Trends in Ecology & Evolution

CellPress

Opinion

How Can Evolution Learn?

Q1 Richard A. Watson[1,*] and Eörs Szathmáry[2]

The theory of evolution links random variation and selection to incremental adaptation. In a different intellectual domain, learning theory links incremental adaptation (e.g., from positive and/or negative reinforcement) to intelligent behaviour. Specifically, learning theory explains how incremental adaptation can acquire knowledge from past experience and use it to direct future behaviours toward favourable outcomes. Until recently such cognitive learning seemed irrelevant to the 'uninformed' process of evolution. In our opinion, however, new results formally linking evolutionary processes to the principles of learning might provide solutions to several evolutionary puzzles – the evolution of evolvability, the evolution of ecological organisation, and evolutionary transitions in individuality. If so, the ability for evolution to learn might explain how it produces such apparently intelligent designs.

Trends

A simple analogy between learning and evolution is common and intuitive. But recently, work demonstrating a deeper unification has been expanding rapidly.

Formal equivalences have been shown between learning and evolution in several different scenarios, including: selection in asexual and sexual populations with Bayesian learning, the evolution of genotype–phenotype maps with correlation learning, evolving gene regulation networks with neural network learning, and the evolution of ecological relationships with distributed memory models.

This unification suggests that evolution can learn in more sophisticated ways than previously realised and offers new theoretical approaches to tackling evolutionary puzzles such as the evolution of evolvability, the evolution of ecological organisations, and the evolution of Darwinian individuality.

Learning and Evolution

New insights and new ways of understanding are often provided by analogies. Analogous reasoning is regarded as a core faculty of human cognition [1], and necessary for complex abstract causal reasoning [2]. In biology, analogy is sometimes considered to be the poor cousin Q2 of homology – similar, but not really the same. But in science more generally, analogies can be founded on perfect equivalences, for example, mathematical isomorphisms or algorithmic equivalence, thus enabling the transfer of ready-made results from one system or discipline to another, for example, between quasispecies theory and population genetics [3,4], electromagnetic fields and hydrodynamics [5], and magnetism and neural networks [6]. The previously casual analogy between learning systems and evolution by natural selection has recently been deepened to a level where such transfer can begin.

How Intelligent is Evolution?

Evolution is sometimes likened to an active problem solver, seeking out ingenious solutions to difficult environmental challenges. The solutions discovered by evolution can certainly appear ingenious. Mechanistically, however, there appear to be good reasons to doubt that cognitive problem solving and evolution are equivalent in any real sense. For example, cognitive problem solving can utilise past knowledge about a problem domain to 'anticipate' future outcomes and direct exploration of solutions, whereas evolutionary exploration is myopic and dependent on undirected variation. Intelligent problem solvers can also form high-level or modular representations of a problem, making it easier to reuse partial solutions in new contexts, whereas evolution merely plods on, filtering random replication errors.

Yet, this is not the whole story. Whilst genetic variation might be undirected, the pattern of phenotypic variation is shaped and biased by the processes of development. Moreover, the organisation of developmental processes (from gene regulatory interactions to morphological body plans) is itself, in large part, a product of past evolution. This affords the possibility that random genetic changes might produce phenotypic changes that are 'informed' by past selection [7–9]. This can direct phenotypic variation into different or higher-level morphological dimensions and/or modularise phenotypic features and redeploy them in new contexts [8,10,11]. The question thus arises: is evolution by natural selection (e.g., by adapting the

[1]Department of Computer Science/
Institute for Life Sciences, University
of Southampton, Highfield,
Southampton, SO17 1BJ, UK
[2]The Parmenides Foundation,
Kirchplatz 1, 82049, Pullach/Munich,
Germany

*Correspondence:
r.a.watson@soton.ac.uk (R.A. Watson).

http://dx.doi.org/10.1016/j.tree.2015.11.009

organisation of developmental processes) able to facilitate subsequent adaptation in the same way that a learning system can exploit knowledge from past experience? If so, evolution might be a ‘smarter’ problem solver than generally appreciated [12] and learning theory could explain how.

Of course, at the time when Darwin sought a mechanistic explanation for evolutionary adaptation, the theory of **algorithms** (see Glossary) did not exist as we know it now and an analogy with learning would not have been illuminating. A century later, when Turing provided the first formal framework of computation, it was immediately used to propose an algorithmic account of learning and intelligence [13]. The well-developed understanding of learning algorithms that we have now vastly expands the space of mechanistic possibilities that might be used to answer Darwin's question.

This opinion paper discusses how (i) recent work shows that the link between learning and evolution is a mathematical equivalence; (ii) accordingly, knowledge from the theory of learning can be converted and redeployed in evolutionary theory; and (iii) this offers exciting opportunities to address fundamental evolutionary puzzles in new ways.

Unifying Learning and Evolution

A system exhibits learning if its performance at some task improves with experience [14]. Reusing behaviours that have been successful in the past (**reinforcement learning**) is intuitively similar to the way selection increases the proportion of fit phenotypes in a population [15–18].

In fact, evolutionary processes and simple learning processes are formally equivalent. In particular, learning can be implemented by incrementally adjusting a probability distribution over behaviours [e.g., **Bayesian learning (Bayesian updating)**] or, if a behaviour is represented by a vector of features or components, by adjusting the probability of using each individual component in proportion to its average reward in past behaviours (e.g., Multiplicative Weights Update Algorithm, MWUA [19]). Harper [20] and Shalizi [21] showed that the former is mathematically equivalent to soft selection on genotypes in asexual populations, and Chastain *et al.* [19] have very recently shown that the latter is equivalent to selection acting on individual alleles at linkage equilibrium in sexual populations [22,23]. Evolution thus acquires information from past selection in the same principled way that simple learning systems acquire information from past experience (see also [24,25]). These results can be seen within the integrative framework provided by Valiant, who shows how formal limits on what can be learned can be transferred to characterise formal limits on what can be evolved [26,27].

Can Evolutionary Systems ‘Anticipate’ Future Outcomes?

A key feature of learning systems that seems disanalogous to evolutionary systems is their ability to anticipate actions that will confer future benefits. But learning systems, just like evolutionary systems, are not really able to ‘see the future’ – they cannot learn from benefits that have not yet occurred. Learning systems are, however, able to extrapolate or generalise from past experience. To move beyond repeating behaviours by rote, generalisation requires an appropriate model – an indirect, usually compact, way of representing behaviours. Learning proceeds simply by incrementally improving the fit of a model to past experience, and new behaviours can then be generated from this model.

