Thesis: Power, D.A. (2019) "Distributed associative learning in ecological community networks", University of Southampton, Electronics and Computer Science, PhD Thesis

UNIVERSITY OF SOUTHAMPTON

Distributed associative learning in ecological community networks

by

Daniel A Power

A thesis submitted for the degree of
Doctor of Philosophy

in the
Faculty of Engineering and the Environment
Computational Engineering and Design

March 2019

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND THE ENVIRONMENT
COMPUTATIONAL ENGINEERING AND DESIGN

Doctor of Philosophy

DISTRIBUTED ASSOCIATIVE LEARNING IN ECOLOGICAL COMMUNITY NETWORKS

by Daniel A Power

Ecological communities are complex, self-organising systems of interacting species exhibiting important and intricate functions. Yet as evolution by natural selection operates within individual species, and not on community structure as a whole, it is not clear whether or how natural selection organises community-level functions. A long-standing open question thus persists: Are there alternative organising mechanisms that would enable us to understand and predict the complex collective behaviours exhibited by natural comunities?

One intriguing possibility is that we might better understand community organisation, not through the principles of evolution (because selection doesn't operate at the community level), but through the principles of learning, driven by selection acting at lower levels of organisation (i.e. coevolution amongst a community's component species). Specifically, in this thesis we show conditions where natural selection on ecological interactions is functionally equivalent to a simple type of distributed associative learning, unsupervised learning. This learning rule is well-known in connectionist learning models of cognitive systems, such as neural-networks, where it can produce many non-trivial collective behaviours.

We build from this result and simulate simple Lotka-Volterra models of communities and show conditions where these self-organising processes result in community structure and assembly dynamics that exhibit non-trivial functional properties at the community level. We use these models to demonstrate how community organisation can be conditioned by past experience in the same sense as connectionist learning models habituate to stimuli. We demonstrate the capabilities of this process in the ecological model by showing how the action of individual natural selection can enable communities to i) form a distributed ecological memory of multiple past states; ii) improve their ability to resolve conflicting constraints among species leading to higher community biomass; and iii) learn to solve complex resource-allocation problems equivalent to difficult computational puzzles like Sudoku.

We identify distributed associative learning as a mechanism by which natural selection contributes to community organisation, resulting in a range of adaptive behaviours at the community level. This mechanism is not a result of community-level selection (it cannot be; we prevent group selection effects), and is emergent only from Darwinian processes at lower levels (i.e. individual-level selection).

Contents

List of Figures

List of Tables

Acronyms

SS	Stable state
CSP	Constraint satisfaction problem

Declaration of Authorship

I, Daniel A Power, declare that this thesis entitled Distributed associative learning in ecological community networks and the work presented in it are my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly while in candidature for a research degree at this University;
2. No part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as: Power, D. A., Watson, R. A., Szathmary, E., Mills, R. , Powers, S.T., Doncaster, C.P. and Czapp, B. (2015) What can ecosystems learn? Expanding evolutionary ecology with learning theory. Biology Direct 10(69).

Signed: DocuSigned by: Daniel Power
BE198D16040940F...

Date: 07 June 2019 | 13:35:13 BST

List of Publications

- Power, D. A., Watson, R. A., Szathmary, E., Mills, R. , Powers, S.T., Doncaster, C.P. and Czapp, B. (2015) What can ecosystems learn? expanding evolutionary ecology with learning theory. Biology Direct 10(69).
- Power, D. A., Watson, R. A. (*in preparation*) Can ecological communities learn to solve Sudoku? Evolutionary feedbacks facilitate constraint satisfaction in community assembly dynamics.
- Watson, R.A., Mills, R., Buckley, C.L., Kouvaris, K., Jackson, A., Powers, S.T., Cox, C., Tudge, S., Davies, A., Kounios, L. and Power, D. (2016) Evolutionary connectionism: algorithmic principles underlying the evolution of biological organisation in evo-devo, evo-eco and evolutionary transitions. Evolutionary Biology, 43(4), pp.553-581.

Acknowledgements

I would like to thank Richard Watson for encouraging and enabling me to explore this fascinating topic. Thanks too, to my cosupervisors Patrick Doncaster and James Dyke for their feedback and guidance. Special thanks to thank everyone who was part of the Institute for Complex Systems Simulation Doctoral Training Centre at the University of Southampton, including Seth Bullock, Jason Noble, Hans Fangor, Nicki Lewin, Tracey Cantlie, plus Richard and James from my supervisory team, and all my fellow students. Thank you for creating such an invigorating and supportive research environment. Thanks to the Engineering and Physical Sciences Research Council (EPSRC) for supporting my work (EP/K503150/1). Thank you to my family, for all their support throughout. Finally, very special thanks to Helen, who never doubted that I would finish this one day.

Chapter 1

Introduction

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 in the present day, and how any such conditioning may shape these systems' responses to new pressures. To do this it is necessary to understand how the adaptive pressure on individuals affects the dynamical properties of the communities they inhabit. Specifically, we need to understand how coevolution of a community's parts affects behaviour of the whole. Yet, as communities are not units of selection, we currently lack a framework linking adaptive pressures on individuals and populations to the dynamical properties of the systems they inhabit. This leaves us unable to address issues such as whether coevolved communities differ from coevolutionary naive populations in the strength or nature of their interactions; whether coevolved and coevolutionary-naive populations differ in their trajectories towards single or multiple stable states; and whether the number and location of dynamical attractors are affected by the organisation of ecological interactions acquired over evolutionary time (Diamond, 1975; Connor and Simberloff, 1979; Thompson et al., 2001; Schulze and Mooney, 2012).

In this thesis we investigate how systems above the Darwinian levels of selection (Maynard Smith and Szathmáry, 1995) may evolve collective behaviours, and observe a deep homology between these and 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. We then use the results of this analysis in numerical simulation and examine the implications for: the evolution of ecological memory; the abilities of coevolving community networks to resolve conflicting constraints among species leading to higher community biomass; and the abilities of coevolving community networks to learning to solve complex resource-allocation problems.

1.1 Community assembly rules

Understanding how assembly rules operate in complex communities and the subsequent effects on the collective organisation and function of ecosystems is a long standing challenge in community ecology (Diamond, 1975; Connor and Simberloff, 1979; Thompson et al., 2001; Schulze and Mooney, 2012). Yet, as evolution by natural selection operates within individual species and not on community structure as a whole, the extent to which natural selection shapes assembly processes and contributes to community-level functions remains unanswered. Without this, we are not able to adequately address fundamental issues such as: how community dynamics are shaped by selective processes; whether evolutionary coevolved communities differ from naive populations in the strength or nature of their interactions; and whether coevolved and coevolutionarily-naive communities differ in their trajectories toward single or multiple alternative stable states (Thompson et al., 2001; Beisner et al., 2003). It is therefore a priority to develop theory on how macroscopic properties such as ecosystem function are affected by microevolutionary processes at lower levels within these systems, as species evolve in response to each other and their environment (Levin, 1998).

Macro-evolutionary studies in phylogenetic community ecology seek to determine how patterns of interspecific relatedness shape community assembly and subsequent effects on community composition, network structure and function (Webb et al., 2002; Cavender-Bares et al., 2009; Nuismer and Harmon, 2015). These studies can be traced back to observations by early ecologists, that closely related species are likely to have similar resource requirements and thus compete for similar resources (Cavender-Bares et al., 2009). A difficulty in these theories is that the pattern of community composition that emerges depends on how specific species' traits affect interspecific interactions and cooccurrence (Srivastava et al., 2012), as well how they evolve over time. Three general patterns can emerge: 1) Phylogenetic over-dispersion, in cases where closely related taxa tend to have similar niches across multiple dimensions, leading to competitive exclusion (Webb et al., 2002; Mayfield and Levine, 2010). (As an example, Graves and Gotelli (1993) showed that foraging flocks in the Amazon rarely contain congeners (also discussed in Webb et al., 2002)) (Figure 1.1a). 2) Phylogenetic clustering resulting from, e.g., environmental filtering on shared physiological tolerances (trait conservatism) (Webb et al., 2002; Mayfield and Levine, 2010) (Figure 1.1b). And 3), phylogenetically random patterns (Nuismer and Harmon, 2015). The sensitivity of community composition to the exact nature of the interactions between community members prevents predictions based on phylogenetic relatedness alone, and we still lack a mechanistic, process-based framework for predicting when phylogenies will have explanatory power (Nuismer and Harmon, 2015). To do this we need to incorporate micro-evolutionary processes and models of ecological interactions.

