

RESEARCH

Open Access



What can ecosystems learn? Expanding evolutionary ecology with learning theory

Daniel A. Power^{1*}, Richard A. Watson², Eörs Szathmáry³, Rob Mills⁴, Simon T. Powers⁵, C. Patrick Doncaster⁶ and Błażej Czapp⁶

Abstract

Background: The structure and organisation of ecological interactions within an ecosystem is modified by the evolution and coevolution of the individual species it contains. Understanding how historical conditions have shaped this architecture is vital for understanding system responses to change at scales from the microbial upwards. However, in the absence of a group selection process, the collective behaviours and ecosystem functions exhibited by the whole community cannot be organised or adapted in a Darwinian sense. A long-standing open question thus persists: Are there alternative organising principles that enable us to understand and predict how the coevolution of the component species creates and maintains complex collective behaviours exhibited by the ecosystem as a whole?

Results: Here we answer this question by incorporating principles from connectionist learning, a previously unrelated discipline already using well-developed theories on how emergent behaviours arise in simple networks. Specifically, we show conditions where natural selection on ecological interactions is functionally equivalent to a simple type of connectionist learning, 'unsupervised learning', well-known in neural-network models of cognitive systems to produce many non-trivial collective behaviours. Accordingly, we find that a community can self-organise in a well-defined and non-trivial sense without selection at the community level; its organisation can be conditioned by past experience in the same sense as connectionist learning models habituate to stimuli. This conditioning drives the community to form a distributed ecological memory of multiple past states, causing the community to: a) converge to these states from any random initial composition; b) accurately restore historical compositions from small fragments; c) recover a state composition following disturbance; and d) to correctly classify ambiguous initial compositions according to their similarity to learned compositions. We examine how the formation of alternative stable states alters the community's response to changing environmental forcing, and we identify conditions under which the ecosystem exhibits hysteresis with potential for catastrophic regime shifts.

Conclusions: This work highlights the potential of connectionist theory to expand our understanding of evo-eco dynamics and collective ecological behaviours. Within this framework we find that, despite not being a Darwinian unit, ecological communities can behave like connectionist learning systems, creating internal conditions that habituate to past environmental conditions and actively recalling those conditions.

Reviewers: This article was reviewed by Prof. Ricard V Solé, Universitat Pompeu Fabra, Barcelona and Prof. Rob Knight, University of Colorado, Boulder.

Keywords: Evolutionary ecology, Alternative stable states, Lotka-Volterra dynamics, Theoretical ecology, Community assembly, Network structures, Ecological memory, Associative learning, Regime shifts, Community matrix

*Correspondence: dap1e12@soton.ac.uk

¹Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, UK

Full list of author information is available at the end of the article

Background

With ever-increasing anthropogenic pressure on natural systems, it is vital to understand how the ecosystems we depend upon have been conditioned by evolutionary processes in historical environments which may have been very different from those they experience in the present day, and how any such conditioning may shape these systems' responses to new pressures. However, as ecosystems are not typically units of selection, we currently lack a framework linking adaptive pressures on individuals and populations to the dynamical properties of the systems they inhabit. In this article we investigate how systems above the Darwinian levels of selection may evolve collective behaviours, and observe a deep homology with emergent properties well understood in connectionist models of learning. We use this homology to develop theoretical analysis of emergent properties of natural selection in ecosystems, and explore the implications for community dynamics through Lotka-Volterra simulation.

Connections and collective behaviours in ecosystems

The structure and organisation of ecological interactions within biological communities causes them to exhibit many complex behaviours that are not straight-forwardly attributable to the summative behaviour of the individuals they contain [1–5]. For example, the structure of the network of interactions in an ecosystem [6, 7] affects many of the system's dynamical behaviours including succession dynamics and community assembly rules [8, 9], the stability, resilience and adaptive capacity of a community [10–14], the presence of alternative stable states [15–17], and the system's susceptibility to regime shifts [18].

From some points of view these system-level behaviours exhibit the appearance of design and/or characteristics in common with organismic functions such as development and complex phenotypes [19–23]. However, an ecological community is not, in most cases, an evolutionary unit [4, 23, 24]; it is an assemblage of species each individually adapted to their biotic and abiotic conditions. Thus the complexity that an ecosystem exhibits is not the product of Darwinian adaptation at the community level [25]. Furthermore, at present we lack general organisational principles that can help us understand and predict how system-level organisation and function results from the many individualistic adaptations on which they depend [3, 19, 23, 26–30], in particular, the reciprocity between the ecological dynamics on the network and the evolutionary changes to the nodes, and hence, connections of the network [30–37]. In short, we do not know how the coevolution of the parts affects the organisation and subsequent behaviour of the whole, i.e. the ecosystem's dynamical properties such as the location and number of its dynamical attractors; the trajectories it takes towards its these attractors (assembly rules); its stability during

assembly and/or succession; and its sensitivity to initial conditions during assembly.

Characterising how evolution and coevolution of the parts affects community-level properties is vital to understanding the responses of ecological communities to changes in environment at all scales. This issue is particularly acute in microbial community research, including medical applications in gut flora, where rapid evolution [38] has the potential to alter the function of those communities we depend upon most intimately, and where there is significant interest in how parental effects create a footprint of community composition that may be remembered throughout life [39]. Coevolutionary processes in gut microbiota have shaped multiple alternative stable states (termed enterotypes) [40], but it remains unclear how the historical conditioning of different communities' networks of interactions evolve in response to environmental changes in cases such as the use of antibiotics [41] and societal changes in diet [42], or how these changes affect the emergent properties of community networks [43] given the alternate enterotypes that act as attractors for these systems. At the macroscopic scale, Case et al. speculate that co-evolutionary processes maintain the distinct bird assemblages on the islands of Bali and Lombok, either side of Wallace's line [44–46]. Although birds are relatively unimpeded by the short stretch of sea that has separated terrestrial species, each island maintains distinct avian communities, and the conjecture is that long periods of coevolution within each community has created 'coevolved' biogeographic provinces; each network maintains a stable state resistant to invasion by members of the other [44]. Yet, without a framework linking microevolutionary changes in interactions between species pairs to dynamical behaviours of whole communities, it remains unclear as to whether a network of coevolved interactions could be the explanation for the observed dynamical stability.

The need to characterise the evolutionary and historical determinants of ecological processes is identified as an important frontier in ecological research [29]. Understanding the evolution and adaptability of ecological interactions is necessary, for example, to characterise the response of an ecosystem to climate change or other perturbations [14, 15, 47, 48] and, more generally, to understand how the number and location of dynamical attractors (alternative stable states) are affected by the organisation of ecological interactions acquired over evolutionary time [17].

These issues connect deeply with the phenomenon of *ecological memory* [29, 49–51] defined by Thompson et al. [29] as "the result of past environmental conditions and subsequent selection on populations [which] is encoded in the current structure of biological communities and reflected in the genetic structure of species".

As an illustrative example, consider the phenomenon of character displacement, [52, 53] in which niche divergence between pairs of isomorphic competitors leads to (genetic) trait divergence and increased likelihood of future coexistence; subsequent communities' distributions are shaped by the "memory" of past competition. Within ecosystems these genetic memories, distributed amongst multiple populations, influence assembly processes, stability, and resilience of the developing and mature ecosystem [29]. However, there remain many open questions about exactly how the microevolutionary modification of interspecific relationships by natural selection shapes any macroevolutionary memory at the community level [29, 54]. For example:

1. How do changes to interactions evolved in past environmental conditions alter the response of the community to future changes or perturbations in environmental conditions?
2. Is ecological memory merely a passive memory (like an imprint in clay) where the persistent effects of the most recent ecological states over-write or blend with those of older states, or can an ecological memory retain information about multiple distinct past states without just averaging them?
3. Can the assembly rules and succession dynamics of a community be systematically organised by selection in past environmental states?
4. How does the formation of an ecological memory affect the possibility of alternative ecological stable states, and regime shifts under subsequent environmental forcing?

The lack of a theoretical framework that links individual adaptations to collective behaviours leaves such questions unanswered. Our aim in this paper is thus to introduce such a framework. We do this by converting and exploiting theory that is already well-developed in another domain, namely, connectionist models of memory and learning. Below, we discuss the characteristics of connectionist models and their relationship to eco-evo dynamics. We then show a formal equivalence between these systems. We conclude that community-level organisation does not require community-level selection. The organisation of a community can be conditioned by past experience (collectively habituated to past environmental conditions) in the same sense, and with the same consequences for collective behaviours at the community level, as connectionist models of memory and learning. In order to demonstrate how making this general link between these disciplines leads to new insights about specific ecological behaviours, we then simulate mathematical models derived from this framework to address the above open questions about ecological memory.

Connectionist models of memory and learning

Connectionism is an approach to modelling cognition, in particular using neural networks, that explains how complex system-level behaviours can arise via the appropriate organisation of many simple components. The first important contribution of these models is to show that although each unit in a network might be very simple (e.g. the activation level of a neuron is simply a non-linear sum of the weighted connections from other neurons [55–57]), if appropriately organised/connected, a network of such units can provide many remarkable collective behaviours, including: a) forming a distributed memory for one or more configurations; b) pattern recognition from partial stimulus; c) the removal of noise from corrupted compositions; and d) classification of ambiguous inputs [55–58]. It has been noted in many different domains that the collective behaviours that can be exhibited by neural networks are not exclusive to neural models and can be exhibited by other types of dynamical systems (e.g. gene regulation networks, immune systems, multi-agent systems, economic systems and social networks) [59–65]. This includes ecological networks (where the growth rate of a species is modelled as a non-linear sum of the weighted fitness-interactions with other species) [64–66].

A deficit in the analogy between neural networks and ecosystems is that whereas neural networks acquire the organisation necessary for their collective behaviours through learning mechanisms designed for that purpose, ecological connections are modified by individual-level natural selection with no such system-level purposes in mind. Although there are similarities with activation dynamics in neural networks, there has not been any reason to expect that both systems may be organised in a similar manner. However, connectionist models also show that network organisations sufficient for many collective behaviours can be generated via learning mechanisms that modify the strength of connections according to only very simple and local reinforcement principles – even by mechanisms that do not require any system-level reward or performance-based feedback. The full significance of this for the evolution of ecological networks has not been previously appreciated [67].

Learning mechanisms in neural network models have two basic types [67]. *Supervised* learning utilises an external reward signal, or error function, to direct incremental changes to connections. We have recently demonstrated a formal equivalence between supervised learning and the evolution of connections in a network that is selected (at the system level) to produce a particular target phenotype or phenotypes [60, 67]. However, in the absence of a group selection mechanism there is no "target" phenotype directing selection at lower levels within ecological communities; supervised learning does not occur at this scale.

The other type of learning in these systems is *unsupervised* learning (Fig. 1), which operates without a reward signal. This may seem counter-intuitive but, when learning correlations or associations, learning what things “often go together” has many useful properties that can be attained without a supervisory signal to indicate what things “should go together” [56, 57], (Watson, R.A., Szathmáry, E.: “Learning in evolution”, opinion article. *Trends in Ecology & Evolution*, forthcoming), [68]. Thus, whereas supervised correlation learning reinforces correlations that are *good* according to some external reward signal, unsupervised correlation learning changes connections simply to reinforce correlations that are *frequent*. Hebbian learning [69] is the simplest unsupervised correlation learning mechanism and is well-understood in neural network models of memory and knowledge representation [56, 57, 68]. Under Hebbian learning, the change in strength of a synaptic connection, $\Delta\omega_{ij}$, is proportional to the co-activation of the neurons it connects: i.e. $\Delta\omega_{ij} = r x_i x_j$, where $r > 0$ is a learning rate, and x_k is the activation level of node k . This type of learning instantiates a very simple positive feedback principle between behaviour and connections, often paraphrased as “neurons that fire together wire together”. The effect of such changes is that *correlation becomes causation*, i.e. variables that happen to be both active at the same time (e.g., because they are stimulated by the same external conditions) become causally related by connections internal to the system, and thus their behaviour becomes more correlated in future. In this manner the network habituates to the perturbations it experiences by internalising information about the pattern of perturbation it has experienced into the organisation of its connections.