The clever part of learning methods concerns how behaviours are parameterised in this model space. In a good model space, desirable future behaviours should be similar (nearby) to behaviours that were useful in the past. For example, perhaps ‘eating apples’ should be close to ‘eating pears’ but far from ‘eating red things’.

Glossary

Algorithm: a self-contained step-by-step set of instructions describing a process, mechanism, or function. An algorithmic description of a mechanism is sufficiently abstract to be ‘multiply realisable’ – i.e., it may be instantiated or implemented in different physical substrates (e.g., biological, computational, mechanical) whilst producing the same results. For example, Darwin's account of evolutionary adaptation (via repeated applications of variation, selection, and inheritance) is fundamentally algorithmic and hence encompasses many possible instantiations (e.g., including the molecular details unknown at the time).

Associative learning/memory: learning correlations between inputs and outputs, or learning what features co-occur in the input [6,26,35,43]. Associative memory is an ability to recall a pattern from a stimulus, for example, ‘Darwin’ → ‘Evolution’, ‘Hebb’ → ‘Learning’ (heteroassociative memory), or to recall a complete pattern from a noisy or partial stimulus, for example, ‘Cha-les -ar-in’ → ‘Charles Darwin’, ‘-ona-d H-b-’ → ‘Donald Hebb’ (autoassociative memory) [6]. Analogue of, for example, evolving the mapping between genotype and phenotype, or the correlations among phenotypic features governed by developmental interactions [32,33,82].

Bayesian learning (Bayesian updating): a learning method using Bayes rule as a principled way to incorporate new information with past experience. Analogue of selection in asexual population (replicator equation) [20,29].

Bivariate model: a model that captures pairwise interactions between features (also known as correlation model restricted to pairwise correlations).

Correlation learning: see **associative learning**.

Deep learning: learning high-level representations by learning correlations on top of correlations, etc. Levels can be learned simultaneously [29], or one at a time (deep belief networks) [30].

Evo-devo: evolutionary developmental biology [7,43]. Here, we are particularly interested in the evolution of developmental organisations that change the

Box 1. Learning (and Evolution) as Model Fitting

Many learning methods operate by incrementally adjusting the parameters of a model to improve the fit with a set of example data (training set) [14,26,29]. Consider a sample of points, for example, phenotypes, characterised by two features or traits, some of which belong to a particular class ('**+**'), for example, high-fitness phenotypes. Learning which feature values are fit on average implicitly represents the class by a region in this 2D space (i). This enables a limited sense of generalisation, for example, novel combinations of fit features generate new points in the same region (e.g., new combinations of fit alleles). However, such a model might be unable to represent the class accurately (**underfitting**), as depicted (e.g., the large region includes many unfit points, and the small region excludes approximately half the fit points). The quality of generalisation can be improved by representing the class in a parameter space or model space (~genotype space) that is different from the feature space (~phenotype space). A basic spectrum of model types is depicted. (i) Representing a class by an average value for each individual dimension or trait is a univariate model. (ii–iii) A **bivariate** or **associative model** can represent pairwise positive or negative correlations among features. Evolutionarily, this can be captured as developmental mapping between genotypes and phenotypes that introduces phenotypic correlations. (ii) A linear correlation model (like linear genotype–phenotype mapping [33]), for example, representing that trait 1 works well only when trait 2 has a similar value, can improve the fit to some extent. (iii) However, a nonlinear correlation model is the simplest model capable of representing multimodal distributions [32], for example, representing that high fitness is conferred only when the two traits are both high or both low. The latter is particularly important because a multivariate model can be constructed by layering one nonlinear model onto the outputs of another (hence deep learning [30]). (iv) In general, multivariate models can represent any data arbitrarily accurately [29]. However, fitting a multivariate model by incremental improvement (learning or evolution) can be troublesome if it is unnecessarily complex. One fundamental problem is overfitting, where fitting the idiosyncrasies of the training data results in a model that fails to generalise well, excluding some potentially desirable points (triangle) (Figure I). Q18

By separating model space from feature space, learned models can be used to generate or identify novel examples with similar structural regularities, or (particularly relevant to evolution) to improve problem-solving or optimisation ability by changing the representation of solutions or reducing the dimensionality of a problem [46,65,79].

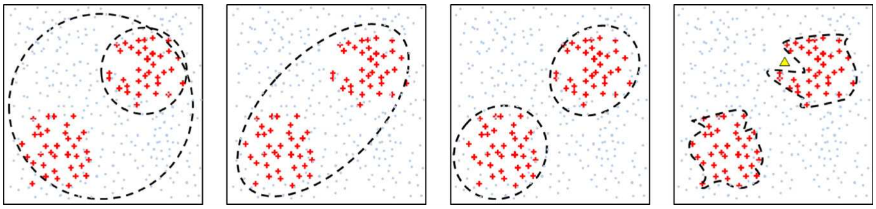

(i) Univariate models, too simple (underfitting)

(ii) Linear correlation model, poor fit (underfitting)

(iii) Nonlinear correlation model (good fit)

(iv) Multivariate model, too complex (overfitting)

Trends in Ecology & Evolution

Figure I. A Basic Spectrum of Model Types (i) to (iv).

In the asexual and sexual populations mentioned earlier, the implicit model space is simply a point in genotype frequency space or allele frequency space, respectively. The latter is a compact way of representing a distribution over genotypes at linkage equilibrium (a **univariate model**, Box 1). This allows a limited sense of generalisation in that new combinations of alleles can be generated from this distribution (i.e., by recombination). In fact, sexual reproduction constitutes a surprisingly efficient trade-off between exploiting alleles that were fit on average in past examples and sampling alleles in new combinations [19]. This simple type of generalisation is ideal when alleles are actually independent (absent of epistasis) whereas asexual reproduction is logical if genotypes cannot be decomposed into independently fit components. Although assuming features are independent is often a pragmatic first approximation and, conversely, assuming complete interdependence covers all eventualities, in most learning tasks neither of these naive extremes is ideal.

For example, in a modular problem, where features in different modules are approximately independent but features in the same module are not, then effective generalisation would be

covariance of phenotypic traits (analogue of correlation learning) [32].

Evo-eco: evolutionary ecology [39,42,54,55,64]. Here, we are particularly interested in the evolution of ecological relationships that change the co-selection of species (analogue of unsupervised correlation learning) [69].

Evo-ego: the evolution of Darwinian individuality [70,71,73,83]. We propose the term 'evo-ego' [34] to refer to the evolution of organisations (reproductive structures) that change the evolutionary unit – i.e., the level of biological organisation that exhibits heritable variation in reproductive success [40]. Here, we are particularly interested in the evolution of reproductive relationships that change the coinheritance of fitness differences [76]. This includes new modes of reproduction modifying the heritability of collectives [40,78] (e.g., vertical transmission of symbionts, as in the origin of eukaryote organelles [83,84]), the origin of chromosomes (via physical linkage of previously independently replicating genetic material [85]), changing reproduction from migrant pool reproduction to group fissioning [71], or encapsulation in compartments (e.g., cell membranes, as in evolutionary transition from replicators on a surface to replicators in compartments) [72,84].

