

## 1 *Trends in Ecology & Evolution: Opinion article*

## *How can evolution learn?*

Richard A. Watson<sup>1</sup> and Eörs Szathmáry<sup>2</sup>

<sup>1)</sup> Dept. Computer Science/ Institute for Life Sciences, University of Southampton,  
Highfield, Southampton, SO17 1BJ, U.K..

<sup>2)</sup> The Parmenides Foundation, Kirchplatz 1, 82049, Pullach / Munich, Germany.

Corresponding author: Watson, R.A. (r.a.watson@soton.ac.uk)

## 8

## 9      Abstract:

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.

20 **Learning and evolution**

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

31 How intelligent is evolution?

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

41 Yet, this is not the whole story. Whilst genetic variation might be undirected, the pattern of  
42 phenotypic variation is shaped and biased by the processes of development. Moreover, the  
43 organisation of developmental processes (from gene-regulatory interactions to morphological  
44 body-plans) is itself, in large part, a product of past evolution. This affords the possibility that

45 random genetic changes might produce phenotypic changes that are ‘informed’ by past  
46 selection [7-9]. This can direct phenotypic variation into different or higher-level  
47 morphological dimensions and/or modularise phenotypic features and re-deploy them in new  
48 contexts [10,8,11]. The question thus arises; is evolution by natural selection (e.g., by adapting  
49 the organisation of developmental processes) able to facilitate subsequent adaptation in the  
50 same way that a learning system can exploit knowledge from past experience? If so, evolution  
51 might be a ‘smarter’ problem solver than generally appreciated [12] and learning theory could  
52 explain how.

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

60 This opinion paper discusses how: a) Recent work shows that the link between learning and  
61 evolution is a mathematical equivalence, b) Accordingly, knowledge from the theory of  
62 learning can be converted and redeployed in evolutionary theory, c) This offers exciting  
63 opportunities to address fundamental evolutionary puzzles in new ways.

64

## 65 **Unifying learning and evolution**

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

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

82 Can evolutionary systems ‘anticipate’ future outcomes?

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

91 The clever part of learning methods concerns how behaviours are parameterised in this  
92 *model space*. In a good model space, desirable future behaviours should be similar (nearby) to

93 behaviours that were useful in the past. For example, perhaps ‘eating apples’ should be close  
94 to ‘eating pears’ but far from ‘eating red things’.

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

107

108 [Box 1]

109

110 For example, in a modular problem, where features in different modules are approximately  
111 independent but features in the same module are not, then effective generalisation would be  
112 provided by new combinations of modules. Genetically, free recombination would disrupt  
113 modules and asexual reproduction would fail to exploit the independence of one module from  
114 another. An appropriate compromise is provided by an intermediate level of recombination,  
115 such as when nucleotides within genes do not recombine, but genes do. Given intragenic  
116 epistasis but not intergenic epistasis, the generalisation this provides explains a significant

117 advantage for sex [28]. However, this relies on an *a priori* correspondence between the physical  
118 linkage of components and their epistatic dependencies [28].

119 Can evolution learn like neural networks learn?

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

133 Again, mathematical equivalences with evolution have recently been shown [26,32].  
134 For evolution, learning of this type requires separating phenotypes from genotypes and  
135 evolving the parameters of a mapping between them. When there is heritable variation in this  
136 mapping that affects phenotypic correlations, natural selection inevitably favours changes that  
137 adhere to correlation learning principles [26,32]. A minimal example is the evolution of a single  
138 ‘relational’ allele, causing subsequent mutations to produce correlated variation in two  
139 phenotypic traits [33] (e.g. via pleiotropy). Pavlicev et al. showed that selection on relational  
140 alleles increases phenotypic correlation if the traits are selected together and decreases it if they  
141 are selected antagonistically (*Hebbian learning*) [33,32]. This simple step from evolving traits

142 to evolving correlations between traits is crucial; it moves the object of natural selection from  
143 fit phenotypes (which ultimately removes phenotypic variability altogether), to the control of  
144 phenotypic variability.

145

146 [Box 2]

147 In larger biological networks, this principle has the same effect as it does in larger neural  
148 networks (hence *evolutionary connectionism* [37]). In the Hopfield network [6], for example,  
149 this type of learning is sufficient for simple cognitive behaviours such as forming an associative  
150 memory capable of storing and recalling multiple distinct activation patterns, and effective  
151 generalisation in numerous recognition and classification tasks [32,78]. Watson et al.  
152 demonstrated conditions where evolved gene-regulation networks produce exactly the same  
153 behaviours [32]; forming a distributed ‘developmental memory’ of multiple phenotypes  
154 selected in the past, and generalising by producing new combinations of phenotypic modules  
155 (Fig. 1).

156

157 [Figure 1]

158

159 These results, and others [37,26,40,41], demonstrate that evolution and learning are not  
160 merely analogous processes but (different instantiations of) the same algorithmic principles.  
161 Transfer of specific models and results between these intellectual domains is already proving  
162 productive. Whilst it is important to apply analogies critically, learning theory is not just one  
163 thing – the issue is not so much to determine where the analogy breaks down, but to find the  
164 right kind of learning theory for each of the biological phenomena that are in need of  
165 explanation. Below we discuss three examples where learning theory makes sense of biological

166 ideas that are currently confusing, suggesting predictions that arise by transferring well-known  
167 learning results.

168 **Future Prospects: Understanding how evolution transforms itself**

169 Learning theory offers new concepts and theoretical tools for tackling several important  
170 puzzles in contemporary evolutionary biology. We identify specific learning models that  
171 inspire new approaches to key open questions in evolutionary developmental biology (evo-  
172 devo), evolutionary ecology (evo-eco) and evolutionary transitions in individuality (or “evo-  
173 ego” [37]) (**Outstanding Questions, Box 3**). Each of these areas is challenging for  
174 evolutionary theory because they involve feedbacks where the *products* of evolution modify  
175 the *mechanisms* of the evolutionary process that created them (Fig. 2) [33,45-49]. Although it  
176 is clear that the processes of variation, selection and reproduction underpinning evolutionary  
177 adaptation are not constants in natural populations, theoretical treatments of ‘modifier alleles’  
178 that enable selection to act on these processes are currently very limited. There is growing  
179 recognition that an integrated framework that puts such feedbacks front-and-centre is  
180 desireable [42-44]. Learning theory is precisely the study of processes that change over time  
181 as a function of past experience [14,31,34]. It is thus ideally suited to describing, not just how  
182 variation, selection and inheritance adapt phenotypes, but how natural selection modifies  
183 variation, selection and inheritance over time. We note that feedbacks on these three processes  
184 result in correlations or co-variance between components that were previously independent [37]  
185 (Fig. 2). Learning theory has well-understood models for each case.