Micro-evolutionary approaches to understanding how evolution affects species distribution tend to focus on pair-wise relationships and how these impact species coexistence. Studies of competition are a core part of evolutionary ecology (e.g. Gause, 1934; Diamond, 1975; MacArthur and Levins, 1967; Finke and Snyder, 2008; Dieckmann et al., 2004) yet theories on pairwise relationships have limited explanatory power in species-rich communities where complex interdependencies mean that niche shifts, due e.g., to character displacement (Beans, 2014), are likely to be influenced by multiple aspects of community composition. If we are to understand the impact of pairwise changes in the structure of ecological networks and whether (as in other complex

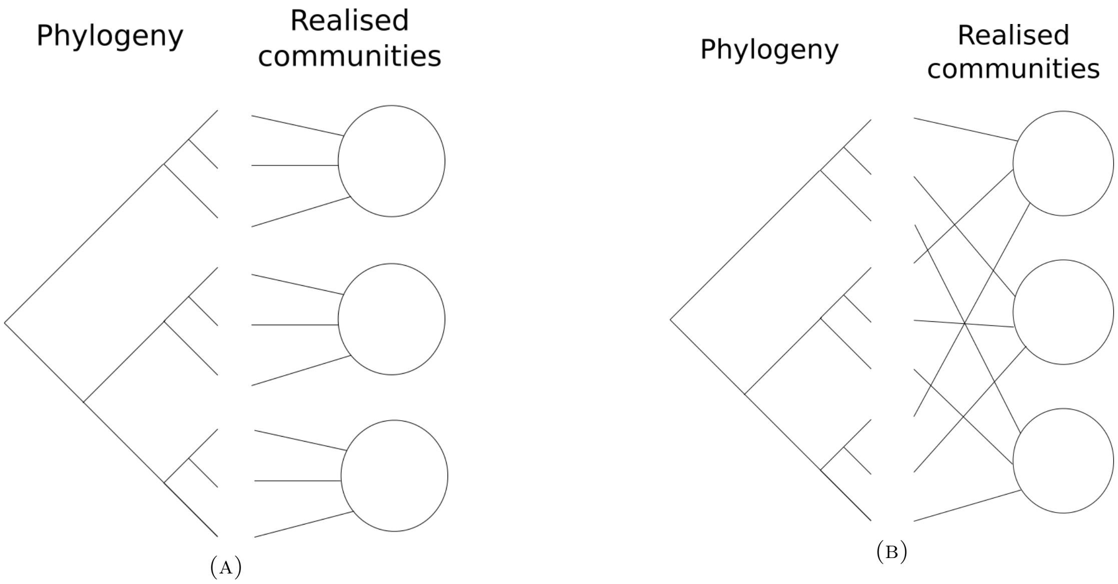

Figure 1.1: Patterns of phylogenetic composition in ecological communities. Phylogenetic clustering (a) occurs when closely related species tend to be found in the same communities or niches. Phylogenetic overdispersion occurs when the pattern of cooccurrance is less clustered than would be expected by chance. Adapted from Cavender-Bares et al. (2009).

adaptive systems; Levine, 1999) changes to network interactions drive any collective community behaviours, it will be necessary to bridge the gap between micro- and macro-scale theories. To do this we must work towards understanding how local interactions are influenced by the patterns of environmental filtering and phylogenetic history that shape species co-occurrence, incorporating theory on how these interactions evolve and their collective effect on community networks.

Difficulties in developing theory on how fine-scale eco-evo feedbacks on local or pairwise interactions affect collective community functions are threefold: 1) communities are not Darwinian units, and are therefore not directly shaped by Darwinian processes at this level; 2) there are good arguments concluding that evolution by natural selection cannot drive any organisation above the level of Darwinian units (Williams, 1966); and 3) identifying trends in the multitude of complex and evolving interactions within ecosystems can seem an impossibly complex task (Levin, 2011).

1.2 Evolutionary determinants of community functions

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 (Levin, 1998, 2005; Anand et al., 2010; Whitham et al., 2006; Ulanowicz, 1986). For example, the structure of the network of interactions in an ecosystem (Neill, 1974; Proulx et al., 2005) affects many of the system's dynamical behaviours including succession dynamics and community assembly rules (Clements, 1916; Weiher and Keddy, 2001), the stability, resilience and adaptive capacity of a community (Holling, 1973; Gallopín, 2006; de Ruiter et al., 1995; Folke, 2006; Staniczenko et al., 2010), the presence of alternative stable states (Higgins et al., 2002; Cropp and Gabric, 2002; Beisner et al., 2003), and the system's susceptibility to regime shifts (Scheffer et al., 2001).

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 (Lenton and van Oijen, 2002; Ulanowicz, 1990; Lenton, 2004; Richardson, 1980; McIntosh, 1995). However, an ecological community is not, in most cases, an evolutionary unit (Whitham et al., 2006; McIntosh, 1995; Levin, 2011); 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 (Smith et al., 1985). 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 (Anand et al., 2010; Lenton and van Oijen, 2002; McIntosh, 1995; Gleason, 1926; Lawton, 1999; Milne, 1998; Thompson et al., 2001; Matthews et al., 2011), in particular, the reciprocity between the ecological dynamics on the network and the evolutionary changes to the nodes, and hence, connections of the network (Post and Palkovacs, 2009; Matthews et al., 2011; Gross and Sayama, 2009; Paperin et al., 2011; Schoener, 2011; Turcotte et al., 2011; Metz et al., 1996; Fussmann et al., 2007). 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 if we are to understand 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 (Harcombe, 2010) 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 (Romano-Keeler and Weitkamp, 2014). Coevolutionary processes in gut microbiota have shaped at least three alternative stable states (termed enterotypes) (Faust and Raes, 2012), but it remains unclear how different communities' networks of interactions evolve in response to environmental changes in cases such as the use of antibiotics (Forslund et al., 2013) and societal changes in diet (David et al., 2014), or how these changes affect the emergent properties of community networks (Shade et al., 2012) 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 Wallaces line (Case et al., 2005; Mayr, 1944; Wallace, 2011). 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 (Case et al., 2005). (Mulga woodland are another interesting example of communities that can exist in alternative stable states; see Folke et al., 2002) Yet, without a framework linking microevolutionary changes in interactions between species pairs to dynamical behaviours of whole communities, it remains moot as to whether a network of coevolved interactions could be the explanation for the observed dynamical stability.

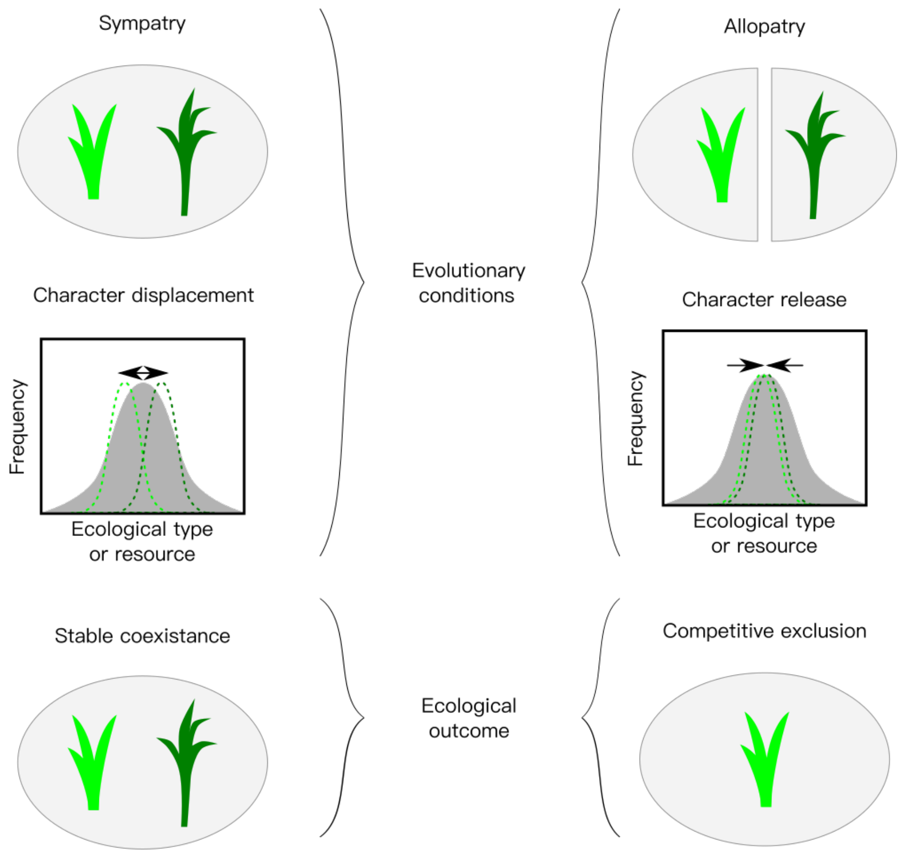

Figure 1.2: Character displacement, in which species that cooccur, evolve to become more likely to cooccur in future, and the converse phenomena, character release. Left images: Character displacement, occurring when species evolve in coexistence (sympatry) with each other results in ecological divergence and future compatibility. Right images: Character release occurs when species evolve allopatrically, resulting in convergent evolution and competitive exclusion. This thesis will draw an analogy between evolutionary processes such as these, where species that evolve together become more compatible, and a distributed associative learning process (Hebb's rule; Hebb, 1949) in which neurons that 'fire together, wire together'.

The need to characterise the evolutionary and historical determinants of ecological processes is identified as an important frontier in ecological research (Thompson et al., 2001). 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 (Higgins et al., 2002; Staniczenko et al., 2010; Angeler et al., 2010; Carpenter et al., 1999) 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 (Beisner et al., 2003).