Evolutionary connectionism: a developing theory for the evolution of biological organisation based on the hypothesis that the positive feedback between network topology and behaviour, well understood in neural network models (e.g., Hebbian learning), is common to the evolution of developmental, ecological, and reproductive organisations [32,34,65,68,79].

Hebbian learning: learning that occurs by altering the strength of synaptic connections between neurons [6,14,29]. For example, 'neurons that fire together wire together' is a Hebbian learning principle that strengthens the connection between two neurons when they are activated at the same time or by the same stimulus. Pavlicev *et al.* [33] showed that the action of natural selection adheres to Hebbian principles when acting on heritable variation that affects correlations (e.g., gene regulatory connections [32]). Power *et al.* [69]

provided by new combinations of modules. Genetically, free recombination would disrupt modules and asexual reproduction would fail to exploit the independence of one module from another. An appropriate compromise is provided by an intermediate level of recombination, such as when nucleotides within genes do not recombine, but genes do. Given intragenic epistasis but not intergenic epistasis, the generalisation this provides explains a significant advantage for sex [28]. However, this relies on an *a priori* correspondence between the physical linkage of components and their epistatic dependencies [28].

Can Evolution Learn Like Neural Networks Learn?

More advanced learning requires more flexible model types (Box 1) that alleviate a dependence on the original feature space; enabling items that appear to be different (far apart in feature space) to be represented as nearby points in model space. A minimal example is a correlation model – a simple way of representing interactions between features. The representation of associations or correlations has the same fundamental relationship to learning as transistors have to electronics or logic gates to computation (and synapses to neural networks). Although mechanisms to learn a single correlation between two features can be trivial, these are also Q5 sufficient, when built up in appropriate networks, to learn arbitrarily complex functions [29]. This type of learning can be implemented by incrementally adjusting the parameters of a correlation model in the direction that reduces error (**supervised learning**) or maximises reward (reinforcement learning) (Box 2, i). For example, this is the basis of neural network learning models (operating by adjusting synaptic connection strengths, hence connectionist learning) which have been extraordinarily successful in numerous learning applications [29–31].

Box 2. Supervised and Unsupervised Correlation Learning and the Level of Selection

(i) The Delta rule is a supervised learning rule that modifies model parameters so as to improve the output (or reduce the error between the current output and the ‘desired’ output)

$$\Delta W_{ij} = r d_i X_j$$

where Δw_{ij} is the change in the interaction coefficient between input j and output i, x_j is the value of the input j, r is a learning rate, and d_i is the desired change in the output (given by the error between the desired and actual outputs). Intuitively, given heritable variation in correlations, natural selection for a target phenotype will evolve correlations in the same direction as the Delta rule, that is, to improve the output [32].

(ii) Hebb's rule, often paraphrased as ‘neurons that fire together wire together’, is an unsupervised learning rule (operating without an external ‘teacher’ to define desired outputs) that modifies model parameters in the direction that amplifies the current output:

$$\Delta W_{ij} = r X_i X_j$$

where x_i is the sign of the current output of unit i.

Whereas supervised correlation learning reinforces correlations that are good, unsupervised correlation learning merely reinforces correlations that are frequent. Nonetheless, this is sufficient for interesting system-level behaviours, such as forming an associative memory of past states [6] building low-dimensional models of high-dimensional data, and in some cases, improving system-level optimisation despite the absence of a global reward function [46,65]. When the current output has the same sign as the desired or locally optimal output (i.e., $x_i = d_i$), Hebb's rule and the Delta rule change interactions in the same direction and hence produce the same dynamical consequences for the behaviour of the system [32]. In other cases, when the current output is not optimal, unsupervised learning reinforces the current output regardless of its value. Selection for robustness, for example, might be analogous to unsupervised learning.

(iii) An interesting parallel exists between unsupervised learning and evolutionary selection on individuals within a collective. Specifically, when individual-level selection causes individuals to adopt behaviours that do not maximise collective fitness (as per any social dilemma), the effect of individual selection is not equivalent to supervised learning for the collective (i.e., $x_i \neq d_i$). Yet, if each individual has adopted a state that is locally fit for them, then individual selection on interactions will act to stabilise that state [65] (like selection for robustness at the collective level). This reinforces the current system configuration (without regard to its effect on collective welfare) as per the action of unsupervised correlation learning [65]. Accordingly, even when the collective is not a unit of selection, such as an ecological community, unsupervised learning behaviours can be produced at the system level [69].

show Hebbian learning in the evolution of ecological interactions.

Hopfield network: a simple type of neural network model where each neuron is (potentially) connected to every other neuron bidirectionally [6,31]. The Hopfield network has been used as a mathematical model for (non-neural) dynamical systems and emergent collective behaviours in many different domains including gene regulation networks and ecological networks.

Inductive bias: because, in principle, there are many general concepts that are consistent with a given set of examples, learning from examples always involves inductive bias (i.e., that *a priori* favours a given class of generalisations). However, there are rather generic inductive biases that, Q17 although fallible in principle, prove extremely effective in practise. Occam's razor is one such bias – favouring simple models over more complex models that explain the same data [14]. See also **parsimony pressure**.

Major evolutionary transitions: evolutionary innovations that have changed the evolutionary unit (the level of biological organisation that exhibits heritable variation in reproductive success): from self-replicating molecules, to chromosomes, to simple cells, to multiorganelle eukaryote cells, to multicellular organisms, to social groups [72,73] (see **evo-ego**).

Multivariate model: a model that captures high-order correlations (greater than pairwise interactions) among features [29].

Overfitting: the tendency of a learning algorithm to perform well on the training set but poorly on the test set resulting from fitting idiosyncrasies of the training set. Failure to generalise correctly [14]. Analogous to securing fitness benefits in current selective environment (robustness) at the expense of potential fitness benefits in future environments (evolvability).

Parsimony pressure: a technique used in learning that penalises model complexity to favour simple models over complex ones. Simple models often produce superior generalisation by alleviating overfitting [14].