This simple principle is capable of producing many remarkable collective behaviours elucidated over more than 50 years of neural network research [56, 57, 68]. Famously, this includes the ability to develop of a *distributed associative memory* which can store and recall multiple patterns of activation in the organisation of synaptic connections [57], facilitating the use of these networks use in pattern recognition, noise reduction and classification (Fig. 1). A main contribution of this paper is to show that in ecological communities, given heritable variation in ecological relationships and certain conditions on ecological constraints, these positive feedback principles obtain from the action of individual natural selection [70]. Table 1 sets out the full analogy we make between connectionist learning in neural networks and eco-evo dynamics in ecosystems, starting with the previously recognised dynamical equivalence (Table 1a-f).

The biological evidence for positive feedback between ecological and evolutionary dynamics is entirely intuitive and already recognised in many areas where populations shape their future selective pressures (e.g. niche

construction [71] and refs. within) but the full implications of this feedback have not been realised [72]. For example, this feedback is part of the backstory involved in “invasional meltdown” [73–75] where species that have been in prolonged contact with one another in one environment facilitate one-another’s invasion into a new environment because they “have had a long evolutionary time to develop a cosy relationship with each other” [73]. Notice the simple positive feedback involved; species that occur in high density at the same time and under the same environmental conditions coevolve to become less competitive with each other over time. In turn, this reduction in niche overlap makes it more likely that they will coexist in high-density together in future. Our first key result is to formalise this principle with population genetics and show its equivalence with unsupervised correlation learning (Table 1g). Then, to demonstrate how this opens-up a transfer of concepts and results between these domains, we use numerical simulation to show conditions where an ecosystem can acquire, hold and recall distributed information about past environmental conditions – i.e. form a distributed ecological memory. This demonstrates several phenomena that are well-understood in connectionist models of memory and learning (Table 1h-m).

The formal link between the disciplines does not depend on the specific scenarios relevant to investigating ecological memory nor on the simplifications that are necessary for the simulation models; for this reason we divide our work into two parts: I) an analytic model and results concerning the general equivalence, II) simulation methods and results concerning ecological memory.

Methods

Methods part I: a model

Ecological dynamics

We model an ecosystem state as a vector of population densities over all N species, $X = \{x_1, x_2, \dots, x_N\}$, ($x_i \geq 0$), and an interaction network, or “community matrix” [6], Ω , where each element, ω_{ij} , represents the fitness effect of species j on species i relative to i on itself ($\omega_{ii} = -1$). We assume $\forall i, j : \omega_{ij} \leq 0$, e.g. competitive (non-trophic) interactions such as via competition for resources. A Lotka-Volterra competition equation (Eq. 1) defines the rate of change of density of a species as a function of its intrinsic growth rate and a weighted sum of interactions with all other species [66]:

$$\frac{dx_i}{dt} = \frac{m_i x_i}{k_{ie}} + \left(k_{ie} + \sum_{j=1}^N \omega_{ij} x_j \right) \quad (1)$$

where x_i is the density of species i , m_i is the intrinsic net growth rate of species i , k_{ie} is the carrying capacity of species i in environment e (i.e. its density

Box 1: Hopfield networks and unsupervised learning using Hebb's rule.

The Hopfield network model [56] originated from the hypothesis that it is the structure of connectivity between units in the central nervous system, rather than differences between the units themselves, that is most important in understanding the brain's complex behaviours. These simple models are fully connected networks of identical units. As units are identical it is solely differences in connections between units that determines the behaviour of each network. Despite this simplicity, these systems display complex behaviours, including the capacity to form multiple distributed memories; indeed, they are the simplest systems that have this capacity. Hopfield networks (and neural networks in general) are able to form multiple memories of configurations because, for each memory they store *correlations* between units, rather than the *states* (or outputs) of units. An effect of this architecture is that the structure of a Hopfield network can be updated to enable it to learn new patterns without over-writing and destroying pre-existing memories (figure 1). Hebb's rule is an unsupervised learning technique that can be used to train Hopfield networks to form memories for one or more configurations. For each training pattern Hebbian learning alters connections between units in the direction that reinforces the correlations between those unit's current outputs. If two units have the same state in multiple patterns they will become strongly correlated; Hebb's rule causes units that 'fire together' to become 'wired together'. Hebb's rule is an 'unsupervised' process because it does not utilise quality functions on the data used to train a network (whereas a supervised process might, for example, use a quality metric to scale changes made to the network). That is, Hebbian learning only acts to reinforce 'frequent' correlations in the training data, rather than correlations that are 'optimal' according to some metric of system performance (as in supervised learning methods).




Fig. 1 Training a Hopfield network using Hebb's rule. Network training: Unsupervised learning processes as used to train a Hopfield network to store two configurations, patterns A and B. Each unit in the Hopfield network corresponds to a pixel in the image display. Six units are highlighted to illustrate the changes to connections during training in pattern A. Hebb's rule alters connections between units such that units of the same sign (1:1 or -1:-1) become more correlated (blue lines) and units of opposite signs (1:-1 or -1:1) become more anti-correlated (red lines). Network behaviour: Training the network on both patterns results in a network with attractors (a.k.a. memories) for these patterns and system dynamics result in all initial conditions converging to one of the trained patterns (**a**). This behaviour enables these systems to be used for a variety of functions, including: (**b**) recovery of complete composition from partial input; (**c**) noise reduction; and (**d**) classification (the input image is a closer match for the plane configuration than the bird configuration)

before interspecific competition), and N is the number of species in the network. The dynamical equivalence between models of this sort, where populations experience exponential growth asymptotically approaching a threshold, and those models used in neural networks of excitation/inhibition between neurons is well recognised ([64] and refs. within). We now turn our attention to the selective pressures on interspecific interactions, and make

new comparisons with unsupervised learning in neural networks.

Evolution of interactions under individual selection

Each interaction coefficient summarises how a variety of structural, physiological and behavioural traits affect the degree to which one species impacts the population growth of another. Although subject to bio-physical

Table 1 Mechanistic equivalence between evo-eco dynamics and learning neural networks, and a map for the comparisons and analogies made in this paper

	Unsupervised correlation learning	Coevolution
	Activation dynamics	Population dynamics
a)	Neural activation level	Species density, x_i
b)	Neural activation pattern	Ecological state, $X = \{x_1, x_2, \dots, x_N\}$
c)	Synaptic connection strength, ω_{ij}	Inter-species fitness interaction, ω_{ij}
d)	Neural network (weight matrix, W).	Ecological network (community matrix, Ω).
e)	Neural activation dynamics: a non-linear weighted sum of inputs from other neurons (and external inputs).	Ecological population dynamics (Eq. 1): species growth is a non-linear function of the sum of weighted fitness interactions from other species (and environmental changes to carrying capacities).
f)	External input patterns (aka. 'training set').	Environmental forcing (in multiple environmental conditions)
	Correlation learning (unsupervised)	Evolution of interactions (individual selection only)
g)	Positive feedback between activation strengths and connection strengths – aka. <i>neurons that fire together wire together</i> . Unsupervised correlation learning mechanism, Hebb's rule: $\Delta\omega_{ij} = r x_i x_j$, where $r > 0$ is a learning rate.	Positive feedback between ecological densities and connections – or <i>species that occur together wire together</i> . Direct effects of individual natural selection on interactions: $v_{ij} = r x_i x_j$, where $r = \frac{m_i}{K_{ie}} \mu$ describes the available mutation (Eq. 3).
	Collective behaviours in neural networks (arising from e.g., Hebbian learning, Fig. 1)	Collective behaviours in ecosystems (arising from individual selection acting upon interspecific correlations)
h)	<i>Memory formation</i> (Fig. 1, top panel) Hebb's rule organises synaptic connections to reinforce the state of the system, decreasing sensitivity to changes in input.	<i>Ecological memory formation</i> (Fig. 3): natural selection organises ecological relationships in a manner that reinforces the current ecological state, decreasing sensitivity to changes in environmental conditions. (Attractors due to environmental variables become attractors of community dynamics [17].)
i)	<i>Distributed associative memory facilitates a memory of multiple patterns</i> (Fig. 1a): the capacity to store multiple patterns of activation in the organisation of synaptic connections and recall patterns from any initial conditions via activation dynamics.	<i>Formation of alternative stable states</i> (Fig. 5a): the creation of a distributed ecological memory in the network of species interactions results in a system with attractors that mimic past ecological states.
j)	<i>Pattern reconstruction</i> (Fig. 1b): the recall of a complete pattern from a partial stimulus.	<i>Ecological assembly dynamics</i> (Fig. 5b): reconstruction of a particular community composition, from a subset of that community.
k)	<i>Error correction</i> (Fig. 1c): the ability to remove noise from a pattern, repair imperfections and restore a complete pattern.	<i>Ecological resilience</i> (Fig. 5c): the ability to recover from perturbations in species densities and restore the complete community.
l)	<i>Recognition</i> or classification of an input or stimulus (Fig. 1d): return the nearest attractor from ambiguous initial conditions.	<i>Ecological sensitivity</i> to initial conditions (Fig. 5d): the switch-like change in response to small variation in initial species densities.
m)	<i>Holding state in dynamics</i> : Hopfield networks and other recurrent networks have an internal state that allows them to display temporal dynamics (independent of input).	<i>Ecosystems hold state in population dynamics</i> (Fig. 6): in systems with multiple attractors this results in a communities capable of hysteresis with tipping points between states.

a-f) The basic components of the analogy made in the introduction to this paper. g) The main contribution of this paper (discussed in part I) – the equivalence of individual natural selection acting on inter-species interactions with a simple associative learning rule such as Hebbian learning. Thus ecological networks *evolve* like neural networks *learn* (Fig. 4). h-m) From this the phenomenology shown in our experiments follows (simulation results, Figs. 3 and 5)

constraints, e.g., stoichiometric constraints on resources, these interactions can often be modified by the evolved characteristics of the constituent species, e.g. traits that alter the overlap of habitat preference or resource utilisation profiles [76] or the time, effort or energy expended on a particular ecological resource or relationship.

We assume that only individual-level selection acts on these interactions. We do not model selection on whole ecosystems (e.g., via a population of ecosystems), nor on species. Thus only changes to traits that directly affect the growth rate of an individual compared to the rest

of the individuals in the species can be selected. Individual selection acts to decrease the competitive effects from others by changing ω_{ij} ; but note that an individual has no intrinsic interest in altering the growth rate of others by changing ω_{ji} . Changes that decrease the density of a competitor, for example, cannot be selected for under individual selection as (in the absence of group selection) such changes benefit all individuals in a species [77]. Therefore any changes to a species growth rate that occur as a side-effect of altering the density of some other species (e.g., via changes to ω_{ji} or via ecological

trade-offs below) are not affected by individual selection (Appendix A).