These issues connect deeply with the phenomenon of *ecological memory* (Thompson et al., 2001; Hendry and McGlade, 1995; Peterson, 2002; Golinski et al., 2008). The term 'ecological memory' is used in different ways in ecological literature. In some usages it refers to 'biological legacies', such as the existence of rare species or seed banks in an ecosystem, or persistent changes to soil composition, 'recording' previous growth or species invasions (Hendry and McGlade, 1995). In this thesis we are specifically interested in the effects of genetic changes in traits that affect species interactions with one another; how past evolution affects future ecological behaviour. We employ the term ecological memory in the same sense as Thompson et al. (2001): '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, (Brown and Wilson, 1956; Dayan and Simberloff, 2005) 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 (Figure 1.2). Within ecosystems these genetic memories, distributed amongst multiple populations, influence assembly processes, stability, and resilience of the developing and mature ecosystem (Thompson et al., 2001). 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 (Thompson, 2005; Thompson et al., 2001). These include issues relating to 1) how changes to interactions evolved in past environmental conditions affect the response of the community to future changes or perturbations in environmental conditions; 2) the nature of ecological memory; whether it is capable of retaining information about multiple distinct past states without just averaging them; 3) whether the assembly rules and succession dynamics of a community can be systematically organised by selection in past environmental states; and 4) how the formation of an ecological memory affects the possibility of alternative ecological stable states, and regime shifts under subsequent environmental forcing (Box 1).

Box 1. **Dynamical landscapes of community networks.**

For many properties of interest, ecological communities can be usefully modelled as a set of dynamic state variables and interactions defined by parameters (e.g. Lotka Volterra models of population dynamics examine the effect of interspecific interactions (parameters) on population densities (variables)). If the system of equations describing the transformation of variables is non-linear, such as in an ecological network of interacting species, variables may persist in one of a number different, locally stable compositions (stable states), but where local stability does not imply global stability (Lewontin, 1969). The community returns to the same state following small perturbation but may shift to an alternative stable state if subjected to a larger disturbance.

A useful heuristic for thinking about community dynamics is a ball and cup analogy, where possible states are represented by a surface or landscape, and the current state of the system represented by a ball moving on the surface (Figure 1.3). Without outside influence community assembly rules cause the ball to roll downhill to a stable state composition. The basin of attraction for each stable state is defined as the set of community compositions that converge (assemble) to that state; stable states with large basins of attraction are those that are assembled from a high number of state compositions. The stability of each stable state can be measured through its amplitude (the size of perturbation the community can withstand and still return to the same stable state) and elasticity (speed of recovery from disturbance) (Westman, 1978).

Under this framework the system can shift state in two ways. If parameters are fixed, the system can only shift from a stable state through adjustment to state variables, such as through perturbation due to outside forces. This is the case in ecological regime shifts, where environmental forcing shifts the system state into a different basin of attraction, in which case an equal but opposite force may not necessarily not return the system to the original basin of attraction (hysteresis, see e.g. Carpenter, 2001). But, if system parameters can change, changes to interactions between system variables can alter the dynamical landscape, potentially altering the location of stable states. Beisner et al. (2003) refers to these two perspectives as the ‘community perspective’ (fixed parameters; fixed number and location of stable states) and the ‘ecological perspective’ (variable parameters; indeterminate number and location of stable states). The fundamental difference between perspectives is timescale. The community perspective considers shorter periods of time in which parameters such as interspecific interactions are assumed to be constant. Under the ecological perspective these parameters do change (or evolve, in the case of species interactions), with consequent effects on the dynamical landscape.

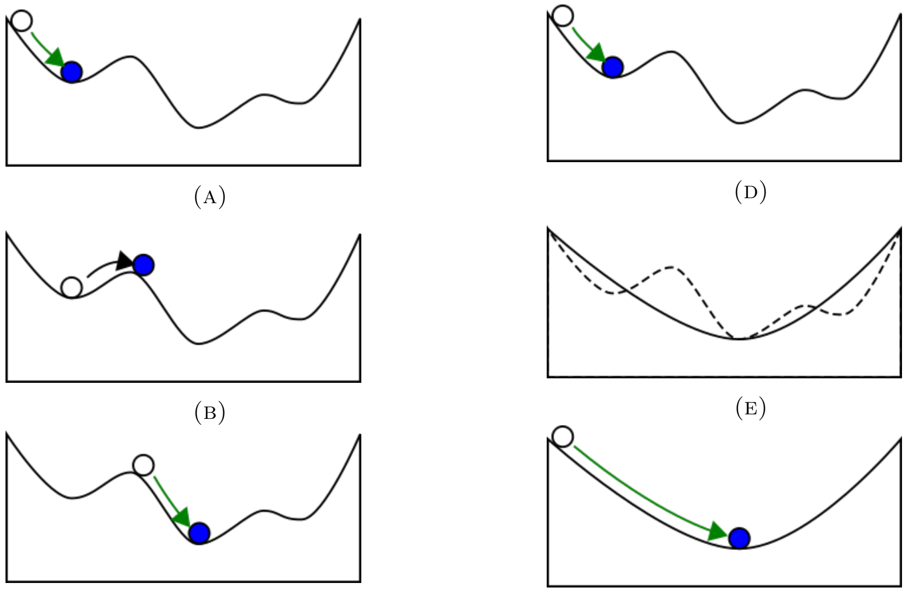

(a)

(d)

(b)

(e)

(c)

(f)

Figure 1.3: Ball and cup representation of differing perspectives community dynamics. Left hand side: Community perspective; fixed parameters define an unchanging dynamical landscape. In a multimodal landscape the system can exist in a number of stable states, with sensitivity to initial conditions determining the state reached during assembly (green arrow). Under the community perspective, once the system has converged to a stable state (a) the system can only shift to a different attractor if subjected to perturbation (black arrow) sufficiently large to move it into the basin of attraction for another stable state (b) & (c). Right hand side: Ecosystem perspective; parameters governing interactions between variables can change, resulting in a shifting dynamical landscape in which attractors may be created and destroyed. From a system in which a set of initial conditions assemble to a certain stable state (d), changes to the dynamical landscape (e) can result in the same initial conditions assembling to a different state composition (f). (Another common way of representing the changing dynamics of the ecosystem perspective is with 3D figures with time as the third axis (Dieckmann et al., 2004; Baedke, 2013).)

1.2.1 Understanding the effects of eco-evo feedbacks on community function

Perhaps one reason we lack theory on how microevolution of interactions affects network structure is the understandable and long-running controversy around the potential of selective processes to have effects at higher levels of organisation. In older literature, particularly in studies of animal behaviour, sometimes individuals are said to behave 'for the good of the species.' For example, V. C. Wynne-Edwards argued for animal behaviour that involved individual sacrifice for species-level benefit (Wynne-Edwards, 1964). Viewpoints such as these largely persisted until Williams (1966), who show showed that higher levels of biological organisation, including species and ecological communities, cannot be organised by natural selection to perform collective functions because selective processes do not (except in very special cases) operate at the these levels (i.e. they are not Darwinian units), and the due to inevitable appearance of cheaters that would destabilise any harmony (tragedy of the commons) (Axelrod and Hamilton, 1981). While group selection theory (or multi-level selection theory) has since made a resurgence under Wilson (1980) and many others (see esp. Okasha, 2006), the general view is that while it may be possible for Darwinian units to exist at the community level, the conditions for this are considered to be rare (e.g. Faust and Raes, 2012), likely to have limited effect, and explanations based on the lowest level of organisation that will explain the observed phenomena are preferable (Williams, 1966).

Under Williams' (and subsequently Dawkins, 1976) unit-centric framework, interactions between components (e.g. community members) may evolve, but explanations of consequences can be wholly derived from understanding changes at each unit. Consequently, while the prevailing perspective acknowledges that communities have an interaction structure, a tendency to seek reductionist explanations has the effect that the potential consequences of natural selection organising that interaction structure in a consistent manner could be overlooked. This potential oversight would leave researchers reliant upon sum-of-the-parts explanations for system-level functions; attempting to explain all properties of a community solely through summing the evolutionary pressures on the Darwinian units in that system, rather than as a complex adaptive system (Levin, 1998) where dynamics exhibit nonlinearity, leading to historical dependency and assembly rules have multiple possible outcomes.

Explanations of community function build on sum-of-the-parts type explanations can be complemented by considering the consequences of individual selection acting on interactions - which can be even more important than the species themselves for understanding how the whole system is functioning and organised (Jordán and Scheuring, 2004). For example, Leigh and Vermeij (2002) seek to understand how ecosystems maximise their biomass, and find that this can result from each selfish component of a system being selected so as to maximise its individual production. If we additionally consider (for example) species' ability to evolve the sensitivity of their growth pattern with respect to the ecological context provided by other community members, then changes in community biomass are also influenced by the extent to which such individually evolved assembly rules create combinations of species that avoid competition with each other (Zuppinger-Dingley et al., 2014). By incorporating the effects of eco-evo feedbacks on community interactions, reductive explanations of community function are enhanced by understanding of the effects of natural selection acting on the very many interspecific links within a community network structure. This has potential to enhance our treatment of issues such as: 1) how natural selection alters the dynamical attractors of community assembly rules; 2) whether mechanisms exist by which the evolution of interspecific interactions results in reduced community-level constraints; 3) whether conditions exist that facilitate evolution of constraint satisfaction behaviour; and 4) the consequences of collective constraint satisfaction for community behaviour and function.

1.2.2 Approach

The lack of a theoretical framework that links individual adaptations to collective behaviours leaves the above sets of questions on ecological memory and collective constraint satisfaction behaviour difficult to answer. This thesis addresses these issues by converting and exploiting theory that is already well-developed in another domain, where networks of simple components can be organised so as to produce powerful collective behaviours. Specifically, we explore the relationship between connectionist models of memory and learning, and find isomorphisms between these models and models of evolving ecological networks. This facilitates a transfer of tools and understanding that sheds new light on the behaviour of ecological communities.