Reinforcement learning: trial and error learning based on an evaluative or reward signal, providing the learner with a measure of the value or quality of a given solution or

Again, mathematical equivalences with evolution have recently been shown [26,32]. For evolution, learning of this type requires separating phenotypes from genotypes and evolving the parameters of a mapping between them. When there is heritable variation in this mapping that affects phenotypic correlations, natural selection inevitably favours changes that adhere to **correlation learning** principles [26,32]. A minimal example is the evolution of a single 'relational' allele, causing subsequent mutations to produce correlated variation in two phenotypic traits [33] (e.g., via pleiotropy). Pavlicev and colleagues showed that selection on relational alleles increases phenotypic correlation if the traits are selected together and decreases it if they are selected antagonistically (**Hebbian learning**) [32,33]. This simple step from evolving traits to evolving correlations between traits is crucial; it moves the object of natural selection from fit phenotypes (which ultimately removes phenotypic variability altogether) to the control of phenotypic variability.

In larger biological networks, this principle has the same effect as it does in larger neural networks (hence **evolutionary connectionism** [34]). In the **Hopfield network** [6], for example, this type of learning is sufficient for simple cognitive behaviours such as forming an **associative memory (learning)** capable of storing and recalling multiple distinct activation patterns, and effective generalisation in numerous recognition and classification tasks [32,35]. Watson *et al.* demonstrated conditions where evolved gene regulation networks produce exactly the same behaviours [32], forming a distributed 'developmental memory' of multiple phenotypes selected in the past, and generalising by producing new combinations of phenotypic modules (Figure 1).

These results, and others [26,34,36,37], demonstrate that evolution and learning are not merely analogous processes but (different instantiations of) the same algorithmic principles. Transfer of

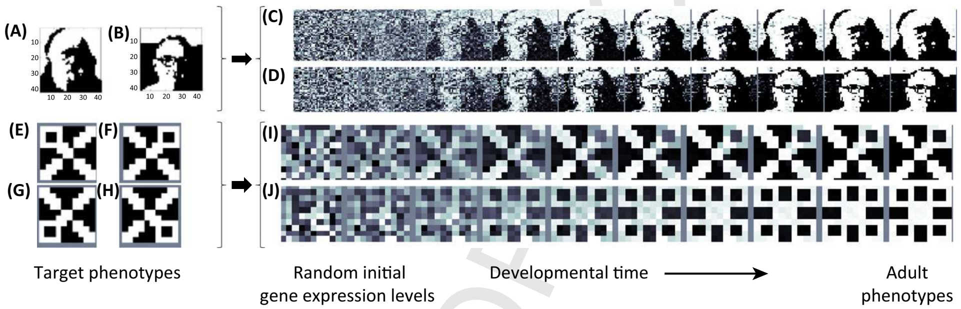

Trends in Ecology & Evolution

Figure 1. A Recurrent Gene Regulation Network (GRN) Evolved in a Varying Environment Exhibits Associative Learning Behaviours. See [32] for details. When a Hopfield network is trained on a set of patterns with Hebbian learning it forms an associative memory of the patterns in the training set. When subsequently stimulated with random excitation patterns, the activation dynamics of the trained network will spontaneously recall the patterns from the training set or generate new patterns that are generalisations of the training patterns [6,31,82]. Here the evolution of connections in a GRN is shown to follow such Hebbian learning principles. The evolved GRN thus forms an associative memory of phenotypes that have been selected for in the past, spontaneously recreating these phenotypes as attractors of development with the GRN and also producing new phenotypes that are generalisations of them. (A–D) A GRN is evolved to produce first one phenotype and then another in an alternating manner [8,49]: A = Charles Darwin, B = Donald Hebb (who first described Hebbian learning). The resulting phenotype is not merely an average of the two phenotypic patterns that were selected in the past (as per a univariate model or free recombination of phenotype pixels). Rather, different embryonic phenotypes (e.g., random initial conditions C and D) develop into different adult phenotypes with this evolved GRN match either A or B (one initial phenotype that falls into each developmental attractor is shown). These two phenotypes can be produced from genotypes that are a single mutation apart [32]. (E–J) In a separate experiment, selection iterates over a set of target phenotypes (E–H). In addition to developing phenotypes that match patterns selected in the past (e.g., I), this GRN also generalises to produce new phenotypes that were not selected for in the past but belong to a structurally similar class, for example, by creating novel combinations of evolved modules (e.g., developmental attractors exist for a phenotype with all four 'loops' [32], J) – see also [8]. This demonstrates a capability for evolution to exhibit phenotypic novelty in exactly the same sense that learning neural networks can generalise from past experience [32].

behaviour, used to amplify successful and reduce unsuccessful behaviours [14]. Analogue of selection.

Supervised learning: learning that changes parameters of a behaviour in the direction that reduces the error (i.e., error = desired output – actual output). Sometimes implies an external teacher that knows the desired output – but in practice, it usually means simply modifying parameters by gradient descent on an error function (rather than by trial and error) [14]. The combination of random variation and selection can affect the same changes in a model given the same gradient, and thus an analogue of selection for a particular target phenotype or phenotypes [32] (see also selection in varying environments or for a set of target phenotypes [8,49]).

Test set: data used to test the accuracy of a model once built (future performance). To test generalisation, the test set includes points that were not presented during training [14]. Analogue of future/novel selective environments.

Training set: data used to build a model (past experience) [14]. Analogue of past selective environments.

Trial and error learning: learning by trialling behaviours at random until a solution is discovered. If each new behaviour that is sampled is a small random modification of the previous behaviour, this becomes a form of incremental adaptation.

Underfitting: the condition that a learned model has failed to accurately fit the training set. Contrast to overfitting [14].

Univariate model: a model that treats each parameter as independent (unable to represent correlations).

Unsupervised learning: learning that aims to optimise a task-independent criterion function based on current output only (e.g., stability or robustness of the output). Occurs without knowledge of a desired output function/external teacher, for example, by reinforcing the current output regardless of its quality. In particular, unsupervised correlation learning, where correlations that are already frequent in the training data are reinforced (rather than correlations that are good with regard to a task, as in reinforcement/supervised correlation learning). The aim of unsupervised learning is to specific models and results between these intellectual domains is already proving productive. Whilst it is important to apply analogies critically, learning theory is not just one thing – the issue is not so much to determine where the analogy breaks down, but to find the right type of learning theory for each of the biological phenomena that are in need of explanation. Below we discuss three examples where learning theory makes sense of biological ideas that are currently confusing, suggesting predictions that arise by transferring well-known learning results.