Analysis of individual-level natural selection acting on ecological interactions

We analyse the rate of accumulation of favourable mutations, ν , in each interaction coefficient, ω_{ij} . In order to study the dynamical interaction between evolutionary and ecological dynamics, we are particularly interested in how the evolution of ω_{ij} is sensitive to the current species densities. The qualitative picture is as follows: Occasionally, mutants arise in species i that are identical to i except for the modification of an interaction coefficient with another species j' in the ecosystem. The origin and establishment of such a mutant can be modelled by applying population genetics theory [78, 79] to the particular case. From the ecological dynamics it follows that the selective coefficient, s , conferred by the change, g , in the interaction coefficient ω_{ij} , is the change in the invasion rate per capita of a mutant type of species i relative to the growth rate per capita of species i without the mutation:

$$s = \frac{m_i}{k_{ie}} \left(k_{ie} + \sum_{j=1}^N \omega_{ij} x_j + g x_{j'} \right) - \frac{m_i}{k_{ie}} \left(k_{ie} + \sum_{j=1}^N \omega_{ij} x_j \right) = \frac{m_i}{k_{ie}} g x_{j'} \quad (2)$$

(Simplified as $s = \frac{m_i}{k_{ie}} g x_{j'}$ henceforth). Since m , k and x are positive, a favourable mutation requires only $g > 0$. Qualitatively, this means that a mutation to an individual of one species, e.g. a change in its habitat or resource usage, is selected for if the mutation reduces the negative influence of another species on its growth rate. We assume that in all species such mutations occur at rate μ per individual per generation. In general, the rate of accumulation of such mutations is equal to the product of the number of individuals, x_i , the beneficial mutation rate, μ , and the average probability that a single new mutation will ultimately fix, $\bar{P}$, such that: $\nu = x_i \mu \bar{P}$ [79]. In large sexual populations with linked loci, $\bar{P}$ will depend on ν , and in different ways depending on the type of recombination, recombination rate, population size, the mutation rate and magnitude of mutations [78, 79]. For unlinked loci, in small populations, or under strong selection and weak mutation where mutations occur serially, $\bar{P}$ is proportional to the selection coefficient, s [79]. Since the effects we want to investigate do not depend on the effects of sexual recombination it is sufficient for our purposes to model the rate of adaptation in this simple manner. In this case, the rate of adaptation, ν_{ij} , in an interaction coefficient, ω_{ij} , is given by:

$$\nu_{ij} = x_i \mu s = \frac{m_i}{k_{ie}} g \mu x_i x_j \quad (3)$$

In more complex cases, where there is interference between alleles at different loci, ν_{ij} may not be linearly proportional to $x_i x_j$ as it is in Eq. 3, but in all cases, the rate of evolutionary change in an interaction coefficient increases with the product of x_i and x_j since mutations must be both created and selected in order for an interaction coefficient to evolve. This is robust to the choice of underlying model (Appendix B). This is entirely intuitive: a) if suitable heritable variation in relationships is available, natural selection always acts to reduce the negative effects of others, and b) the rate of adaptation of the interaction coefficient between two species, e.g. by character displacement, is driven by their co-occurrence [54].

This is our first key result, describing how selection acts on inter-species relationships as a function of the current ecological state (Table 1g). Equation 3 tells us that the rate of adaptation on inter-species relationships is proportional to the co-occurrence of the species involved: Hence, *species that occur together* (arise in high density at the same time and under the same conditions), *“wire” together* (and there will be selection for changes to interactions that makes those species more likely to co-occur in future) - as per the principle of unsupervised correlation learning. (Correlation learning can be produced either by a *reduction* in negative interactions, as here, or by an *increase* in positive interactions, with the same effect on system dynamics, i.e. either will increase the future co-occurrence of the species that have co-occurred in the past.) We now investigate the consequences of this finding for collective behaviours in an ecological community, using ecological memory as a case study.

Methods part II: simulation

In general, Eq. 1 may exhibit unstable or even chaotic behaviour. In the following investigations we restrict our simulations to interactions that are symmetric ($\forall i, j : \omega_{ij} = \omega_{ji}$) as per competition for shared resources or for competition coefficients estimated from utilisation functions [80], in which case the dynamics have only fixed point equilibria [81]. During simulation we allow the ecological dynamics to equilibrate at each time step (over τ iterations of Eq. 1). Then all interaction coefficients are updated according to the direct effect of natural selection in proportion to the rate of adaptation (Eq. 3). Then ecological constraints are applied to these interactions as follows, and the process is repeated.

Ecological constraints/evolutionary trade-offs on changes to ecological interactions

In ecosystems where niche space is saturated, the capacity of natural selection to alter interactions is subject to inevitable ecological constraints and evolutionary trade-offs that prevent selection from eliminating all competition. Individuals with traits that cause them to avoid

competition with one species may be forced to compete more with others. Thus the interaction between two species is more generally governed by a) the evolvable characteristics of the species as described by Eq. 3, and b) evolutionary trade-offs or ecological constraints applied by the physical properties of the environment (e.g., energy spent on exploiting one resource cannot also be spent exploiting another). Here these trade-offs are represented by normalisation conditions that conserve the sum of interactions to and, by symmetry, from each species. Specifically, for all species i and j ($j \neq i$), $\sum_{j=1}^N \omega_{ij}(t) = Q_i$ and $\sum_{i=1}^N \omega_{ji}(t) = Q_j$, where $Q_i = Q_j < 0$ is a constant (the sum of interaction terms in row/column i at time $t = 0$) (see Appendix C). Such normalisation represents ecological niches that resist change in width more than change in location, e.g. individuals can more easily change which resources they depend on than how dependent they are overall [82].

Although natural selection always acts to reduce competitive impacts from others, the fact that the rate of adaptation is greater for some competitive interactions than others (Eq. 3), together with these normalising evolutionary trade-offs, will mean that the competition between some species will increase. When the interaction, ω_{ij} , from some species j to a given species i is, for example, made *less* competitive (decreased in magnitude) by the evolution of heritable traits, all the other interactions involving i , i.e., ω_{ih} ($h \neq j$) and ω_{hi} ($h \neq i$), are made more competitive by these normalising evolutionary trade-offs. This, in turn, leaves all interactions not involving i relatively less competitive. Self-interactions are not modified by either evolutionary or normalisation mechanisms ($\omega_{ii} = -1$).

Environmental forcing

To investigate ecological memory we are interested in how the evolution of ecological interactions is influenced by past ecological states. To model the evolution of an ecosystem under varying environmental conditions that

force or drive the ecosystem to adopt different ecological states, we define two environmental conditions, E_1 and E_2 , that have differing effects on the carrying capacities of the species (Appendix C). Relative to a default environment E_0 , environment E_1 increases the carrying capacity of some species and decreases others, whilst in E_2 , a different subset of species is increased/decreased. E_1 and E_2 may represent hot-dry and cold-wet climates, for example; or high/low levels of some key broadly-utilised resource such as phosphorous input rates for a lake habitat [48]. Given that individuals from each species experience both conditions over evolutionary timescales, these conditions could vary in space (e.g. geographic localities, [33]), rather than in time (e.g. seasonal change). To make the effects of these two conditions on community composition easily identifiable we utilise environmental forcing patterns corresponding to two arbitrary but easily identifiable pictograms (Fig. 2). Here the *hot* and *cold* pictograms describe two different configurations of species densities representing, for example, hot dry savannah and cold wet/temperate ecological states, respectively. The environment is switched between E_1 and E_2 every T evolutionary updates.

Model parameters of the simulations and methods used for assessing ecological attractors are described in Appendix C.

Results

We use the series of four open questions concerning ecological memory listed above to exemplify some of the implications of our general result.

i) Changes to interactions evolved in past ecological states “canalise” the response of the community to subsequent changes in environmental conditions or future perturbations

Experiment 1 investigates how evolution in fixed environmental conditions changes the ecological dynamics of the




Fig. 2 The carrying capacities of 400 species are affected differently by two different environmental conditions, E_1 and E_2 . For our purposes, the specific patterns of carrying capacities for the two conditions are arbitrary and can thus be depicted by 20×20 pixel ‘hot’ (a) and ‘cold’ (b) pictograms where a black pixel at location x, y indicates an elevated carrying capacity, $(k_0 + \alpha)$, and a white pixel a depressed carrying capacity, $(k_0 - \alpha)$, for the $(20x + y + 1)^{th}$ species in the ecosystem, Appendix C (the two-dimensions of the pictogram are abstract and do not imply any spatial arrangement of the species)




Fig. 3 Ecological dynamics before and after evolution in E_1 . **a–b** Before evolution of interactions, when forcing is applied, some species densities increase, others decrease. **a** Four species responding differently to E_1 and E_2 (H='high', L='low'). **b** Vectors of all species population densities are displayed in a pixel array as per Fig. 2. Under a given pattern of environmental forcing (top row), an initially random pattern of species densities (middle row), equilibrates at a pattern of species densities (after τ timesteps) (bottom row). Initial species densities do not alter the attractor attained (5 independent examples). **c** After evolution of interactions in E_1 , equilibrium states are governed by that past pattern of environmental forcing and not by the current environment. This ecological memory is a stable attractor, reached from any initial pattern of species densities, regardless of the pattern of environmental forcing (some distortion is visible under E_2 forcing)

community. Before the evolution of interactions, during the ecological phase of simulation, the ecosystem arrives at a stable equilibrium corresponding to the pattern prescribed by the current environmental forcing (Fig. 3a and b). Inter-species interactions are then evolved in environment E_1 , i.e., without changes to the environmental forcing during evolution. The process is repeated for 800 ecological and evolutionary cycles. We then assess how evolved interactions have altered the sensitivity of the ecosystem to subsequent environmental forcing. We find that the ecosystem now arrives at a stable equilibrium corresponding to the E_1 pattern (the pattern it experienced when interactions were evolving) even when the environmental forcing is subsequently changed to E_0 or E_2

(Fig. 3c). Experiment 1 thus shows that the effect of evolving ecological interactions by individual natural selection under fixed environmental conditions is to create a stable attractor for the specific ecological state experienced in that past environment, reducing the responsiveness of the ecosystem to respond to subsequent environmental forcing, and increasing the adaptive capacity of the system to withstand changes to population densities. This behaviour demonstrates the basics of an ecological memory, but only a memory of one pattern. Even passive systems can remember one pattern, e.g. an imprint in clay, but connectionist models show that a dynamical network is capable of storing and recalling multiple patterns.

ii) Ecological memory can retain and recall information about multiple distinct past states

In Experiment 2 varying environmental conditions are applied to cause the ecosystem to adopt two different ecological states (E_1 and E_2) repeatedly whilst inter-species interactions are evolving. The effect of these evolved changes plus normalising evolutionary trade-offs are illustrated in Fig. 4. We see that their evolution is identical to Hebbian learning (again this is robust to the choice of underlying model, Appendix D).

