local stresses in those structures [23,24,29,32]. When the internal organisation of a system is incompatible with, or stressed by, the pattern of forces acting on the system from the environment (or from its own dynamical behaviour), that organisation is caused to change by those forces (Figure 1B). Such systems bend, deform, or give way in the direction that accommodates to the specific stresses created by that pattern of forcing. This can result in a sort of memory (or engram) of that forcing that remains in the structural organisation of the system and can thereby influence the system's subsequent dynamics in a manner that reconstitutes or 'recalls' the past pattern. A beautifully simple example is shown in paper folding, where up-folds or down-folds become easier as the paper yields to the forcing it has experienced, leaving a memory that can recreate a given input–output relation or classification [33].

To make sense of what this means, let us start simple. In simple form, a physical memory is trivial. A simple univariate memory (i.e., modelling each variable independently), or imprint, can be shown by something as ordinary as a uniform plastic material like a bed of clay. The features in the pattern of forcing have a one-to-one correspondence with the elements of the material. (The biological analogue of this is a direct one-to-one genotype–phenotype mapping where selection acts on each genetically specified trait independently without pleiotropy). However, this can only remember one pattern at a time, e.g., the deformation of the clay to a new object overwrites the memory of any previous object. At best, the history of patterns it has been exposed to is reduced to an average or consensus pattern. In contrast, when the deformation in a system occurs in an internal structural organisation that affects linkages or relationships between features, as in these examples, this is capable of a higher-order or associative memory [23,29,34–36]. The deformation in such a system can store and recall multiple patterns with high fidelity [35,37]. In this case, it is possible for the internal organisation to represent a set or class of patterns, or to learn a non-linear functional relationship between input and output variables. Such physical learning thereby shows competence in some conventional learning tasks such as classification, function learning, or reinforcement learning [23,24,29]. Learning in a simple neural network can also be understood as an example of a system with connections that deform under stress. Hebb's rule [38], a simple method of updating neural connections often described as 'neurons that fire together wire together', is equivalent to energy minimisation on the weights of the network given the current (e.g., forced) pattern of the states [35,39–41]. Like a neural network, a physical system with this property can also *generalise* to recognise, classify, or produce novel patterns that belong to the same general class of patterns even though these specific cases have not been experienced previously.

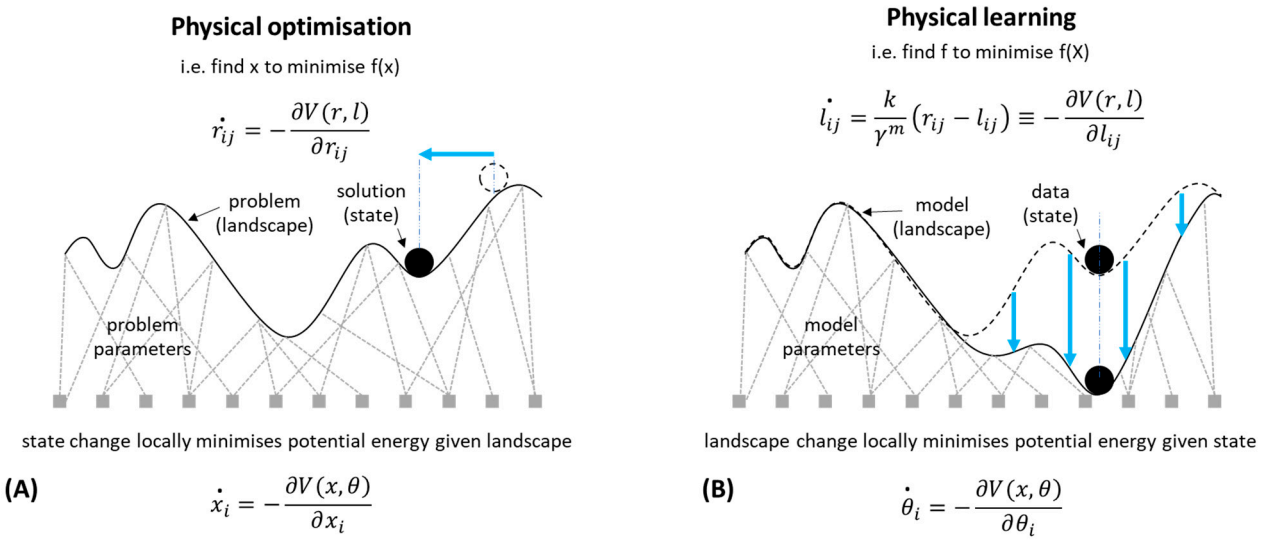

Figure 1. Optimisation and learning in physical systems are complementary processes. (**A**) Physical optimisation is described by the change in state ($x = \{x_i\}$) given the set of parameters ($\theta = \{\theta_i\}$) defining a potential function, V. (**B**) Physical learning, in contrast, is described by the change in the parameters of a model given some state or distribution of states. We consider optimisation and learning in a system of masses connected by springs, $m\ddot{r}_{ij} = -\frac{\partial V(r,l)}{\partial r_{ij}} - \gamma\dot{r}_{ij}$, and for simplicity, when this system is overdamped (e.g., with unit viscosity, $\gamma = 1$, and when m is small), it is well described by the first-order ODE, $\dot{r}_{ij} = -\frac{\partial V(r,l)}{\partial r_{ij}}$. For optimisation, the change in the separation, r_{ij}, between two masses i and j, is given by the derivative of potential energy, V (a function of all states, $r = \left\{r_{ij}\right\}$, and all natural lengths, $l = \left\{l_{ij}\right\}$), wrt r_{ij}. For learning, we show that the change in the natural length of a spring, l_{ij}, between two masses i and j is given by the derivative of potential energy, V, wrt l_{ij}. Thus, in physical optimisation, the states relax to the parameters, and in complement to this, in physical learning, the parameters accommodate to the states. Optimisation is exemplified by a ball that changes position (blue arrow), rolling down a landscape with a fixed shape, and learning is exemplified by a landscape that changes shape (blue arrows), giving way under the weight of the ball with a fixed position. Usually, these two effects are studied in isolation, without feedback.

To understand the potential adaptive properties of physical systems further, let us now begin to relate this kind of learning to energy minimisation and optimisation. Consider the dynamics of a physical system described by the local minimisation of an energy function. A simple example is a ball moving downhill, a fluid filling a mould, or perhaps a bag of sticks yielding and rearranging to gravity or some other forcing, finding a configuration that locally resolves the physical frustration that exists between components, as one stick pushes another into new positions. (A biological example would be a tissue yielding under mechanical stress, or a chemical reaction network finding equilibrium, among many others). Hopfield and Tank [42] illustrate how this can be interpreted as an optimisation process that can solve non-trivial optimisation problems. That is, if the interactions between components of the system correspond to the constraints of a problem, then the natural behaviour of those components is to change state in the direction that relaxes. That is, to change in the direction that reduces any stresses caused by those interactions, thereby finding states that resolve the frustrations between state variables caused by violations of those constraints. Thus, a physical system, by following the local gradient of its energy function, moves toward better solutions for the constraint problem (Figure 1A). This proceeds until a local minimum is reached in the energy landscape, corresponding to a locally optimal solution to the system of problem constraints [42].

The physical learning examples described above can now be understood as a special case of this very fundamental energy-minimisation principle. That is, learning can be understood as the local optimisation of model parameters to data, and this can be implemented

in a physical system when the internal organisation of a system gives way under stress. The difference with physical optimisation is that in a simple optimisation process, we consider how the state of the system accommodates to the landscape of the problem, whereas in a physical learning system, we consider the reverse, i.e., how the landscape of the model accommodates to the state of the data. Optimisation is a process provided by a dynamical system described by some state variables (solution variables) and some parameters describing the interactions between them (the problem parameters). Learning is a process provided by a dynamical system described by some data and some variables describing the relationships between them (the model parameters). In both cases, the energy function can be the same kind of function—a product of state variables and interaction variables. Learning and optimisation are thus complementary processes (Figure 1); the variables that change in one case are the variables that do not change (the system parameters) in the other and in both cases, the dynamical behaviour is simply local energy minimisation—either directly in the (ordinary) state variables for an optimisation process or in the interaction variables for a learning process (Figure 1).

1.3. Induction and Deduction

Since learning and optimisation are both described by local gradient-following processes, simply applied to different subsets of variables within a dynamical system, there is a level of abstraction where they can be understood to be examples of the same principle; however, it is a mistake to believe there is no conceptual difference between them. A crucial distinction is that optimisation involves a set of state variables (external or observable), and learning involves the interactions between them (internal or unobservable). The latter can be understood as *second-order* variables that control the interactions between *ordinary* (first-order) state variables. In Figure 1, the relational aspect of the problem/model is represented by a network of interactions that define the shape of the landscape (shown in the style of Waddington's 'epigenetic landscape' determining the dynamics of the developmental process [43]). (Clay, in contrast, might be represented by vertical ties that connect straight down from each point on the surface to a corresponding point on the base).

The conceptual import of this mechanical distinction is that learning systems intrinsically involve *induction*. Induction is the process of inferring general rules from specific examples [44–46]. For example, *this swan is white, that swan is white* → *all swans are white*. Unlike deductive inference (which can draw specific conclusions from general rules), we immediately notice that inductive conclusions are not logically valid (you cannot conclude with certainty that *all* swans are white if you have not observed all swans). Inductive conclusions are not deductively supported by the evidence—they go beyond the data. Induction is nonetheless necessary for learning with generalisation; the ability to perform well on previously unseen cases or, relatedly, the ability to generate novel examples from the model. Generalisation necessarily requires inferences that go beyond that which is supported by the data (by definition), and without generalisation, we have only memory, not genuine learning.

Whilst induction in learning systems can take many forms, the principle is simple (We are not intending to advance or augment a general understanding of induction or inductive inference in this paper; 'natural induction' merely aims to acknowledge the importance of inferences that are under-determined by observations, and their necessity for generalisation and adaptation, in the effect we are modelling). In particular, although the learning process results in a particular model (e.g., "all swans are white") there may be many models that are equally compatible with the data (e.g., "all animals are white", "the first two swans are white and all others are pink", or "swans are always one colour"). In a physical system, this means that there are potentially many energy landscapes that are consistent with the pattern of forcing the system has experienced in the past. This means that its internal organisation is necessarily *under-determined* by its experience (just as general conclusions are not deductively entailed by specific observations). Nonetheless, when the model deforms, a particular landscape is obtained (rather than all the possible

landscapes consistent with the data). It thereby represents a particular conclusion that, although consistent with past experience, is neither assured by past experience nor is it the only conclusion compatible with past experience. This is why the distinction between ordinary state variables and second-order state variables is important to learning, i.e., because the latter are under-determined by observations on the former and thereby afford the possibility of generalisation.

Optimisation of state variables on this model landscape (finding a local minimum) constitutes the generative recall of a pattern from the model. This can include the recall of specific states from past experience but also the production of novel instances via generalisation. How the landscape generalises is under-determined by the data and instead depends on the specific nature of the model space (i.e., internal architecture) [47]. The under-determination of the model and the consequently specific nature of the distribution generated from the model is what makes the optimisation of model parameters in physical learning different from the optimisation of (ordinary) external or observable state variables in physical optimisation. Whilst the increments to the model parameters are also determined by the current data point and the current model parameters, the latter are internal variables (general rules) that are not directly controlled by the application of the external forcing (specific instances). These internal variables may depend in principle on the entire history of experiences and may involve subtle symmetry-breaking dynamics that involve the whole system and dominate the influence of its initial conditions [47].

To further clarify, the difference between learning and optimisation tends to zero only if each internal variable has a one-to-one correspondence with an external state variable (like the bed of clay); in this case, the model space is degenerate precisely because the model parameters cannot represent any interactions or associative relationships between features of the data. It is this non-relational property that prevents clay from representing underlying structural regularities in the data, and this prevents it from being able to generalize (Note that evolution by natural selection is often characterised in this way, i.e., as independent selective coefficients acting on a vector of independent alleles or non-pleiotropic traits. In this case, selection is not capable of inducing a generalised model of past selective conditions; however, when selection acts on the parameters of a developmental process, with complex pleiotropic interactions, it is possible to store and recall multiple fit phenotypes in a single genotype and for generalisation in this model space to produce novel phenotypes from the same class [35,48]). All the examples of physical learning therefore involve re-arrangements to internal structural variables that are changed by the pattern of forcing applied but under-determined by the true causal processes that produced the patterns and correlations in that forcing.

1.4. What Is Adaptation?

The quote from Darwin at the beginning of the Introduction Section draws our attention to the exquisite and particular organisation of relationships between the parts of a system. (We would not use the word "perfected" in modern thinking, but there has to be something special about the arrangement of parts).

A system that is deformed or distorted by the forcing applied to it, or simply relaxes under the stresses it experiences, is not (normally) what we mean by adaptation (regardless of whether that deformation involves observable variables or internal structure). Even learning a function or a generalised class of patterns, though it might be analogous to simple forms of cognitive learning, is not the same sense of adaptation that we seem to refer to in biological evolution. Accordingly, the physical optimisation and physical learning examples above have not been claimed to be new sources of adaptation, although learning principles can deepen our understanding of adaptation by natural selection [34,49]. Thus, what do we really mean by adaptation and how does it relate to learning and optimisation?

Definitions of adaptation and adaptive organisation that are tied to natural selection, organismic reproduction, or Darwinian fitness are not useful for answering this question. The 'appearance of design' concept [4] is closer to what we need because it is not necessarily

tied to Darwinian processes, but it is not easily quantifiable. In fact, defining exactly what we mean by adaptation is an open problem in biology. Note that Darwinian fitness, the number of offspring produced, (or inclusive fitness, for that matter), is part of an explanation of biological adaptation, not the explanandum itself (and besides, the relationship between natural selection and maximisation of these quantities is notoriously fraught [50,51]). It is not a high number of offspring that needs explaining, and believing that the 'goodness of fit' between an organism and environment is defined as reproductive output fails to engage with the special properties of the biological organisations that facilitate this. To assess whether there can be any sources of adaptation other than natural selection, it is necessary to let go of the idea that survival and reproduction are the ultimate assessors of adaptation. Whilst these are undeniably important to biological organisms, they are part of the Darwinian explanation for how adaptation happens. In contrast, the goodness of fit between an organism and its environment and the appearance of design are much more general notions and more conceptually apt to understanding what adaptation is. However, the former is too easily satisfied in a physical sense (e.g., an imprint in clay) and the latter is rather subjective.

How can we make a general notion of adaptation more objective and quantifiable? Our approach is to define adaptation as a process that provides an optimization or constraint-solving capability that is non-trivial, i.e., superior to a local hill-climber, and more specifically, one that improves its problem-solving capability with experience, i.e., learns to find better solutions (without assuming that it finds optimal or perfect solutions). In evolutionary computation, the adaptive capability of natural selection is demonstrated through its ability to find good solutions to optimisation problems [52]. Despite the controversy attached to viewing adaptation as a problem-solving process, we find it useful to conceptualise adaptation in this way [53] A less value-laden conception of evolution as a dynamical process, without a problem to be solved and without a pre-existing niche to be occupied, is acknowledged [54,55]. However, if adaptation is construed as merely 'whatever happens as a result of natural selection', for example, then it remains tied to natural selection. Adaptation needs to address a design-like property (independent of natural selection), and problem-solving is one way to characterise this. At the least, a process that can provide a non-trivial problem-solving competency is a high bar for assessing adaptation. To formalise this, we begin with a notion of searching for points in some configuration space that optimises some quantity, i.e., a process of function optimisation defined over some given space of variables [50,53] (Figure 1A). An optimisation ability does not require that a process finds the globally optimal (best possible) configuration [50]; this would be too strict. Natural selection, for example, does not provide optimal solutions. An adaptation built by natural selection need only be a good solution to a problem, not an optimal one. Conversely, neither do we want to define adaptation in a manner that is trivially satisfied [47]. If we adopt a very simple notion of problem-solving, such as that provided by a hill-climbing or local gradient ascent process, that merely finds a local optimum of an objective function, it would be trivially satisfied by a physical system as discussed—and therefore not as satisfying as a concept of adaptation. The behaviour of any physical system that can be described by the local minimisation of an energy function can be interpreted as an optimisation process in the limited sense that it finds locally optimal solutions to its implicit energy-minimisation 'problem' [42] (Figure 1A).