186 **a) Evo-devo: the evolution of evolvability and correlation learning**

187 The evolution of developmental biases and constraints, accumulated over past selection, *might*  
188 improve the distribution of phenotypes explored in the future [52,45,54,7]. But the core issue

189 in the evolution of evolvability [7,51,52,45] is that selection cannot favour traits for benefits  
190 that have not yet been realised [9,53].

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

197 Whilst generalisation is not always easy, it does not require clairvoyance – it simply requires  
198 the ability to find structural regularities that are deep enough to be invariant over time [26].  
199 Accordingly, the possibility that evolution can learn from experience to favourably bias future  
200 exploration need not be any more mysterious than the basic result that learning from a training  
201 set can produce good generalisation on an unseen test set [55]. This also sheds light on the  
202 tension between robustness and evolvability. Here the problem is that adapting variation  
203 mechanisms so that they are *less* likely to produce *deleterious* variants (e.g. via canalisation)  
204 is often more immediately advantageous than adapting them so that they are *more* likely to  
205 produce *adaptive* variants [9,56]. Learning theory understands this tension extremely well.  
206 Specifically, *over-fitting* occurs when learning improves performance on training data but  
207 worsens performance on test data. To avoid this, the complexity of a model can be limited (e.g.  
208 by applying a *parsimony pressure*) to prevent memorisation of unnecessary details and force  
209 solutions to capture deeper regularities (Box 1). This explains why a cost of connections  
210 increases evolved modularity and improves evolvability [55,57].

211 Using past experience to favourably direct future behaviour is a hallmark of intelligence.  
212 By showing that incremental adjustment in the parameters of an appropriate model is sufficient  
213 to achieve this, learning theory puts this behaviour within reach of evolution by natural

214 selection, and identifies conditions where it can learn to favourably direct future exploration  
215 (Box 3, prediction 1).

216 **b) Evo-eco: ecological organisation and unsupervised correlation learning**

217 Organisms can modify their biotic and abiotic environment and thereby alter the selective  
218 pressures that act on themselves [43,46,48,58-61]. By modifying the network of ecological  
219 dependencies with other species, this *might* result in ecological organisations that increase the  
220 self-regulation of ecosystem variables, the resilience of ecological networks, or the efficiency  
221 of resource utilisation [62-64]. But since ecosystems are not, in most cases, evolutionary units  
222 [49,50], such feedbacks could also result in effects that are destructive in the long term, e.g.  
223 making an ecosystem more brittle or susceptible to catastrophic regime change, decreasing  
224 total biomass, etc. At present, however, we have no general organising principles for  
225 understanding how the structural organisation of ecological networks changes over  
226 evolutionary time, nor how this affects ecological functions and dynamics [58,49,64-67].

227 A different type of learning is relevant here. *Unsupervised* learning mechanisms do not  
228 depend on an external reward signal. By reinforcing correlations that are *frequent*, regardless  
229 of whether they are *good*, unsupervised correlation learning can produce system-level  
230 behaviours without system-level rewards (Box 2.ii). This can be implemented without  
231 centralised learning mechanisms as in *connectionist* models of intelligence [31,6,37] or  
232 distributed multi-agent systems [35] (simple forms of *collective intelligence* [35,68-70]).

233 Recent theoretical work shows that selection acting only to maximise *individual* growth rate,  
234 when applied to inter-specific competition coefficients within an ecological community,  
235 produces unsupervised learning at the system level [38] (Box 2.iii). This is an exciting  
236 possibility because it means that, despite not being a unit of selection, an ecological community  
237 might exhibit organisations that confer coordinated collective behaviours, e.g. a distributed  
238 ecological memory that can recall multiple past ecological states [38].

239 Learning theory shows that incremental adjustment in the parameters of individual  
240 behaviours is sufficient to achieve such collective behaviours [35], putting them within reach  
241 of individual-level selection. Accordingly, learning theory describes conditions where  
242 individual-level natural selection might facilitate ecological organisation and collective  
243 behaviour (Box 3, prediction 2).

244 **c) Evo-ego: The evolution of individuality and deep correlation learning**

245 In the major evolutionary transitions [47,71-73] “entities that were capable of independent  
246 replication before the transition can replicate only as part of a larger whole after the  
247 transition”[74,73]. These transitions in individuality [75,47,71] involve the evolution of new  
248 mechanisms of inheritance or reproductive codispersal (e.g., vertical genetic transmission,  
249 compartmentalisation, reproductive linkage) [76,74,73,77] that create new evolutionary units.  
250 But there’s a catch: If individual and group interests are aligned then selection applied at the  
251 group level doesn’t alter evolutionary outcomes, and if individual and group interests are not  
252 aligned then individual-level selection will oppose the creation and maintenance of adaptations  
253 that enforce selection at the group level [47]. Given this, how can evolution at one level of  
254 biological organisation systematically create reproductive organisations that facilitate non-  
255 trivial adaptation at a higher level of organisation before that level of adaptation exists?