After evolution we find that, in the absence of further environmental forcing, the ecological dynamics have two stable attractors corresponding to E_1 and E_2 , reached from any initial species densities (Fig. 5a, Appendix C). An ecological memory can thus retain information about multiple distinct past states without just averaging them or blending them (for example, the system does not have an attractor for the union of both patterns).

iii) The assembly rules of a community can self-organise to recreate past environmental states

After evolution in the varying environment (Experiment 2) either of the two patterns can be completely recalled or assembled from an initial subset of species. That is, when the initial species densities have just a few species present in a density that matches one of the previous patterns, the ecological dynamics act to recreate the full pattern to which that ‘partial stimulus’ belonged (Fig. 5b). This experiment also reveals more about how the stability and resilience of the community is affected by the presence of multiple memories. When the initial conditions are ‘corrupted’ versions of a previous pattern, the complete pattern is restored, even when the corruption is severe (Fig. 5c.) (thus maintaining/re-creating the current ecological pattern), and when the initial species densities partially resemble both patterns, the population dynamics ‘break symmetry’, causing all species to adopt the pattern to which the initial conditions are closest (thus ‘choosing’ between two ecological states - not blending them).

iv) Ecological memory can create multiple ecological stable states, and may exhibit critical transitions between them under subsequent environmental forcing

Figure 6 examines the response of the ecosystem to patterns of environmental forcing that change linearly from E_1 to E_2 . Before the evolution of interactions, the response of the ecosystem is proportional to the environmental forcing applied (Fig. 6a). In contrast, after the evolution of interactions (Experiment 2), the response of the system is discrete or switch-like and exhibits significant hysteresis (Fig. 6b). That is, as the pattern of environmental forcing moves incrementally from E_1 to E_2 , the response of the system is to stay on E_1 considerably past the mid-point and then suddenly switch to E_2 . Conversely, when the

environmental forcing is reversed, the ecosystem retains a configuration matching E_1 considerably past the mid-point before switching back to E_1 . The dynamics underlying this hysteresis loop are shown by the vector field of species densities changes (Figs. 6c, 9a). This also shows that the response of the population dynamics to perturbations in species densities slows down near the critical transition (consistent with early-warning signals for a tipping point [83, 84]). Figure 6d. shows how the response of the ecosystem to forcing changes over evolutionary time. Around generation 470, the ecosystem exhibits non-linear but non-catastrophic transitions [18]. The catastrophic regime change is not a general instability property of the evolved system – it only occurs when the environmental forcing is similar to a past state that is remembered by the ecosystem – unstructured forcing results in a linear response (Appendix E).


These observations demonstrate a conversion of one type of ecological alternate stable state into another. Beisner et al. [17] describe the “ecosystem” perspective on alternate stable states, which involves changes driven by abiotic environmental conditions, and the “community” perspective, which involves multiple attractor states that can exist under fixed environmental conditions. Figure 6 shows a system that converts alternate “ecosystem states” into alternate “community states”; thus converting patterns from past environmental states into ecological memories.

Figure 10 shows a ‘bestiary’ of ecological attractors changing over evolutionary time, showing some diversity before settling down to the two-attractors shown in Figs. 5 and 6. During long-term simulation we find that, as the forcing used to switch the system between attractors is of fixed value, while the effect of the evolved changes to interactions is ever-increasing, eventually the level of forcing applied is insufficient to shift the system from its current attractor. When this happens, the system becomes ‘stuck’ at one attractor, effecting a breakdown in observed behaviour (Appendix F, Figs. 10 and 11).

Discussion

Our results formalise the intuitive idea that individual selection on ecological interactions produces positive feedback on species co-occurrence. By recognising an equivalence between this feedback and principles of unsupervised correlation learning we are able to use concepts from connectionist models to understand and illustrate the consequences of these changes for system-level behaviours. This makes several contributions to our understanding of evo-eco interactions:

Evo-eco dynamics have predictable consequences for collective behaviours. The worked example developed in our simulations converts informal notions about the evolutionary and historical determinants of ecological



processes into a model that makes specific predictions about how past ecological conditions alter the selective pressures on the component species and hence modifies their future ecological behaviours. This presents a specific model for non-trivial ecological memory that can be empirically tested (Appendix G). From this model we better-understand the necessary conditions for a distributed ecological memory to form, such as the presence of evolutionary trade-offs that cause species to become *more dependent* on other species [85] rather than just becoming *less competitive* with them (Appendix H).

Ecological communities can exhibit organised collective behaviours. Under certain conditions, memories of past ecological states can be stored in a distributed way in the organisation of evolved ecological relationships. Such memories are not simply the summative or average result of multiple species each with individually alternate stable states. The connections that produce these behaviours are organised not by evolutionary adaptation at the community level, but rather by evolutionary adaptation at the individual level and 'past experience' of historic environmental conditions. The organisation of the system is thus conditioned by past environmental conditions, causing it to collectively habituate to the patterns of perturbation it has experienced [47].

Community assembly rules can be organised to re-assemble specific past states. The assembly of complete and specific past ecological states may be triggered by partial environmental cues or a small number of founders (as in invasional meltdown) (Fig. 5b) or, similarly, the system can recover each specific state from corruptions of that state (Fig. 5c). The learned assembly rules result in a system that 'classifies' initial compositions according to their similarity to past ecological states and will return community composition to the state that most closely resembles initial conditions (Fig. 5d). Ecological memories can thus direct subsequent community assembly to

recreate multiple past ecological states in a complex and collective, but predictable, manner.

Stability and resilience tends to increase but instability and regime shifts are also predictable. We find that evolutionary pressures on ecological interactions tend to increase ecosystem resilience (recovery after perturbations to species densities) (Fig. 3) and adaptive capacity (robustness to environmental forcing) (Figs. 3 and 6) [10, 11, 13, 16, 47]. However, if the evolutionary history of an ecosystem has included a multi-modal distribution of environmental conditions, then this can result in alternative stable states (rather than universal stability) and may exhibit critical transitions in changing from one state to another [10]. This switch-like change in the community (Fig. 6) is only exhibited when the forcing that is applied is similar to past forcing – when arbitrary forcing is applied the response may remain linear (Appendix E, Fig. 9). Critical transitions between alternate stable states may thus reflect memories of specific past states and are not necessarily arbitrary non-linear responses to the current forcing pattern. Past experience of distinct environmental conditions (e.g. temperatures) may thus make future responses to related forcing (e.g. climate change) more likely to exhibit discrete changes in ecological states, critical transitions or tipping points [15, 18]. This suggests that critical transitions are not necessarily the arbitrary failure of an ecological community but can be a matter of 'recalling' alternate states familiar from past conditions.

Future work should explore the ultimate equilibrium of these evo-eco feedbacks (Appendix F, Figs. 10 and 11), and investigate relaxation of some of the simplifying assumptions utilised in the memory behaviours illustrated here (Appendix H). In particular, this paper has not investigated the effect of evo-eco feedbacks on trophic (e.g. predator-prey) interactions or mutualisms, nor have we investigated the ecological analogues of other behaviours that can be produced by unsupervised learning in more




Fig. 5 Ecological dynamics after evolution in varying environmental conditions. The evolved ecosystem exhibits two attractor states (rightmost frames) that are reached from various initial species densities (leftmost frames). **a** Random initial species densities develop into one of two possible attractors corresponding to the patterns of forcing experienced in the evolutionary past. **b** Initial configurations that resemble a small part of E_1 (i and ii) or E_2 (iii and iv) develop into equilibria that fully recreate E_1 and E_2 respectively. **c** Initial configurations that are partially randomised versions of E_1 (i. 20 %, ii. 80 %) or E_2 (iii. 80 %, iv. 20 %) develop into equilibria that ‘repair’ the corresponding state. **e** For initial conditions between E_1 and E_2 , ($E_1 : E_2$ ratio = i.80:20, ii.55:45, iii.45:55 iv.20:80) the dynamics ‘recognise’ the pattern that is resembled most closely

general neural networks (e.g. with multi-layered or asymmetric connection structures). Some intriguing further research directions are also suggested:

Do brains learn in the same way that ecosystems evolve? We have shown that ecosystems evolve in the same way that brains learn, but recognising evo-eco dynamics and connectionist learning models as different instantiations of the same underlying adaptive mechanisms also sheds light in the other direction, i.e. on cognitive processes [86, 87].

Can an ecosystem gain from experience? The idea of *sequential selection*, where non-arbitrary organisation arises in a system over evolutionary time without selection being applied at the system level [21, 88] suggests that a biological community “may gain from experience” by using “a system ‘memory’ carried in the gene pool” [19]. Our work in other domains has shown that the more specific sense of system memory demonstrated in the current paper can improve the ability of an adaptive network [32] to solve constraint problems or optimise resource




Fig. 6 Response to environmental forcing before and after evolution in varying environmental conditions. **a-b** Population dynamics under slowly-changing environmental forcing, changing first from E_1 to E_2 (middle row, left-to-right), and then back again from E_2 to E_1 (bottom row, right-to-left). **a** Before evolution of interactions, changes in species densities are proportional to forcing. **b** After evolution (Experiment 2), species densities show an abrupt switch between attractors with hysteresis. **c** Vector field for the population dynamics. The unstable equilibrium is revealed at the boundary of the shaded region indicating where species densities move away from E_1 . Points near the critical transition (solid circle) have slower population dynamics than points far from critical transition (dashed circle). **d** Evolution of two-attractor system. Initially, change in species densities is proportional to environmental forcing. Around generation 470 non-linear but non-catastrophic transitions are observable. Finally, two stable attractors with a catastrophic transition and hysteresis

allocation problems (without an external reward signal) [58]. This suggests that adaptation at the ecosystem level is possible in a formal sense without group selection; not adaptation in a Darwinian sense, but rather in the same sense and by the same mechanism as connectionist models of organismic adaptation [60].

Similarly, demonstration that ecosystems exhibit collective adaptive behaviours without being units of selection prompts inquiry as to whether these systems are capable of more complex computational tasks. For example, natural ecosystems are under very many constraints that limit species abundance (e.g. phosphate availability). Does selection on individuals improve a system's ability to resolve these constraints? Hopfield networks are

known to be able to solve complex constraint satisfaction problems [89]. Do these abilities translate to ecological networks?

Conclusions

We have introduced the framework of connectionist learning as a tool to expand our understanding of eco-dynamics and collective ecological behaviours. Within this framework we find that, despite not being an evolutionary unit, an ecological community can behave like an (unsupervised) learning system, creating internal organisations that collectively habituate to past environmental conditions, and actively recalling past responses to those conditions.

Previously there have only been two choices in how to interpret collective behaviours in ecosystems – i.e. either they have no system-level organisation or some mechanism of group selection must be involved. Our findings demonstrate that there is a third possibility. Ecological organisations that produce collective behaviours can arise from the positive feedback of individual natural selection and ecological population dynamics without invoking group selection. Specifically, given the presence of evolutionary trade-offs, the effect of individual-level natural selection acting on interspecific relationships is dynamically equivalent to a mechanism of unsupervised correlation learning and ecosystems can thereby exhibit organised collective behaviours via the same principles of connectionist learning that apply to neural networks. What is it that ecosystems learn? We find that they have the potential to learn *where to go* (i.e. evolved ecological attractors recreate past ecological states, where an attractor may be the climax community resulting from a successional process [90]), *how to get there* (i.e. the successional or assembly process) and *how to stay there* (i.e. the relationships that increase the resilience and stability of those mature ecological states). Of course, interpreting evo-eco dynamics as a connectionist learning system is not obligatory. A description in terms of individual natural selection and ecological population dynamics only is entirely compatible – indeed, we have provided this level of description for all the results in this paper. But recognising the equivalence with connectionist models enables us to convert and exploit well-understood concepts and results from this discipline to understand the organisation of ecological communities in new ways, and thereby to recognise the potential for predictable collective behaviours.

Reviewers' comments

Reviewer report 1: Prof. Ricard V Solé Universitat Pompeu Fabra, Barcelona

Reviewer 1: This is a rather intriguing paper, suggesting a novel view of ecosystems as a special class of neural-like networks. The underlying idea is that several relevant features of ecological networks, as well as the assembly dynamics of ecosystems and their resilience against different classes of perturbations, can be understood in terms of the attractor dynamics associated to unsupervised neural networks. Such a connection would make possible to rethink community dynamics and the emergence of complex ecosystems as connectionist objects.