1.3 Connectionism

Connectionism is the study of how systems of individually-simple components can exhibit collective behaviours when interactions (connections) between units are appropriately organised (e.g. how complex cognitive processes arise from the organisation of individually-simple neurons in the brain; Rumelhart and McClelland, 1987). 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 McCulloch and Pitts, 1943; Hinton et al., 1999; Hopfield, 1982), if appropriately organised (or connected), a network of such units can provide many remarkable collective behaviours. This includes memory-like behaviours resulting from the ability of these systems to develop multiple, specific attractors relating to experienced or trained states (McCulloch and Pitts, 1943; Hinton et al., 1999; Hopfield, 1982; Watson et al., 2011a) as well as constraint satisfaction behaviours, resulting from the ability of these systems to 'self-learn' attractors for low-competition states, without having experienced these states (Watson et al., 2011a).

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; Farmer, 1990; Watson et al., 2011b; Mikhailov et al., 1990; Fernando et al., 2009; Vohradsky, 2001; Noonburg, 1989; Poderoso and Fontanari, 2007). 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; Wilson, 1992; Noonburg, 1989; Poderoso and Fontanari, 2007).

A deficit in the analogy between neural networks and ecosystems is that whereas neural networks acquire the organisation necessary for their collective behaviours through distributed learning mechanisms designed for that purpose, ecological connections are modified by individual-level natural selection with no such system-level purposes in mind. Although ecological networks may have population dynamical similarities with neural 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 distributed 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 (Watson et al., 2014).

Learning mechanisms in neural network models have two basic types (Watson et al., 2014). *Supervised* learning utilises an external reward signal, or error function, to direct incremental changes to connections. Watson et al. (2011b, 2014) 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. 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, which operates without a reward signal. It 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' (Hinton et al., 1999; Hopfield, 1982; Watson and Szathmáry, 2016; Ackley et al., 1985). 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.* The simplest unsupervised correlation learning mechanism is Hebbian learning (Hebb, 1949), well-understood in neural network models of memory and knowledge representation (Hinton et al., 1999; Hopfield, 1982; Ackley et al., 1985).

Hebb's rule, paraphrased as neurons that 'fire together, wire together', modifies connections in a very simple manner: adjusting the strength of connections in the direction that amplifies or reinforces the current output of each unit (Hebb, 1949). It is a fully distributed learning mechanism, in that the change in connections is function of only the two connecting nodes. Specifically, the change in strength of a synaptic connection, $\Delta\omega_{ij}$, is proportional to the co-activation of the neurons it connects: i.e. $\Delta\omega_{ij} = rx_ix_j$, where $r > 0$ is a learning rate, and x_k is the activation level of node k. 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.

Despite its relative simplicity, Hebbian learning in distributed networks is known to produce a range of complex system-level computational processes, including forming 1) a distributed associative memory which can store and recall multiple patterns of activation in the organisation of synaptic connections (Hopfield et al., 1986; Hopfield, 2008), facilitating the use of these networks use in pattern recognition, noise reduction and classification (Box 2, Figure 1.4); and

2) and forming attractors for states that minimise inter-unit constraints, including the ability to ‘self-model’ enabling them to solve complex constraint satisfaction problems (Watson et al., 2011b; Hinton et al., 1999; Hopfield, 1982; Ackley et al., 1985) (Figure 1.5).

Box 2. **Hopfield networks and unsupervised learning using Hebb's rule.**

The Hopfield network model (Hopfield, 1982) 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 1.4). Hebb's rule is an unsupervised learning technique (Prugel-Bennett and Shapiro, 1993) 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).

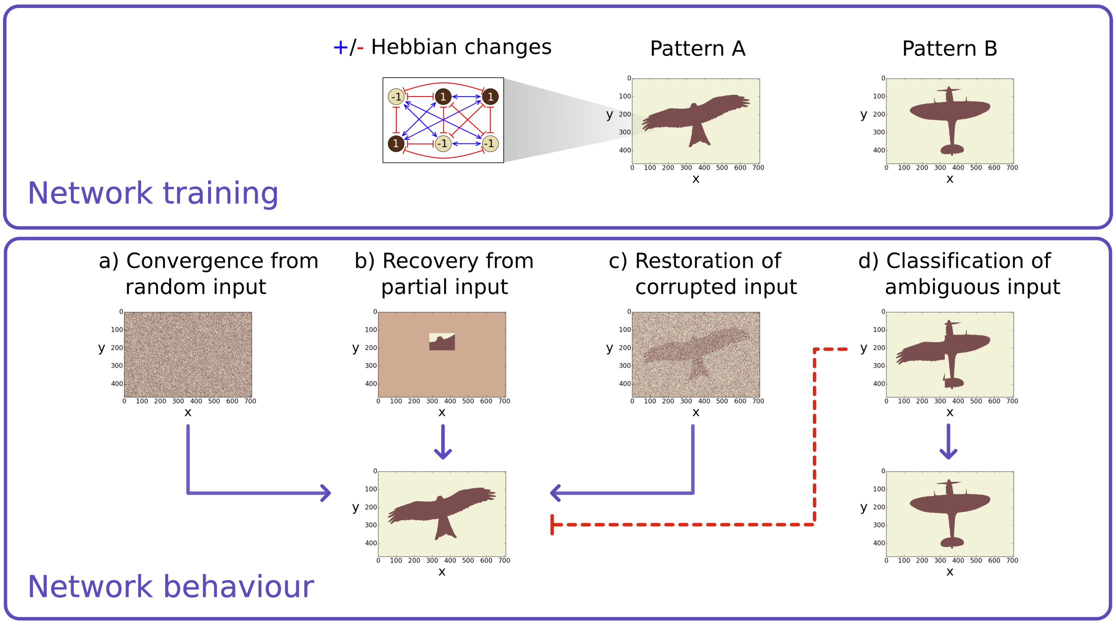

FIGURE 1.4: Training a Hopfield network using Hebb's rule.
Network training: Unsupervised learning processes as used to train a Hopfield network to store two configurations, patterns A and B. Each unit in the Hopfield network corresponds to a pixel in the image display. Six units are highlighted to illustrate the changes to connections during training in pattern A. Hebbs 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).

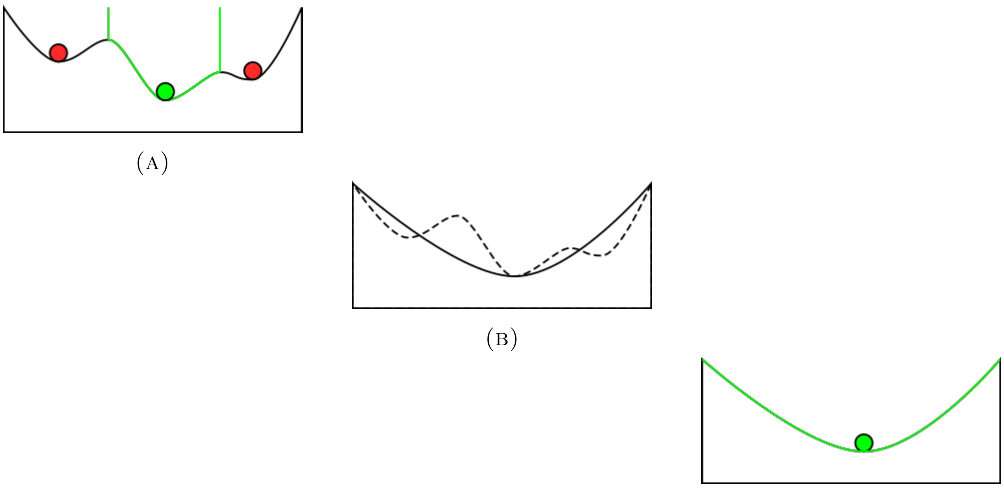

(A)

(B)

(C)

Figure 1.5: Learning to solve constraint satisfaction problems. Constraint satisfaction problems (CSPs) are computational problems where the challenge is for a dynamical system to learn (develop an attractor for) a configuration that maximally satisfies constraints in a network. a) Typically, the initial system state in a CSP consists of a dynamical landscape where the global optimum for constraint satisfaction (green ball) is very hard to find due to the presence of very many local optima (red balls in the diagram); there may be initial conditions that converge to the maximally satisfying state, but there are many more initial conditions that converge to sub-optimal states for constraint satisfaction. b) During the learning process changes to interactions between units alter the dynamical landscape. c) The system has *learned to solve* the CSP if it is able to sufficiently enlarge the basin of attraction corresponding to the maximally satisfying state, such it can reliably find that state (i.e. solve every time). The diagram shows a simplified system that initially has three attractors. Computationally challenging CSPs, such as the eight queens problem or Sudoku, have many thousand sub-optimal attractors, with the basin of attraction for the global optimum usually being sufficiently small that is cannot be found in reasonable time through stochastic search methods.

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; Laland et al., 2014, and refs. within) but the full implications of this feedback have not been realised (Laland et al., 2009). For example, this feedback is part of the backstory involved in invasional meltdown (Gallardo and Aldridge, 2015; Heimpel et al., 2010; Simberloff and Von Holle, 1999) 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' (Gallardo and Aldridge, 2015). 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. (See also the fascinating example of utricularia, and its associated periphyton and zooplankton grazers (Ulanowicz, 1995), in which a three-way autocatalytic loop emerges and no species can be understood without considering the other two.)