It is not very useful to adopt a definition of adaptation that is satisfied by a literal ball rolling down a literal hill (even if it is functionally equivalent to the process described by natural selection, see Discussion). Accordingly, the local optimisation behaviour of physical systems is not sufficient to constitute adaptation. So, instead, we seek a non-trivial problem-solving competency—not necessarily optimal, but not trivial, either [47]. Physical systems can do better than the trivial case. Simulated annealing is a famous example of a computational optimisation process that finds solutions better than local optima [56], and actual annealing (e.g., in a cooling metal) occurs spontaneously. This shows that the spectrum of possibilities for natural adaptation is non-empty—but still, we aspire to more

than cooling metal as a natural example of adaptation. In what follows, we show that the ability of a physical system to learn offers the possibility of a physical system that *learns to solve problems better with experience*. Put differently, we classify the ordinary (first-order) local energy minimisation behaviour of physical systems as trivial, but optimisation that improves with experience or *second-order optimisation* is algorithmically interesting.

In the previous examples of physical optimisation, the quality of a solution improves over a given state trajectory, but it only reaches a local optimum and its ability to find solutions of good quality does not change with experience. Lower-energy states, constituting better solutions to the frustrated state variables in the system, may exist but are not obtained. The optimisation ability does not change over time—restarting the system results in the same locally optimal outcome if it restarts from the same initial position. In the examples of physical learning, a system incrementally improves the fit of an (implicit) internal model to a conventional learning task such as classification or representing a function, e.g., [29]. A problem-solving ability is not demonstrated (except in the same sense of finding a locally optimal fit of model parameters to the data). Both these behaviours (learning and optimisation) are very natural and do not involve particularly limiting assumptions about the system. In both cases the behaviour is determined by local energy minimisation of the same energy function—the only difference is which variables give way under stress and which are held constant. In optimisation, the state gives way to the problem, and in learning, the model gives way to the data (Figure 1); however, so far we have considered these two processes in isolation: the outputs of the learning process do not affect the optimisation landscape, and the outputs of the optimisation process do not affect the learning data. This limiting assumption means that learning responds to data that is given by external conditions (fixed training data), and optimisation responds to a landscape that is given by the problem definition (fixed problem). These effects have been studied separately, but in general dynamical systems both these effects will happen at the same time. What happens to the structure and dynamics of the system in the general case where there is feedback between the two is much more interesting.

1.5. Change in State and Change in Interaction Structure

What happens when both the state and the structure of a system are variable? If the structure of the system determines the system's behaviour (i.e., its state dynamics), then changing the structure has the effect of changing the behaviour of the system. But how can a system change *its own* behaviour, let alone 'improve itself'? We are specifically interested in cases where the state influences the change in structure and structure influences the change in the state through dynamical feedback. These kinds of systems are studied under other various names. "Meta-dynamical systems" are a general concept that describes systems where the parameters of a dynamical system are, in fact, slow-changing variables [57]. Examples discussed include evolving gene-regulation networks, neural networks and immune systems. Note that a meta-dynamical system is, from the perspective where all dynamical variables are lumped together, just a dynamical system that behaves. But when these variables are separated (into state variables and structural variables), it becomes reasonable to describe it as a system that 'changes its behaviour', i.e., has a behaviour (on a fast timescale) and it changes its own behaviour (on a slow timescale). This is not an arbitrary separation of variables but depends on timescales and also what is observable and what is internal or (not observable)—or, for our purposes, which variables are directly forced and which ones are indirectly forced or induced. In "self-organised" systems, the interaction between state and structure is of various kinds and the emphasis is on the spontaneous organisation of system components, or 'order for free' [58]. In "adaptive networks", the dynamics of behaviour on a network (e.g., utility-based replication of game strategies) are influenced by the network topology and, reciprocally, changes to the network topology are influenced by the behaviour on the network [59]. This can result in structures that alter state dynamics in predictable ways, e.g., resulting in an increase in equilibrium levels of cooperation [60].

Ashby [61] also studied systems with dynamical state variables as well as dynamical interactions, or wiring parameters, to describe mechanisms of "ultra-stability" and homeostasis. Random re-organisation of the wiring is triggered by stress until such time as this brings the essential variables back into their viable range. "Adaptive improvisation", "sequential selection", and the principle of "least rattling" [32,62,63] develop and extend similar ideas; i.e., keep re-organising until an organisation that does not trigger further re-organisation is obtained. These works demonstrate principles similar to those shown here insomuch as there is a dynamical feedback between observed state variables and (internal) interaction variables that control the relationships between the state variables. They demonstrate the ability of such dynamics to satisfice criteria of order and stability but they do not show a second-order optimisation ability or the ability to find solutions systematically better than local optima.

Since we know that ordinary physical systems can exhibit both optimisation behaviour (given variable state) and learning behaviour (given variable structure), can a system with both flexible state and flexible structure exhibit an increase in problem-solving competency spontaneously?

1.6. Adaptation by Natural Induction (a Physical Model)

Here we examine systems which have the conditions for both physical optimisation and physical learning—systems where both the state and the structure give way under stress, and these changes feedback on each other (Figure 2). We show that a dynamical system described by a network of viscoelastic connections can spontaneously learn to solve an optimisation problem better with experience under natural conditions. The conditions for this effect are that the connections are viscoelastic, meaning that they give way slightly under stress, and that the state configuration of the system is occasionally disturbed, e.g., subjected to shocks or perturbations. For example, a network of masses connected by springs and subject to disturbances meets these conditions if the springs are not perfectly elastic, i.e., if the springs are slightly plastic, as all physical springs are. The specific self-conditioned organisation of the springs that obtains constitutes an adaptive organisation insomuch as it causes the system to generate state configurations that are particularly high-quality solutions to difficult combinatorial optimisation problems. The quality of the solutions discovered can be extremely rare compared to those found by local gradient methods in the same problem, even if such first-order optimisation has repeated attempts.

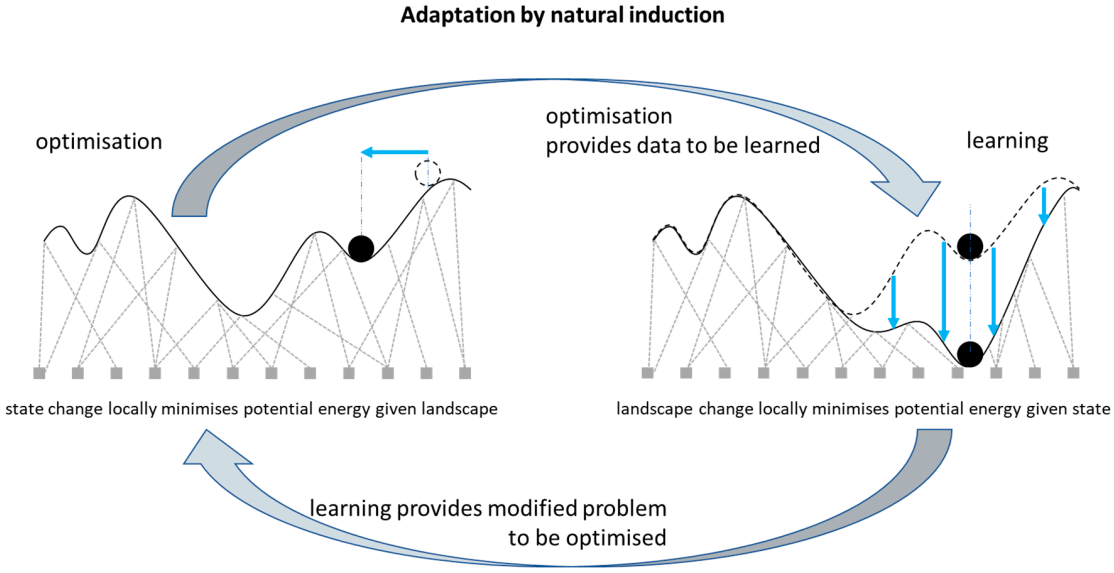

Figure 2. Adaptation by natural induction results from positive feedback between optimisation and learning. Adaptation by natural induction occurs naturally in physical systems described by a network of viscoelastic connections (subject to disturbances). This involves dynamical feedback:

energy minimisation on state variables given problem parameters, and energy minimisation on those parameters given the states visited. The initial shape of the problem landscape causes changes in state (optimisation, left). This state at any point in time provides a data point that causes (relatively slow) changes in the problem parameters (learning, right). This provides a modified set of problem parameters to optimise, and so on. Disturbances (randomising the state variables) occur infrequently enough that the system spends most of its time at energy minima (locally optimal solutions) and frequently enough that the deformation of the parameters occurs over a distribution of such optima.

We call the spontaneous feedback between learning and optimisation (occurring without selection or design) *natural induction*, and the ability to improve optimisation ability with experience (i.e., finding solutions better than local optima) *adaptation by natural induction* to reflect its close association with inductive learning processes and to emphasise its complementarity with natural selection. (There is some possibility that 'abduction' or 'transduction' are more accurate terms but we feel that induction captures the basic flavour sufficiently). Just as Darwin's natural selection is a process that requires no literal intelligent oversight, but mirrors the intentional selection of the breeder, so our natural induction requires no literal intelligent oversight, but nonetheless mirrors the way in which an intelligent agent can inductively model the behaviour of a system to discover solutions that resolve its constraints.

The underlying algorithmic concept is as follows: In each trajectory of an optimisation process, the state changes until it reaches a local equilibrium—corresponding to a locally optimal solution. This optimising process is subject to repeated shocks or perturbations, which effectively randomise the state variables (but not the interaction variables). Each local equilibrium state (achieved in between these disturbances) is thus a point drawn from a distribution of such locally optimal solutions. Meanwhile, the problem parameters give way slightly in a manner that accommodates to the current state. As the system spends time at each locally optimal solution visited, the problem parameters model this distribution of state configurations; that is, the problem parameters are model parameters and the distribution of locally optimal solutions is the 'training data'. The result is that the system models the outcomes of its own behaviour (hence a "self-modelling" dynamical system [40,64]). In so doing, this changes the dynamics of the optimising system, making it more likely to visit solutions that have already been visited. This changes the distribution of equilibria discovered, which changes the model/problem again, and so on (Figure 2). This affects dynamical feedback between learning and optimisation—the state variables give way in a manner that accommodates to the current system parameters (optimisation), and meanwhile, the system parameters give way in a manner that accommodates to the current state variables (learning). This results in an increased optimisation ability, in particular in the ability to discover solutions of exceptionally rare quality [40]. In principle, natural induction may occur in any dynamical system where internal structures give way or reorganise slightly under the stress or forcing a system experiences; the worked example illustrated in this paper uses a network of viscoelastic connections, i.e., springs, that can stand in for many types of network connections in different (biological and non-biological) substrates. (The extent to which this results in effective adaptation, and optimisation better than a local hill-climber, will depend on the suitability of the inductive bias implicit in the system architecture, i.e., whether the causal geometry of the system doing the learning is like that of the environment it is learning about—See Discussion).

In other settings, the underlying 'learning to optimise' principles involved in this effect have been demonstrated in other kinds of self-modelling dynamical systems, including neural networks [40,41,65], gene-regulation networks [35,48,66], social networks [67,68], and ecological networks [11,65]. In all previous cases, however, there was, at some level or another, either a mandated learning process [40], a selection process [11,35], or a utility-maximisation process [67,68]. Uniquely, in this paper, we show that the same adaptive principle is exhibited naturally and spontaneously in a physical system—and it is therefore an adaptive process independent of natural selection. Crucially, there is no differential

survival or reproduction involved either for the network as a whole, its component parts, or the connections between them; i.e., we demonstrate that natural induction occurs without natural selection (natural selection is neither involved at 'run time', i.e., as the adaptation occurs, nor in establishing the initial conditions, the 'set-up', or the construction of the system).

In the following experiments, we detail some example dynamical systems (in two scenarios) each described by a system of masses connected by springs. The first scenario is generic—a network of viscoelastic connections. This illustrates a system that gets better at minimising its own energy function (by changing the organisation of its interactions). The second scenario describes a system that finds good solutions to an independent system of problem parameters (or an external 'environment') via changes to its organisation (by changing the coupled dynamics of the system and the environment). This is illustrated for both continuous problem variables (Scenario 2a) and binary problem variables (Scenario 2b). In all three cases, the mechanism of adaptation by natural induction is the same. We present numerical simulations of these systems to illustrate their optimisation capabilities and demonstrate that they constitute adaptation (by our stringent problem-solving criterion). To conclude, we discuss how natural induction differs from natural selection, and point to some of the important implications.

Our purpose in using a physical model built from masses and springs is not because we are interested in a mechanism depending on these specific physical components but as an illustration of a general dynamical property in networks of various kinds (biological and non-biological). The purpose of describing this as a system in this physical way is to be as generic as possible and to ensure we are not using assumptions that derive from biological systems with properties that depend on prior natural selection. Darwin's description of evolution by natural selection, though embedded in detailed biological observations [1], was also what we might now describe as an *algorithm*, independent of this biological substrate or any particular instantiation [3,4,52,69]. In the same way, we aim to use the models that follow as an illustration of an adaptive algorithm that may be instantiated in multiple different substrates (specifically, dynamical systems described by networks of viscoelastic connections, including biological networks of many kinds).

The relevant conceptual territory of this paper thus lies at the interface of biological thinking, physical systems, and general dynamical systems. Like the other work in this area [32,58,61], we aim to make a small contribution to triangulating the space of possibilities for spontaneous adaptation in natural systems—the aim of this paper is not to displace existing theories but to illustrate additional possibilities and potentialities. This places the work naturally at the cusp of physics and biology [58], and contributes to understanding their interaction [70]. It is now understood that evolution depends on the interaction of natural selection with the self-organisation properties of non-living matter [70–72]. Natural induction adds to this the possibility that dynamical feedbacks among unorganised components are capable of non-trivial adaptation absent natural selection.

This concept space, in particular the use of energy landscapes and attractors in biological processes, has long been a part of evolutionary thinking [43,73]. Unlike much of this prior work, we also use the theoretical framing of computer science and optimisation to provide a rigorous test for non-trivial adaptation (second-order optimisation). This helps to distinguish mechanisms that provide adaptation from other types of complex systems phenomena. Whereas other authors have sought to understand how natural selection might be implemented in primitive physical systems [8], we aim to demonstrate that primitive physical systems correspond to different principles of adaptation familiar to machine learning [74], in particular, adaptation that is not natural selection, nor a weakened form of natural selection.