Authors: *We thank Prof. Solé for a clear summary of our work.*

Reviewer 1: The authors make a detailed comparison between both classes of networks and the key dynamical features shared by both. Some of the reported similarities are rather generic. They include the presence of attractor

dynamics or the robustness against stochastic changes. These are almost inevitable components of most complex adaptive systems. Due to their dissipative behaviour, attractors are expected to occur, and given the nature of the interactions among neurons or species (described by connectionist models sharing several structural properties) multiple attractors are inevitable.

Authors: *Yes, the existence of multiple attractors is indeed a generic property of dynamical systems. However, it is not the attractors per se that we claim are interesting – it is their location with respect to the history of the system. The emergence of attractor states that resemble past ecological states is not a property of complex adaptive systems in general.*

Reviewer 1: The list of mechanistic equivalences given in table I is systematic, and several specific ingredients are very valuable and non-trivial. This is the case, for example, of the weights of trophic (or mutualistic) interactions among species. Could these connections incorporate some class of learning driven by evolutionary forces?

Authors: *This point concerns the symmetry of connections – addressed separately below.*

Reviewer 1: Some other items in the list could be collapsed into a single one: points h-k all correspond to the expected properties of associative memory.

Authors: *We retain the separate table entries to emphasise how each of these associative memory properties is equivalent to distinct properties within natural ecosystems, and to indicate the separate experiments we have used to demonstrate these phenomena.*

Reviewer 1: There is one very important assumption that might limit the potential generality of this work. At some point the authors choose a symmetric connectivity matrix. This is certainly appropriate to develop theoretical models, as successfully shown by the classic work of John Hopfield. It guarantees the existence of stable solutions (fixed points) that are the minima of a multidimensional landscape and this is the reason such choice of Hebbian learning rules has been studied in detail. Under these assumptions, a large class of nonlinear networks (neural or not) can be fully understood in terms of a potential-derived description of the dynamics. Powerful results, such as the Cohen-Grossberg theorem, can be easily applied to a broad range of multidimensional models including both neural and ecological ones. But we should not forget that such a situation needs to be considered as an exception, not the rule. Symmetry allows us to assume a landscape with multiple valleys, each one harbouring a stable attractor with a more or less large basin. Unfortunately, as soon as we abandon this assumption, we get into a much more complex universe plagued by oscillatory and chaotic solutions, far away from the simple attractor metaphor grounded on stable fixed points. Both neural and ecological dynamics are known to exhibit

complex fluctuations, which makes proper analysis (and a comparative one) difficult.

Authors: The second reviewer, Professor Knight, makes similar comments regarding symmetry in his review below. We respond to both reviewers' comments on this topic together (see below).

Reviewer 1: A robust response is parameter dependent in both classes of systems, but some qualitative traits make them considerably different and these differences should be taken into account. Ecological systems are strongly driven by the fact that they are dissipative structures connected to energy flows. Here, energy and matter dominate over information, whereas the latter is a crucial component. Perhaps for this reason the network motifs (and other structural patterns) differ between neural nets and food webs. Ecosystems also display some special properties that are difficult to map into a neural counterpart. Species can be located in different levels along food chains, and some species are specially relevant when dealing with network fragility against species removals. Such keystone species, if removed or reduced in population, can trigger deep rearrangements in the network, affecting multiple levels or even causing extinction. These are important features and connect our understanding of ecosystems with their evolutionary and ecological dynamics. Although both neural nets and ecological webs can experience shifts and catastrophes, the implications and the nature of these transitions are rather different.

Authors: We agree – these are good points. Certainly, there will be cases where energy and matter flows dominate the dynamics of ecological communities. This has been the only way to think about ecological dynamics thus far. We think it is significant, therefore, to demonstrate that there are any conditions where ecological dynamics become information-based, as shown. More work is needed to investigate the exact nature of the conditions that are important in distinguishing when energy and matter dominate information, and vice versa. We think that the presence of normalising ecological constraints, having the effect of forcing the evolution of fitness 'dependencies' between species, is one of the important features of such conditions (Appendix H).

Reviewer 1: If we consider as valid the concept of Red Queen dynamics, we should also expect the propagation of changes resulting from pairwise coevolution through the food web. It has been shown that indirect effects have a relevance at least as large as direct links. The possibility that learning, in a Hebbian sense, might be difficult to be sustained under this indirect, changing effects needs to be seriously considered. Other components of ecological organization, such as the stabilizing effects of space, might also need to be taken into account.

Despite the previous criticisms, I still find the proposed comparison very intriguing. I think it deserves

serious inspection and in the future it might play an important role in understanding a very important class of complex ecosystems: those associated to the microbiome. Here we have a system dominated by mutualistic interactions, which might be more suitable to be modeled (on a first approximation, and with some caution) with nonlinear neural networks. In this context, we do know that coevolutionary forces have shaped the interactions between microbiome partners and the host niche. Moreover, the microbiome might not escape from a description grounded in a complex system that adapts and learns. More importantly, engineered approaches targeting special components of the microbiome could benefit from the picture provided in this paper.

Authors: We thank Prof Sole for his generous review.

Reviewer report 2: Prof. Rob Knight, University of Colorado, Boulder

Reviewer 2: This interesting manuscript uses the connectionist paradigm from artificial intelligence to provide a new theoretical framework describing how ecosystems can maintain "memory", or nonlinear interactions between current inputs and past states. It addresses an important need in the field for a better theoretical understanding of how complex microbial communities are assembled and behave in the face of a changing environment, including ever-changing hosts.

Authors: We thank Prof Knight for this encouraging introduction. We agree that while our results can be applied in a range of scenarios (including plant and animal communities) they may be of particular interest in microbial communities.

Reviewer 2: The model is based on a generalization of the classic Lotka-Volterra model, updated to include the possibility of allelic evolution. Although this model is sufficiently general to support many kinds of dynamic behaviour, the simulations are restricted to the specific case where interactions are symmetric. Although this makes the simulations easier to interpret, it is unlikely to be a good model for real microbial communities where many interactions are known to be asymmetric (e.g. production of an antimicrobial by species A where species B does not have a retaliation strategy against that particular species, or differential responses to depletion of the same nutrient). However, it could reasonably be argued that relaxing this constraint should be a topic of future work.

Authors: Professor Sole has made similar comments on symmetry in his review above. We aggregate our response to both reviewers together (see below).

Reviewer 2: However, even in this simplified model, the simulation results are interesting, for example providing instances of canalization into alternative states.

Additionally, the demonstration of “ecological memory”, where the response to a current perturbation depends on past states of the community, is convincing. It would be useful to see a more detailed exploration of how much evolution is required for this (i.e. can you still get the same results from ecological changes even without evolution, and, for example, would heritable gene expression changes e.g. through epigenetic modification have the same effect?). Again, these points could be discussed briefly rather than necessarily adding a lot of new experiments to the present manuscript. It would also be useful to connect this work to the literature on microbial co-occurrence and what the implications are for our ability to detect meaningful ecological interactions using currently used techniques, and perhaps to point the way to what kinds of technical developments would be needed in future to do this right.

Authors: *When changes to interaction coefficients are provided by genetic evolution, the separation in timescales with ecological population dynamics is clear (as is necessary for the perspective taken in this paper - see Beisner 2003 for discussion on fast and slow variables). If changes in interaction coefficients were provided by non-genetic inheritance this could facilitate similar behaviours, in principle, but only if they persisted on timescales that were significantly longer than population dynamics. We appreciate the suggestion to connect our work with that on microbial co-occurrence. Appendix G discusses some empirical tests for distributed learning that could be applied in that context.*

Reviewer 2: Overall, this is an interesting theoretical development and provides an exciting new way to look at ecosystem assembly and evolution.

Authors: *We thank Prof Knight for this enthusiastic review.*

Response to reviewer comments on symmetry

Authors: *Both reviewers correctly note that our simulations are carried out under the assumption of symmetrical interactions, and we agree that this is a limitation of our computational models. However, the analysis, in Methods Part 1, where we show that evolved changes in interactions are Hebbian, does not rely upon this assumption. Moreover, whilst the computer simulations do employ the assumption of symmetrical interactions to facilitate analysis of emergent dynamical behaviours, it is not the case that interesting dynamics in neural networks in general depend upon this symmetry. Very many important collective behaviours take place in non-symmetric neural networks, many of which also incorporate Hebbian learning. We strongly agree with both reviewers that relaxing this assumption during future modelling is an important direction for future research.*

Appendices

Appendix A: Individual selection in ecosystems

A mutation to an individual in species i that decreases the competitive effect, ω_{ij} , of species j on species i directly affects the fitness of the individual carrying the mutation and not other individuals in species i , and can thus be favoured by individual selection. It is only changes to traits that directly affect the growth rate of an individual compared to the rest of the individuals in the species can be affected by individual-level selection. Traits that increase the growth rate of all individuals in the species equally have no differential individual benefit (despite conferring benefit to the species as a whole). In particular, a mutation to a trait in an individual in species i that changes its competitive effect, ω_{ji} , on some other species j , e.g. decreasing the density of a competitor species, may thereby indirectly increase the growth rate of species i . But this will benefit all individuals in species i , not just the mutant, and therefore has no differential selective benefit to the individual that bears the mutation [77]. Likewise, the competitive effect of species j on species i may, by virtue of normalising ecological constraints, be decreased as a side-effect of increasing the competitive effect of species k on species i . But again this would not be favoured by individual selection as the benefit is felt by all individuals in species i (conversely, changes to ω_{ij} could be selected under individual selection even though, as a result of indirect effects through changes in density of other species or through normalising ecological constraints, their net effect is to decrease the density of their own species). It is therefore only direct effects on individual fitness that are taken into account by the selection coefficient described here; i.e., Eq. 2. evaluates the change in growth rate of individuals in species i due to changes in ω_{ij} and not ω_{ji} , and furthermore, only changes to ω_{ij} caused by positive selection coefficients, not those caused by indirect normalisation effects. This correctly disregards any changes to a species growth rate that occurs as an indirect side-effect of altering the density of some other species.

Appendix B: The relationship between rate of adaptation and product of species densities in more complex cases

In the main text, the rate of adaptation, v_{ij} , of each inter-specific interaction coefficient is modelled with Eq. 3 corresponding to the case where there is no interference between simultaneously segregating alleles at different loci. In large sexual populations with linked loci, the rate of adaptation will depend on the type of recombination, recombination rate, population size, the mutation rate and magnitude of mutations. Here we compare the rate of adaptation of an interaction coefficient for three different models. In each case, the rate of adaptation, v_{ij} , of an inter-specific interaction coefficient describing the fitness effect

of species j on species i , is $v_{ij} = x_i \mu \bar{P}$, where $x_i \mu$ is the rate with which beneficial mutations arise in species i , and $\bar{P}$ is the average probability that a single new mutation will ultimately fix (see main text). In all cases, $\bar{P}$ is a function of the selection coefficient $s_i = \frac{m_i}{k_{ie}} g x_j$ (Eq. 2, main text) where m_i is the intrinsic net growth rate of species i , k_{ie} is the carrying capacity of species i in environment e , and g is the change in the interaction coefficient due to an individual mutation. Here we write $s_i = \beta x_j$, for clarity of the comparisons that follow.