256 In neural networks, *deep learning* [29] exploits correlation learning at multiple-scales to  
257 build multivariate models (Box 1). *Deep belief nets* [29], an exciting recent development  
258 igniting renewed interest in neural networks, achieve this in a bottom-up fashion, ‘freezing’  
259 each layer before adding the next. This creates the need to infer low-level representations that  
260 are useful for learning higher-order representations before the higher level of representation  
261 exists. Unsupervised learning provides a solution. By reducing the effective dimensionality of  
262 the data it ‘primes’ good performance at the next layer even though it is not informed by what  
263 the data will be used for at the next level [78]. In evolutionary systems, selection at one level

264 of organisation can operate like unsupervised learning at a higher level of organisation (Box  
265 2.iii)[38]. Abstract models incorporating these features show that individual-level selection can  
266 thus prime the systematic formation of adaptive higher-level evolutionary units without pre-  
267 supposing selection at the higher level [79,80]. New optimisation methods based on these  
268 principles demonstrate problem-solving capabilities that cannot be achieved with single-level  
269 adaptation [79,36]. We think this suggests that such *deep optimisation* principles might explain  
270 how evolutionary transitions facilitate *deep evolution*<sup>i</sup>, i.e. the evolution of adaptive biological  
271 complexity through successive levels of biological organisation [37,47,71-73] (Box 3,  
272 prediction 3).

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

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

## 283 **Conclusions**

284 Learning and evolution share common underlying principles both conceptually and formally  
285 [16,18-22,26,32,37,38,41]. This provides access to well-developed theoretical tools that have  
286 not been fully exploited in evolutionary theory (and conversely, suggests opportunities for

---

<sup>i</sup>Thanks to Seth Bullock for suggesting the term “deep evolution”.

287 evolutionary theory to expand cognitive science [81,82]). Learning theory is not just a different  
288 way of describing what we already knew about evolution. It expands what we think evolution  
289 is capable of. In particular, it shows that via the incremental evolution of developmental,  
290 ecological or reproductive organisations natural selection is sufficient to produce significant  
291 features of intelligent problem-solving.

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

299  
300

## 301 Acknowledgements

302 Thanks to Daniel Power, Gunter Wagner, Louis Kounios, Kostas Kouvaris, Chrisantha Fernando,  
303 Daniel Weinreich, Rob Mills, Paul Ryan, Adam Jackson, Simon Powers, Simon Tudge, Chris Cox,  
304 Markus Brede, Samir Okasha, Seth Bullock and Chris Buckley for formative discussion.

305

## 306 Glossary:

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

313 **Associative learning/memory:** Learning correlations between inputs and outputs, or learning what features co-  
314 occur in the input [6,42,78,26]. Associative memory is an ability to recall a pattern from a stimulus e.g.

315 “Darwin” → “Evolution”, “Hebb” → “Learning” (hetero-associative memory), or to recall a complete pattern  
316 from a noisy or partial stimulus, e.g. “Cha-les –ar-in” → “Charles Darwin”, “-ona-d H-b-” → “Donald Hebb”  
317 (auto-associative memory) [6]. Analogue of e.g. evolving the mapping between genotype and phenotype, or  
318 the correlations among phenotypic features governed by developmental interactions [33,39,32].

319 ***Bayesian learning (Bayesian updating)***: A learning method using Bayes rule as a principled way to incorporate  
320 new information with past experience. Analogue of selection in asexual population (replicator equation)  
321 [20,31].

322 ***Bivariate model***: A model that captures pairwise interactions between features (aka correlation model restricted  
323 to pairwise correlations).

324 ***Correlation learning***: see ***associative learning***

325 ***Deep learning***: Learning high-level representations by learning correlations on top of correlations, etc. Levels can  
326 be learned simultaneously [31], or one at a time (deep belief networks) [29].

327 ***Evo-devo***: Evolutionary developmental biology [7,42]. Here we are particularly interested in the evolution of  
328 developmental organisations that change the co-variance of phenotypic traits (analogue of correlation  
329 learning) [32].

330 ***Evo-eco***: Evolutionary ecology [67,58,46,59,49]. Here we are particularly interested in the evolution of ecological  
331 relationships that change the co-selection of species (analogue of unsupervised correlation learning) [38].

332 ***Evo-ego***: The evolution of Darwinian individuality [83,74,71,72]. We propose the term “evo-ego” [37] to refer to  
333 the evolution of organisations (reproductive structures) that change the evolutionary unit – i.e. the level of  
334 biological organisation that exhibits heritable variation in reproductive success [47]. Here we are particularly  
335 interested in the evolution of reproductive relationships that change the co-inheritance of fitness  
336 differences [77]. This includes new modes of reproduction modifying the heritability of collectives [47,80]  
337 (e.g. vertical transmission of symbionts, as in the origin of eukaryote organelles [83,84]), the origin of  
338 chromosomes (via physical linkage of previously independently-replicating genetic material [85]), changing  
339 reproduction from migrant pool reproduction to group-fissioning [72], or encapsulation in compartments  
340 (e.g. cell membranes, as in evolutionary transition from replicators on a surface to replicators in  
341 compartments) [84,73].

342 ***Evolutionary connectionism***: A developing theory for the evolution of biological organisation based on the  
343 hypothesis that the positive feedback between network topology and behaviour, well-understood in neural  
344 network models (e.g. Hebbian learning), is common to the evolution of developmental, ecological and  
345 reproductive organisations [37,70,35,36,32].

346 ***Hebbian learning***: Learning that occurs by altering the strength of synaptic connections between neurons  
347 [14,6,31]. For example, ‘neurons that fire together wire together’ is a Hebbian learning principle that  
348 strengthens the connection between two neurons when they are activated at the same time or by the same  
349 stimulus. Pavlicev et al. [33] showed that the action of natural selection adheres to Hebbian principles when

350        acting on heritable variation that affects correlations (e.g. gene-regulatory connections [32]). Power et al.  
351        [38] show Hebbian learning in the evolution of ecological interactions.

352        **Hopfield network:** A simple type of neural-network model where each neuron is (potentially) connected to every  
353        other neuron bi-directionally [6,30]. The Hopfield network has been used as a mathematical model for (non-  
354        neural) dynamical systems and emergent collective behaviours in many different domains including gene-  
355        regulation networks and ecological networks.

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

361        **Major evolutionary transitions:** Evolutionary innovations that have changed the evolutionary unit (the level of  
362        biological organisation that exhibits heritable variation in reproductive success): from self-replicating  
363        molecules, to chromosomes, to simple cells, to multi-organelle eukaryote cells, to multicellular organisms, to  
364        social groups [74,73] (see evo-ego).

365        **Multivariate model:** A model that captures high-order correlations (greater than pairwise interactions) among  
366        features [31].