2. Methods

In the models that follow, we assume that the initial energy function of a dynamical system defines an objective function, i.e., the quality of solutions to a problem [42]. Initially,

then, the system finds locally optimal solutions to its own 'energy-minimisation problem' [42]. However, we define adaptation as doing better than this, i.e., finding solutions better than locally optimal solutions. Logically, this requires that in order for a physical system to go somewhere different from the nearest locally optimal solution in configuration space, the energy function (describing the dynamics of the system) must be different from the objective function (describing the quality of solutions to the problem) [47]. In the models that follow, the dynamics of the system change over time through physical learning, causing the behaviour of the system later in time to deviate from its initial behaviour. To exhibit adaptation, the new behaviour (defined by a new energy function) must cause the system to arrive at configurations that are superior solutions, i.e., better than locally optimal solutions in the original energy function (objective function). A learning system can, for example, accumulate information from experience (i.e., from multiple samples of points in an objective function) about regularities or underlying structure in the problem space, and use this to optimise better with experience [40]. This experience is not provided by an external 'teacher', or by any source with privileged information about the problem structure or the location of good solutions, it is 'unsupervised' learning [34]. It is acquired through 'self-modelling' [40], provided by its natural energy minimisation dynamics and the state disturbances, together with the natural generalisation ability of the physical learning process.

We construct a system of N masses constrained in a 2D plane (located at x_i, y_i) connected by Hookean (linear) ideal springs and viscous dampers. Typically, in such systems, pairs of masses are connected by a single spring and damper. With these elements alone the force between each pair of masses is

$$F_{ij} = -k\left(r_{ij} - l^0\right) - \gamma \dot{r_{ij}}$$

where $r_{ij} = \sqrt{\left(x_i - x_j\right)^2 + \left(y_i - y_j\right)^2}$ is the distance between the masses, $\dot{r_{ij}}$ is their relative velocity, k is the spring constant, γ is a globally defined damping constant, and l^0 is the natural length of each spring. In damped systems, we expect all the kinetic energy to dissipate until only elastic potential energy remains, as given by

$$\mathrm{V} = \sum_{ij \in S} \frac{1}{2} \mathrm{k} \left(r_{ij} - l^0\right)^2$$

where $\{S\}$ is a set of springs. However, real springs are not perfectly elastic, they are imperfectly elastic, slightly plastic, or viscoelastic, meaning that they give way when they are stressed (as is familiar when a spring is stretched too far or held in a stretched position for a long time). A simple example is 'creep', where the natural length of a spring increases slightly when it is held under tension or decreases under sustained compression. This slow change in natural length can be modelled by adding a second viscous damper in series with the ideal spring. See Figure 3A, which introduces the properties of *Maxwell material* [75].

Specifically, in this configuration, it can be shown (Appendix A) that the force between pairs of masses is

$$F_{ij} = -k\left(r_{ij} - l_{ij}\right) - \gamma \dot{r_{ij}}$$

where l_{ij} is the natural spring length, $l_{0,}$ plus a viscous deformation of the spring and evolves as

$$\dot{l_{ij}} = \frac{k}{\gamma_m}\left(r_{ij} - l_{ij}\right) \tag{1}$$

and γ_m is the damping constant of the second spring and we have assumed the timescale of this deformation is slower than viscous drag on spring displacements, $\gamma_m \gg \gamma$ (i.e., changes in natural length are slow compared to changes in displacement of the masses). Notice that this deformation of the springs is directional; from Equation (1) we can see that under tension, or compression, the spring will lengthen or shorten, respectively.

This deformation described in Equation (1) can also be derived from considering local minimisation energy with respect to spring length (Appendix A), which concurs with our previous work on Hopfield networks [40]. Specifically, the potential energy is

$$V = \sum_{ij \in S} \frac{1}{2} k (r_{ij} - l_{ij})^2$$

and thus,

$$\dot{l}_{ij} = -\frac{\partial V}{\partial l_{ij}} \equiv \frac{k}{\gamma_m}(r_{ij} - l_{ij});$$

that is, the normal energy-minimisation state dynamics of the system (with perfect springs) describe changes in the positions of the masses given a set of spring parameters, and conversely, the viscoelastic change in the springs is described by the minimisation of the same energy function given the positions of the particles. Thus, the particle positions relax given the frustrations created by the current spring parameters, and the spring parameters (slowly) relax given the frustration created by the current positions of the particles (Figures 1 and 2).

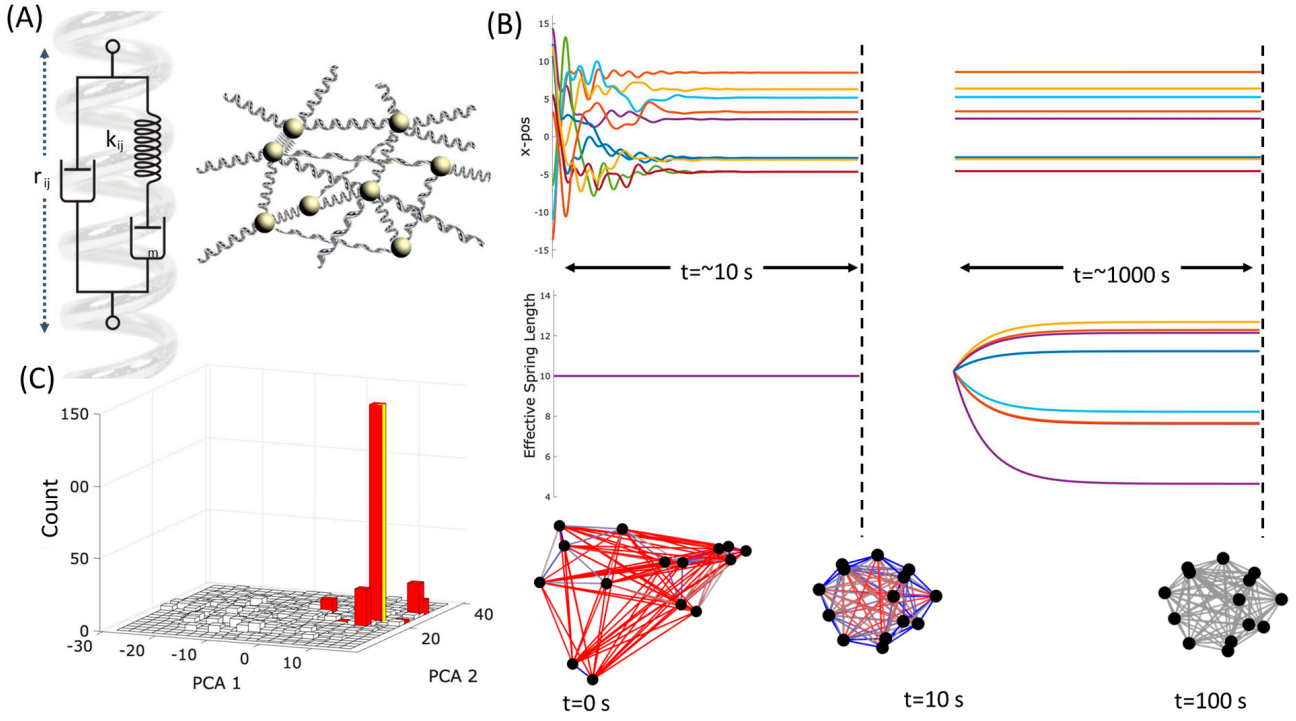

Figure 3. A viscoelastic mass-spring-damper system forms a memory of a previously visited configuration. (**A**) (**Left**) A deformable spring is modelled as an ideal spring in series with a viscous damper (γ_m), which in turn is in parallel with a second viscous damper (γ) (a Maxwell configuration). (**Right**) We simulate a set of frictionless masses in a 2D plane connected by deformable springs. (**B**) (**Left**) The dynamics of the spring positions in the x-direction settle to stable equilibria over a period of ~10 s (**top**); there is no appreciable change in the natural lengths of the spring over this period (**middle**). (**Right**) Over a longer timescale, the positions of the masses are fixed (**top**) but the natural lengths deform (**middle**). (**Bottom row**) The network of connected massed in a 2D plane. Red and blue indicate springs under extension or compression, respectively. The initial condition of the masses in the plane (**left**) over short timescales decays to a stable configuration, which is under frustration (**middle**). Over long timescales, this frustration is released as the natural length deforms (**right**). (**C**) Histograms of the first two principal components of the pairwise difference between spring positions in the equilibrium configurations reached starting from 1000 initial conditions. Before the deformation of springs, the final configurations are widely distributed (white bins). After

the deformation, most initial conditions converge to one of very few configurations (red bins). The most visited configuration aligns with the configuration at which the deformation took place (yellow line). Note: to enable visibility of small values, the tallest red bin, which extends beyond the vertical limit of the plot, is truncated.

As described thus far, the system does nothing very interesting—the fast state dynamics fall to the nearest local optimum and then the slow structural dynamics accommodate to this state, resulting in a new landscape with one memory—at the position of the arbitrary local optimum that was initially found. This therefore affords nothing more than local optimisation (with memory); however, the system comes alive when we introduce a 'pulse'—repeated disturbances applied to the state variables. This makes the feedback between physical optimisation and physical learning much more interesting. The disturbances effect a sort of ebb and flow or push and pull dynamic; i.e., the weights push the states, then the states push the weights, and so on, as per the two directions of influence (arcing arrows) in Figure 2. This does not require a mechanism to explicitly interleave these two directions; it is achieved simply by letting the weights change slowly all the time and subjecting the states to occasional disturbances. When the states are in a high-energy configuration, they are pushed around by the weights until they reach a local equilibrium and are unable to change any further, and then, whilst they are resting at a local optimum for a long time, the states slowly push on the weights, and then the next disturbance, and so on. Thus, the two directions of influence are implicitly taking turns and the result of the feedback is not simply a single equilibrium state but an organised change in the distribution of a dynamical behaviour.

To the basic setup, we introduce occasional disturbances to the state variables (particle positions, not spring parameters). We use three scenarios that differ in the initial arrangement of the system (the ways that particles can move and the arrangement of springs connecting them) to illustrate adaptation by natural induction applied to (1) solving an intrinsic energy minimisation problem and (2) solving an external optimisation problem (over continuous and discrete variables, 2a and 2b, respectively).

3. Experiments and Results

3.1. Adaptation by Natural Induction—Generic Case

We start by demonstrating the dynamics of a small system and then the ability of such a system to retain information about the configuration of past states. For illustrative purposes, we initialise a small system with $N = 15$ masses and connect pairs of masses with springs with probability $p_c = 90\%$, with uniform natural length $l^0 = 10$, and uniform spring constants $k = 10$.

Over a relatively short timescale (≈10 s), the masses transiently oscillate and settle to an equilibrium (Figure 3B, upper left); however, note that at this equilibrium, not all the springs achieve their natural lengths and some springs are under compressive or extensive forces (Figure 3B, bottom middle network, blue and red, respectively), and the system is frustrated. Like a neural network with symmetric connections, a physical system of masses connected by springs has only fixed-point attractors (the local minima of the energy function) [37]. In general, such a system may have many such attractors, attaining a variety of equilibria with different amounts of frustration (total energy at equilibrium, V).

Over much longer timescales (≈100 s) the spring lengths change slowly, see Figure 3B (right), and the effective natural length, $l_{ij}^{eff} = l_{ij}$, of each spring deforms in the direction that reduces this tension, see Figure 3B (bottom right network). Since the spring lengths are now different, the state dynamics will be altered. But in what way? How will the attractors of the system with the new spring values relate to those of the original system?

To examine how the plastic deformation of springs has affected the stable configurations of the system, we plot the distribution of final equilibria (sampled from the same distribution of random starting points) before and after the springs deform. In Figure 3C, we plot a histogram of the equilibrium states visited from different random initial condi-

tions. To plot this in two dimensions, we use the first two principal components of the distances between pairs of masses at equilibrium. The equilibrium reached in any one trajectory depends on the initial conditions (i.e., the starting point of the trajectory). Before the springs deform, we observe a broad distribution of stable equilibrium configurations (Figure 3C, blue distribution). One such endpoint, chosen arbitrarily, is indicated by the yellow column. The system is allowed to rest at this configuration whilst the springs slowly deform. After the springs deform, the system has a new energy function, and its trajectories and equilibria are different (even when using the same initial conditions). We see that the distribution is much more narrowly concentrated (Figure 3C, red distribution) and from most initial conditions (>98%) the masses converge to the yellow equilibrium, i.e., the configuration the system was in when the springs deformed.

Note that this is not simply a system finding stability (in particles and springs), but a system changing its dynamics (by altering its slow variables) in such a way that it can recreate a particular pattern of its fast variables. This is explained by noting that the effect of spring deformation at any particular state is twofold: it lowers the energy of this particular state (because the springs accommodate to this state and thus resist it less), but more importantly, it increases the size of the basin of attraction for this state (i.e., increases the number of random initial conditions for trajectories that arrive there). This is functionally analogous to the formation of memory in, for example, a Hopfield network under Hebbian adjustment to connections [37,40,76]. The changes to the springs thus change the relative size of the attractor basins in the energy function (not just their depth) such that this configuration is visited much more often than others. We interpret this as a memory of a past state configuration—or we might say that the slow system variables have taken an associative 'imprint' of its own past state (hence a 'self-modelling' dynamical system).

We now examine what happens when springs deform over a distribution of equilibria rather than a single arbitrary equilibrium. The sampling of random initial conditions is provided by the random disturbances, i.e., shocks or perturbations that randomise the positions of the particles. If a system is subject to disturbances of this kind, with intervals of relaxation in between, then the system will visit many different equilibria in proportion to the size of its attractor basins. We assume the time for the fast variables to reach equilibrium is much less than the interval between disturbances, which in turn is much less than the time for the slow variables (spring lengths) to reach equilibrium. Under these conditions, the system spends most of its time at local minima in the energy function but visits many such equilibria on the timescale where the springs deform.

The effect of this is interesting when the number of different equilibria in the initial system dynamics is large. We simulate a system of $N = 300$ with lower connectivity $p_c = 50\%$, with periodic disturbances, see Figure 4A. Since the system spends most of its time at equilibrium configurations (not on the transients), the computation is approximated by running the system until it settles to equilibrium and then updating the spring lengths once for each equilibrium visited (rather than updating them at every time step on the transients as well). This is a good approximation in the limit of a slow deformation rate, with sufficiently long periods spent close to equilibrium.

Although the springs potentially form a memory of many state configurations, we find the dynamics of the system with the deformed springs consistently settle to one of very few equilibria as before (Figure 4B). This is expected because the system is forming a memory of its own behaviour with positive feedback—the more it visits a state, the more that state is memorised, and the more that state is sampled in future, and so on. However, is there anything special about the particular equilibrium (or small set of equilibria) it converges to? Are they good-quality solutions or arbitrary configurations?

As before, the deformation of the springs has the effect of lowering the energy of state configurations it has visited previously, and widens their basins, making it more likely that it will visit them again. To assess their quality as solutions to the original energy minimisation problem, we need to examine the configurations found with the new system dynamics and report the energy that these configurations had in the *original* system

dynamics. Since the springs now have different lengths, it is possible that the new attractor states may not have been attractor states in the original system. Hence, instead, we find which attractor (of the original system) it is closest to, or more exactly, which equilibrium state it is attracted to; that is, we use the final state configuration found in the new system (with the new energy function) to set the initial conditions of the original system and let this system settle (in the original energy function). (For the purposes of taking these statistics, the spring lengths are unaffected by this side assay).