Case a) No interference

In simple cases when there is no interference between simultaneously segregating alleles at different loci (e.g. where genes are under weak selection per locus, free recombination and the linkage disequilibria among alleles sweeping to fixation are negligible), the probability of fixation, $\bar{P} = s_i$. Thus, as per Eq. 3 main text:

$$v_{ij} = \beta \mu x_i x_j \quad (4)$$

where x_i is the density of species i , x_j is the density of species j and μ is the beneficial mutation rate.

Case b) Linked genes on a linear genome

Weissman & Barton [79] consider the effects of interference between linked genes on a linear genome. Here the genomic rate of fixation of beneficial mutations is ([79] Eq. 7):

$$v = \frac{v_0}{1 + 2v_0/R} \quad (5)$$

where, v_0 is the genomic rate of fixation of beneficial mutations in the absence of interference and R is the total genetic map length in Morgans. The authors use the approximation $v_0 = 2x\mu s$, where x is species density and

s is the selection coefficient. With $s_i = \beta x_j$ as before, this gives the rate of adaptation on an interaction:

$$v_{ij} = \frac{2\beta \mu x_i x_j}{1 + 4\beta \mu x_i x_j / R} \quad (6)$$

Case c) Occasional outcrossing

Neher et al. [78] study the rate of adaptation in unlinked loci in facultative sexuals where the rate of outcrossing is very small. Whereas Weissman and Barton examine the case of obligately sexual populations, this case represents occasionally/facultatively sexual populations (e.g. plants). On condition that $r^2/s^2 \gg 4x\mu$, the rate of accumulation of beneficial mutations in this case is given by ([78] Eq. 12b):

$$v \approx x\mu s^2 \left(1 - \frac{4x\mu s^2}{r^2} \right) \quad (7)$$

where r is the outcrossing rate. With $s_i = \beta x_j$ as before, this gives the rate of adaptation on an interaction:

$$v_{ij} \approx x_i \mu (\beta x_j)^2 \left(1 - \frac{4x_i \mu (\beta x_j)^2}{r^2} \right) \quad (8)$$

Comparison of the three cases

Figure 7 plots the rate of adaptation v_{ij} as a function of x_i and x_j for these three different cases. We observe that case a, where the rate of adaptation is directly proportional to the product $x_i x_j$ as modelled in our simulations, and the two more complex cases (b and c) are all qualitatively similar. Although in some cases the absolute rate of adaptation is more strongly limited by the recombination rate than the mutation supply or the strength of selection, for example [79], the relative rates of adaptation are still determined largely by the product of x_i and x_j . More specifically, all three cases have the essential characteristic that the rate of adaptation is zero when either x_i or x_j




Fig. 7 Rate of adaptation v_{ij} as a function of x_i and x_j for three different models. **a** Eq. 3 from main text. **b** Eq. 7 from Weissman & Barton [79], **c** Eq. 12b from Neher et al. [78]. We observe that **(b)** is very close to a linear scaling of **(a)** and, although **(c)** shows slight qualitative differences in the shape of the function, it maintains the essential qualitative characteristic. In all cases, the rate of evolutionary change in an interaction coefficient increases with the product of x_i and x_j . (As per our simulation experiments: $k_{ie} = 10$, $m_i = 0.5$, $g = 0.1$, $\mu = 1.0 \times 10^{-5}$. For case **(b)**, the map length, $R = 1$. For case **(c)**, the out-crossing rate, $r = 0.01$)

is zero, and otherwise, the greater the value of one, the greater the rate of increase with the other. Thus, although the shape of the alternate functions differs from ours, the essential behaviour is preserved. Intuitively, mutations must be both created and selected for an interaction coefficient to evolve.

Appendix C: Additional methods: Normalisation, variable environments, measuring ecological attractors and model parameters

Normalisation methods

In each evolutionary step, all interaction terms in $\Omega(t)$ are updated by natural selection according to Eq. 3 to produce $\Omega'(t)$ and then renormalised to produce $\Omega(t+1)$. Renormalisation preserves the conditions that for each species i and all other species j ($j \neq i$), $\sum_{j=1}^N \omega_{ij}(t) = Q_i$, and $\sum_{j=1}^N \omega_{ji}(t) = Q_i$, where $Q_i < 0$ is a constant for each species. Specifically, an iterative row and column normalisation (below) is applied to $M(k=0) = \Omega'$, until the values of M converge within a specified accuracy, i.e. $(\sum_{ij} (m_{ij}(k+1) - m_{ij}(k))^2 < 10^{-5}$, where k is the iteration counter, as follows:

$$M(k+1) = \text{column_norm}(\text{row_norm}(M(k))) \quad (9)$$

where $\text{row_norm}(m_{ii}) = m_{ii}$, $\text{column_norm}(m_{ii}) = m_{ii}$, i.e. self-interactions are unaffected, and

$$\text{row_norm}(m_{ij(i \neq j)}) = \frac{m_{ij}(k)}{\sum_{j=1(i \neq j)}^N m_{ij}(k)} \quad (10)$$

and

$$\text{column_norm}(m_{ij(i \neq j)}) = \frac{m_{ij}(k)}{\sum_{i=1(i \neq j)}^N m_{ij}(k)} \quad (11)$$

Variable environments

We investigate the effect of variable environments as follows. The carrying capacity of the i^{th} species in a default ecological environment, E_0 , is k_{i0} . For simplicity in our simulations we let $k_{i0} = k_0$, for all i , where k_0 is a constant. Prior to the evolution of interactions, this causes all species to equilibrate at the same density. To model the evolution of an ecosystem under varying environmental conditions that force or drive the ecosystem to adopt different ecological states, we define two other environmental conditions that alter carrying capacities. The pattern of equilibrium species densities under one environmental condition, E_1 , increases the carrying capacity of some species to $k_0 + \alpha$ and decreases others to $k_0 - \alpha$, where $\alpha = 0.1$. In E_2 , a different subset of species is increased/decreased in a similar manner. See Fig. 2 main text.

Measuring ecological attractors

We examine the ecological attractors in the ecosystem by Monte Carlo sampling, i.e., by repeatedly setting the

species densities to random initial conditions and running to an equilibrium. To measure the inherent attractors induced by evolutionary changes, this sampling is carried out in the absence of environmental forcing – i.e., in E_0 . In some experiments we also investigate the amount of environmental forcing required to push the ecosystem out of equilibrium in one pattern of species densities and into the attractor basin of another stable equilibrium. Whenever, as here, interactions control the correlation of species densities and not their absolute densities, the complement of any attractor pattern is also necessarily an attractor [57, 58, 67]. However, so long as initial conditions are more similar to the past states experienced during evolution than the opposite of those past states these unnatural attractors are precluded. Accordingly, we examine initial conditions, x , satisfying the condition $(|x - E_1| < |x - E'_1|)$ and $(|x - E_2| < |x - E'_2|)$ where E' is the inverse of E (i.e. $E' = 2\bar{E} - E$).

Model parameters

$N = 400$, number of species.

$m_i = 0.5$, growth rate of all species.

$s(t=0) = 0.1$, initial species densities.

$k_0 = 10$, a parameter governing the extrinsic component of carrying capacity in E_0 .

$\alpha = 0.1$ increment/decrement of particular carrying capacities in environments E_1 and E_2 .

$T = 1$, number of evolutionary changes applied in each environment before switching.

$\tau = 5000$, number of ecological timesteps (Eq. 1) between 'initial' and 'final'.

$g = 0.01$, constant of proportionality in selection-limited evolution (Eq. 3)

Interaction coefficients are initialised as follows:

$$\omega_{ij}(t=0) = \begin{cases} -1, & \text{if } i = j \text{ (i.e. self interactions)} \\ -0.2, & \text{otherwise} \end{cases}$$

$Q_i = \sum_{j=1(i \neq i)}^N \omega_{ij}(t=0)$, normalisation constant (the sum of the non-self interactions in any one row/column remains equal to their sum at time $t=0$).

The quantitative values of these parameters will naturally have quantitative effects on the behaviour of the eco-evolutionary dynamics that we simulate. Since the simulations are a phenomenological model of ecosystem evolution, what matters is the relative rather than absolute rates of adaptation on different interaction coefficients – in particular, which interactions increase, which decrease and which remain largely unchanged. This pattern, and its sensitivity to different modelling choices, is investigated in Appendix D.

Appendix D: Equivalence of Hebbian and evolved changes in more complex cases

In the main text the rate of adaptation of each inter-specific interaction coefficient is modelled with Eq. 3 corresponding to the case where there is no interference between simultaneously segregating alleles at different loci. Appendix B shows that the characteristics of the rate of adaptation in more complex cases is qualitatively similar although they are quantitatively different. Here we simulate evolution using these alternative models and incorporating normalising ecological constraints. Figure 8 shows that the quantitative differences in the three equations do not alter the pattern of positive, negative and neutral changes that are produced in the evolving interaction matrix. Specifically, the pattern of changes in interactions have the same direction as the Hebbian model in all cases. Accordingly, there will be parameter ranges where they produce the same distributed memory phenomena in the ecosystem. Investigations of quantitative differences remain for future work.

Appendix E: Response to environmental forcing that is *not* similar to environments experienced during evolution

Figure 9 shows that an ecosystem can exhibit a non-catastrophic response when forced in arbitrary directions (b) and simultaneously exhibit hysteresis and catastrophic

regime shifts when forced in directions that have been experienced previously over evolutionary time (a). This emphasises that the evolved ecological memory causing the switching behaviour is conditioned by the systems' evolutionary history, and thus causes recall (or recognition) of a specific point in a multi-dimensional space of species densities, rather than a general stability/instability property resulting from unorganised or arbitrary evolutionary changes.

Appendix F: Development and breakdown of multiple attractors over long evolutionary timescales

Figure 10 shows how the attractors of the ecosystem change over evolutionary time in Experiment 2. Interestingly, we see that in the long term the two-attractor state is unstable because, rather than reinforcing the ecological patterns that are 'forced' by the external environment, the system begins to reinforce its own patterns of behaviour [58], and positive feedback causes one (slightly stronger) attractor to outcompete the other (Fig. 11).

Appendix G: Empirical tests for distributed learning in ecosystems

The dynamical behaviours we observe in the evolved ecosystem are consistent with ecological memory, alternate ecological states, succession dynamics, assembly





Fig. 8 Evolved changes to interactions are Hebbian in more complex population conditions. Change in interactions between the first 16 species are shown under evolution in a changing environment. Rate of adaptation is controlled by our equation (top row), that from Weissman & Barton [79] (middle row) and Neher et al. [78] (bottom row). **a-c** the change in interactions due to direct selection effects (see Fig. 4 main text). **d** When normalising ecological constraints are taken into account, some interactions are decreased, some left unchanged, and others are increased. The resulting direction of change is the same in all three cases and identical to Hebbian changes (Fig. 4e. main text). ($k_{ie} = 10$, $m_i = 0.5$, $g = 0.1$, $\mu = 1.0 \times 10^{-5}$, $\alpha = 3.5$. For case ii, the map length, $R = 1$. For case iii, the out-crossing rate, $r = 0.01$). For visualisation, the magnitude of changes in **(d)** are multiplied by 5



rules, regime changes and founder effects observed in natural ecosystems. These behaviours follow from simple component principles (i.e. the availability of heritable variation in inter-specific interactions, and the presence of ecological constraints or evolutionary trade-offs) and direct evidence for these behaviours is testable. For example, consider the evolution of a small microbial community. Given a culturable community with stable coexistence dynamics, we could first test whether it has i) one or ii) alternative stable states. This requires sampling many different initial species compositions and allowing species densities to equilibrate. i) If a single state, we can then force the system into a different state ('alternate ecosystem state', [17]) – e.g. by changing temperature, nutrient influx – and hold it there for evolutionary time. Then remove the forcing and retest for multiple attractors ('alternate community states'). If a memory has been conditioned by this forcing then a new attractor will be exhibited. ii) If the system initially has more than one attractor state, then we can estimate the basin size for each attractor by counting the number of different initial conditions that arrive at one or the other. By leaving the system in one attractor over evolutionary time this should increase the relative basin size in proportion to the time spent in that attractor. Next we need to assess the extent to which such a memory is collective or merely the sum of individual memories. This can be done by swapping-in evolved species for species in the original community one-by-one and assessing the relative contribution of individual and collective genetic changes on the dynamical behaviour of the system.