367        **Over-fitting:** The tendency of a learning algorithm to perform well on the training set but poorly on the test set  
368        due to fitting idiosyncrasies of the training set. Failure to generalise correctly [14]. Analogous to securing  
369        fitness benefits in current selective environment (robustness) at the expense of potential fitness benefits in  
370        future environments (evolvability).

371        **Parsimony pressure:** A technique used in learning that penalises model complexity to favour simple models over  
372        complex ones. Simple models often produce superior generalisation by alleviating over-fitting [14].

373        **Reinforcement learning:** Trial and error learning based on an evaluative or reward signal, providing the learner  
374        with a measure of the value or quality of a given solution or behaviour, used to amplify successful and  
375        reduce unsuccessful behaviours [14]. Analogue of selection.

376        **Supervised learning:** Learning that changes parameters of a behaviour in the direction that reduces the error (i.e.  
377        error = desired output-actual output). Sometimes implies an external teacher that knows the desired output  
378        – but in practice, it usually means simply modifying parameters by gradient descent on an error function  
379        (rather than by trial and error) [14]. The combination of random variation and selection can effect the same  
380        changes in a model given the same gradient; thus an analogue of selection for a particular target phenotype  
381        or phenotypes [32] (see also selection in varying environments or for a set of target phenotypes [8,51]).

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

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

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

389     ***Under-fitting***: The condition that a learned model has failed to accurately fit the training set. Contrast, over-  
390         fitting [14].

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

392     ***Unsupervised learning***: Learning that aims to optimise a task-independent criterion function based on current  
393         output only (e.g. stability or robustness of the output). Occurs without knowledge of a desired output  
394         function/external teacher, e.g. by reinforcing the current output regardless of its quality. In particular,  
395         unsupervised correlation learning, where correlations that are already frequent in the training data are  
396         reinforced (rather than correlations that are good w.r.t. a task, as in reinforcement/supervised correlation  
397         learning). The aim of unsupervised learning is to discover categories, clusters or regularities inherent in the  
398         training samples and hence reduce the effective dimensionality of the data [14] (see Box 2).

399

## 400     References

- 401     1. Hofstadter, D. (2001). Analogy as the Core of Cognition, in Dedre Gentner, Keith Holyoak, and Boicho  
402             Kokinov (eds.) *The Analogical Mind: Perspectives from Cognitive Science*, Cambridge, MA: The MIT  
403             Press/Bradford Book, 2001, pp. 499–538
- 404     2. Penn, D. C., & Povinelli, D. J. (2007). Causal cognition in human and nonhuman animals: A comparative,  
405             critical review. *Annu. Rev. Psychol.*, 58, 97-118.
- 406     3. Eigen, M. (1971). Selforganization of matter and the evolution of biological macromolecules.  
407             *Naturwissenschaften*, 58(10), 465-523.
- 408     4. Wilke, C. O. (2005). Quasispecies theory in the context of population genetics. *BMC evolutionary  
409             biology*, 5(1), 44.
- 410     5. Feynman, R. P., Leighton, R. B., & Sands, M. L. (1963). *The Feynman Lectures on Physics: Mainly  
411             Mechanism, Radiation and Heat*. Addison Wesley.
- 412     6. Hopfield (1982) Neural networks and physical systems with emergent collective computational abilities,  
413             *Proceedings of the National Academy of Sciences USA*, 79 (8) 2554-2558.
- 414     7. Gerhart, J., & Kirschner, M. (2007). The theory of facilitated variation. *Proceedings of the National  
415             Academy of Sciences USA*, 104(1), 8582-8589.
- 416     8. Parter, M., Kashtan, N., & Alon, U. (2008). Facilitated variation: how evolution learns from past  
417             environments to generalize to new environments. *PLoS Computational Biology*, 4(11), e1000206.
- 418     9. Pigliucci, M. (2008). Is evolvability evolvable?. *Nature Reviews Genetics*, 9(1), 75-82.

- 419 10. Kashtan, N., & Alon, U. (2005). Spontaneous evolution of modularity and network motifs. *Proceedings*  
420 *of the National Academy of Sciences USA*, 102(39), 13773-13778
- 421 11. Wagner G.P., Pavlicev, M. & Cheverud (2007) The road to modularity. *Nature Reviews Genetics* 8: 921-  
422 931.
- 423 12. Fernando, C. (2013). Design for a Darwinian brain: part 1. philosophy and neuroscience. In *Biomimetic*  
424 *and Biohybrid Systems* (pp. 71-82). Springer Berlin Heidelberg.
- 425 13. Turing, A. M. (1950). Computing machinery and intelligence, *Mind*, 59, 236, 433-460
- 426 14. Mitchell, T. (1997). *Machine Learning*, McGraw Hill
- 427 15. Maynard Smith, J. (1986) *The Problems of Biology*. Oxford Univ. Press.
- 428 16. Frank, S. A. (1996) The Design of Natural and Artificial Adaptive Systems, Ch. 14 in *Adaptation*, Eds.  
429 M.R.Rose & G.V. Lauder. Academic Press.
- 430 17. Skinner, B. F. (1953). *Science and human behavior*. Simon and Schuster.
- 431 18. Bateson, G. (1979). *Mind and nature: A necessary unity* (pp. 98-100). New York: Dutton.
- 432 19. Chastain, E., Livnat, A., Papadimitriou, C., & Vazirani, U. (2014). Algorithms, games, and  
433 evolution. *Proceedings of the National Academy of Sciences*, 201406556.
- 434 20. Harper, M. (2009). The replicator equation as an inference dynamic. *arXiv preprint arXiv:0911.1763*.
- 435 21. Shalizi, C. R. (2009). Dynamics of Bayesian updating with dependent data and misspecified  
436 models. *Electronic Journal of Statistics*, 3, 1039-1074.
- 437 22. Hofbauer, J., & Sigmund, K. (1998). *Evolutionary games and population dynamics*. Cambridge University  
438 Press
- 439 23. Barton, N. H., Novak, S., & Paixão, T. (2014). Diverse forms of selection in evolution and computer  
440 science. *Proceedings of the National Academy of Sciences*, 111(29), 10398-10399.
- 441 24. Frank, S. A. (2009). Natural selection maximizes Fisher information. *Journal of Evolutionary  
442 Biology*, 22(2), 231-244.
- 443 25. Frank, S. A. (2012). Natural selection. V. How to read the fundamental equations of evolutionary change  
444 in terms of information theory. *Journal of evolutionary biology*, 25(12), 2377-2396
- 445 26. Valiant, L. (2013). *Probably Approximately Correct: Nature's Algorithms for Learning and Prospering in a  
446 Complex World*. Basic Books.
- 447 27. Valiant, L. G. (2007). Evolvability. In *Mathematical Foundations of Computer Science 2007* (pp. 22-43).  
448 Springer Berlin Heidelberg.
- 449 28. Watson, R. A., Weinreich, D. M., & Wakeley, J. (2011). Genome structure and the benefit of  
450 sex. *Evolution*, 65(2), 523-536.
- 451 29. Hinton, G. E. (2007). Learning multiple layers of representation. *Trends in cognitive sciences*, 11(10),  
452 428-434.
- 453 30. Hopfield & Tank (1986) Computing with neural circuits: A model. *Science* 233: 625-633.
- 454 31. Rumelhart, D.E., James L. McClelland and the PDP Research Group (1986) *Parallel Distributed  
455 Processing*. Vols. I and II. MIT Press.