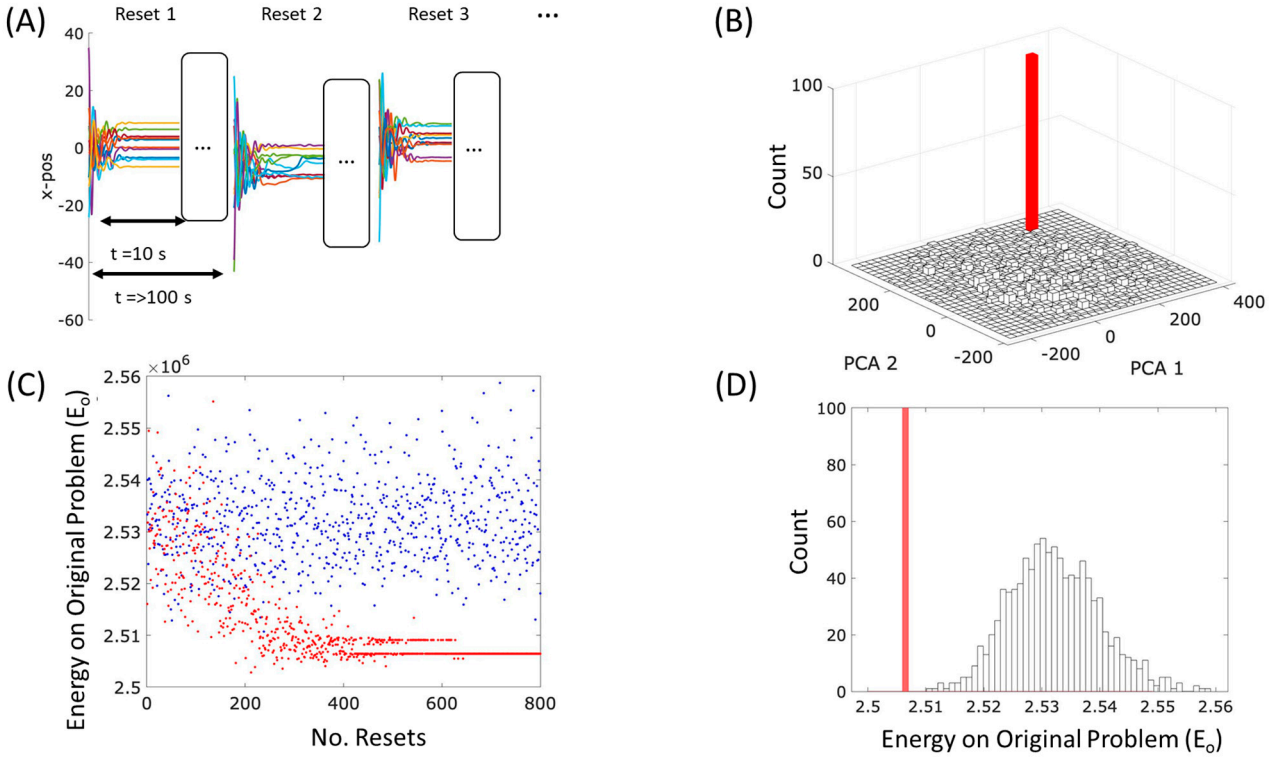

Figure 4. Adaptation by natural induction discovers exceptionally low-energy configurations (Scenario 1). (**A**) The protocol is described in terms of the x-displacement of springs over multiple disturbances (resets). The system is disturbed every ~1000 steps and settles to a stable configuration after which the springs deform. (**B**) Histograms of the principal components of the pairwise difference between spring positions in the equilibrium configurations reached, starting from 1000 random initial states. Before deformation, there is a broad distribution of final configurations (white bins); after the protocol in (**A**) all initial conditions converge to a single stable configuration (note: the tallest red bin is truncated and extends beyond the figure). (**C**) After each reset, we rerun the original spring system initialised at the final learned configuration. The plot shows E_o verses the number of resets. Without deformation, E_o (bule dots) is widely distributed but after many resets, the system consistently finds the basin of attraction of low E_o (red dots). (**D**) The distribution of E_o before (white) and after (red).

We plot the equilibrium energy of the original system (V_o) obtained using this procedure over time, as measured by the number of disturbances, thus tracking how this energy changes as the springs deform (Figure 4C, see blue dots). For reference, we also plot the energy of equilibria in the original system when starting from the same initial conditions i.e., caused by identical disturbances (Figure 4C, red dots). We observe that as the springs deform, the system has to find lower and lower energy configurations, finally converging on a small number of very low energy configurations (Figure 4C, blue dots). In Figure 4D, we plot a histogram of stable equilibria found from random initial conditions before and after spring deformation (resettled with the original spring values), and blue and red histograms, respectively. We note that the new distribution converges on states

that are particularly low energy—lower than any of those found in the original behaviour of the system (from the same distribution of initial conditions).

We stress that this is not because it is finding the same arbitrary configuration and lowering its energy, but because it is finding different configurations. These low-energy configurations were present as attractors in the original system dynamics but their basins of attraction were very small and were thus very rarely found with the original dynamics (indeed, not found at all in the number of samples shown). The new system dynamics has enlarged the attractor basins for these particularly good quality solutions, even though we would not even expect them to be visited, even once, under the original dynamics with this number of samples. How is this possible?

The explanation has two parts. First, in systems built from a large number of pairwise interactions, the lower-energy equilibria tend to have larger basins of attraction (by virtue of limits on the slope of the energy function arising from being the sum of many forces) [40]. It is thus expected that the attractors that are better quality solutions to the energy minimisation problem are visited more often than low-quality solutions (this is essentially why gradient methods, like simulated annealing, have some success in general energy landscapes/objective functions). With the positive feedback on the system dynamics provided by spring deformation, we might expect the system to converge on the largest attractors, and these tend to be the best attractors. By virtue of the fact that this makes a system visit good quality solutions more reliably over time, this is already a type of adaptation by our strict criteria, i.e., a system that learns to optimise better with experience; however, this simple kind of reinforcement learning is not the whole explanation and the results are significantly more interesting than this reasoning suggests. Specifically, the system finds solutions that are lower in energy than any solution found by the original dynamics. This means that the system is not just forming a memory of low-energy configurations it had already visited, it is visiting configurations that are novel (and even lower in energy). This is possible because an associative memory can generalise—it can generate novel patterns from the same class, not just patterns it has been trained on (i.e., already visited) [40]. This is possible because the new spring organisations are not just a memory of past states but an induced model that goes beyond the training data to generate new configurations with similar structural regularities.

To quantify this and explore the robustness of the results to different realisations, we quantify how unlikely the learned energy was to be found in the unlearned system as the number of standard deviations (STDs) between the mean energy of equilibria found in the original system and the minimum final learnt energy. Over 10 runs, we found that 80% had more than 3 STDs from the mean and the median number was 3.4 STDs (we calculated this using the log energy, which is better represented by a normal distribution).

Note that if the original spring parameters represent the constraints of an optimisation problem, then the original state dynamics just find locally optimal solutions to that problem (in proportion to how likely they are to be sampled) [42]. In contrast, the state dynamics of the new system, deformed under these conditions, finds configurations that are much better than any of the locally optimal solutions to the problem found in the same number of samples from the same initial conditions; in other words, through a natural stress-reduction process arising spontaneously from the deformation of the springs, the system learns to solve the problem better with experience.

3.2. Solving 'External' Problems via Natural Induction

In Scenario 1, the effect of adaptation by natural induction is demonstrated under general and natural conditions, i.e., for any network of viscoelastic connections. Scenario 2 examines cases that demonstrate the same effect under different conditions, changing the way the problem is encoded into the system and the way that a solution is 'read off'. In all cases, the mechanism of adaptation by natural induction that finds the solutions works in the same way.

In Scenario 1, the springs that represent the problem and the springs that deform are the same springs (we compare the system behaviour before and after deformation using a copy of the original springs). It is a system that gets better at minimising its own energy function by changing the organisation of its interactions. (The quality of a solution, therefore, cannot be directly read off from the final state configuration in Scenario 1 because the original spring parameter values that represent the problem are lost as the springs change. This is why it was necessary to take each of the new configurations it finds and assess their quality by finding the nearest attractor in the original system, as described above). Can natural induction solve a problem that is external to the learning network? That is, a scenario more analogous to the conventional (though artificial) separation of organism and environment? In this second scenario, we examine a system that finds good solutions to an independent system of problem parameters or an external 'environment', via changes to its organisation (by changing the coupled dynamics of the system and the environment).

In Scenario 2, the system has two different types of springs—springs that are deformable and springs that are not. The non-plastic springs represent the problem constraints (or external environment), and the springs that are plastic affect inductive learning (the adaptive system). Thus, the plastic springs modify the total system dynamics in the same way as before, finding configurations that solve the constraints defined by the problem springs. Eventually, the plastic springs 'melt away' completely, leaving only the unchanged problem springs to determine the final attractor state. In this way, this scenario uses two types of springs and demonstrates the same increase in problem-solving competence without needing to do the side assay where we transferred the new state back to the old spring values and re-settled the system. This setup also offers an interpretation where the problem springs represent a problem external to the system (in the environment) and the deformable springs represent an adaptive system (organism or agent). These are coupled together through the shared state variables (an interface or phenotype) to produce combined dynamics (organism and environment in interaction). The effect of connecting these two networks together is that the environmental dynamics induce a model of itself into the agent's internal structure (the deformable springs), and the generalised nature of this model's dynamics has the effect of steering or chaperoning the interface variables into increasingly superior solution states (until leaving them in a solution state as their influence attenuates).

To model this, we introduce a second type of spring and distinguish between problem springs (P-springs) and learning springs (L-springs). Specifically, P-springs are perfectly elastic and do not deform, i.e., see Figure 5A, left. These springs define the energy landscape of the problem to be solved, i.e., we interpret the problem as one to find the lowest energy configuration of the P-springs. The L-springs have the Maxwell configuration previously described (see Figure 5A, right). These springs are highly connected (almost fully), weaker, and uniform in natural length and spring constant (as before). The intuition is that the energy landscape of the initial system (with P-springs and initial values of L-springs) will still be dominated by the stronger P-springs. To confirm this, we plot histograms of the energies of equilibria found when only P-springs are present against and those found by first running on both P-springs plus L-springs but then removing L-springs and letting it resettle (as per the previous assay). These histograms were almost identical (see Appendix B). We can thus think of L-springs as the addition of an initially neutral field of plastic material that (initially!) does not alter the problem or the physical optimisation dynamics, and also will eventually melt away to nothing, leaving only the bare problem; But their transient influence in between leaves the system in a special state configuration.

In the following experiments, we follow the previous protocol of periodic disturbances and to this, we add an additional long settling period at the end of the experiment (Figure 5B). During this final period, the effective natural length of the L-springs converges to the distance between pairs of masses at the equilibrium of P-springs only, i.e., the L-springs exert no forces at equilibrium.

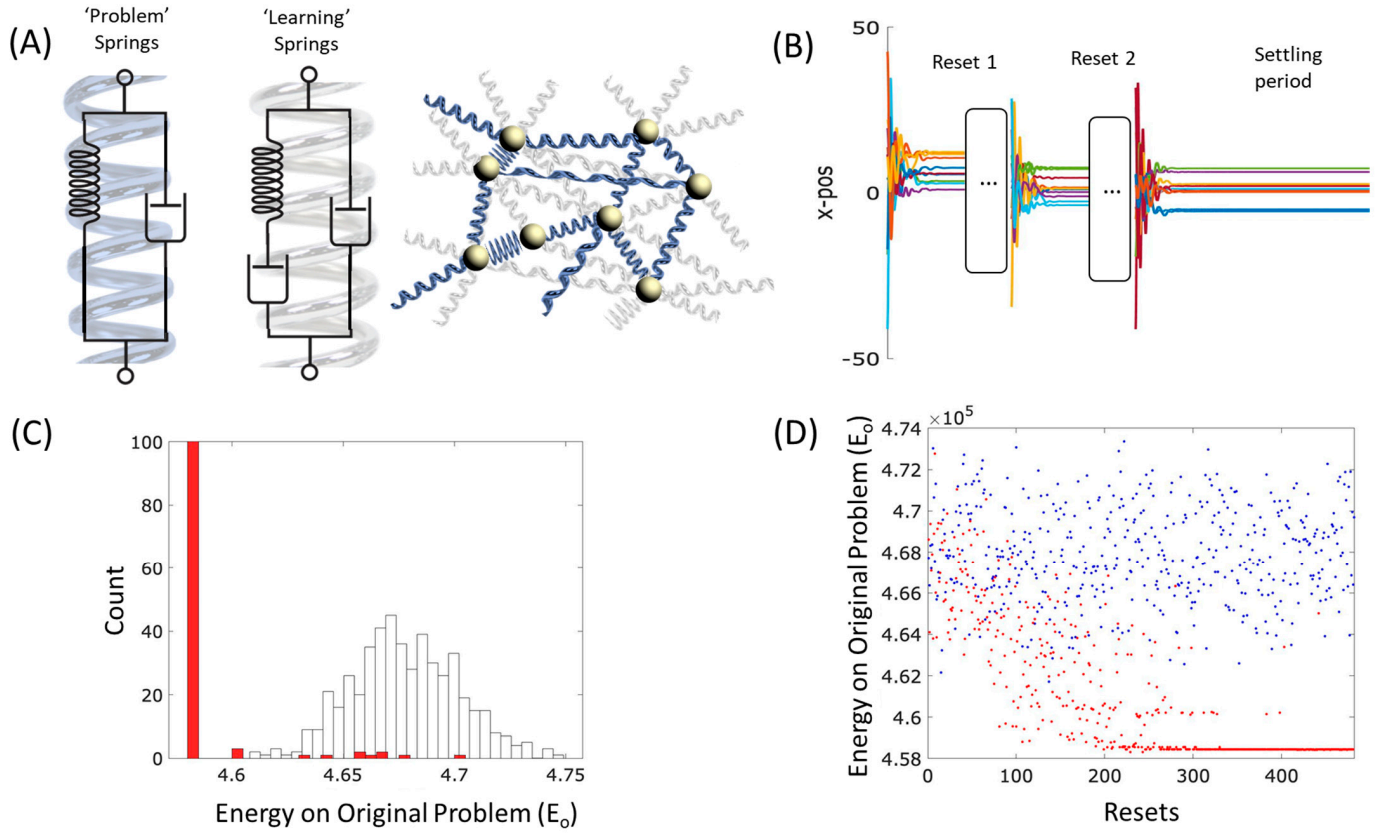

Figure 5. A viscoelastic network discovers low energy configurations of a network of 'problem' springs (Scenario 2). (**A**) A material comprising a set of highly connected viscoelastic springs (leftmost spring configuration and blue) and a relatively sparse set of perfectly elastic springs (middle spring configuration and grey). (**B**) The same protocol is described in Figure 2 but the system is allowed to settle after disturbances stop. After this period, all tension is released in the viscoelastic springs and they no longer contribute to the energy. (**C**) The distribution of energies starting from random initial conditions of the problem springs only (white) and after the settling period (red). (**D**) Energy after the settling period versus the number of resets starting from random initial conditions before deformation (blue) and as deformation progresses (red).

We examine two different kinds of external problems; specifically, a problem defined over continuous states (Scenario 2a) and a combinatorial optimisation problem defined over binary states (Scenario 2b).