Appendix H: Asymmetric interactions, the importance of normalising ecological constraints, and other future work

One important aspect of evo-eco dynamics that is highlighted by this model is the importance of normalising ecological constraints or evolutionary trade-offs for collective behaviours. These constraints prevent a species A from benefiting from the presence of species B without also becoming dependent on B. That is, it is not just the case that A grows faster in the presence of B, but that A's growth is slower when B is absent. Under these conditions, changes to interactions do not merely increase the growth of each species in a manner that is sensitive to its ecological context, but more specifically, they modify *correlations* between species densities. We assume in the present model that an adaptation that, for example, decreases the niche overlap with one species increasing the niche overlap with others. But the extent to which species evolve dependencies rather than just (context-sensitive) individual advantages in natural ecosystems is an empirical matter – and from this work we recognise it as a matter that is centrally important to the possibility of collective behaviours that are more than the sum of the individual behaviours.

This paper has investigated only competitive interactions and has not investigated mutualistic interactions or asymmetric interactions such as characteristic of trophic, e.g. predator-prey, relationships. The observation that selected changes to interactions are Hebbian does not depend on them being symmetric (or competitive). That is, Eq. 3 is not sensitive to any assumptions about the initial values of interaction coefficients, e.g., whether ω_{ij}




Fig. 10 'Bestiary' of ecological attractors changing over evolutionary time. From an array of different random initial species densities (left-most column), the ecological states reached in the population dynamics changes over evolutionary time. **a** Initially, all initial conditions lead to the same ecological attractor (with all species at the same carrying capacity). **b** New attractor states begin to appear and become established. **c** In the long term, the two-attractor state is unstable and positive feedback causes one of the attractor states to 'out-compete' the other. **d** Eventually the one remaining attractor breaks down as only the strongest species (those that were high density in both patterns) take over [65]




Fig. 11 In the long term the two-attractor system breaks down. Monte Carlo sampling of the ecological attractor states from random initial species densities during evolutionary time. Initially, all attractor states contain species densities that are only minor deviations from the default attractor (E_0) in Euclidean distance. The signed pattern of the attractor state, i.e. in terms of $+/ -$ with respect to the mean species densities, either matches E_1 (blue) or E_2 (green). As the two-attractor state emerges, at around generation 525 (a classic pitchfork bifurcation, but the unstable fixed point is not shown), the magnitudes (as well as signs) of the attractor states closely match the two targets. In the long term, one of the attractors, in this case E_1 , outcompetes the other and becomes the only attractor. Eventually (after ~ 575 generations), this attractor also degrades, i.e. the equilibrium magnitudes no longer match the original target closely (Fig. 10)

and ω_{ji} are equal or even have the same sign, and therefore applies to predator-prey relationships as well as symmetric competitive interactions. Eq. 3 also shows that the selective pressures on *changes* to interactions are symmetric i.e., $\Delta\omega_{ij} = \Delta\omega_{ji}$ (except for the influence of individually-varying carrying capacities), so there is no systematic reason for interactions to become asymmetric over evolutionary time. In the examples investigated in this paper the interaction coefficients are initialised symmetrically and, accordingly, they remain approximately symmetric. The evolutionary model could be applied to asymmetric interactions, but asymmetric interactions introduce the possibility of non-fixed point attractors, e.g. cycles, that complicate the behaviour of the eco-evolutionary dynamics and their measurement considerably. (We note that where ω_{ij} and ω_{ji} differ, the addition of multiple symmetric changes through natural selection will make them less asymmetric over evolutionary time, i.e., bring the ratio of these terms closer to 1, and could evolve them to take the same sign even when they started out with opposite signs. This implies that the effect of evolutionary change would be to increase the stability of the ecological dynamics and reduce or remove chaotic or cyclic attractors over time).

We have assumed that each interaction coefficient is independently modifiable whereas in natural populations traits may affect many interactions simultaneously. Here we chose to investigate scenarios where none of the interaction coefficients reach zero or go positive (which is possible in principle despite the normalisation employed). The equations used exhibit unstable behaviour in this case and a different approach to modelling would be required to handle mutualistic interactions. In natural populations one member of a population can gain selective advantage by changing its relationship to other members of its own species, but our simulations have fixed self-interactions at -1 and have investigated only the evolution of interactions with members of other species.

A key technical distinction between the recent work on associative memory in gene networks [67] and the models utilised here is that the Lotka-Volterra equations represent unsigned (positive) state variables, as is natural for species densities, rather than signed (positive and negative) state variables representing under- or over-expressed gene activity (compared to some normal level). Although it is possible and common to model interesting dynamical behaviours using either signed or unsigned state variables in neural networks, the use of unsigned variables means that Hebb's rule, or natural selection, will only alter interactions in one direction, i.e., the product $x_i x_j$ is always positive (although crucially it may have different magnitudes). The assumption of normalising constraints that cause some interactions to become more competitive as a side effect of others becoming less

competitive is thus important to the results that we have shown.

In particular, as mentioned above, the assumption of these normalising constraints means that changes to interactions, although motivated by increases in individual growth rates, have the effect of (also) altering the *dependency* of one species on another. Without these constraints, the effect of unconstrained changes to interactions is to make high density species fitter in all conditions, rather than making them dependent on the simultaneous high density state of specific species (and hence less fit in some conditions). It is therefore important for future work to investigate how different ways of modelling such constraints impact the behaviours illustrated here. For example, rather than a Lotka-Volterra model, a stoichiometric model of species interactions may alleviate the need for an explicit normalisation mechanism.

Assuming that ecological dynamics (i.e., changes in species density) are much more rapid than evolutionary changes (i.e., genetic changes affecting the coefficients of inter-species fitness dependencies) [91], most evolution occurs whilst ecological dynamics are at or near equilibrium, as modelled here. The behaviour of evo-eco dynamics when these processes have more similar timescales [35] deserves attention. However, the fact that we model varying ecological conditions, causing the ecosystem to visit more than one ecological equilibrium, means that the interaction of ecological and evolutionary dynamics is non-trivial even though their timescales are kept separate in our simulations (following [24]). Moreover, any model assuming a single ecological attractor will overlook the interesting behaviours modelled here, regardless of whether the timescales are separated or similar.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

RAW, RM and STP conceived and developed the original hypothesis; RAW, RM and BC prototyped the models; RAW, DAP and ES designed the study and the model specification; DAP developed and performed the modelling and analysis; RAW and DAP wrote the first draft of the manuscript; ES, CPD, RM and STP contributed substantially to development of arguments. All authors read and approved the manuscript.

Acknowledgements

Thanks to Mike Lewis for investigations using preliminary models and to C.L. Buckley, James Dyke, Kostas Kouvaris, Adam Jackson, Chris Cox, Alexandra Penn, Eva Jablonka, Ferenc Jordán, Bob Holt, Ellen Clarke and Ran Feldesh for conceptual and technical discussion. This work has been supported by the EPSRC: EP/K503150/1.

Author details

¹Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, UK. ²Institute for Life Sciences/Electronics and Computer Science, University of Southampton, Southampton, UK. ³The Parmenides Found, Center for the Conceptual Foundations of Science, Pullach, Germany. ⁴Department of Informatics, Faculty of Sciences, University of Lisbon, Lisbon, Portugal. ⁵Department of Ecology & Evolution, University of Lausanne, Lausanne, Switzerland. ⁶School of Biological Sciences, University of Southampton, Southampton, UK.