- 456 32. Watson, R.A., Wagner, G. P., Pavlicev, M., Weinreich, D. M., & Mills, R. (2014). The evolution of  
457 phenotypic correlations and “developmental memory”. *Evolution*. 68(4) 1124-1138.
- 458 33. Pavlicev, M., Cheverud, J. M., & Wagner, G. P. (2011). Evolution of adaptive phenotypic variation  
459 patterns by direct selection for evolvability. *Proceedings of the Royal Society B: Biological  
460 Sciences*, 278(1713), 1903-1912.
- 461 34. Watson, R.A., Buckley, C.L., Mills, R. (2010) Optimisation in ‘Self-modelling’ Complex Adaptive Systems.  
462 *Complexity*. 16(5):17-26.
- 463 35. Watson, R.A., Mills, R., Buckley, C.L. (2011) Global Adaptation in Networks of Selfish Components.  
464 *Artificial Life*. 17(3):147-66.
- 465 36. Watson, R.A., Mills, R., Buckley, C.L. (2011) Transformations in the scale of behavior and the global  
466 optimization of constraints in adaptive networks, *Adaptive Behaviour* 19(4): 227-249.
- 467 37. Watson R.A., Mills, R., Buckley, C.L., Kouvaris, K., Jackson, A., Powers, S.T., Cox, C., Tudge, S., Ryan, P.,  
468 Davies, A., Kounios, L., Power, D. (to appear) Evolutionary Connectionism: Algorithmic Principles  
469 Underlying the Evolution of Biological Organisation in Evo-Devo, Evo-Eco and Evolutionary Transitions.  
470 *Evolutionary Biology* (also ECS Technical report, University of Southampton, *Soton eprints*, April 2015).
- 471 38. Power, D.A., Watson, R.A., Szathmáry, E., Mills, R., Powers, S.T. Doncaster, C.P. & Czapp, B. (to appear)  
472 What can ecosystems learn? Expanding evolutionary ecology with learning theory. *Biology Direct*; (also  
473 *arXiv preprint 1506.06374*)
- 474 39. Watson, R.A., Buckley, C.L., Mills, R.M. & Davies, A. (2010) Associative Memory in Gene Regulation  
475 Networks. *Artificial Life XII*. 194-202.
- 476 40. Börgers, Tilman, and Rajiv Sarin. "Learning through reinforcement and replicator dynamics." *Journal of  
477 Economic Theory* 77.1 (1997): 1-14
- 478 41. Friston, K. (2010). The free-energy principle: a unified brain theory?. *Nature Reviews Neuroscience*,  
479 11(2), 127-138.
- 480 42. Müller (2007) Evo–devo: extending the evolutionary synthesis. *Nature Reviews Genetics* 8(12) 943-949
- 481 43. Laland & Sterelny (2006) Perspective: seven reasons (not) to neglect niche construction. *Evolution* 60(9)  
482 1751-62.
- 483 44. Pigliucci, M. (2007). Do we need an extended evolutionary synthesis?. *Evolution*, 61(12), 2743-2749.
- 484 45. Wagner G.P. & L. Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution*  
485 50:967–976.
- 486 46. Post, D.M. & Palkovacs, E.P. (2009). Eco-evolutionary feedbacks in community and ecosystem ecology:  
487 interactions between the ecological theatre and the evolutionary play. *Philos. Trans. R. Soc. B, Biol. Sci.*,  
488 364, 1629–1640.
- 489 47. Okasha, S. (2006). *Evolution and the Levels of Selection*. Clarendon Press
- 490 48. Laland, K. N., Sterelny, K., Odling-Smee, J., Hoppitt, W., & Uller, T. (2011). Cause and effect in biology  
491 revisited: is Mayr’s proximate-ultimate dichotomy still useful? *Science*, 334(6062), 1512-1516