Future Prospects: Understanding How Evolution Transforms Itself

Learning theory offers new concepts and theoretical tools for tackling several important puzzles in contemporary evolutionary biology. We identify specific learning models that inspire new approaches to key open questions in evolutionary developmental biology (**evo-devo**), evolutionary ecology (**evo-eco**), and evolutionary transitions in individuality (or **'evo-ego'** [34]) (see Outstanding Questions). Each of these areas is challenging for evolutionary theory because they involve feedbacks where the products of evolution modify the mechanisms of the evolutionary process that created them (Figure 2, Key Figure) [33,38–42]. Although it is clear that the processes of variation, selection, and reproduction underpinning evolutionary adaptation are not constants in natural populations, theoretical treatments of 'modifier alleles' that enable selection to act on these processes are currently very limited. There is growing recognition that an integrated framework that puts such feedbacks front-and-centre is desirable [43–45]. Learning theory is precisely the study of processes that change over time as a function of past experience [14,29,46]. It is thus ideally suited to describing, not just how variation, selection, and inheritance adapt phenotypes, but how natural selection modifies variation, selection, and inheritance over time. We note that feedbacks on these three processes result in correlations or covariance between components that were previously independent [34] (Figure 2). Learning theory has well-understood models for each case.

Learning Theory Approach A. Evo-Devo: The Evolution of Evolvability and Correlation Learning

The evolution of developmental biases and constraints, accumulated over past selection, might improve the distribution of phenotypes explored in the future [7,38,47,48]. But the core issue in the evolution of evolvability [7,38,47,49] is that selection cannot favour traits for benefits that have not yet been realised [9,50].

Learning theory offers a solution. First, a memory of phenotypes that have been selected in the past (e.g., Figure 1) can facilitate faster adaptation whenever these phenotypes are selected again in the future [8,33]. Second, and more importantly, because learned models can generalise (e.g., Figure 1J), an evolved memory can, as illustrated by Parter *et al.* [8], also facilitate faster adaptation to new targets. In short, evolvability is to evolution as generalisation is to learning.

Whilst generalisation is not always easy, it does not require clairvoyance – it simply requires the ability to find structural regularities that are deep enough to be invariant over time [26]. Accordingly, the possibility that evolution can learn from experience to favourably bias future exploration need not be any more mysterious than the basic result that learning from a **training set** can produce good generalisation on an unseen **test set** [51]. This also sheds light on the tension between robustness and evolvability. Here the problem is that adapting variation mechanisms so that they are less likely to produce deleterious variants (e.g., via canalisation) is often more immediately advantageous than adapting them so that they are more likely to produce adaptive variants [9,52]. Learning theory understands this tension extremely well. Specifically, **overfitting** occurs when learning improves performance on training data but worsens performance on test data. To avoid this, the complexity of a model can be limited (e.g., by applying a **parsimony pressure**) to prevent memorisation of unnecessary details and

discover categories, clusters, or regularities inherent in the training samples and hence reduce the effective dimensionality of the data [14] (Box 2).

Key Figure

Challenges in Current Evolutionary Theory – Caused When the Products of Evolution Modify the Mechanisms of Evolution

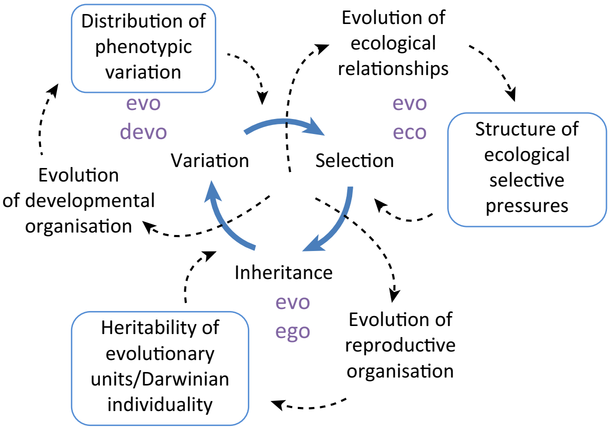

Trends in Ecology & Evolution

Figure 2. Key components of evolution by natural selection – variation, selection, and inheritance [61] – are defined by structures (boxed) that are themselves modified by products of the evolutionary process (dotted arrows). Evo-devo: the evolution of developmental interactions modifies the distribution of phenotypic variation. Evo-eco: the evolution of ecological interactions modifies the structure of selective pressures. Evo-ego: the evolution of reproductive interactions (e.g., vertical transmission of symbionts, or transition from replicators on a surface to replicators in compartments) that modify evolutionary individuality by changing mechanisms of inheritance. These feedbacks are difficult to accommodate in evolutionary theory but are well studied in learning systems. We note that each of these feedbacks results in correlations or covariance between components that were previously independent: (i) the evolution of phenotypic correlations mean that traits do not vary independently; (ii) the evolution of ecological dependencies mean that selection pressures on one species are not independent of the selective pressures on another; and (iii) the evolution of new reproductive mechanisms mean that evolutionary units are not inherited independently. But, in evo-devo, correlations evolve within a single evolutionary unit; in evo-eco, correlations evolve between multiple evolutionary units, and; in evo-ego, correlations change the evolutionary unit (such that multiple, previously separate units become a new single unit at a higher level of organisation) [34]. Learning theory has models that correspond to each of these cases.

force solutions to capture deeper regularities (Box 1). This explains why a cost of connections increases evolved modularity and improves evolvability [51,53].

Using past experience to favourably direct future behaviour is a hallmark of intelligence. By showing that incremental adjustment in the parameters of an appropriate model is sufficient to achieve this, learning theory puts this behaviour within reach of evolution by natural selection, and identifies conditions where it can learn to favourably direct future exploration (see Outstanding Questions, prediction 1).

Learning Theory Approach B. Evo-Eco: Ecological Organisation and Unsupervised Correlation Learning

Organisms can modify their biotic and abiotic environment and thereby alter the selective pressures that act on themselves [39,41,44,54–57]. By modifying the network of ecological dependencies with other species, this might result in ecological organisations that increase the self-regulation of ecosystem variables, the resilience of ecological networks, or the efficiency of resource utilisation [58–60]. But since ecosystems are not, in most cases, evolutionary units [42,61], such feedbacks could also result in effects that are destructive in the long term, for example, making an ecosystem more brittle or susceptible to catastrophic regime change, decreasing total biomass, etc. At present, however, we have no general organising principles for understanding how the structural organisation of ecological networks changes over evolutionary time, nor how this affects ecological functions and dynamics [42,54,60,62–64].

A different type of learning is relevant here. **Unsupervised learning** mechanisms do not depend on an external reward signal. By reinforcing correlations that are frequent, regardless of whether they are good, unsupervised correlation learning can produce system-level behaviours without system-level rewards (Box 2, ii). This can be implemented without centralised learning mechanisms as in connectionist models of intelligence [6,29,34] or distributed multiagent systems [65] (simple forms of collective intelligence [65–68]).