Scenario 2a: Solving a general (continuous-value) external problem.

We simulate $N = 300$ randomly distributed masses connected by P-springs $P_c^P = 50\%$, with uniformly distributed natural length $l_{ij}^P \in [0, 1]$ and uniform spring constants, $k^P = 1$. To this, we add densely connected L-springs (connectivity $P_C^L \sim 99\%$ to avoid instabilities caused by symmetries of fully connected systems), with uniform natural length $l_{ij}^L \in [0, 1]$ and uniform, much weaker, spring constants $k_{ij}^L = 0.1$.

We plot the energy of the final equilibria of the learned system versus the number of disturbances (Figure 5C,D (red)) and, for reference, plot this on a background of the energy of equilibria found when only P-springs are present (blue). Again, we find the system converges to low energy configurations of the P-springs. This is confirmed by a histogram of the energy of equilibria from P-springs only (Figure 5C (blue)), and P-springs plus final L-springs (Figure 5C (red)). Again, to quantify the robustness, we calculate the number of STDs between the final energy found and the mean energy of the problem springs over 10 independent realisations. Again, we find that 80% of runs are over 3 STDs from the

mean, and the average number is 3.0, indicating that the obtained solutions of this quality were extremely rare without the influence of the deforming springs.

Scenario 2b: Solving a discrete combinatorial optimisation problem.

The previous examples demonstrated that, under relatively broad assumptions, a mass-spring-damper system can exhibit adaptation in the strict sense of learning to solve a problem better with experience. In the previous cases, the problem space is a continuous state space. Can this kind of adaptation by natural induction solve the kind of combinatorial optimisation problems that are more familiar in computer science? Combinatorial optimisation problems are a very general class, including MAX-SAT, TSP, etc. The main departure from Scenario 2a is that the problem space in Scenario 2b is binary, not continuous. For this, we need to constrain the initial system architecture further to approximate binary states. This requires some additional mechanistic constraints (see Appendix C for details of Scenario 2b), but the natural induction mechanism is the same as before (Scenario 2a), i.e., using constant problem springs and a separate uniform network of deformable learning springs. We use the problem springs to define a spin glass system. Finding the lowest energy configuration of this system corresponds to solving the MaxCut graph partition problem and is NP-hard in general [77].

As before, we find that adaptation by natural induction finds good quality solutions that are very rarely found without learning. We ran 10 runs and found that in 8 of these, the binarized solution energy in the presence of learning was over 4 STDs away from the mean found without learning, demonstrating a significant increase in the quality of solutions found.

4. Discussion

4.1. The Relationship between Natural Induction and Natural Selection

Adaptation by natural induction and adaptation by natural selection share a number of features: Both involve the incremental accumulation of small changes over time, both can result in the increased fit of an adaptive system to a system of constraints, and both involve simple gradient-following principles. They are also different in the mechanisms that they require (their necessary and sufficient conditions), how they work (their algorithmic principles) and, consequently, in their adaptive competencies. Whereas natural selection depends on the differential survival and/or reproduction of entities, natural induction operates by the differential easing of frustrated relationships between entities (The accommodation of internal connections to a state configuration has some similarity to the over-production of neural connections and their differential retention or reinforcement by 'selective stabilization', which may also result in increased response to or memory of activation patterns. The emphasis of natural induction on the differential easing of frustrated interactions agrees with an emphasis on network integration or 'survival of the fitted', in contrast to 'survival of the fittest'. However, natural induction is a physical model without any population or selection process, whereas selective stabilisation depends on an over-produced population of connections, and in terms of algorithmic competence, natural induction demonstrates an increase in adaptative capabilities). Moreover, the raison d'etre of selection as a theory of biological evolution is to avoid dependence on variation that is directed toward adaptive outcomes, whereas natural induction exploits the fact that variation directed toward easing frustrated interactions is normal in physical systems and adaptively significant (Darwin suggested that some variation was developmentally environmentally directed but was not specific about its adaptive significance. The nature of developmental bias and phenotypic plasticity, and the potential of these and other factors to influence genetic evolution adaptively, is an active topic [6,78]. These mechanistic differences are important, but it is perhaps the difference in their algorithmic principles and adaptive competencies that is more important. Even though both include gradient-following principles, they are different algorithms [79], and this is clear because they have different problem-solving capabilities.

A common interpretation of adaptation by natural selection, i.e., characterised as a population following fitness gradients to a local peak in a fitness landscape [73,80], aligns well

with local gradient-following principles of (first-order) physical optimization. It is acknowledged that biological evolution is not necessarily a good optimiser [54], the idea of natural selection climbing gradients in a static fitness landscape is a serious over-simplification [55], identifying any quantity that natural selection maximises is problematic [50,51], and the conception of natural selection as a problem-solving process has been criticised [51,54,55,81]. However, these are mostly arguments that weaken the interpretation of natural selection as an optimisation process, i.e., natural selection is, at least sometimes, less effective at optimisation than a local gradient process. Accordingly, doing better than local optimisation is a conservative criterion for adaptation.

At a suitable level of abstraction, both natural selection and physical optimisation can be described as processes that follow local gradients to a local optimum. Clearly, though, they are mechanistically different ways of implementing this process. Specifically, natural selection can be described as a process of random generation and selective retention, whereas in the Newtonian model of a physical system, a ball, for example, is deterministically caused to move in a directional fashion by the reaction to the slope, i.e., to roll downhill, not uphill. No population of balls, random variation nor selection process need be involved. It is also possible, however, to think about a statistical mechanics process where the position of the ball is represented by a probability distribution of possible future positions, which is updated to amplify positions that are lower in energy compared to those that are higher in energy. So, does it really matter whether it is a statistical–mechanics (variation and selection) process or a Newtonian (directed learning) process? If the gradient-following outcome is the same, does the mechanism matter?

In some contexts, the mechanism seems to matter a lot to biological thought; Evolution by natural selection can only occur if there is a population, suitable variation, and selection—anything else is not natural selection. On the other hand, it is common to equate evolutionary adaptation with a hill-climbing process, and to the extent that the change in an evolving population is non-arbitrary, this level of functional equivalence seems to capture its adaptive competence. However, if that is true, it has a curious implication; Although it is traditional for evolutionary thought to consider processes that go uphill (in a fitness landscape), whereas physical models go downhill (in an energy function), this does not make either one cleverer, i.e., a more effective optimiser, than the other (regardless of whether they are implemented in a statistical manner or a Newtonian manner). That would suggest that a literal ball rolling down a literal hill would also constitute adaptation. If local hill-climbing is sufficient to produce the biological adaptation we observe, this would mean that the genius of Darwin's theory is just that it provides a hill-climbing process capable of operating in the appropriate organismic 'design space'. Alternatively, the mechanistic details of evolution by natural selection, or the context in which it occurs, may matter to its adaptive competence—making it different from a hill climber [79]. It seems likely that the details do matter, but to the extent that evolution by natural selection is formally characterised as a simple (substrate-independent) hill-climbing process, it excludes these potentially important aspects of the actual biological process.

The models presented here can be interpreted as literal physical systems, but they can also be interpreted as models that stand in for other natural optimisation processes, including those involving the local gradient-following capacity of evolution by natural selection. Adaptation by natural induction is instantiated as a Newtonian process in the models we have illustrated—with forces and (directional) reactions rather than a statistical mechanics process (random variation and selection) and because they are different mechanisms, it means that natural induction has different necessary and sufficient conditions to natural selection, and may apply in cases where natural selection does not. However, in terms of adaptive competence, that is not the important difference between natural induction and natural selection. The problem-solving competence of natural induction is not the same as that of natural selection because they are different algorithms (not because they are different mechanisms for implementing the same algorithm). Adaptation by natural induction finds better solutions than a local gradient-following optimisation process, a.k.a.

a hill-climber. If biological evolution is algorithmically equivalent to a hill-climber (first-order optimisation only), then its adaptive competence is inferior to adaptation by natural induction. Conversely, if biological evolution is a more sophisticated optimiser than a hill climber, then the substrate-independent algorithm of random variation and selection does not describe it.

It might be appropriate to conceive biological evolution as a simple gradient process when natural selection acts on a simple vector of genes or vector of phenotypic traits individually determined by corresponding genes in a one-to-one fashion [66]. In this case, the action of evolution by natural selection is analogous to the bed of clay—a univariate model; that is, each selective coefficient is responsible for the change in frequencies of alleles at one locus. Other work has illustrated, however, that an evolutionary process operating on heritable variation in the connections of a dynamical network can exhibit the same kind of adaptive competence as natural induction. This includes gene-regulation networks [35,48,66]. In this case, the gene network constitutes a dynamical developmental process that generates phenotypes indirectly (and 'disturbances' are provided by a lifecycle that resets developmental states to a neutral, undifferentiated, epigenetic state). This gives it the possibility of representing relational interactions between traits (i.e., pleiotropic interactions in the genotype–phenotype map of such a developmental process may constitute an associative model of selected phenotypes [34]). Previous work shows that under these circumstances, the outcome is superior to local gradient optimisation [66]; that is, evolution by natural selection acting on the parameters of a developmental process (second-order) can solve problems that evolution by natural selection acting on a directly encoded phenotype (first-order) cannot.

One interpretation of this is that natural selection can exhibit the same adaptive competence as natural induction, after all. However, adaptation by natural induction is not the algorithm that Darwin described (evolution by natural selection does not describe a process that provides model induction, a network of viscoelastic connections, an ability to learn to adapt better with experience, nor the significance of 'pulse' or disturbances, amongst other things). So, even though the basic gradient-following process might, in biological cases, be provided by a variation and selection process, in cases where natural selection operates on the parameters of a developmental process constituted by a network of interactions [35,48,66], it might be more correct to attribute the adaptive competence to natural induction and not to natural selection. After all, whereas the theory of evolution by natural selection focuses on the mechanism of variation and selection, this paper demonstrates (using purely Newtonian processes) that this superior adaptive competence arises from the dynamical feedback between model induction and optimisation, and not from random variation and selection. Whereas natural selection depends on the differential survival and reproduction of things, natural induction fundamentally depends on the differential easing of frustrated relationships between things. Although adaptation by natural induction is therefore fully compatible with a Darwinian model of evolutionary change (in suitable network contexts), these are different adaptive algorithms with different necessary conditions, different algorithmic principles, and different adaptive competencies.

We also note that the relationship between evolution and learning has been recognised and developed by many. At a suitable level of abstraction, evolution and some learning methods appear to be the same algorithm (i.e., 'trial and error' plus reinforcement equates to random variation and selection). Both can be understood as processes that optimise a function by, to a first approximation, following local gradients. It has often been noted that reinforcement learning and evolution by natural selection are closely analogous [34,82], and indeed, the replicator equation (an abstraction of biological evolution under natural selection) and Bayesian updating (a learning optimisation process) have been shown to be formally equivalent [83,84]. See also [17,48,49,85–87] for the relationship between learning and evolution. These works expand and deepen our understanding of the adaptation provided by natural selection; however, note that adaptation by natural induction involves a two-way feedback between an optimisation process and an inductive learning process—

the latter on its own is simply an optimisation process in model parameters, and not sufficient to demonstrate an improvement in problem-solving competency.

The results in this paper thus demonstrate that a dynamical system described by a network of viscoelastic connections, and subject to occasional disturbances, exhibits adaptation in the more stringent sense of learning to optimise better with experience or improving its problem-solving competency over time—and this is not the same mechanism, algorithm, or competency as natural selection.

Where might networks of suitable connections occur naturally? By far the most familiar examples of viscoelastic networks are in fact biological ones. Ecological networks, protein networks, cytoskeletal networks, metabolic networks, bio-electrical networks, social networks, and the biosphere as a whole are all networks at least partially characterised by linkages that are likely to give way under stress and are subject to at least occasional perturbations. Although these networks all involve biological individuals and materials, most of them are not (always) evolutionary units, so natural selection does not straightforwardly apply. This suggests that the interaction of natural selection and natural induction may be complex and possibly widespread. Outside of systems that we already recognise as biological, another obvious candidate where natural induction may be important is the origins of life and origins of evolution [88,89]. To the extent that a pre-biotic chemical network has an internal conformation structure that gives way under stress, we speculate that it has the potential to induce a model of its past experience that can anticipate and generalise without having properties sufficient to be a bonafide evolutionary unit.

Whilst we wish to make the case that the conditions for this are quite natural and not onerous (i.e., do not require selection or design), we do not claim that these conditions are ubiquitous or even frequently or commonly met. Bear in mind that the conditions for evolution by natural selection, namely self-replicating systems with heritable variation in reproductive success, are hardly ubiquitous in the physical world (even though they may be realised in various substrates [90,91]) and their origin is not known. Instead, we claim that the contrary assumption, that natural selection is the only possible naturally occurring mechanism of spontaneous adaptation, is not correct. Who would make such a claim? Actually, this assumption is quite widely adopted, usually implicitly, with very wide-reaching and important implications [4].

Reasoning in many biological domains often depends on this assumption implicitly or explicitly. For example, the assertion that loose ecological communities (or the biosphere as a whole) cannot possibly be adapted because they are not evolutionary units (or members of a population) is an argument that presupposes natural selection to be the only possible source of adaptation. Biological systems are replete with dynamical systems described by networks in many different substrates (e.g., chemical reaction networks, metabolic networks, protein networks, gene-regulatory networks, bio-electric networks, ecological networks, and social networks). Linkages in many networks give way under stress (perfect elasticity is an idealisation) and biological systems are rarely isolated from shocks or disturbances. To the extent that these come together in natural systems, natural induction has the potential to provide a mechanism of spontaneous adaptation that is relevant to biological systems in many domains. This includes those where natural selection does not apply and those where natural selection also applies. How natural induction interacts with natural selection, and the broader implications for biological thought and the origins of adaptive complexity, will be analysed in future work.

4.2. Limitations

Natural induction does not occur unless particular conditions are met. It requires a set of state variables and a set of structural parameters such as a network of connections (interaction terms in the dynamical system). These connections need to give way slightly under stress, and the system needs to be subject to disturbances (or episodic stress). These conditions are not uncommon in natural networks but are, of course, not universal. We have made a number of specific assumptions in the particular systems we have illustrated

(spring constants, timing parameters, connectivity, etc.), but the central claim of this paper is robust to these choices and other details; that is, natural selection is not the only possible source of spontaneous adaptation. The aim of our illustrations is not to be the 'last word' or definitive case on such issues, but merely to open up the discussion and fuel debate that is productive.

Nonetheless, in order for natural induction to produce adaptation (i.e., an increase in optimisation capability), some general conditions are required. These correspond to the conditions for good generalisation in a learning system. The problems solved here are constructed from pairwise constraints. These capture a large class of problems, but not all. It is of course easy to construct optimisation problems that natural induction cannot solve. The state dynamics need to spend most of their time at configurations that are better than average (better than random). This condition is easily met and the condition that they spend most of their time at local optima, as in our illustrations, is probably not necessary but not yet investigated. Likewise, we suspect that it is not essential that interactions are symmetric (which guarantees only fixed-point attractors) but asymmetric interactions and non-fixed-point dynamics are not investigated in this paper. Additionally, the system must be subject to disturbances—not too frequent to prevent the system from spending most of its time at good configurations, but not too infrequent that it does not visit a representative sample of good configurations. Disturbances do not necessarily need to be complete resets of the state configuration in order for some induction to occur, but we anticipate that partial resets will limit the independence of the samples from which associations are being learned. This is a complicated matter, however, because partial resets to state variables can sometimes act in a similar manner to updates in interaction parameters. Together, these conditions correspond to the fact that learning works well only when the training data is representative of the class that must be learned—you cannot learn a general class from a single example (or an impoverished distribution of samples) and that is why the disturbances are needed. Learning rate also matters (particularly in online learning where there is positive feedback between what is learned and the data that is learned from) because you do not want learning to take unnecessarily long nor for the system to converge on the first (arbitrary) state it experiences.