1.4 Species that fire together, wire together

The dynamical equivalence between many mathematical models of neural networks and ecological models used to simulate the growth of interacting populations, is well recognised and not new (Noonburg, 1989). A neural network can be considered as a network of interconnected units in which activity levels at each unit are inhibited or excited by the activity of the other units. This is logically equivalent to an interacting set of populations with the density of each population affected negatively or positively by its competition or cooperation with the other populations. With constant coefficients, however, the Lotka-Volterra model does not have the capability for modelling adaptation, which is a vital component in any neural network or associative learning model.

Our observation however, that there exists equivalence, not just in dynamics, but in how these systems self-organise invites the intriguing possibility that these systems may share other similar properties, with consequences for our understanding of natural communities. If natural communities can self-organise in the same way as neural networks, there may be conditions in which they exhibit the fascinating behaviours of these dynamically-similar systems. This has potential to help us understand how low-level selective pressures in natural communities natural communities may contribute to system-level community functions, improving our understanding of these system's ability to form ecological memories and resolve internal constraints.

Table	**Unsupervised correlation learning**	**Coevolution in ecological networks**
2.1	Activation dynamics.	Population dynamics.
2.2	Changes in connections due to unsupervised correlation learning.	Evolution of interactions due to individual selective pressure (Chapter 2).
3.1	Ability of a neural network to form a distributed associative memory.	Collective organisation of community networks resulting in a distributed ecological memory (Chapter 3).
4.1	Conditions required for neural networks to learn to solve constraint satisfaction problems.	Conditions required for ecological community networks to evolve assembly rules with enhanced constraint satisfaction behaviour (Chapter 4).
5.1	Ability of a neural network to learn to solve complex constraint satisfaction problems by augmenting its dynamical behaviour with correlational information from its own attractors.	Ability of an ecological community network to evolve its assembly rules so as to solve complex resource-allocation problems (Chapter 5).

Table 1.1: Distributed associative learning mechanisms in neural networks and ecological community networks. A roadmap for this thesis, and the first in a series of tables we use to highlight the equivalence between these systems. This equivalence stems from consistent dynamical behaviours between these systems, previously identified (Noonburg, 1989) and described in more detail in Table 2.1. Table 2.2 lays out a primary contribution of the thesis, identifying that the selective pressure on interspecific interactions in a community network has a natural isomorphism with simple associative learning rules from connectionist models, demonstrating that natural communities *evolve*, like neural networks *learn*. We build from this isomorphism and show conditions where these self-organising processes result in community structure and assembly dynamics that exhibit non-trivial functional properties at the community level, demonstrating: how distributed learning processes enable each system to form distributed associative memories of past states (Table 3.1); how equivalent conditions in each system enable these systems to learn to resolve conflicting constraints (Table 4.1); and how each system is able to learn to solve complex combinatorial problems by augmenting its dynamical behaviour with correlational information from its own attractors (Table 5.1).

1.4.1 Thesis structure

This thesis uses computational models to demonstrate the potential for distributed associative learning mechanisms to occur in ecological community networks. Recognising this enables us to leverage theory from neural networks, where the effects of distributed associative learning mechanisms are more familiar and better understood.

We begin by showing that species evolve their interactions in proportion to their cooccurrence, and observe that the selective pressure on interactions is isomorphic to the modification of interactions in a neural network trained with Hebb's rule (Chapter 2). We build from this result and simulate ecological networks, demonstrating conditions in which changes in interspecific interactions in these networks are 'Hebbian' (Chapter 3). Recognising the equivalence between these systems enables transfer of theory between disciplines and enables us to make several predictions regarding the consequences of distributed associative learning in ecological communities. This includes: i) The ability of an ecological community to form a distributed associative memory of multiple past states (Chapter 3). ii) The ability of an ecological network to resolve internal constraints, in the same conditions that are required for a neural network to do the same (Chapter 4). iii) The ability of an ecological network to perform 'predictive' self modelling and augment their dynamical behaviour using correlational information from their own attractors to learn how to solve complex resource allocation problems. Each of these behaviours occurs in the same manner as in neural networks and can be anticipated by understanding the isomorphism between these two systems. Table 1.1 sets out the relationship between the components of the analogy.

Chapter 2

Analysis of individual-level natural selection acting on ecological interactions

2.1 Introduction

In order to understand the historical determinants of community functions, it is necessary to develop a framework linking the effects of natural selection on community members to the assembly rules of the communities they inhabit. In this chapter we develop such a framework, building from well understood Lotka-Volterra models, and incorporating the effects of natural selection on interspecific interactions. We then make new comparisons between these processes and unsupervised learning processes in neural networks.

The dynamical equivalence between Lotka-Voltera models, where populations experience exponential growth asymptotically approaching a threshold, and those models used in neural networks of excitation/inhibition between neurons is well recognised (Noonburg, 1989, and refs. within). Units within each system (species densities vs. neurons) respond in an isomorphic manner to the network of interactions between them (see Table 2.1 for a full breakdown of this analogy). Yet until now, the analogy between these systems has not extended further than this. Typical usage of Lotka-Volterra models involves ecological dynamics over relatively short timescales in which interspecific interaction coefficients do not change. Without the ability to change network interactions the Lotka-Volterra model does not have the capability for modelling adaptation, a necessary component in any neural network or learning model.

Species experience selective pressure on interspecific interactions due to niche overlap with other community members. High-density competitors can drive competition avoidance which causes species to modify their niches, expanding them into less utilised space (for example, by switching dependency to less-utilised resources). In this chapter we examine the selective pressures on interspecific interactions in the Lotka-Volterra model and derive formula for the rate of adaptation of interspecific interactions in these systems. We find that the adaptation of interspecific

	Unsupervised correlation learning	**Coevolution in ecological networks**
2.1.0	Activation dynamics.	Population dynamics.
2.1.1	Neural activation level.	Species density, x_i.
2.1.2	Neural activation pattern.	Ecological state, $X = \{x_i, x_2, ...x_N\}$.
2.1.3	Synaptic connection strength, ω_{ij}.	Inter-species fitness interaction, ω_{ij}.
2.1.4	Neural network (weight matrix, W).	Ecological network (community matrix, Ω).
2.1.5	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).

Table 2.1: The dynamical equivalence between activation dynamics in a neural network and population dynamics in ecological community networks (previously recognised; Noonburg, 1989).

interactions in a community network is isomorphic to the adaptive dynamics in neural networks trained with a well known unsupervised learning rule.

We demonstrate this behaviour by training a neural network, but substituting the usual learning rule with the function describing the changes in interactions in an evolving ecological network. We demonstrate that this generates the same interesting behaviours as in a neural network trained with Hebb's rule (Hebb, 1949). Finally, we test the generality of this finding by examining rates of adaptation in more complex ecological scenarios, and find this isomorphism is robust to the underlying choice of model.

2.2 Methods

2.2.1 Lotka Volterra model

We model an ecosystem state as a vector of population densities over N species, $X = \{x_1, x_2, ..., x_N\}$, $(\forall i : x_i \geq 0)$, and an interaction matrix representing the community network (Neill, 1974) Ω, in which each element, ω_{ij}, represents the effect of species j on the growth rate of species i (but not vice versa) relative to i on itself (ω_{ii} = -1). In all interaction matrices we assume competitive (non-trophic) interactions such as via competition for resources, such that $\forall i, j : \omega_{ij} \leq 0$. We use a Lotka-Volterra competition equation to define 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 (Wilson, 1992):

$$\frac{dx_i}{dt} = \frac{m_i x_i}{k_i}\left(k_i + \sum_{j=1}^{N} \omega_{ij} x_j\right) \tag{2.1}$$

where x_i, m_i and k_i are the density, growth rate and carrying capacity of species i respectively. We now turn our attention to the selective pressures on interspecific interactions.

2.2.2 Evolution of interactions under individual selection

Each interaction coefficient in Equation 2.1 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 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 (Hutchinson, 1965) 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, and we do not incorporate drift in our model. Thus the only changes to interactions are those selected for by individual-level natural selection, and only those changes to traits that directly affect the growth rate of an individual compared to the rest of the individuals in the species.

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 (Wilson, 1980). Therefore any changes to the growth rate of a species that occur as a side-effect of altering the density of some other species (e.g., via changes to ω_{ji}) are not affected by individual selection.

Only mutations to an individual in species i that decrease the competitive effect, ω_{ij}, of species j on species i directly affecting the fitness of the individual carrying the mutation and not other individuals in species i, can thus be favoured by individual selection. Examples of mutations with this property include those that cause an individual to spend more of its time in habitats away from a high density competitor, or those that reduce its dependency on a resource shared with a competitor.