Recent theoretical work shows that selection acting only to maximise individual growth rate, when applied to interspecific competition coefficients within an ecological community, produces unsupervised learning at the system level [69] (Box 2, iii). This is an exciting possibility because it means that, despite not being a unit of selection, an ecological community might exhibit organisations that confer coordinated collective behaviours, for example, a distributed ecological memory that can recall multiple past ecological states [69].

Learning theory shows that incremental adjustment in the parameters of individual behaviours is sufficient to achieve such collective behaviours [65], putting them within reach of individual-level selection. Accordingly, learning theory describes conditions where individual-level natural selection might facilitate ecological organisation and collective behaviour (see Outstanding Questions, prediction 2).

Learning Theory Approach C. Evo-Ego: The Evolution of Individuality and Deep Correlation Learning

In **major evolutionary transitions** [40,70–72] 'entities that were capable of independent replication before the transition can replicate only as part of a larger whole after the transition' [72,73]. These transitions in individuality [40,70,74] involve the evolution of new mechanisms of inheritance or reproductive codispersal (e.g., vertical genetic transmission, compartmentalisation, reproductive linkage) [72,73,75,76] that create new evolutionary units. But there is a catch: if individual and group interests are aligned then selection applied at the group level does not alter evolutionary outcomes, and if individual and group interests are not aligned then individual-level selection will oppose the creation and maintenance of adaptations that enforce selection at the group level [40]. Given this, how can evolution at one level of biological organisation systematically create reproductive organisations that facilitate nontrivial adaptation at a higher level of organisation before that level of adaptation exists?

In neural networks, **deep learning** [30] exploits correlation learning at multiple scales to build **multivariate models** (Box 1). Deep belief nets [30], an exciting recent development igniting renewed interest in neural networks, achieve this in a bottom-up manner, 'freezing' each layer before adding the next. This creates the need to infer low-level representations that are useful for learning higher-order representations before the higher level of representation exists. Unsupervised learning provides a solution. By reducing the effective dimensionality of the data it 'primes' good performance at the next layer, even though it is not informed by what the data will be used for at the next level [35]. In evolutionary systems, selection at one level of organisation can operate like unsupervised learning at a higher level of organisation (Box 2, iii) [69]. Abstract models incorporating these features show that individual-level selection can thus prime the systematic formation of adaptive higher-level evolutionary units without presupposing selection at the higher level [77,78]. New optimisation methods based on these principles demonstrate problem-solving capabilities that cannot be achieved with single-level adaptation [77,79]. We think this suggests that such deep optimisation principles might explain how evolutionary transitions facilitate deep evolution [†], that is, the evolution of adaptive biological complexity through successive levels of biological organisation [34,40,70–72] (see Outstanding Questions, prediction 3).

Efficiently reducing a problem by rescaling a search process at a higher level of representation is another hallmark of intelligent problem solving. Again, learning theory places this within reach of evolution by showing how incremental adaptation, in the right model, can achieve this.

Taken together, correlation learning, unsupervised correlation learning, and deep correlation learning thus provide a formal way to understand how variation, selection, and inheritance, respectively, might be transformed over evolutionary time (Figure 2). We do not claim that evolvability, ecosystem organisation, or the level of evolutionary unit will always increase – on the contrary, we argue that learning theory can be used to characterise the conditions when it will and when it will not.

Concluding Remarks

Learning and evolution share common underlying principles both conceptually and formally [16,18–22,26,32,34,37,69]. This provides access to well-developed theoretical tools that have not been fully exploited in evolutionary theory (and conversely suggests opportunities for evolutionary theory to expand cognitive science [80,81]). Learning theory is not just a different way of describing what we already knew about evolution. It expands what we think evolution is capable of. In particular, it shows that via the incremental evolution of developmental, ecological, or reproductive organisations natural selection is sufficient to produce significant features of intelligent problem solving.

In current evolutionary theory, it seems impossible that natural selection can anticipate what is needed in novel selective environments, that ecological organisation can occur without community-level selection, or that new levels of individuality could emerge systematically from selection on lower-level units. We argue that specific types of learning provide concrete models for such phenomena and suggest predictions that might be tested. We think this offers the potential to better explain how the process of random variation and selection results in the apparently intelligent designs it produces.

Q7 Uncited references

[83–85].

Acknowledgments

We thank Daniel Power, Gunter Wagner, Louis Kounios, Kostas Kouvaris, Chrisantha Fernando, Daniel Weinreich, Rob Mills, Paul Ryan, Adam Jackson, Simon Powers, Simon Tudge, Chris Cox, Markus Brede, Samir Okasha, Seth Bullock, and Chris Buckley for formative discussions.

Q8 References

1. Hofstadter, D. (2001) Analogy as the core of cognition. In *The Analogical Mind: Perspectives from Cognitive Science* (Gentner, D. *et al.*, eds), pp. 499–538, MIT Press
2. Penn, D.C. and Povinelli, D.J. (2007) Causal cognition in human and nonhuman animals: a comparative, critical review. *Annu. Rev. Psychol.* 58, 97–118

Outstanding Questions

Evo-Devo, 2. Evo-Eco, and 3. Evo-Ego Research and Potential Impact of Taking a Learning Theory Approach (see Future Prospects A, B, and C, respectively)

1. Is evolvability evolvable?

Evolutionary problem: Developmental organisations change over evolutionary time in response to the short-term fitness benefits such as from phenotypic robustness. But: *How could the evolution of developmental organisations favour variability that facilitates long-term evolvability?*

Learning theory insight: Learning theory describes conditions where incremental reward-based adaptation can result in successful generalisation to previously unseen situations.

Example prediction: Short-term selection can increase long-term evolvability if it benefits from an appropriate **inductive bias**, for example, that the genotype–phenotype map is complex enough to represent structure (epistatic interactions) in the selective environment but simple enough to avoid overfitting that structure [51,53].

2. Can ecosystem functions be adapted without ecosystem selection?

Evolutionary problem: The organisation of ecological relationships in an ecosystem changes over evolutionary time as a result of individual selection within each component species. But: *Given that an ecological community is not a Darwinian unit, how can ecological organisations be anything other than the arbitrary consequence of happenstance contingencies?*

Learning theory insight: Unsupervised learning can be produced by very simple component-level reinforcement mechanisms (e.g. ‘neurons that fire together wire together’) without a system-level reward function but can nonetheless result in nontrivial system-level behaviours (e.g., associative memory).

Example prediction: Individual selection within an ecological community can result in coordinated behaviours for the ecosystem as a whole if ‘species that fire together wire together’ (e.g., species that are frequently in high

---

[†] Thanks to Seth Bullock for suggesting the term ‘deep evolution’.