- 492 49. Whitham, T. G., Bailey, J. K., Schweitzer, J. A., Shuster, S. M., Bangert, R. K., LeRoy, C. J., ... & Wooley, S.  
493 C. (2006). A framework for community and ecosystem genetics: from genes to ecosystems. *Nature  
Reviews Genetics*, 7(7), 510-523.
- 495 50. Lewontin (1970) The Units of Selection, *Annual Review of Ecology and Systematics*, 1. 1-18.
- 496 51. Kashtan, N., Noor, E. & Alon, U. (2007) Varying Environments Can Speed up Evolution. *Proceedings of  
the National Academy of Sciences USA* , Vol. 104, No. 34 (Aug. 21, 2007), pp. 13711-13716.
- 498 52. Kirchner, M. & Gerhart, J. (1998). Evolvability. *Proceedings of the National Academy of Sciences USA*,  
499 95:8420-8427.
- 500 53. Sniegowski, P.D., Murphy, H.A. (2006). Evolvability. *Current Biology* 16:R831-R834.
- 501 54. Toussaint, M., & von Seelen, W. (2007) Complex adaptation and system structure, *BioSystems* 90: 769–  
502 782
- 503 55. Kouvaris, K., Clune, J., Kounios, L., Brede, M. Watson, R.A. (2015) How Evolution Learns to Generalise:  
504 Principles of under-fitting, over-fitting and induction in the evolution of developmental organisation.  
505 arXiv:1508.06854 preprint.
- 506 56. Clune, J., Misevic, D., Ofria, C., Lenski, R. E., Elena, S. F., & Sanjuán, R. (2008). Natural selection fails to  
507 optimize mutation rates for long-term adaptation on rugged fitness landscapes. *PLoS Computational  
508 Biology*, 4(9), e1000187.
- 509 57. Clune, J. Mouret J.B. & Lipson H. (2013) The evolutionary origins of modularity. *PRSB*, 280(1755).
- 510 58. Matthews, B., Narwani, A., Hausch, S., Nonaka, E., Peter, H., Yamamichi, M., Sullam, K.E. & Turner, C. B.  
511 (2011). Toward an integration of evolutionary biology and ecosystem science. *Ecology Letters*, 14(7),  
512 690-701.
- 513 59. Schoener, T. W. (2011). The newest synthesis: understanding the interplay of evolutionary and  
514 ecological dynamics. *Science*, 331(6016), 426-429
- 515 60. Wright, J. P., & Jones, C. G. (2006). The concept of organisms as ecosystem engineers ten years on:  
516 progress, limitations, and challenges. *BioScience*, 56(3), 203-209.
- 517 61. Powers, S.T., Penn, A.S., Watson, R.A. (2011) The Concurrent Evolution of Cooperation and the  
518 Population Structures that Support it. *Evolution*. 65(6):1527–1543.
- 519 62. Woods, H. A., & Wilson, J. K. (2013). An information hypothesis for the evolution of homeostasis. *Trends  
520 in Ecology & Evolution*, 28(5), 283-289
- 521 63. Lenton, T. M., & van Oijen, M. (2002). Gaia as a complex adaptive system. *Philosophical Transactions of  
522 the Royal Society B: Biological Sciences*,357(1421), 683-695.
- 523 64. Cropp, R., & Gabric, A. (2002). Ecosystem adaptation: Do ecosystems maximize resilience?. *Ecology*,  
524 83(7), 2019-2026.
- 525 65. Holling, C. S., & Gunderson, L. H. (2002). Resilience and adaptive cycles. *Panarchy: Understanding  
526 transformations in human and natural systems*, 25-62.
- 527 66. Lawton, J. H. (1999). Are there general laws in ecology?. *Oikos*, 84(2), 177-192.
- 528 67. Levin, S. A. (1998). Ecosystems and the biosphere as complex adaptive systems. *Ecosystems*, 1(5), 431-  
529 436.

- 530 68. Bettencourt, L. (2009). The rules of information aggregation and emergence of collective intelligent  
531 behavior. *Topics in Cognitive Science*, 1(4), 598-620.
- 532 69. Panzarasa, P., & Jennings, N. R. (2006). Collective cognition and emergence in multi-agent systems.  
533 *Cognition and Multi-Agent Interaction*, 401.
- 534 70. Farmer, J. D. (1990). A Rosetta stone for connectionism. *Physica D: Nonlinear Phenomena*, 42(1), 153-  
535 187.
- 536 71. Michod, R. (1999), *Darwinian Dynamics, Evolutionary Transitions in Fitness and Individuality*
- 537 72. Michod, R. (2007) Evolution of individuality during transition from unicellular to multicellular life.  
538 *Proceedings of the National Academy of Sciences USA*, 104 8613-8618
- 539 73. Szathmáry, E. (2015). Toward major evolutionary transitions theory 2.0. *Proceedings of the National  
540 Academy of Sciences USA*, 201421398.
- 541 74. Maynard Smith, J. & Szathmary, E. (1995) *Major Transitions in Evolution*. W. H. Freeman
- 542 75. West, S. A., Fisher, R. M., Gardner, A., & Kiers, E. T. (2015). Major evolutionary transitions in  
543 individuality. *Proceedings of the National Academy of Sciences USA*, 201421402
- 544 76. Jablonka, E., & Szathmáry, E. (1995). The evolution of information storage and heredity. *Trends in  
545 ecology & evolution*, 10(5), 206-211.
- 546 77. Ryan, P.A., Powers, S.T. and Watson, R.A. (accepted) Social niche construction and evolutionary  
547 transitions in individuality, *Biology and Philosophy*, to appear.
- 548 78. O'Reilly, R. C. & Munakata, Y. (2000) *Computational Explorations in Cognitive Neuro-science:  
549 Understanding the Mind by Simulating the Brain* (MIT Press, Cambridge, MA).
- 550 79. Mills, R., Jansen, T. & Watson, R.A. (2014) Transforming Evolutionary Search into Higher-Level  
551 Evolutionary Search by Capturing Problem Structure. *IEEE Trans. on Evolutionary Comp.* 18(5) 628-642.
- 552 80. Watson, R. A., Palmius, N., Mills, R., Powers, S. T., & Penn, A. (2009) Can Selfish Symbioses Effect  
553 Higher-level Selection? In *Advances in Artificial Life. Darwin Meets von Neumann* (2011) (pp. 27-36).  
554 Springer Berlin Heidelberg.
- 555 81. Adams, P. (1998). Hebb and Darwin. *Journal of theoretical Biology*, 195(4), 419-438.
- 556 82. Fernando, C., Szathmáry, E., & Husbands, P. (2012). Selectionist and evolutionary approaches to brain  
557 function: a critical appraisal. *Frontiers in computational neuroscience*, 6.
- 558 83. Margulis (1993) Origins of Species: Acquired Genomes and Individuality, *BioSystems*, Vol. 31(2-3)121-5.
- 559 84. Sigmund & Szathmáry (1998) On merging lines and emerging units. *Nature* 392, 439-441.
- 560 85. Maynard Smith & Szathmáry (1993) The origin of chromosomes I. Selection for linkage. *Journal of  
561 Theoretical Biology*, 164(4), 437-446.
- 562