2.3 Results

2.3.1 Analysis of individual-level natural selection acting on ecological interactions

We analyse the rate of accumulation of favourable mutations, v, 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 (Neher et al., 2010; Weissman and Barton, 2012) to the particular case. From the ecological dynamics it follows that the selective coefficient, s, conferred by the change, g, in the interaction coefficient $\omega_{ij'}$, 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:

$$s = \frac{m_i}{k_i}\left(k_i + \sum_{j=1}^{N}\omega_{ij}x_j + gx_{j'}\right) - \frac{m_i}{k_i}\left(k_i + \sum_{j=1}^{N}\omega_{ij}x_j\right) = \frac{m_i}{k_i}gx_{j'} \tag{2.2}$$

Henceforth simplified as:

$$s = \frac{m_i}{k_i}gx_j \tag{2.3}$$

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: $v = x_i\mu\bar{P}$ (Weissman and Barton, 2012). In large sexual populations with linked loci, $\bar{P}$ will depend on v, and in different ways depending on the type of recombination, recombination rate, population size, the mutation rate and magnitude of mutations (Neher et al., 2010; Weissman and Barton, 2012). 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 (Weissman and Barton, 2012). 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, v_{ij}, in an interaction coefficient, ω_{ij}, is given by:

$$v_{ij} = x_i\mu s = \frac{m_i}{k_{ie}}g\mu x_i x_j \tag{2.4}$$

2.3.2 Equivalence between rate of adaptation and unsupervised learning in a neural network

Equation 2.4 tells us that the rate of adaptation on inter-species relationships is proportional to the co-occurrence of the species involved. Hence, as with Hebb's rule (Hebb, 1949), *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). This is exactly as per the principle of unsupervised correlation learning, which 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. (Note that change in interactions is not directly affected by current interaction rates; it is primarily a function of unit cooccurrence. Thus competitive interactions can, in theory, become mutualistic, just as negative interactions can become positive in Hopfield networks trained with Hebb's rule.)

We demonstrate the equivalence between Equation 2.4 and Hebbian learning by comparing the results of training a neural network with each of these functions. We show how training with either function enables a Hopfield network to build an associative memory of distinct states, enabling consequent behaviours including pattern recognition, noise correction and classification (Figure 2.1). The fundamental property that enables this equivalent result is that both functions modify connections between units in proportion to the cooccurrence of those units.

2.3.3 The relationship between rate of adaptation and product of species densities in more complex cases.

Thus far, we have only considered the case where there is no interference between simultaneously segregating alleles at different loci. However, 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. We now compare the rate of adaptation of an interaction coefficient for two additional models. In each case, the rate of adaptation, v_{ij}, of an interspecific 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. In all cases, $\bar{P}$ is a function of the selection coefficient $s_i = \frac{m_i}{k_i} g x_j$ (Equation 2.3) where m_i is the intrinsic net growth rate of species i, k_i is the carrying capacity of species i, 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.

2.3.3.1 Case a) No interference

We have already explored the simple case 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),

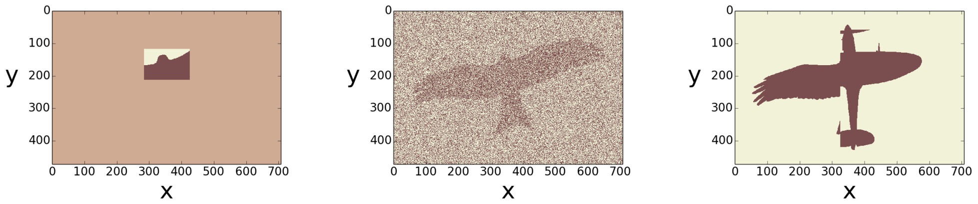

(A) Input patterns

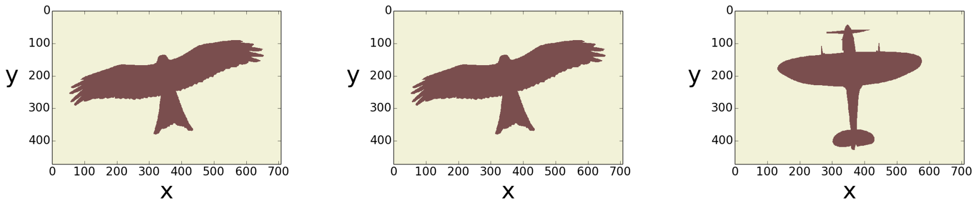

(B) Outputs from Hopfield network trained with Hebbs rule

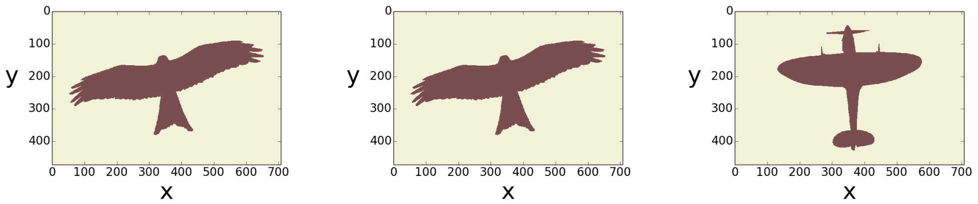

(C) Outputs from Hopfield network trained with Equation 2.4 from this thesis

FIGURE 2.1: Equivalent dynamical behaviours resulting from training a Hopfield network using Hebb's rule and the equations discussed in this chapter (Hopfield, 1982; Hebb, 1949). Row A: Input patterns to Hopfield network. Patterns specified so as to demonstrate the ability perform pattern reconstruction, error or noise correction and classification of an ambiguous input. Row B: Corresponding outputs from a Hopfield network trained to recall bird and plane patterns using Hebbian learning. Row C: Outputs from a hopfield network trained to recall bird and plane patterns using Equation 2.4. The Hopfield network trained with Equation 2.4 is able to perform all of the behaviours exhibited by the Hopfield trained in the conventional manner (with Hebbian learning).

the probability of fixation, $\bar{P} = s_i$. Thus, as per Equation 2.4, substituting for s_i as above:

$$v_{ij} = x_i \mu s = \beta \mu x_i x_j \tag{2.5}$$

2.3.3.2 Case b) Linked genes on a linear genome

Weissman and Barton (2012) consider the effects of interference between linked genes on a linear genome. Here the genomic rate of fixation of beneficial mutations is:

$$v = \frac{v_0}{1 + 2v_0/R} \tag{2.6}$$

(Weissman and Barton, 2012, Equation 7)

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:

$$v_{ij} = \frac{2\beta\mu x_i x_j}{1 + 4\beta\mu x_i x_j / R} \tag{2.7}$$

2.3.3.3 Case c) Occasional outcrossing

Neher et al. (2010) study the rate of adaptation in unlinked loci in facultative sexuals where the rate of outcrossing is very small. Whereas Weissman and Barton (2012) 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:

$$v \approx x\mu s^2 \left(1 - \frac{4x\mu s^2}{r^2}\right) \tag{2.8}$$

(Neher et al., 2010, Equation 12b)

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

$$v_{ij} \approx x_i \mu (\beta x_j)^2 \left(1 - \frac{4 x_i \mu (\beta x_j)^2}{r^2}\right) \tag{2.9}$$

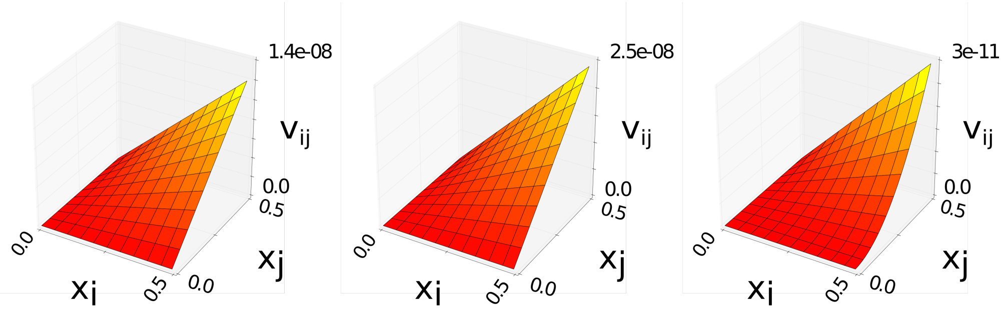

(a)
Our model

(b)
W&B2012
eq 7

(c)
N&al.2010
eq 12b

Figure 2.2: Rate of adaptation v_{ij} as a function of x_i and x_j for three different models. a) Equation 2.4 from this thesis. b) Equation 7 from Weissman and Barton (2012), c) Equation12b from Neher et al. (2010). 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. Parameters set as per the experiment we demonstrate in Chapter 3: $k_i = 10$, $m_i = 0.5$, $g = 0.1$, $\mu = 1.0 \times 10^{-5}$. For case b), the map length, $R = 1$. For case c), the out-crossing rate, $r = 0.01$.)

2.3.3.4 Comparison of the three cases

We compare the the rate of adaptation v_{ij} as a function of x_i and x_j for these three different cases in Figure 2.2. We observe that case a, where the rate of adaptation is directly proportional to the product $x_i x_j$ as in Equation 2.3, 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 (as in Weissman

	Unsupervised correlation learning	**Coevolution in ecological networks**
2.2.0	Changes in connections due to unsupervised correlation learning.	Evolution of interactions due to individual selective pressure.
2.2.1	Positive feedback between activation strengths and connection strengths. Also known as *neurons that fire together, wire together.*	Positive feedback between ecological densities and interspecific interactions, or *species that occur together, wire together.*
2.2.2	Unsupervised correlation learning mechanism, Hebb's rule: $\Delta\omega_{ij} = rx_ix_j$, where $r > 0$ is a learning rate.	Direct effects of individual natural selection on interactions: $v_{ij} = rx_ix_j$, where $r = \frac{m_i}{k_i}g\mu$ describes the available mutation (Equation 2.4).

Table 2.2: The natural isomorphism between learning processes in unsupervised neural networks and the evolution of interactions in ecological community networks. Interactions in each network change as a function of units' cooccurrence.

and Barton, 2012), 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 is zero, and otherwise, the greater the value of one, the greater the rate of increase with the other. Although the shape of the alternate functions differs, this essential behaviour is preserved; intuitively, mutations must be both created and selected for an interaction coefficient to evolve. As already shown for Equation 2.3, we can demonstrate that each of these functions are able to produce the same necessary pattern of change in a neural network to form a distributed associative memory (Appendix, Figure A.1.)