Finally, natural induction is an inductive learning process and there is an important fundamental limitation to any inductive learning process—the need for a suitable *inductive bias*. Generalisation cannot occur without induction, but any general rule is necessarily under-determined by past experience. Over the set of all possible general rules that are compatible with the data, any prediction is possible (e.g., "all swans are white except the next one, which is pink"); therefore, any inductive learning process can only produce a prediction (right or wrong) by preferring particular generalisations over others. This preference is not determined by the data (by definition) and is known as inductive bias. An inductive bias can be as simple as a preference for simple models over complex models (a.k.a. parsimony pressure, or regularisation in machine learning).

The simplest kind of model capable of representing associations, and thus capable of non-trivial generalisation, is a correlation model. This is the inductive bias underlying our results; that is, in the examples presented here (and others [11,34,35,48,66]), the model space in which induction occurs is built from pairwise interactions—like it is in neural networks. This works well in many learning problems because it is as simple as possible but not more so (generalisation with a univariate model, like the bed of clay, is limited to simple similarity measures [34]). For our worked examples above (and in the previous work), correlation learning is a good inductive bias because the constraints that determine the structure of the problem are also built from pairwise interactions. This is why generalising over the distribution of some particular local optima is able to predict the location of (i.e., enlarge the attractor for) other local optima that have not been previously visited. In general terms, the implicit inductive bias will be suitable whenever the model and the problem are built from a similar causal geometry (in this case, a network of pairwise interactions)—this is natural when a system is learning by adjusting its own connections [40].

The acknowledgement of an inductive bias is not to suggest in any way that the model was somehow given the solution in advance. All learning requires induction, and a learning process does not know answers in advance, it acquires this information from experience. Although it is not usually presented as such, all optimisation is really a task that requires induction. That is, an optimisation process must predict the location of (hard-to-find) good solutions from (easy-to-find) samples. If, conversely, an adaptive process had already visited the location of good solutions, then the problem is already solved. Even a simple hill-climbing process or an evolutionary process is no better than random guessing if it is not employing a suitable inductive bias—in this case, the assumption of local smoothness [92]. To visualise this, imagine a natural selection process on a truly random fitness landscape with no auto-correlation—here the location of any solution sampled in the past provides no information whatsoever about the location of potentially better solutions, and thus natural selection has no optimisation ability.

In adaptation by natural induction, we are simply exploiting a similar auto-correlation bias but in a slightly deeper representation instead of the original features. Natural selection depends on the assumption that good solutions have state *values* that are similar to the values found in other good solutions. In natural induction, the implicit assumption is that good solutions have *correlations* among their state values that are similar to the correlations found in other good solutions. In difficult problems, the simple value-based assumption is limited and the correlation-based assumption provides a little more competence. Problems that are even more difficult, if they have any learnable structure at all, require deeper models and the principles exploited by natural induction can be extended in this direction [93]. Recognising that adaptation requires learning, and learning requires generalisation, and generalisation requires an inductive bias, helps us to understand how adaptation really works and what is required. Without inductive bias, all adaptation (including natural selection) would be magical.

We have not yet investigated systems with hidden states. Although the problem springs can represent a problem that is external to the learning springs, in the examples illustrated thus far, all the state variables are shared. Inducing a model of a complex system that has hidden states can be much more difficult, and reflexively, a learned model that contains hidden states (or a 'deep' representation) can express relationships that a shallow model cannot [41,94–96].

In the particular system of masses and springs used here, where the parameters of the problem and the induced model are embodied in spring lengths, we show that the change in the parameters is given by the differential of the potential energy function (see Appendix A). This seems intuitively natural—forcing the state variables causes the internal structure to give way in the direction that reduces the stress in the system and lowers the potential energy of that state. This also enlarges the dynamical attractor of that state configuration. The generality of this concept for other physical systems and other structural variables is not known, although other possibilities have been demonstrated. For example, earlier work modelled changes in interaction strengths analogous to spring constants rather than spring lengths [11,40,68]. (If springs only weaken, and never increase in strength, this describes an update rule that is equivalent to one 'half' of Hebb's rule (i.e., the half that *decreases* the magnitudes of weights). The limit of this is that ultimately all interactions (L-springs) go to zero and all states are independent of other states (except for the influence of P-springs). This also depends on the differential reduction of frustrated correlation parameters (without any increases in correlation parameters that are not frustrated) so the initial effect of L-springs cannot be zero. Nonetheless, the current work likewise depends on differential easing and demonstrates that this is sufficient for significant adaptive problem-solving, even when P-springs cannot be changed).

The illustrations of natural induction presented here using masses and springs have all been conducted in a heavily damped regime, which minimises oscillatory behaviour. We have not yet investigated the kind of adaptive algorithm that is instantiated in the oscillating regime. This is obviously a lot more complicated, but we suspect that it might

be quite interesting [97]. There are possibilities of states being represented by phases, interactions being represented by resonance properties, and learning being represented by harmonic phase locking that is natural between coupled oscillators. Preliminary work in another context suggests that phase synchronisation could play an important role in scaling up the adaptive process from one level of organisation to another [41,47,98].

5. Conclusions

It has been argued strongly, and widely assumed across biological thinking, that natural selection is the only possible mechanism capable of producing spontaneous adaptation in natural systems. Here we show that this assumption is false; adaptive organisation occurs spontaneously in physical systems with suitable natural properties, through an effect we call adaptation by natural induction. This occurs in dynamical systems described by a network of viscoelastic connections subject to disturbances. A viscoelastic connection is simply one that 'gives way' slightly under stress, which is a natural property of many physical materials, biological networks, and complex systems more generally. When disturbances cause the system to visit a distribution of locally optimal solutions, the changes to the connections in the network learn a generalised associative model of the local solutions visited, which causes the system to adapt in the rigorous sense of improving its problem-solving competency; that is, it provides optimisation ability that is superior to local optimisation and improves with experience. The simplicity of natural induction, and its necessary and sufficient conditions, offer a solution to both the chicken and the egg problems—i.e., natural selection is not involved at run time nor in the construction or setup of the system. This has important implications for our understanding of biological evolution and adaptation in other complex systems—not least that adaptation can occur spontaneously in systems that are not units of selection.

Author Contributions: Conceptualization, C.L.B., B.M., A.T. and R.A.W.; Methodology, C.L.B., B.M., A.T. and R.A.W.; Software, C.L.B.; Validation, C.L.B. and R.A.W.; Formal analysis, C.L.B.; Investigation, C.L.B. and R.A.W.; Writing—Original draft, C.L.B. and R.A.W.; Writing—Review & editing, C.L.B., T.L., M.L., B.M., A.T. and R.A.W.; Visualization, C.L.B.; Supervision, R.A.W. All authors have read and agreed to the published version of the manuscript.

Funding: Financial support is gratefully acknowledged from grants #62230 and #62220 * (R.A.W) and #62212 (M.L.) from the John Templeton Foundation (the opinions expressed in this paper are those of the authors and not those of the John Templeton Foundation), and BBRSC grant BB/P022197/1 and the UKRI Horizon Europe Guarantee scheme as part of the METATOOL project (C.L.B). (* including APC).

Data Availability Statement: Given the synthetic/theoretical nature of the study, no new empirical data were created (the specifications for the numerical simulations are given in the paper). Further inquiries can be directed to C.L.B.

Acknowledgments: The authors thank Chrisantha Fernando, Adam Davies, Dave Prosser, Freddy Nash, Jamie Caldwell, Tobias Uller, Kostas Kouvaris, Christoph Thies, Tazzio Tissot, Jonathon Young, Lucas Mathieu, Samuel Lennard, Eva Jablonka, Charlie Munford, Hod Lipson, Josh Bongard, Joshua Knowles, Christopher Congleton and Luval Clejan for discussion.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A

Here we derive the equations of motion for the Maxwell spring by resolving forces in the normal way. First, we define the displacement of each element; thus,

$$x = x_s + x_m = x_2 \tag{A1}$$

where x is the length of the entire spring and x_m and x_2 are the lengths of the ideal spring, and damper in series and parallel, respectively (see Figure A1).

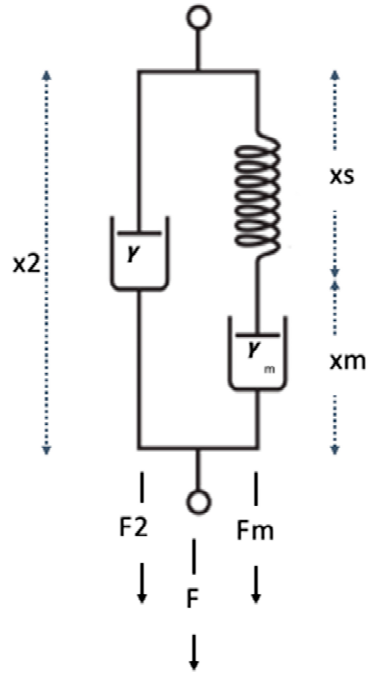

Figure A1. A maxwell configuration comprising of a spring in series with a damper in series with a second damper.

The force on the unit mass is given by

$$F = F_m + F_2 \text{ and } F_m \equiv F_s \tag{A2}$$

where F is the total force and F_s, F_m, and F_2 are forces in the ideal spring, and damper in series and parallel, respectively. Writing

$$F_s = -k(x_s - l_o),$$

$$F_m = -\gamma_m \dot{x}_m$$

$$F_2 = -\gamma \dot{x_2}$$

we can we replace $\dot{x_2}$ with $\dot{x} = v$ because we can see from Equation (A1) that this is just the velocity of the combined system. Taking the time derivative of Equation (A1) and substituting it in Equation (A2), we arrive at

$$\dot{x} = -\frac{\dot{F_s}}{k} - \frac{F_m}{\gamma_m}$$

Rearranging this equation and using Newton's law yields the equations of motion for the unit mass as

$$\dot{v} = F_m - \gamma v$$

$$\dot{x} = v$$

$$\dot{F_m} = -\frac{kF_m}{\gamma_m} - kv$$

We can replace $F_m \equiv F_s$. To convince ourselves that this is correct, we consider the case where the dashpot is stiff and effectively rigid, i.e., $\gamma_m \to \infty$. Integrating $\dot{F_m}$ and taking initial conditions, we return a simple damped spring as

$$\dot{v} = -k\left(x - l^0\right) - \gamma v$$

$$\dot{x} = v$$

We can show that this is equivalent to the 'learning' rule in Equation (A1) in the main text by again integrating $\dot{F_m}$

$$\dot{v} = -k(x - l_l) - \gamma v$$
$$\dot{x} = v$$

where we have made the substitution

$$l_l = l^0 - \int_0^t \frac{F_m}{\gamma_m} \tag{A3}$$

Now, differentiating both sides of Equation (A3) and substituting F_m with the force on the spring, i.e., $F_m = -k(x - l_l)$, we get

$$\dot{l_l} = \frac{k(x - l_l)}{\gamma_m}$$

with learning rate $\frac{k}{\gamma_m}$.

Another way to arrive at this learning rule is to take the derivative of the energy with respect to the parameter. Specifically, considering the energy of the above system,

$$V = \frac{1}{2}k(x - l_l)^2$$

and defining an update that minimises this energy by gradient descent gives

$$\dot{l_l} = -\frac{dV}{dl} \equiv \frac{k(x - l_l)}{\gamma_m}$$

Appendix B

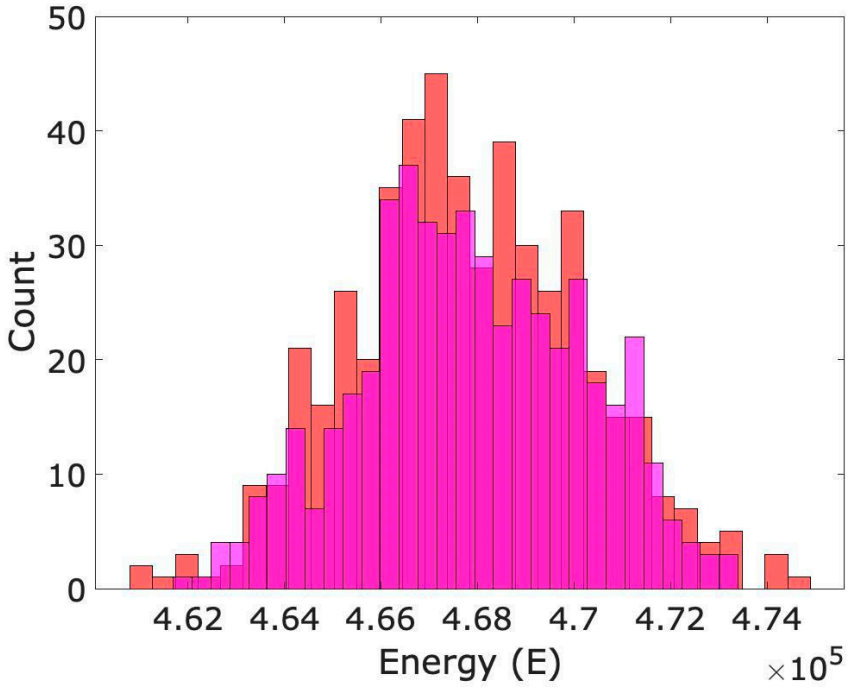

Figure A2. The distribution of found energies on the P-spring alone (red) versus running them with P- and L-springs (before learning) and then removing the L-springs (magenta). The distributions are almost identical. We can thus think of L-springs as the addition of an initially neutral field of plastic material that (initially!) does not alter the problem.

Appendix C

In Scenario 2b, we examine a combinatorial optimisation problem specified as finding the minimal energy state of a spin glass system. We consider a network of N binary spins $s_i \in \{-1, 1\}$ with symmetric interactions, $J_{ij} = J_{ji}$, and $J_{ii} = 0$, but randomly connected $J_{ij} \sim U\{0, -1\}$ with energy defined by $H = -\sum_{ij} J_{ij} s_i s_J$.