563

#### FIGURE LEGENDS

564

565 **Fig. 1: A recurrent gene-regulation network (GRN) evolved in a varying environment exhibits associative  
566 learning behaviours.** See [32] for details. When a Hopfield network is trained on a set of patterns with *Hebbian  
567 learning* it forms an *associative memory* of the patterns in the training set. When subsequently stimulated with  
568 random excitation patterns, the activation dynamics of the trained network will spontaneously recall the  
569 patterns from the training set or generate new patterns that are generalisations of the training patterns [39,6,30].  
570 Here the evolution of connections in a GRN is shown to follow such Hebbian learning principles. The evolved  
571 GRN thus forms an associative memory of phenotypes that have been selected for in the past, spontaneously  
572 recreating these phenotypes as attractors of development with the GRN and also producing new phenotypes  
573 that are generalisations of them. A-D) A GRN is evolved to produce first one phenotype and then another in an  
574 alternating fashion [51,8]: A=Charles Darwin, B=Donald Hebb (who first described *Hebbian Learning*). The  
575 resulting phenotype is not merely an average of the two phenotypic patterns that were selected in the past (as  
576 per a univariate model or free recombination of phenotype pixels). Rather, different embryonic phenotypes (e.g.  
577 random initial conditions C and D) develop into different adult phenotypes with this evolved GRN match either  
578 A or B (one initial phenotype that falls into each developmental attractor is shown). These two phenotypes can  
579 be produced from genotypes that are a single mutation apart [32]. E-J) In a separate experiment, selection  
580 iterates over a set of target phenotypes (E-H). In addition to developing phenotypes that match patterns selected  
581 in the past (e.g. I), this GRN also generalises to produce new phenotypes that were not selected for in the past  
582 but belong to a structurally similar class e.g., by creating novel combinations of evolved modules (e.g.  
583 developmental attractors exist for a phenotype with all 4 ‘loops’ [32], J) – see also [8]. This demonstrates a  
584 capability for evolution to exhibit phenotypic novelty in exactly the same sense that learning neural networks  
585 can generalise from past experience [32].

586

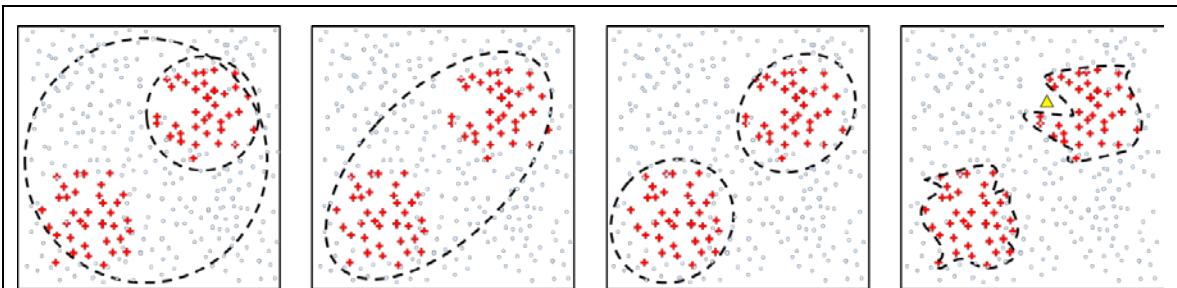
587 **KEY FIGURE. Fig. 2: Challenges in current evolutionary theory – caused when the products of evolution modify**  
588 **the mechanisms of evolution.** Key components of evolution by natural selection – variation, selection and  
589 inheritance [50] – are defined by structures (boxed) that are themselves modified by products of the  
590 evolutionary process (dotted arrows). ***Evo-devo:*** the evolution of developmental interactions modifies the  
591 distribution of phenotypic variation, ***Evo-eco:*** the evolution of ecological interactions modifies the structure of  
592 selective pressures, ***Evo-ego:*** the evolution of reproductive interactions (e.g. vertical transmission of symbionts,  
593 or transition from replicators on a surface to replicators in compartments) that modify evolutionary individuality  
594 by changing mechanisms of inheritance. These feedbacks are difficult to accommodate in evolutionary theory  
595 but well-studied in learning systems. We note that each of these feedbacks results in correlations or co-variance  
596 between components that were previously independent: a) The evolution of phenotypic correlations mean that  
597 traits do not vary independently, b) the evolution of ecological dependencies mean that selection pressures on  
598 one species are not independent of the selective pressures on another, and c) the evolution of new reproductive  
599 mechanisms mean that evolutionary units are not inherited independently. But, in evo-devo, correlations evolve  
600 within a single evolutionary unit; in evo-eco, correlations evolve between multiple evolutionary units, and; in  
601 evo-ego, correlations change the evolutionary unit (such that multiple, previously separate units, become a new  
602 single unit at a higher level of organisation) [37]. Learning theory has models that correspond to each of these  
603 cases.

604

## DISPLAY BOXES

605

606



- i) univariate models, too simple (under-fitting)
- ii) linear correlation model, poor fit (under-fitting)
- iii) non-linear correlation model (good fit)
- iv) multivariate model, too complex (over-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*) [14,31,26]. Consider a sample of points, e.g. phenotypes, characterised by two features or traits, some of which belong to a particular class ('+')<sup>1</sup>, e.g. high-fitness phenotypes. Learning which feature values are fit on average implicitly represents the class by a region in this two-dimensional space (i). This enables a limited sense of generalisation, e.g. 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 (*under-fitting*), 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 a developmental mapping between genotypes and phenotypes that introduces phenotypic correlations. ii) A linear correlation model (like a linear genotype-phenotype mapping [33]), e.g., representing that *trait 1* works well only when *trait 2* has a similar value, can improve the fit to some extent. iii) However, a non-linear correlation model is the simplest model capable of representing multi-modal distributions [32], e.g., 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 non-linear model onto the outputs of another (hence *deep learning* [29]). iv) In general, multivariate models can represent any data arbitrarily accurately [31]. However, fitting a multivariate model by incremental improvement (learning or evolution) can be troublesome if it is unnecessarily complex. One fundamental problem is *over-fitting*, where fitting the idiosyncrasies of the training data results in a model that fails to generalise well, excluding some potentially desirable points (triangle).

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 [34,35,36].