2.4 Discussion

We identify that the rate of change in community interactions due to selective pressure on individual species is proportional to the cooccurrence of the species involved. This is the first key result of this thesis, describing how selection acts on interspecific interactions as a function of the current ecological state. Identifying this relationship is an essential first step in bridging the gap between micro-evolutionary processes and collective community functions.

This finding is robust to the choice of underlying model. In more complex cases, where there is interference between alleles at different loci, v_{ij} may not be linearly proportional to x_ix_j as it is in Equation 2.4, 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 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 (Thompson, 2005).

The second important result of this thesis is in identifying the isomorphism that exists between these dynamics and unsupervised learning processes in neural networks, where neurons that 'fire together, wire together' (Hebb, 1949). This finding has potential to open up transfer of knowledge between these disciplines. That is, given the known equivalence at dynamical level (Table 2.1), and by identifying equivalence in how interactions are modified in these systems

(Table 2.2), we open up the intriguing possibility that natural communities may perform some of the fascinating computational processes exhibited by artificial neural networks.

We begin this transfer of knowledge in the next chapter, where we identify that distributed associative learning can take place in ecological networks, producing similar behaviours as we show in this chapter for neural networks.

Chapter 3

The evolution of ecological memory

3.1 Introduction

Ecological assembly rules are shaped by the genetic structure of species and the evolutionary history of communities, along with the residual effects of ecological events. The evolutionary history of a community encodes in the genetic structure of species a record of past historical conditions that constrains the trajectories of community assembly. This is the phenomena of ecological memory: the genetic changes to community members resultant from from historical environmental conditions and subsequent natural selection (Thompson et al., 2001).

Ecological memories affect how species interact with one another and, consequently, how communities assemble, their stability and their ability to recover from disturbance. However, the mechanisms by which ecological memories in interspecific interactions affects these system-level dynamical effects are not well understood. Specifically, it is not known, in a computational sense, what *type* of memory an ecological memory is. Are ecological memories a simple, passive memories (like an impression in clay), or are these natural communities capable of forming memories with conditional dependencies between units?

Lack of clarity rearding the consequences of ecological memory inhibits our ability to understand the historical determinants of present-day community behaviours. For example, natural communities have the capacity to exist in multiple different, stable configurations, community compositions that are stable over ecologically-relevant timescales (Beisner et al., 2003). Regime shifts, due to ecological (environmental) forcing, may cause communities to transition from one stable state to another. Along with resistance to environmental conditions, multiple different states may be possible under the same conditions (Figure 3.1). However, whilst it is known that communities can have multimodal dynamical landscapes, the relationship between community attractor states and evolutionary history is not well understood. This leaves many open research questions about exactly how evolutionary memory is shaped by selective pressures in past conditions (Thompson, 2005; Thompson et al., 2001). For example:

FIGURE 3.1: Mulga woodland, an example of alternative stable states. Natural communities have the capacity to exist in multiple different, stable configurations, known as alternative stable states (Beisner et al., 2003). Mulga woodlands of can exist in either a grass-rich regime, or a shrub-dominated regime. Regime shifts in this case may be caused by combinations of fire, herbivory, and variable rainfall. Either state can be resilient, dependent upon management. Image from Folke et al. (2002).

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?

Previously, the lack of a theoretical framework linking individual adaptations to collective behaviours meant questions such of these were very difficult to address. This prior lack should not be not surprising, given the focus of much of evolutionary biology on evolutionary units (genes, individuals, populations, etc) rather than the potential of natural selection to modify interactions between units in a consistent manner. The ability of natural selection to consistently modify interactions is important to understand, as appropriately organised networks are capable of producing many important collective effects, including the ability to perform distributed associative learning and form associative memories. Should natural selection in ecological communities modify interspecific interactions (network correlations) in a consistent manner, it may be that these networks are similarly capable of collective learning behaviours.

Chapter 2 demonstrates that the selective pressure on interspecific interactions is proportional to species' cooccurrence. These dynamics have a natural isomorphism with the manner in which interactions in a neural network are adjusted by Hebbian learning. Given 1) the dynamical equivalence between natural communities and simple neural networks (Noonburg, 1989) and 2) equivalent organisational processes, this leads to the suggestion that these ecological networks might exhibit similar dynamical behaviours to neural networks.

When appropriately organised, neural networks form distributed associative memories of experienced or trained states. Organising (training) neural networks involves modifying the strength of the connections (or associations; interactions) between units in the network, in accordance with the training rule applied (such as Hebbian learning, Hebb, 1949). With appropriate training data, modifying interactions causes the system to form attractors for multiple distinct states, thus causing the system to converge to one of these states when activated. The formation of multiple, *specific* attractors enables the use of these networks in several related processes, including: 1) Pattern reconstruction, whereby a complete pattern is recalled from a partial stimulus. 2) Error correction; in which noise is removed from a signal or imperfections are repaired. 3) Pattern recognition; in which ambiguous initial conditions are transformed into the most similar trained state. 4) The ability to hold state in dynamics; to display temporal dynamics independent of input (hysteresis).

In this chapter we investigate whether ecological networks can be organised, just as neural networks can be, so as to form distributed associative memories of specific past states. We use computer simulation of a simple Lotka-Volterra model and find that selective pressure on interspecific interactions in natural communities, isomorphic to learning processes in neural networks, creates similar dynamical behaviours, enabling community networks to store distributes associative memories in the interactions between species. This enables ecological community networks to perform dynamical behaviours equivalent to those listed above for neural networks. These equivalences are set out in Table 3.1, showing the analogy we make between the dynamical consequences for a neural network resulting from these systems' ability to form a distributed associative memory, and the equivalent dynamical behaviours for ecological community networks.

	Unsupervised correlation learning	**Coevolution**
3.1.0	Ability of a neural network to form a distributed associative memory	Collective organisation of community networks resulting in ecological memory formation
3.1.1	*Memory formation* (Figure 1.4, top panel) Hebb's rule organises synaptic connections to reinforce the state of the system, decreasing sensitivity to changes in input.	*Ecological memory formation* (Figure 3.3): 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; Beisner et al., 2003).
3.1.2	*Distributed associative memory facilitates a memory of multiple patterns* (Figure 1.4.a): the capacity to store multiple patterns of activation in the organisation of synaptic connections and recall patterns from any initial conditions via activation dynamics.	*Formation of alternative stable states* (Figure 3.5.a): the creation of a distributed ecological memory in the network of species interactions results in a system with attractors that mimic past ecological states.
3.1.3	*Pattern reconstruction* (Figure 1.4.b): the recall of a complete pattern from a partial stimulus.	*Ecological assembly dynamics* (Figure 3.5.b): reconstruction of a particular community composition, from a subset of that community.
3.1.4	*Error correction* (Figure 1.4.c): the ability to remove noise from a pattern, repair imperfections and restore a complete pattern.	*Ecological resilience* (Figure 3.5.c): the ability to recover from perturbations in species densities and restore the complete community.
3.1.5	*Recognition* or classification of an input or stimulus (Figure 1.4.d): return the nearest attractor from ambiguous initial conditions.	*Ecological sensitivity* to initial conditions (Figure 3.5.d): the switch-like change in response to small variation in initial species densities.
3.1.6	*Holding state in dynamics*: Hopfield networks and other recurrent networks have an internal state that allows them to display temporal dynamics (independent of input).	*Ecosystems hold state in population dynamics* (Figure 3.6): in systems with multiple attractors this results in a communities capable of hysteresis with tipping points between states.

Table 3.1: Ecological community networks exhibit equivalent distributed associative learning processes to those in neural networks. Given certain conditions this enables community networks to form distributed associative memories of specific experienced states, just as neural network can be trained to the same effect.

3.2 Methods

3.2.1 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 B.1). 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, for example, El Nino versus La Nina years, hot-dry and cold-wet climates; or high/low levels of some key broadly-utilised resource such as phosphorous input rates for a lake habitat (Carpenter et al., 1999). Given that individuals from all species experience both environments over evolutionary timescales, these conditions could vary in space (e.g. geographic localities, Paperin et al., 2011), 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 (Figure 3.2). 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 after each complete cycle of ecological dynamics (see below).

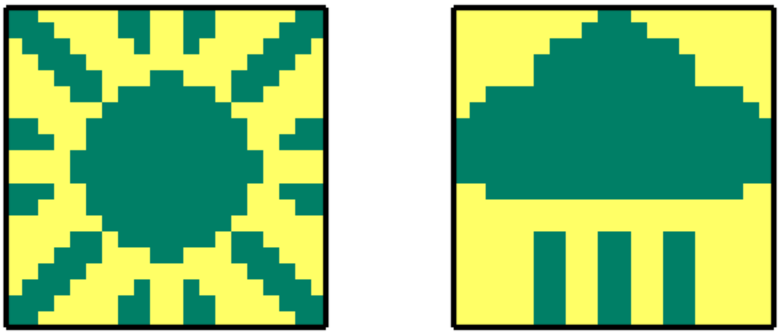

(A) E1 (B) E2

FIGURE 3.2: 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/dry' and 'cold/wet' pictograms where a green pixel at location x, y indicates an elevated carrying capacity, $(k_0 + \alpha)$, and a yellow pixel a depressed carrying capacity, $(k_0 - \alpha)$, for the $(20x + y + 1)^{th}$ species in the ecosystem, Appendix B.1. (The two-dimensions of the pictogram are abstract and do not imply any spatial arrangement of the species).