3. Eigen, M. (1971) Self-organization of matter and the evolution of biological macromolecules. *Naturwissenschaften* 58, 465–523
4. Wilke, C.O. (2005) Quasispecies theory in the context of population genetics. *BMC Evol. Biol.* 5, 44 Q13
5. Feynman, R.P. *et al.* (1963) *The Feynman Lectures on Physics: Mainly Mechanism, Radiation and Heat,* Addison Wesley
6. Hopfield, J.J. (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.* 79, 2554–2558
7. Gerhart, J. and Kirschner, M. (2007) The theory of facilitated variation. *Proc. Natl. Acad. Sci. U.S.A.* 104, 8582–8589
8. Parter, M. *et al.* (2008) Facilitated variation: how evolution learns from past environments to generalize to new environments. *PLoS Comput. Biol.* 4, e1000206
9. Pigliucci, M. (2008) Is evolvability evolvable? *Nat. Rev. Genet.* 9, 75–82
10. Kashtan, N. and Alon, U. (2005) Spontaneous evolution of modularity and network motifs. *Proc. Natl. Acad. Sci. U.S.A.* 102, 13773–13778
11. Wagner, G.P. *et al.* (2007) The road to modularity. *Nat. Rev. Genet.* 8, 921–931
12. Fernando, C. (2013) Design for a Darwinian brain: part 1. Philosophy and neuroscience. In *Biomimetic and Biohybrid Systems* (Lepora, N.F. *et al.*, eds), pp. 71–82, Springer Q9
13. Turing, A.M. (1950) Computing machinery and intelligence. *Mind* 59, 433–460
14. Mitchell, T. (1997) *Machine Learning,* McGraw Hill
15. Maynard Smith, J. (1986) *The Problems of Biology,* Oxford University Press
16. Frank, S.A. (1996) The design of natural and artificial adaptive systems. In *Adaptation* (Rose, M.R. and Lauder, G.V., eds), pp. 451–505, Academic Press Q10
17. Skinner, B.F. (1953) *Science and Human Behavior,* Simon and Schuster
18. Bateson, G. (1979) *Mind and Nature: A Necessary Unity,* Dutton
19. Chastain, E. *et al.* (2014) Algorithms, games, and evolution. *Proc. Natl. Acad. Sci. U.S.A.* 111, 10620–10623
20. Harper, M. (2009) The replicator equation as an inference dynamic. *arXiv arXiv* 0911.1763 Q11
21. Shalizi, C.R. (2009) Dynamics of Bayesian updating with dependent data and misspecified models. *Electron. J. Stat.* 3, 1039–1074 Q14
22. Hofbauer, J. and Sigmund, K. (1998) *Evolutionary Games and Population Dynamics,* Cambridge University Press
23. Barton, N.H. *et al.* (2014) Diverse forms of selection in evolution and computer science. *Proc. Natl. Acad. Sci. U.S.A.* 111, 10398–10399
24. Frank, S.A. (2009) Natural selection maximizes Fisher information. *J. Evol. Biol.* 22, 231–244
25. Frank, S.A. (2012) Natural selection. V. How to read the fundamental equations of evolutionary change in terms of information theory. *J. Evol. Biol.* 25, 2377–2396
26. Valiant, L. (ed.) (2013) *Probably Approximately Correct: Nature's Algorithms for Learning and Prospering in a Complex World*, Basic Books
27. Valiant, L.G. (2007) Evolvability. In *Mathematical Foundations of Computer Science* (Rao, G.S., ed.), pp. 22–43, Springer Q12
28. Watson, R.A. *et al.* (2011) Genome structure and the benefit of sex. *Evolution* 65, 523–536
29. Rumelhart, D.E. *et al.* and the PDP Research Group (1986) *Parallel Distributed Processing (Vol. 1, 2),* MIT Press
30. Hinton, G.E. (2007) Learning multiple layers of representation. *Trends Cogn. Sci.* 11, 428–434
31. Hopfield, J.J. and Tank, D.W. (1986) Computing with neural circuits: a model. *Science* 233, 625–633
32. Watson, R.A. *et al.* (2014) The evolution of phenotypic correlations and 'developmental memory'. *Evolution* 68, 1124–1138
33. Pavlicev, M. *et al.* (2011) Evolution of adaptive phenotypic variation patterns by direct selection for evolvability. *Proc. R. Soc. B Biol. Sci.* 278, 1903–1912
34. Watson, R.A. *et al.* (2015) Evolutionary connectionism: algorithmic principles underlying the evolution of biological organisation in evo-devo, evo-eco and evolutionary transitions. *Evol. Biol.* (in press)
35. O'Reilly, R.C. and Munakata, Y. (2000) *Computational Explorations in Cognitive Neuro-science: Understanding the Mind by Simulating the Brain,* MIT Press
36. Börgers, T. and Sarin, R. (1997) Learning through reinforcement and replicator dynamics. *J. Econ. Theory* 77, 1–14
37. Friston, K. (2010) The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* 11, 127–138
38. Wagner, G.P. and Altenberg, L. (1996) Complex adaptations and the evolution of evolvability. *Evolution* 50, 967–976
39. Post, D.M. and Palkovacs, E.P. (2009) Eco-evolutionary feedbacks in community and ecosystem ecology: interactions between the ecological theatre and the evolutionary play. *Philos. Trans. R. Soc. B Biol. Sci.* 364, 1629–1640
40. Okasha, S. (2006) *Evolution and the Levels of Selection,* Clarendon Press
41. Laland, K.N. *et al.* (2011) Cause and effect in biology revisited: is Mayr's proximate-ultimate dichotomy still useful? *Science* 334, 1512–1516
42. Whitham, T.G. *et al.* (2006) A framework for community and ecosystem genetics: from genes to ecosystems. *Nat. Rev. Genet.* 7, 510–523
43. Müller, G.B. (2007) Evo-devo: extending the evolutionary synthesis. *Nat. Rev. Genet.* 8, 943–949
44. Laland, K.N. and Sterelny, K. (2006) Perspective: seven reasons (not) to neglect niche construction. *Evolution* 60, 1751–1762
45. Pigliucci, M. (2007) Do we need an extended evolutionary synthesis? *Evolution* 61, 2743–2749
46. Watson, R.A. *et al.* (2010) Optimisation in 'self-modelling' complex adaptive systems. *Complexity* 16, 17–26
47. Kirchner, M. and Gerhart, J. (1998) Evolvability. *Proc. Natl. Acad. Sci. U.S.A* 95, 8420–8427
48. Toussaint, M. and von Seelen, W. (2007) Complex adaptation and system structure. *BioSystems* 90, 769–782
49. Kashtan, N. *et al.* (2007) Varying environments can speed up evolution. *Proc. Natl. Acad. Sci. U.S.A.* 104, 13711–13716
50. Sniegowski, P.D. and Murphy, H.A. (2006) Evolvability. *Curr. Biol* 16, R831–R834
51. Kouvaris, K. *et al.* (2015) How evolution learns to generalise: principles of under-fitting, over-fitting and induction in the evolution of developmental organisation. *arXiv arXiv* 1508.06854
52. Clune, J. *et al.* (2008) Natural selection fails to optimize mutation rates for long-term adaptation on rugged fitness landscapes. *PLoS Comput. Biol.* 4, e1000187
53. Clune, J. *et al.* (2013) The evolutionary origins of modularity. *Proc. Biol. Sci.* 280, 20122863
54. Matthews, B. *et al.* (2011) Toward an integration of evolutionary biology and ecosystem science. *Ecol. Lett.* 14, 690–701
55. Schoener, T.W. (2011) The newest synthesis: understanding the interplay of evolutionary and ecological dynamics. *Science* 331, 426–429
56. Wright, J.P. and Jones, C.G. (2006) The concept of organisms as ecosystem engineers ten years on: progress, limitations, and challenges. *BioScience* 56, 203–209
57. Powers, S.T. *et al.* (2011) The concurrent evolution of cooperation and the population structures that support it. *Evolution* 65, 1527–1543
58. Woods, H.A. and Wilson, J.K. (2013) An information hypothesis for the evolution of homeostasis. *Trends Ecol. Evol.* 28, 283–289
59. Lenton, T.M. and van Oijen, M. (2002) Gaia as a complex adaptive system. *Philos. Trans. R. Soc. B Biol. Sci.* 357, 683–695
60. Cropp, R. and Gabric, A. (2002) Ecosystem adaptation: do ecosystems maximize resilience? *Ecology* 83, 2019–2026
61. Lewontin, R.C. (1970) The units of selection. *Annu. Rev. Ecol. Syst.* 1, 1–18
62. Holling, C.S. and Gunderson, L.H., eds (2002) Resilience and adaptive cycles. In *Panarchy: Understanding Transformations in Human and Natural Systems*, pp. 25–62, Island Press