Finding the lowest energy configuration of this system corresponds to solving the MaxCut graph partition problem and is NP-hard in general [77]. To represent this problem in our mass-spring-damper system, we consider the following setup: Each mass is attached to a vertical rod constrained to travel vertically and tied with a spring to a fixed point at height $y = 0$ (Figure A3A), which affects bistability (i.e., the binarisation) in the y-position. Pairs of masses are randomly connected with a set of P-springs as before. To simplify our simulation, we ignore spatial constraints and assume that every pair of rods is a unit length apart. Resolving forces, we can write the potential energy of this system as

$$V = \sum_{ij \in \{S\}} \frac{1}{2} k \left(\sqrt{\left(1 + \Delta y_{ij}^2\right)} - l \right)^2 + \sum_i \frac{1}{2} k_p \left(\sqrt{\left(\epsilon^2 + y_i^2\right\}} - l_p \right)^2 \tag{A4}$$

where y_i is the vertical displacement of each mass with respect to the fixed point, $\Delta y_{ij}^2 = (y_j - y_i)^2$ is the y-distance between masses, and γ is damping effected by friction on the rod. The natural length and spring constant of the P-springs and tie springs are k, l, and k_p, l_p, respectively. The distance of the tie to the rod is given by ϵ. The first term of this equation is the potential energy due to the P-springs, note that by construction there are no net forces in the x-direction. The second term is the contribution to the energy of the tie spring. In the absence of P-springs, and with the tie point close to the rod ($\epsilon \to 0$), the system is driven into either an up or down position, $y_i = \pm l_p$. This bistability is preserved in the presence of P-springs if $k_p > k$, such that the second term in Equation (A4) dominates the first. Adding P-springs, which are much longer than unit length (i.e., longer than the distance between pairs of rods), forces coupled masses to misalign and emulate the behaviour of a negative weight in a binary spin glass system. We can see this by noting that when the natural length of the P-springs is much longer than the separation of rods, $l^0 \gg 1$, the first term in LHS of the energy Equation (A4) is dominated by $-2L|yi - y_j|$, and the masses separate. (It would also be possible to emulate positive weights when the natural lengths of the P-springs are shorter than the distance between rods, i.e., $l^0 \ll 1$).

As before, $N = 300$, with connectivity $p_c^P = 0.1$ and parameters $k = 1$, $l = 10$, $k_p = 1$, $l_P = 1$, γ. The equivalent spin glass problem can be constructed by translating all long springs in $\{S\}$ into negative weights in $J_{ij} = -1$. We run the mass-spring-damper to equilibrium and interpret these as solutions to the spin glass problem by interpreting all rods above and below the mid-point of travel as negative and positive spins, respectively (Figure A3B). We compare the performance of this system on the MaxCut problem against the standard hill-climbing algorithm (HC) where we flip each spin with probability P = 1/N and run it for 20,000 steps. In Figure A3C, we see that the distribution of solutions found (before L-springs deform) is comparable to the solutions found by a hill-climbing algorithm.

To allow the system to adapt, as before, we introduce a highly connected set of L-springs ($P_c^L = 0.99$ and $k^L = 0.1$ and $l^L = 20$). We follow the procedure as in Scenario 2, with periodic disturbances, and after some number of disturbances, we let the system relax for a long period to read out the found solution. In Figure A3C, we have plotted the energy of the found solutions against the number of disturbances (red dots) as well as the energies found with P-springs alone (blue dots), demonstrating again that the system finds better solutions over time. In Figure A3D, we convert the binarised solution states into their equivalent MaxCut solution. We ran 10 runs and found that for 8 of these, the binarized solution energy in the presence of learning was over 4 STDs away from the mean found without learning, again demonstrating a significant increase in the quality of solutions found.

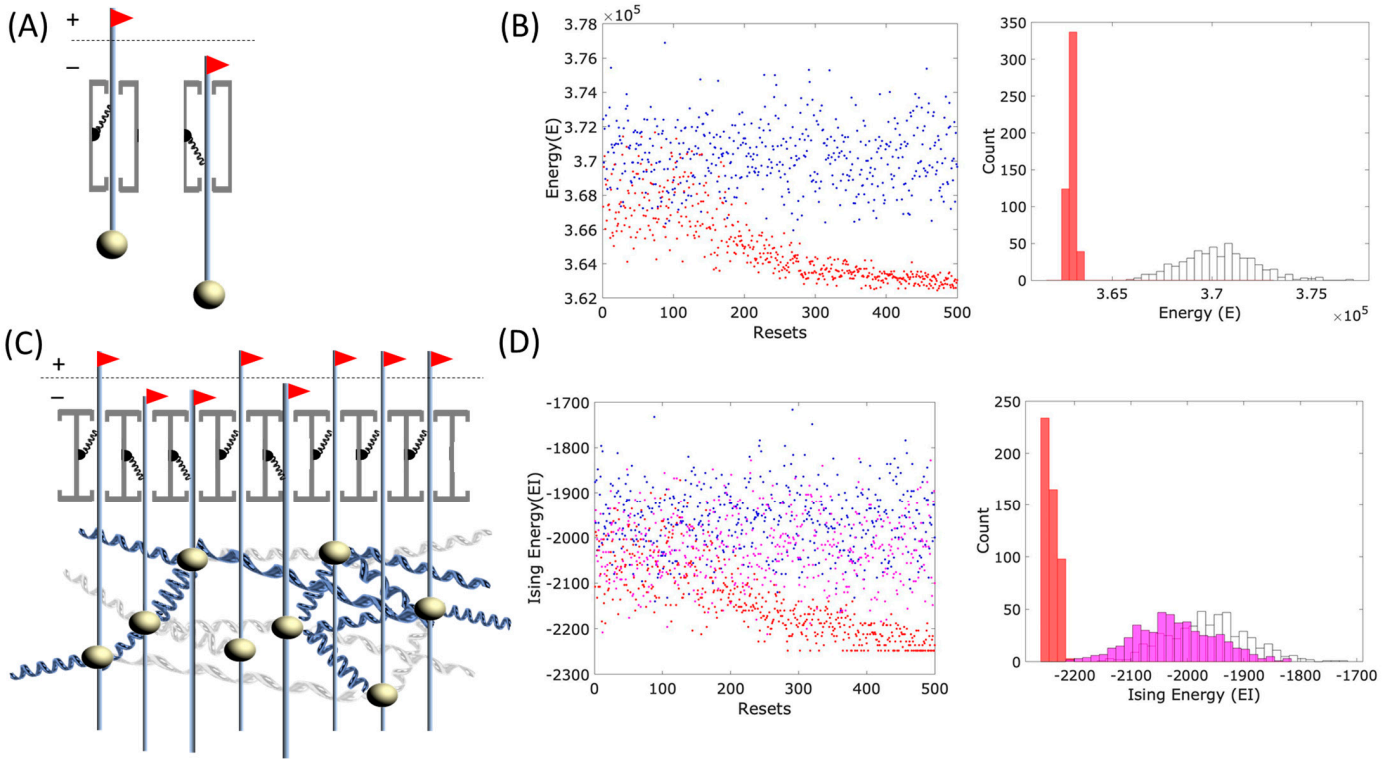

Figure A3. Adaptation by natural induction discovers solutions to a binary constraint problem (Scenario 3). (**A**) A system of masses set on rods constrained to move in the y-direction. Each rod is tied to the y-axis with a short spring, which, in the absence of other springs, yields a bistable resting position in an up/down position. (**B**) Pairs of masses are also connected by a set of elastic long springs ('problem springs', blue). These springs are longer than the distance between the rods in the horizontal plane and they thus have the effect that the masses want to misalign. Masses are also densely connected with set of uniform-length deforming spring ('learning springs', grey). A binary readout of the displacement of the masses is interpreted based on whether the tie spring is oriented upward or downwards (i.e., red flags above/below the midpoint). (**C**) Left: the energy of the final configuration of springs versus the number of resets, without learning springs (blue) and as learning springs deform over time (red). Right: a histogram of the energy of the equilibria found for just the problem springs (white bins) and from different initial conditions for the final L-spring lengths (red bins). The L-springs allow the system to consistently visit low-energy solutions defined by P-springs. Furthermore, the lowest energy state was extremely unlikely to have been visited without the deformed L-springs. (**D**) Same as (**C**) but here the final state is binarized and energy is measured on the underlying Ising energy of the problem. Again, red and white bins are with and without L-springs, respectively. The magenta bins show the distributions of solutions from multiple runs of the hill climber.

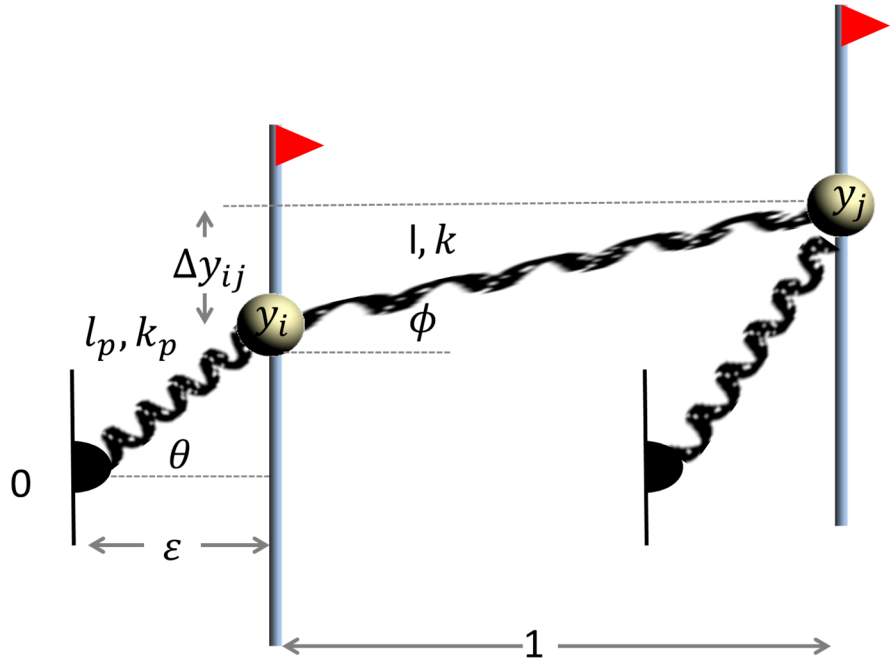

Figure A4. Two masses constrained by metal rods to move in the y-direction. The rods are tied to a pivot point using a tie spring.

Resolving forces in the y-direction for the tie spring we have

$$F^{p}_{y_i} = -k_p(\sqrt{\varepsilon^2 + y_i} - l_p)\,\frac{y_i}{\sqrt{\varepsilon^2 + {y_i}^2}} - \gamma \dot{y_i}$$

where we use $sin\ sin\ \theta = \frac{y}{\sqrt{\varepsilon^2 + {y_i}^2}}$ and assume some friction on the vertical rod. Simplifying, we get

$$F^{p}_{y_i} = -k_p(1 - \frac{l_p}{\sqrt{\varepsilon^2 + {y_i}^2}})\,y_i$$

Similarly, resolving forces for the spring between masses i and I we have

$$F_{y_{ij}} = -k(\sqrt{{\Delta y_{ij}}^2 + 1} - l)\,\frac{\Delta y_{ij}}{\sqrt{{\Delta y_{ij}}^2 + 1}}$$

where $\Delta y_{ij} = y_j - y_j$ and which simplifies to

$$F_{y_{ij}} = -k(1 - \frac{l}{\sqrt{\Delta y_{ij}^2 + 1}})\Delta y_{ij}$$

The total forces on mass y_i is

$$F_{y_i} = F^{p}_{y_i} + \sum_{j \in \{S\}} F_{y_{ij}}$$

and the potential energy of this system is

$$V = \sum_{ij \in \{S\}} \frac{1}{2} k \left(\sqrt{\left(1 + \Delta y_{ij}^2\right)} - l \right)^2 + \sum_i \frac{1}{2}\, k_p \left(\sqrt{\{\epsilon^2 + y_i^2\}} - l_p \right)^2$$

References

1. Darwin, C. *On the Origin of Species by Means of Natural Selection, or, The Preservation of Favoured Races in the Struggle for Life*; John Murray: London, UK, 1859.
2. Futuyma, D.J. *Evolutinary Biology*; Sinauer Associates: Sunderland, MA, USA, 1979.
3. Lewontin, R.C. The units of selection. *Annu. Rev. Ecol. Syst.* **1970**, *1*, 1–18.
4. Dawkins, R. Universal darwinism. In *Evolution from Molecules to Men*; Cambridge University Press: Cambridge, UK, 1983; pp. 403–425.
5. Laland, K.N.; Uller, T.; Feldman, M.W.; Sterelny, K.; Müller, G.B.; Moczek, A.; Jablonka, E.; Odling-Smee, J. The extended evolutionary synthesis: Its structure, assumptions and predictions. *Proc. R. Soc. B Biol. Sci.* **2015**, *282*, 20151019. [CrossRef]
6. Levin, M. Darwin's agential materials: Evolutionary implications of multiscale competency in developmental biology. *Cell. Mol. Life Sci.* **2023**, *80*, 142. [CrossRef] [PubMed]
7. Nowak, M.A.; Ohtsuki, H. Prevolutionary dynamics and the origin of evolution. *Proc. Natl. Acad. Sci. USA* **2008**, *105*, 14924–14927. [CrossRef]
8. Cairns-Smith, A.G.; Hartman, H.; Cairns-Smith, G. *Clay Minerals and the Origin of Life*; CUP Archive: Cambridge, UK, 1986.
9. Maynard Smith, J.; Szathmary, E. *The Major Transitions in Evolution*; Oxford University Press: Oxford, UK, 1997.
10. Lovelock, J.E.; Margulis, L. Atmospheric homeostasis by and for the biosphere: The Gaia hypothesis. *Tellus* **1974**, *26*, 2–10. [CrossRef]
11. Power, D.A.; Watson, R.A.; Szathmáry, E.; Mills, R.; Powers, S.T.; Doncaster, C.P.; Czapp, B. What can ecosystems learn? Expanding evolutionary ecology with learning theory. *Biol. Direct* **2015**, *10*, 69. [CrossRef]
12. Levin, S.A. Ecosystems and the biosphere as complex adaptive systems. *Ecosystems* **1998**, *1*, 431–436. [CrossRef]
13. Wilson, D.S. Two Meanings of Complex Adaptive Systems. In *Complexity and Evolution: Toward a New Synthesis for Economics*; Wilson, D.S., Kirman, A., Eds.; The MIT Press: Cambridge, MA, USA, 2016; pp. 31–46.
14. Ghalambor, C.K.; McKay, J.K.; Carroll, S.P.; Reznick, D.N. Adaptive versus non-adaptive phenotypic plasticity and the potential for contemporary adaptation in new environments. *Funct. Ecol.* **2007**, *21*, 394–407. [CrossRef]
15. West-Eberhard, M.J. *Developmental Plasticity and Evolution*; Oxford University Press: Oxford, UK, 2003.
16. Mitchell, T.M. *Machine Learning*; McGraw-Hill: New York, NY, USA, 1997; Volume 1.
17. Campbell, D.T. The general algorithm for adaptation in learning, evolution, and perception. *Behav. Brain Sci.* **1983**, *6*, 178–179. [CrossRef]
18. Campbell, D.T. Blind variation and selective retentions in creative thought as in other knowledge processes. *Psychol. Rev.* **1960**, *67*, 380. [CrossRef]
19. Edelman, G.M. *Neural Darwinism: The Theory of Neuronal Group Selection*; Basic Books: New York, NY, USA, 1987.
20. Fernando, C.; Goldstein, R.; Szathmáry, E. The neuronal replicator hypothesis. *Neural Comput.* **2010**, *22*, 2809–2857. [CrossRef] [PubMed]
21. Fernando, C.; Szathmáry, E. *Natural Selection in the Brain, in towards a Theory of Thinking: Building Blocks for a Conceptual Framework*; Springer: Berlin/Heidelberg, Germany, 2009.
22. Fernando, C.; Szathmáry, E.; Husbands, P. Selectionist and evolutionary approaches to brain function: A critical appraisal. *Front. Comput. Neurosc.* **2012**, *6*, 24.
23. Stern, M.; Murugan, A. Learning without neurons in physical systems. *arXiv* **2022**, arXiv:2206.05831. [CrossRef]
24. Strong, V.; Holderbaum, W.; Hayashi, Y. Electro-active polymer hydrogels exhibit emergent memory when embodied in a simulated game environment. *Cell Rep. Phys. Sci.* **2024**. [CrossRef]
25. McGregor, S.; Vasas, V.; Husbands, P.; Fernando, C. Evolution of associative learning in chemical networks. *PLoS Comput. Biol.* **2012**, *8*, e1002739. [CrossRef]
26. Parsa, A.; Wang, D.; O'Hern, C.S.; Shattuck, M.D.; Kramer-Bottiglio, R.; Bongard, J. Evolving programmable computational metamaterials. In Proceedings of the Genetic and Evolutionary Computation Conference, Boston, MA, USA, 9–13 July 2022.
27. Venkatesan, T.; Williams, S. *Brain Inspired Electronics*; AIP Publishing LLC: New York, NY, USA, 2022; p. 010401.
28. Kim, K.-H.; Gaba, S.; Wheeler, D.; Cruz-Albrecht, J.M.; Hussain, T.; Srinivasa, N.; Lu, W. A functional hybrid memristor crossbar-array/CMOS system for data storage and neuromorphic applications. *Nano Lett.* **2012**, *12*, 389–395. [CrossRef]
29. Stern, M.; Hexner, D.; Rocks, J.W.; Liu, A.J. Supervised learning in physical networks: From machine learning to learning machines. *Phys. Rev. X* **2021**, *11*, 021045. [CrossRef]
30. Wright, G.; Onodera, T.; Stein, M.M.; Wang, T.; Schachter, D.T.; Hu, Z.; McMahon, P.L. Deep physical neural networks enabled by a backpropagation algorithm for arbitrary physical systems. *arXiv* **2021**, arXiv:2104.13386.
31. Stern, M.; Pinson, M.B.; Murugan, A. Continual learning of multiple memories in mechanical networks. *Phys. Rev. X* **2020**, *10*, 031044. [CrossRef]
32. Chvykov, P.; Berrueta, T.A.; Vardhan, A.; Savoie, W.; Samland, A.; Murphey, T.D.; Wiesenfeld, K.; Goldman, D.I.; England, J.L. Low rattling: A predictive principle for self-organization in active collectives. *Science* **2021**, *371*, 90–95. [CrossRef]
33. Stern, M.; Arinze, C.; Perez, L.; Palmer, S.E.; Murugan, A. Supervised learning through physical changes in a mechanical system. *Proc. Natl. Acad. Sci. USA* **2020**, *117*, 14843–14850. [CrossRef] [PubMed]
34. Watson, R.A.; Szathmary, E. How can evolution learn? *Trends Ecol. Evol.* **2016**, *31*, 147–157. [CrossRef]