Received: 8 April 2015 Accepted: 26 October 2015

Published online: 08 December 2015

References

- Levin SA. Ecosystems and the biosphere as complex adaptive systems. *Ecosystems*. 1998;1(5):431–6.
- Levin SA. Self-organization and the emergence of complexity in ecological systems. *Bioscience*. 2005;55(12):1075–9.
- Anand M, Gonzalez A, Guichard F, Kolasa J, Parrott L. Ecological systems as complex systems: challenges for an emerging science. *Diversity*. 2010;2(3):395–410.
- Whitham TG, Bailey JK, Schweitzer JA, Shuster SM, Bangert RK, LeRoy CJ, et al. A framework for community and ecosystem genetics: from genes to ecosystems. *Nat Rev Genet*. 2006;7(7):510–23.
- Ulanowicz RE. *Growth and Development: Ecosystems Phenomenology*. New York, NY, USA: Springer; 1986.
- Neill WE. The community matrix and interdependence of the competition coefficients. *Am Nat*. 1974;108(962):399–408.
- Proulx SR, Promislow DE, Phillips PC. Network thinking in ecology and evolution. *Trends Ecol Evol*. 2005;20(6):345–53.
- Clements FE. *Plant Succession: an Analysis of the Development of Vegetation* vol. 242. Washington, DC, USA: Carnegie Institution of Washington; 1916.
- Weiher E, Keddy P. *Ecological Assembly Rules: Perspectives, Advances, Retreats*. Cambridge, UK: Cambridge University Press; 2001.
- Holling CS. Resilience and stability of ecological systems. *Annu Rev Ecol Syst*. 1973;4:1–23.
- Gallopin GC. Linkages between vulnerability, resilience, and adaptive capacity. *Glob Environ Chang*. 2006;16(3):293–303.
- de Ruiter PC, Neutel AM, Moore JC. Energetics, patterns of interaction strengths, and stability in real ecosystems. *Science*. 1995;269:1257–60.
- Folke C. Resilience: the emergence of a perspective for social–ecological systems analyses. *Glob Environ Chang*. 2006;16(3):253–67.
- Staniczenko P, Lewis OT, Jones NS, Reed-Tsochas F. Structural dynamics and robustness of food webs. *Ecol Lett*. 2010;13(7):891–9.
- Higgins PA, Mastrandrea MD, Schneider SH. Dynamics of climate and ecosystem coupling: abrupt changes and multiple equilibria. *Phil Trans R Soc London Series B Biol Sci*. 2002;357:647–56.
- Cropp R, Gabric A. Ecosystem adaptation: Do ecosystems maximize resilience? *Ecology*. 2002;83(7):2019–26.
- Beisner BE, Haydon DT, Cuddington K. Alternative stable states in ecology. *Front Ecol Environ*. 2003;1(7):376–82.
- Scheffer M, Carpenter S, Foley JA, Folke C, Walker B. Catastrophic shifts in ecosystems. *Nature*. 2001;413(6856):591–6.
- Lenton TM, van Oijen M. Gaia as a complex adaptive system. *Phil Trans R Soc London Series B Biol Sci*. 2002;357(1421):683–95.
- Ulanowicz RE. Aristotelean causalities in ecosystem development. *Oikos*. 1990;57(1):42–48.
- Lenton TM. In: Schneider SH, editor. *Clarifying Gaia: regulation with or without natural selection. Scientists debate Gaia: the next century*. Cambridge, Massachusetts: MIT Press; 2004. pp. 15–25.
- Richardson JL. The organismic community: resilience of an embattled ecological concept. *BioScience*. 1980;30(7):465–71.
- McIntosh RP. Ha gleason's 'individualistic concept' and theory of animal communities: a continuing controversy. *Biol Rev*. 1995;70(2):317–57.
- Levin SA. Evolution at the ecosystem level: On the evolution of ecosystem patterns. *Contrib Sci*. 2011;7(1):11–16.
- Smith JM, Burian R, Kauffman S, Alberch P, Campbell J, Goodwin B, et al. Developmental constraints and evolution: a perspective from the mountain lake conference on development and evolution. *Q Rev Biol*. 1985;60(3):265–87.
- Gleason HA. The individualistic concept of the plant association. *Bulletin Torrey Botanical Club*. 1926;53(1):7–26.
- Lawton JH. Are there general laws in ecology? *Oikos*. 1999;84(2):177–92.
- Milne BT. Motivation and benefits of complex systems approaches in ecology. *Ecosystems*. 1998;1(5):449–56.
- Thompson JN, Reichman O, Morin PJ, Polis GA, Power ME, Sterner RW, et al. Frontiers of ecology. *BioScience*. 2001;51(1):15–24.
- Matthews B, Narwani A, Hausch S, Nonaka E, Peter H, Yamamichi M, et al. Toward an integration of evolutionary biology and ecosystem science. *Ecol Lett*. 2011;14(7):690–701.
- Post DM, Palkovacs EP. Eco-evolutionary feedbacks in community and ecosystem ecology: interactions between the ecological theatre and the evolutionary play. *Phil Trans R Soc B: Biol Sci*. 2009;364(1523):1629–40.
- Gross T, Sayama H. *Adaptive Networks - Theory, Models and Applications*. Germany: Springer; 2009.
- Paperin G, Green DG, Sadedin S. Dual-phase evolution in complex adaptive systems. *J R Soc Interface*. 2011;8:609–29.
- Schoener TW. The newest synthesis: understanding the interplay of evolutionary and ecological dynamics. *Science*. 2011;331(6016):426–9.
- Turcotte MM, Reznick DN, Hare JD. The impact of rapid evolution on population dynamics in the wild: experimental test of eco-evolutionary dynamics. *Ecol Lett*. 2011;14(11):1084–92.
- Metz JA, Geritz SA, Meszén G, Jacobs FJ, Van Heerwaarden JS. Adaptive dynamics, a geometrical study of the consequences of nearly faithful reproduction. *Stochastic Spatial Struct Dyn Syst*. 1996;45:183–231.
- Fussmann G, Loreau M, Abrams P. Eco-evolutionary dynamics of communities and ecosystems. *Funct Ecol*. 2007;21(3):465–77.
- Harcombe W. Novel cooperation experimentally evolved between species. *Evolution*. 2010;64(7):2166–72.
- Romano-Keeler J, Weitkamp JH. Maternal influences on fetal microbial colonization and immune development. *Pediatr Res*. 2014;77(1-2):189–95.
- Faust K, Raes J. Microbial interactions: from networks to models. *Nat Rev Microbiol*. 2012;10(8):538–50.
- Forslund K, Sunagawa S, Kultima JR, Mende DR, Arumugam M, Typas A, et al. Country-specific antibiotic use practices impact the human gut resistome. *Genome Res*. 2013;23(7):1163–9.
- David LA, Maurice CF, Carmody RN, Gootenberg DB, Button JE, Wolfe BE, et al. Diet rapidly and reproducibly alters the human gut microbiome. *Nature*. 2014;505(7484):559–63.
- Shade A, Peter H, Allison SD, Baho DL, Berga M, Bürgmann H, et al. Fundamentals of microbial community resistance and resilience. *Front Microbiol*. 2012;3:417.
- Case TJ, Holt RD, McPeck MA, Keitt TH. The community context of species' borders: ecological and evolutionary perspectives. *Oikos*. 2005;108(1):28–46.
- Mayr E. Wallace's line in the light of recent zoogeographic studies. *The Q Rev Biol*. 1944;19(1):1–14.
- Wallace AR. *The Geographical Distribution of Animals: with a Study of the Relations of Living and Extinct Faunas as Elucidating the Past Changes of the Earth's Surface* vol. 1. Cambridge, UK: Cambridge University Press; 2011.
- Angeler DG, Trigel C, Drakare S, Johnson RK, Goedkoop W. Identifying resilience mechanisms to recurrent ecosystem perturbations. *Oecologia*. 2010;164(1):231–41.
- Carpenter SR, Ludwig D, Brock WA. Management of eutrophication for lakes subject to potentially irreversible change. *Ecol Appl*. 1999;9(3):751–71.
- Hendry R, McGlade J. The role of memory in ecological systems. *Proc R Soc Lond Ser B: Biol Sci*. 1995;259(1355):153–9.
- Peterson GD. Contagious disturbance, ecological memory, and the emergence of landscape pattern. *Ecosystems*. 2002;5(4):329–38.
- Golinski M, Bauch C, Anand M. The effects of endogenous ecological memory on population stability and resilience in a variable environment. *Ecol Model*. 2008;212(3):334–41.
- Brown WL, Wilson EO. Character displacement. *Syst Zool*. 1956;5(2):49–64.
- Dayan T, Simberloff D. Ecological and community-wide character displacement: the next generation. *Ecol Lett*. 2005;8(8):875–94.
- Thompson JN. *The Geographic Mosaic of Coevolution*. Chicago, IL, USA: University of Chicago Press; 2005.
- McCulloch WS, Pitts W. A logical calculus of the ideas immanent in nervous activity. *Bull Math Biophys*. 1943;5(4):115–33.
- Hinton GE, Sejnowski TJ. *Unsupervised Learning: Foundations of Neural Computation*. Cambridge, MA, USA: MIT Press; 1999.
- Hopfield JJ. Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci*. 1982;79(8):2554–8.
- Watson RA, Buckley CL, Mills R. Optimisation in 'self-modelling' complex adaptive systems. *Complexity*. 2010;16(5):17–26.
- Farmer JD. A rosetta stone for connectionism. *Physica D: Nonlinear Phenomena*. 1990;42(1):153–87.
- Watson RA, Mills R, Buckley CL. Global adaptation in networks of selfish components: emergent associative memory at the system scale. *Artif Life*. 2011;17(3):147–66.

61. Mikhailov A, Mit'kov I, Sveshnikov N. Molecular associative memory. *BioSystems*. 1990;23(4):291–5.
62. Fernando CT, Liekens AM, Bingle LE, Beck C, Lenser T, Stekel DJ, et al. Molecular circuits for associative learning in single-celled organisms. *J R Soc Interface*. 2009;6(34):463–9.
63. Vohradsky J. Neural model of the genetic network. *J Biol Chem*. 2001;276(39):36168–73.
64. Noonburg V. A neural network modeled by an adaptive lotka-volterra system. *SIAM J Appl Math*. 1989;49(6):1779–92.
65. Poderoso FC, Fontanari JF. Model ecosystem with variable interspecies interactions. *J Phys A: Math Theoretical*. 2007;40(30):8723.
66. Wilson DS. Complex interactions in metacommunities, with implications for biodiversity and higher levels of selection. *Ecology*. 1992;73(6):1984–2000.
67. Watson RA, Wagner GP, Pavlicev M, Weinreich DM, Mills R. The evolution of phenotypic correlations and “developmental memory”. *Evolution*. 2014;68(4):1124–38.
68. Ackley DH, Hinton GE, Sejnowski TJ. A learning algorithm for boltzmann machines. *Cogn Sci*. 1985;9(1):147–69.
69. Hebb D. *The Organization of Behavior; a Neuropsychological Theory*. New York, NY, USA: Wiley/Psychology Press; 1949.
70. Ulanowicz RE. Utricularia's secret: the advantage of positive feedback in oligotrophic environments. *Eco Model*. 1995;79(1):49–57.
71. Laland K, Odling-Smee J, Turner S. The role of internal and external constructive processes in evolution. *J Physiol*. 2014;592(11):2413–22.
72. Laland KN, Odling-Smee J, Feldman MW, Kendal J. Conceptual barriers to progress within evolutionary biology. *Found Sci*. 2009;14(3):195–216.
73. Gallardo B, Aldridge DC. Is great britain heading for a ponto–caspien invasional meltdown? *J Appl Ecol*. 2015;52(1):41–9.
74. Heimpel GE, Frelich LE, Landis DA, Hopper KR, Hoelmer KA, Sezen Z, et al. European buckthorn and asian soybean aphid as components of an extensive invasional meltdown in north america. *Biol Invasions*. 2010;12(9):2913–31.
75. Simberloff D, Von Holle B. Positive interactions of nonindigenous species: invasional meltdown? *Biol Invasions*. 1999;1(1):21–32.
76. Hutchinson GE. *The niche: an abstractly inhabited hypervolume. The ecological theatre and the evolutionary play*. New Haven, CT: Yale University Press; 1965, pp. 26–78.
77. Wilson DS. *The Natural Selection of Populations and Communities*. Menlo Park, CA, USA: Benjamin/Cummings Pub. Co.; 1980.
78. Neher RA, Shraiman BI, Fisher DS. Rate of adaptation in large sexual populations. *Genetics*. 2010;184(2):467–81.
79. Weissman DB, Barton NH. Limits to the rate of adaptive substitution in sexual populations. *PLoS Genet*. 2012;8(6):1002740.
80. May RM. Some notes on estimating the competition matrix, α . *Ecology*. 1975;56(3):737–41.
81. Hughes JB, Roughgarden J. Aggregate community properties and the strength of species' interactions. *Proc Natl Acad Sci*. 1998;95(12):6837–42.
82. Roughgarden J. Evolution of niche width. *Am Nat*. 1972;106(952):683–718.
83. Scheffer M, Bascompte J, Brock WA, Brovkin V, Carpenter SR, Dakos V, et al. Early-warning signals for critical transitions. *Nature*. 2009;461(7260):53–9.
84. Dakos V, Scheffer M, van Nes EH, Brovkin V, Petoukhov V, Held H. Slowing down as an early warning signal for abrupt climate change. *Proc Natl Acad Sci*. 2008;105(38):14308–12.
85. Finn JT. Measures of ecosystem structure and function derived from analysis of flows. *J Theor Biol*. 1976;56(2):363–80.
86. Adams P, Hebb and darwin. *J Theor Biol*. 1998;195(4):419–38.
87. Fernando C, Szathmáry E, Husbands P. Selectionist and evolutionary approaches to brain function: a critical appraisal. *Front Comput Neurosci*. 2012;6(24):0.
88. Betts RA, Lenton TM. Second chances for lucky gaia: a hypothesis of sequential selection. *Gaia Circular*. 2008;1(1):4–6.
89. Hopfield JJ, Tank DW. Computing with neural circuits- a model. *Science*. 1986;233(4764):625–33.
90. Law R, Morton RD. Permanence and the assembly of ecological communities. *Ecology*. 1996;77(3):762–75.
91. Roughgarden J. *Theory of Population Genetics and Evolutionary Ecology: an Introduction*. New York, NY, USA: Macmillan; 1979.

Submit your next manuscript to BioMed Central and take full advantage of:

- Convenient online submission
- Thorough peer review
- No space constraints or color figure charges
- Immediate publication on acceptance
- Inclusion in PubMed, CAS, Scopus and Google Scholar
- Research which is freely available for redistribution

Submit your manuscript at
www.biomedcentral.com/submit