607

608 BOX 1: Learning (and evolution) as model fitting

609

- 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  $\Delta w_{ij}$  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, i.e. 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 [34,35]. 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 [35] (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 [35]. 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 [38].

610

BOX 2: supervised and unsupervised correlation learning and the level of selection

611

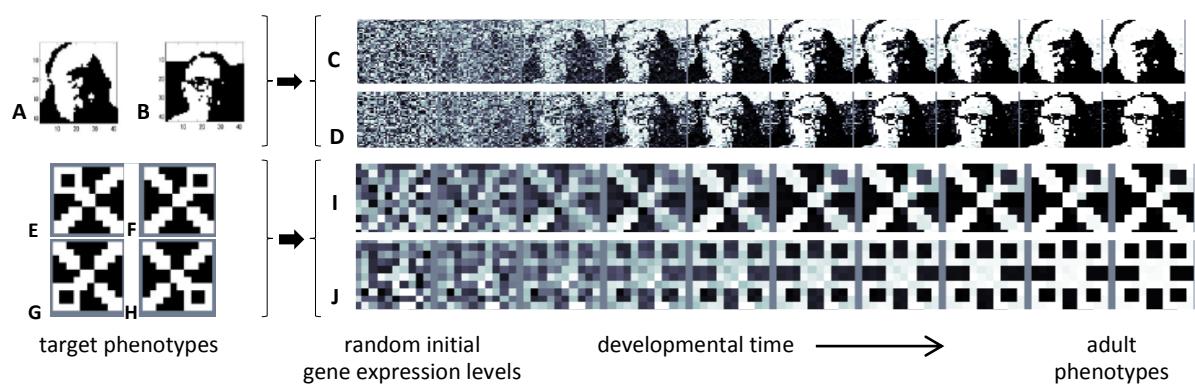
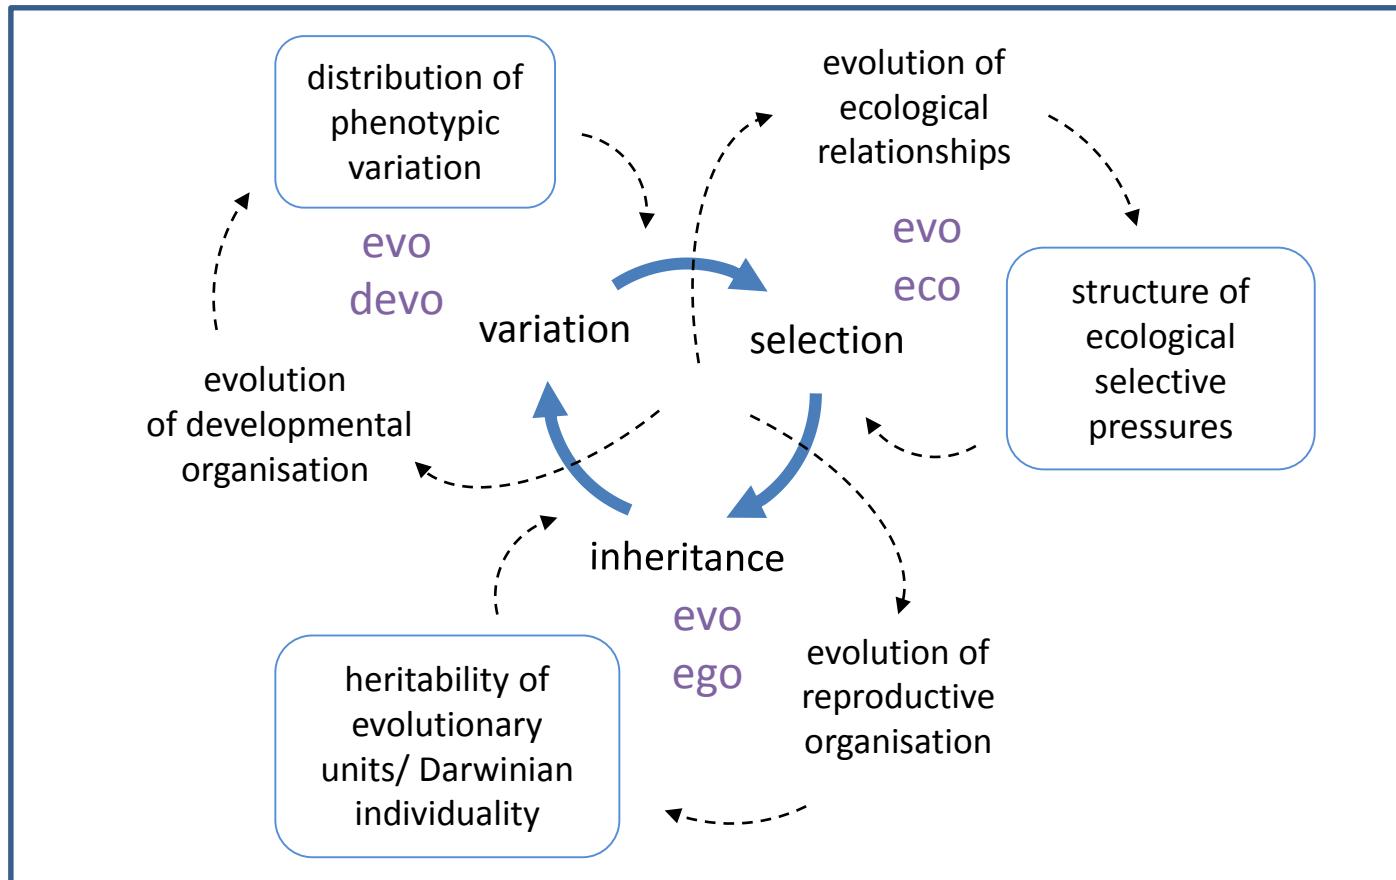


Fig. 1.



### **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, e.g. 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 [57,55].

### **2. Can ecosystem functions be adapted without ecosystem selection?**

**Evolutionary problem:** The organisation of ecological relationships in an ecosystem changes over evolutionary time due to 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 non-trivial 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 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 multi-scale 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 fashion [37,36].

**BOX 3. Outstanding Questions in 1) 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).**

### Trends Box

- 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.