density under the same environmental conditions reduce resource competition or increase mutualism) [38].

3. Is individuality evolvable?

Evolutionary problem: Reproductive mechanisms defining the level of Darwinian individuality change from one level of organisation to another over evolutionary time. But: *How can selection at one level of biological organisation favour reproductive organisations that support individuality at a higher level before that new level of organisation exists?*

Learning theory insight: Deep learning aims to construct hierarchical or multiscale models. This can be achieved in a bottom-up layer-wise manner by using unsupervised learning at one level to reduce the dimensionality of the problem space and thus construct representations that are useful for the next level up, even before that next level exists.

Example prediction: Individual selection on reproductive mechanisms can create new evolutionary units that are adaptive at a higher level of biological organisation (even before that level of organisation exists) if the implicit adaptive problem posed by the environment is 'hierarchically decomposable' in a recursive manner [34,36].

63. Lawton, J.H. (1999) Are there general laws in ecology? *Oikos* 84, 177–192

Q16 64. Levin, S.A. (1998) Ecosystems and the biosphere as complex adaptive systems. *Ecosystems* 1, 431–436

65. Watson, R.A. *et al.* (2011) Global adaptation in networks of selfish components. *Artific. Life* 17, 147–166

66. Bettencourt, L. (2009) The rules of information aggregation and emergence of collective intelligent behavior. *Top. Cogn. Sci.* 1, 598–620

67. Panzarasa, P. and Jennings, N.R. (2006) Collective cognition and emergence in multi-agent systems. In *Cognition and Multi-Agent Interaction* (Son, R., ed.), p. 401, Cambridge University Press

68. Farmer, J.D. (1990) A Rosetta stone for connectionism. *Phys. D Nonlin. Phenom.* 42, 153–187

Q15 69. Power, D.A. *et al.* (2015) What can ecosystems learn? Expanding evolutionary ecology with learning theory. *Biol. Direct.* in press

70. Michod, R.E. (ed.) (1999) *Darwinian Dynamics, Evolutionary Transitions in Fitness and Individuality*, Princeton University Press

71. Michod, R.E. (2007) Evolution of individuality during transition from unicellular to multicellular life. *Proc. Natl. Acad. Sci. U.S.A.* 104, 8613–8618

72. Szathmáry, E. (2015) Toward major evolutionary transitions theory 2.0. *Proc. Natl. Acad. Sci. U.S.A.* 112, 10104–10111

73. Maynard Smith, J. and Szathmary, E. (1995) *Major Transitions in Evolution,* W.H. Freeman

74. West, S.A. *et al.* (2015) Major evolutionary transitions in individuality. *Proc. Natl. Acad. Sci. U.S.A.* 112, 10112–10119

75. Jablonka, E. and Szathmáry, E. (1995) The evolution of information storage and heredity. *Trends Ecol. Evol.* 10, 206–211

76. Ryan, P.A. *et al.* Social niche construction and evolutionary transitions in individuality. *Biol. Philos.* (in press)

77. Mills, R. *et al.* (2014) Transforming evolutionary search into higher-level evolutionary search by capturing problem structure. *IEEE Trans. Evol. Comp.* 18, 628–642

78. Watson, R.A. *et al.* (2009) Can selfish symbioses effect higher-level selection? In *Advances in Artificial Life. Darwin Meets von Neumann* (Kampis, G. *et al.*, eds), pp. 27–36, Springer

79. Watson, R.A. *et al.* (2011) Transformations in the scale of behavior and the global optimization of constraints in adaptive networks. *Adapt. Behav.* 19, 227–249

80. Adams, P. (1998) Hebb and Darwin. *J. Theoret. Biol.* 195, 419–438

81. Fernando, C. *et al.* (2012) Selectionist and evolutionary approaches to brain function: a critical appraisal. *Front. Comput. Neurosci.* 6, 24

82. Watson, R.A. *et al.* (2010) Associative memory in gene regulation networks. In *Proceedings of the Artificial Life Conference XII* (Fellermann, H. *et al.*, eds), pp. 194–202, MIT Press

83. Margulis, L. (1993) Origins of species: acquired genomes and individuality. *BioSystems* 31, 121–125

84. Sigmund, K. and Szathmáry, E. (1998) On merging lines and emerging units. *Nature* 392, 439–441

85. Maynard Smith, J. and Száthmary, E. (1993) The origin of chromosomes I. Selection for linkage. *J. Theoret. Biol.* 164, 437–446