35. Watson, R.A.; Wagner, G.P.; Pavlicev, M.; Weinreich, D.M.; Mills, R. The evolution of phenotypic correlations and "developmental memory". *Evolution* **2014**, *68*, 1124–1138. [CrossRef] [PubMed]
36. Sun, X.; Dong, C.; Levin, B.E.; Caunca, M.; Al Hazzouri, A.Z.; DeRosa, J.T.; Stern, Y.; Cheung, Y.K.; Elkind, M.S.V.; Rundek, T.; et al. Erratum to: Systolic Blood Pressure and Cognition in the Elderly: The Northern Manhattan Study. *J. Alzheimer's Dis.* **2021**, *84*, 915. [CrossRef] [PubMed]
37. Hopfield, J.J. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA* **1982**, *79*, 2554–2558. [CrossRef]
38. Hebb, D. *The Organization of Behavior. A Neuropsychological Theory*; Psychology Press: Hove, UK, 1949.
39. Watson, R.; Buckley, C.L.; Mills, R.; Davies, A. Associative memory in gene regulation networks. In Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems, ALIFE 2010, Odense, Denmark, 19–23 August 2010; pp. 659–666.
40. Watson, R.A.; Buckley, C.; Mills, R. Optimization in "self-modeling" complex adaptive systems. *Complexity* **2011**, *16*, 17–26. [CrossRef]
41. Watson, R.A.; Mills, R.; Buckley, C. Transformations in the scale of behavior and the global optimization of constraints in adaptive networks. *Adapt. Behav.* **2011**, *19*, 227–249. [CrossRef]
42. Hopfield, J.J.; Tank, D.W. "Neural" computation of decisions in optimization problems. *Biol. Cybern.* **1985**, *52*, 141–152. [CrossRef]
43. Waddington, C. *The Strategy of the Genes*; George Allen & Unwin: London, UK, 1957.
44. Skyrms, B. *Choice and Chance: An Introduction to Inductive Logic*; Dickenson Pub. Co.: Los Angeles, CA, USA, 1975.
45. Fisher, R.A. The logic of inductive inference. *J. R. Stat. Soc.* **1935**, *98*, 39–82. [CrossRef]
46. Solmonoff, R. A formal theory of inductive inference. (parts I and II). *Inf. Control.* **1964**, *7*, 224–254. [CrossRef]
47. Watson, R. Agency, Goal-Directed Behavior, and Part-Whole Relationships in Biological Systems. *Biol. Theory* **2024**, *19*, 22–36. [CrossRef]
48. Kouvaris, K.; Clune, J.; Kounios, L.; Brede, M.; Watson, R.A. How evolution learns to generalise: Using the principles of learning theory to understand the evolution of developmental organisation. *PLoS Comput. Biol.* **2017**, *13*, e1005358. [CrossRef] [PubMed]
49. Valiant, L. *Probably Approximately Correct: Nature's Algorithms for Learning and Prospering in a Complex World*; Duke University Press: Durham, NC, USA, 2013; ISBN 978-0465060726.
50. Grafen, A. Formalizing Darwinism and inclusive fitness theory. *Philos. Trans. R. Soc. B Biol. Sci.* **2009**, *364*, 3135–3141. [CrossRef]
51. Birch, J. Has Grafen formalized Darwin? Commentary on Grafen's 'The Formal Darwinism project in outline'. *Biol. Philos.* **2014**, *29*, 175–180. [CrossRef]
52. Holland, J. *Adaptation in Natural and Artificial Systems*; University of Michigan Press: Ann Arbor, MI, USA, 1975; Volume 7, pp. 390–401.
53. Fields, C.; Levin, M. Competency in Navigating Arbitrary Spaces: Intelligence as an Invariant for Analyzing Cognition in Diverse Embodiments. *Entropy* **2022**, *24*, 819. [CrossRef] [PubMed]
54. Gould, S.J.; Lewontin, R.C. The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. Proceedings of the royal society of London. *Ser. B Biol. Sci.* **1979**, *205*, 581–598.
55. Levins, R.; Lewontin, R. *The Dialectical Biologist*; Harvard University Press: Cambridge, MA, USA, 1985.
56. Kirkpatrick, S.; Gelatt, C.D., Jr.; Vecchi, M.P. Optimization by simulated annealing. *Science* **1983**, *220*, 671–680. [CrossRef]
57. Varela, F.J.; Bourgine, P. Introduction: Toward a Practice of Autonomous Systems. In Proceedings of the First European Conference on Artificial Life, Paris, France, 11–13 December 1991; MIT Press: Cambridge, MA, USA, 1992.
58. Kauffman, S.A. *The Origins of Order: Self-Organization and Selection in Evolution*; Oxford University Press: New York, NY, USA, 1993.
59. Gross, T.; Sayama, H. *Adaptive Networks: Theory, Models and Applications*; Springer Publishing Company, Incorporated: New York, NY, USA, 2009.
60. Santos, F.C.; Pacheco, J.M.; Lenaerts, T. Cooperation prevails when individuals adjust their social ties. *PLoS Comput. Biol.* **2006**, *2*, e140. [CrossRef]
61. Ashby, W.R. *Design for a Brain: The Origin of Adaptive Behaviour*; Springer Science & Business Media: Boston, MA, USA, 1952.
62. Soen, Y.; Knafo, M.; Elgart, M. A principle of organization which facilitates broad Lamarckian-like adaptations by improvisation. *Biol. Direct* **2015**, *10*, 68. [CrossRef]
63. Betts, R.A.; Lenton, T.M. *Second Chances for Lucky Gaia: A Hypothesis of Sequential Selection*; Met Office: Exeter, UK, 2008.
64. Zarco, M.; Froese, T. Self-modeling in Hopfield neural networks with continuous activation function. *Procedia Comput. Sci.* **2018**, *123*, 573–578. [CrossRef]
65. Watson, R.A.; Mills, R.; Buckley, C.L.; Kouvaris, K.; Jackson, A.; Powers, S.T.; Cox, C.; Tudge, S.; Davies, A.; Kounios, L.; et al. Evolutionary connectionism: Algorithmic principles underlying the evolution of biological organisation in evo-devo, evo-eco and evolutionary transitions. *Evol. Biol.* **2016**, *43*, 553–581. [CrossRef]
66. Kounios, L.; Clune, J.; Kouvaris, K.; Wagner, G.P.; Pavlicev, M.; Weinreich, D.M.; Watson, R.A. Resolving the paradox of evolvability with learning theory: How evolution learns to improve evolvability on rugged fitness landscapes. *arXiv* **2016**, arXiv:1612.05955.
67. Davies, A.P.; Watson, R.A.; Mills, R.; Buckley, C.L.; Noble, J. "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. *Artif. Life* **2011**, *17*, 167–181. [CrossRef] [PubMed]

68. Watson, R.A.; Mills, R.; Buckley, C.L. Global adaptation in networks of selfish components: Emergent associative memory at the system scale. *Artif. Life* **2011**, *17*, 147–166. [CrossRef]
69. Bickhard, M.H. Variations in Variation and Selection: The Ubiquity of the Variation-and-Selective-Retention Ratchet in Emergent Organizational Complexity, Part II: Quantum Field Theory. *Found. Sci.* **2003**, *8*, 283–293. [CrossRef]
70. Salazar-Ciudad, I.; Jernvall, J.; Newman, S.A. Mechanisms of pattern formation in development and evolution. *Development* **2003**, *130*, 2027–2037. [CrossRef]
71. Forgacs, G.; Newman, S.A. *Biological Physics of the Developing Embryo*; Cambridge University Press: Cambridge, UK, 2005.
72. Newman, S.A. *Self-Organization in Embryonic Development: Myth and Reality, in Self-Organization as a New Paradigm in Evolutionary Biology: From Theory to Applied Cases in the Tree of Life*; Springer: Berlin/Heidelberg, Germany, 2022; pp. 195–222.
73. Provine, W.B. *Sewall Wright and Evolutionary Biology*; University of Chicago Press: Chicago, IL, USA, 1989.
74. Alexander, S.; Cunningham, W.J.; Lanier, J.; Smolin, L.; Stanojevic, S.; Toomey, M.W.; Wecker, D. The autodidactic universe. *arXiv* **2021**, arXiv:2104.03902.
75. Roylance, D. *Engineering Viscoelasticity*; Department of Materials Science and Engineering–Massachusetts Institute of Technology: Cambridge, MA, USA, 2001; Volume 2139, pp. 1–37.
76. Hopfield, J.J.; Tank, D.W. Computing with neural circuits: A model. *Science* **1986**, *233*, 625–633. [CrossRef]
77. Karp, R.M. *Reducibility among Combinatorial Problems*; Springer: Berlin/Heidelberg, Germany, 2010.
78. Livnat, A.; Love, A.C. Mutation and evolution: Conceptual possibilities. *BioEssays* **2024**, *46*, 2300025. [CrossRef]
79. Watson, R.A. Is evolution by natural selection the algorithm of biological evolution? In Proceedings of the ALIFE 2012: The Thirteenth International Conference on the Synthesis and Simulation of Living Systems, East Lansing, MI, USA, 19–22 July 2012; MIT Press: Cambridge, MA, USA, 2012.
80. Wright, S. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In Proceedings of the Sixth International Congress of Genetics, Ithaca, NY, USA, 24–31 August 1932.
81. Birch, J. Natural selection and the maximization of fitness. *Biol. Rev.* **2016**, *91*, 712–727. [CrossRef]
82. Skinner, B.F. Selection by consequences. *Science* **1981**, *213*, 501–504. [CrossRef]
83. Harper, M. The replicator equation as an inference dynamic. *arXiv* **2009**, arXiv:0911.1763.
84. Shalizi, C.R. Dynamics of Bayesian updating with dependent data and misspecified models. *Electron. J. Stat.* **2009**, *3*, 1039–1074. [CrossRef]
85. Chastain, E.; Livnat, A.; Papadimitriou, C.; Vazirani, U. Algorithms, games, and evolution. *Proc. Natl. Acad. Sci. USA* **2014**, *111*, 10620–10623. [CrossRef]
86. Frank, S.A. Natural selection maximizes Fisher information. *J. Evol. Biol.* **2009**, *22*, 231–244. [CrossRef] [PubMed]
87. Vanchurin, V.; Wolf, Y.I.; Katsnelson, M.I.; Koonin, E.V. Towards a Theory of Evolution as Multilevel Learning. *arXiv* **2021**, arXiv:2110.14602. [CrossRef]
88. Pross, A. Causation and the origin of life. Metabolism or replication first? *Orig. Life Evol. Biosph.* **2004**, *34*, 307–321. [CrossRef]
89. Damer, B.; Deamer, D. The hot spring hypothesis for an origin of life. *Astrobiology* **2020**, *20*, 429–452. [CrossRef] [PubMed]
90. Campbell, J.A. *Universal Darwinism: The Path of Knowledge*; Greate Space: Scotts Valley, CA, USA, 2011.
91. Hodgson, G.M. Generalizing Darwinism to social evolution: Some early attempts. *J. Econ. Issues* **2005**, *39*, 899–914. [CrossRef]
92. Wolpert, D.H.; Macready, W.G. No free lunch theorems for optimization. *IEEE Trans. Evol. Comput.* **1997**, *1*, 67–82. [CrossRef]
93. Caldwell, J.; Knowles, J.; Thies, C.; Kubacki, F.; Watson, R. Deep Optimisation: Transitioning the Scale of Evolutionary Search by Inducing and Searching in Deep Representations. *SN Comput. Sci.* **2022**, *3*, 253. [CrossRef]
94. Caldwell, J.; Knowles, J.; Thies, C.; Kubacki, F.; Watson, R. Deep Optimisation: Multi-scale Evolution by Inducing and Searching in Deep Representations. In Proceedings of the International Conference on the Applications of Evolutionary Computation (Part of EvoStar), Seville, Spain, 7–9 April 2021; Springer: Cham, Switzerland, 2021.
95. Caldwell, J.R.; Watson, R.A.; Thies, C.; Knowles, J.D. Deep optimisation: Solving combinatorial optimisation problems using deep neural networks. *arXiv* **2018**, arXiv:1811.00784.
96. Watson, R.; Levin, M. The collective intelligence of evolution and development. *Collect. Intell.* **2023**, *2*, 26339137231168355. [CrossRef]
97. Wang, T.; Roychowdhury, J. OIM: Oscillator-based Ising machines for solving combinatorial optimisation problems. In Proceedings of the Unconventional Computation and Natural Computation: 18th International Conference, UCNC 2019, Tokyo, Japan, 3–7 June 2019; Springer: Berlin/Heidelberg, Germany, 2019.
98. Tissot, T.; Levin, M.; Buckley, C.; Watson, R.A. An Ability to Respond Begins with Inner Alignment: How Phase Synchronisation Effects Transitions to Higher Levels of Agency. *bioRxiv* 2023, *under submission*.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.