3.2.2 Ecological dynamics

We utilise the ecological model described in Chapter 2 with a slight adjustment to allow for variable environmental conditions affecting species' carrying capacities, k. We thus use the

Lotka Volterra equation:

$$\frac{dx_i}{dt} = \frac{m_i x_i}{k_{ie}} \left(k_{ie} + \sum_{j=1}^{N} \omega_{ij} x_j \right) \tag{3.1}$$

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 before interspecific competition). N is the number of species in the network. Each cycle of ecological dynamics consists of τ iterations of Equation 3.1.

3.2.3 Rate of adaptation

Similarly, we use the rate of adaptation in interactions developed in Chapter 2, allowing for variation in k:

$$v_{ij} = x_i \mu s = \frac{m_i}{k_{ie}} g \mu x_i x_j \tag{3.2}$$

where μ is the beneficial mutation rate, and g is the size of the conferred change in ω_{ij}.

3.2.4 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 Equation 3.2, 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 $\sum_{j=1}^{N} \omega_{ji}(t) = Q_j$, where $Q_i = Q_j < 0$ is a constant (the sum of interaction terms in row/column i at time $t = 0$) (Sinkhorn-Knopp normalisation; Knight, 2008, See Appendix B.1). 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 (Roughgarden, 1972).

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 (Equation 3.2), 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 ($\omega_{ii} = -1$).

Evolutionary processes (Equation eq: rate of adaptation memory and Sinkhorn-Knopp normalisation) are applied at the end of every T ecological cycles. Model parameters of the simulations and methods used for measuring the properties of ecological attractors are described in Appendix B.1.

3.3 Results

We use the series of four open questions concerning ecological memory listed above to exemplify implications related to memory that follow from the general result outlined in Chapter 2.

3.3.1 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 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 (Figures 3.3.a and 3.3.b). Interspecific 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 cycles, with evolutionary processes applied after each ecological cycle. 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 (Figure 3.3.c). 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 environmental conditions or perturbations 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.

3.3.2 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 Figure 3.4. We see that their evolution is identical to Hebbian learning. As in Chapter 2, this is robust to the choice of underlying model (Appendix B.2).

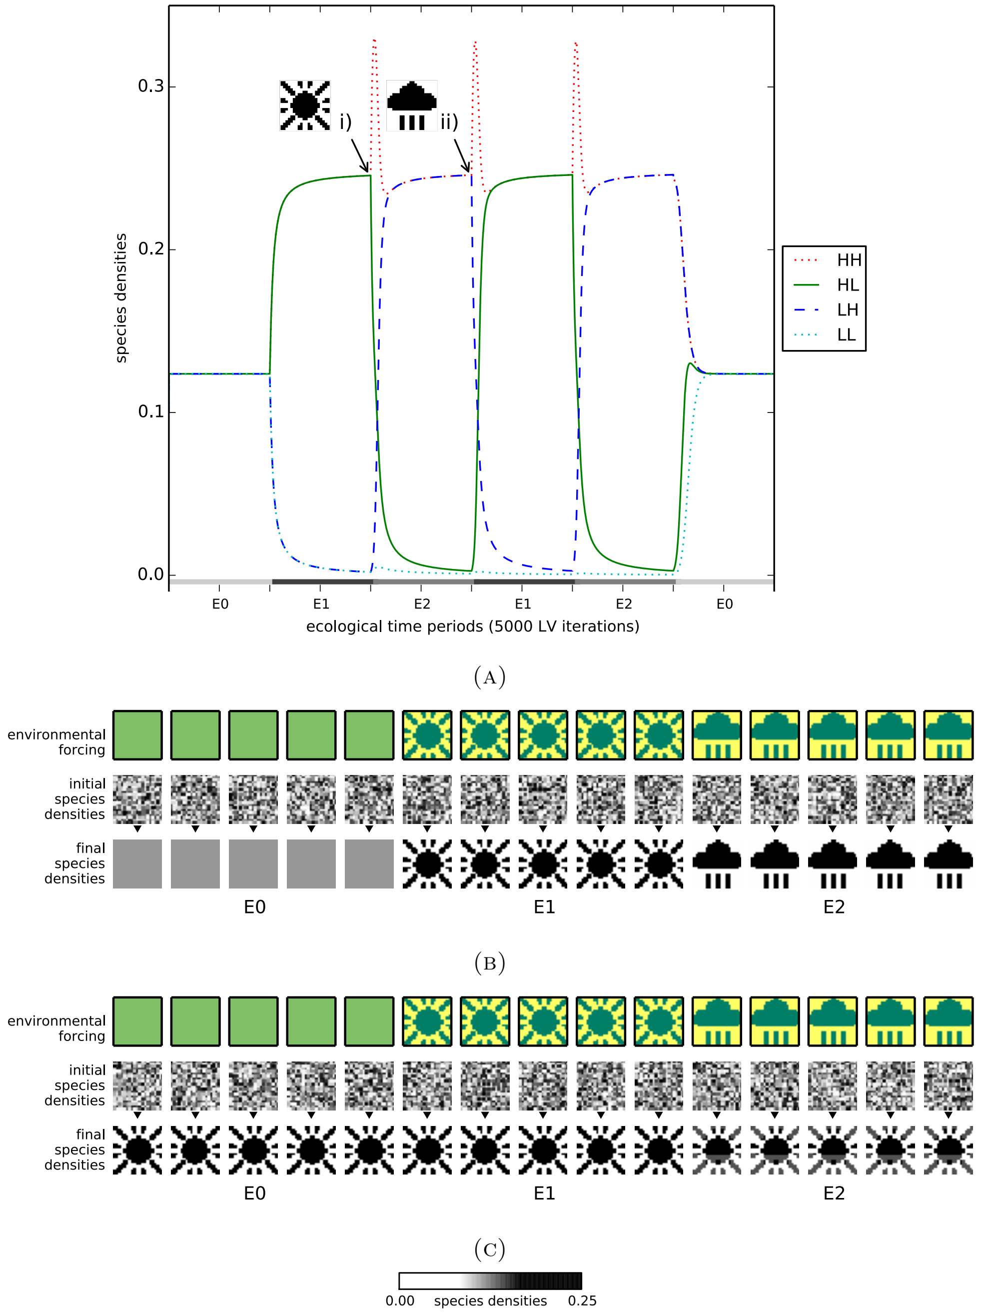

(A)

(B)

(C)

Figure 3.3: Ecological dynamics before and after evolution in E_1 only. a) Ecological dynamics before any evolution, showing the four types of response to forcing in two environments. Lines are labelled according to species carrying capacity in E_1 and E_2. (H='high', L='low'; such that, e.g., HH indicates species that have a high carrying capacity in both environments). Species start at the same initial density and, as there are only two carrying capacities in each environment, there is no variation in species densities of the same label (labels being LL, LH, HL, or HH). b) End states of ecological dynamics before any evolution. Vectors of all species population densities are displayed in a pixel array as per Figure 3.2. 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 for each pattern of forcing). c) End states of ecological dynamics after 500 evolutionary timesteps in environment E_1 only. Layout as in (b). Here, equilibrium states are governed by that past pattern of environmental forcing and not by the current environmental forcing. We observe a community network with an ecological memory for a single 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). The next experiment demonstrates that community networks are capable of forming memories of multiple, distinct states.

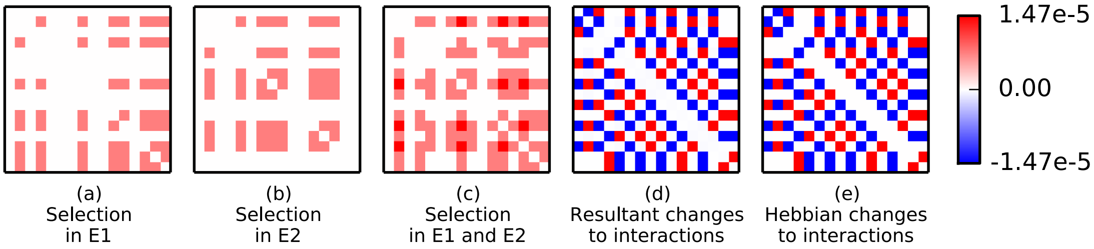

FIGURE 3.4: Evolved interactions are identical to Hebbian interactions. Change in interactions between the first 16 species are shown. a-b) Some of the competitive interaction coefficients are decreased by the direct effects of selection in E_1 and E_2, respectively. c) The combined effect of selection in the two environments is that some interactions are decreased in both environments, some in only one environment and others in neither environment. This depicts the relative rate of change due to direct selection effects. d) When normalising ecological constraints are taken into account, some interactions are decreased, some left unchanged, and others are increased. The resulting changes are identical to (e). e) The result of Hebb's rule applied to the interactions between the first 16 species summed over E_1 and E_2 (r is scaled to give the same mean magnitude as (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 (Figure 3.5.a., Appendix B.1). 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).

3.3.3 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 (Figure 3.5.b). The system has not just formed arbitrary attractors, but specific attractors that recreate the experienced states. 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 (Figure 3.5.c.) (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).

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

Figures 3.6 and 3.7 examine 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

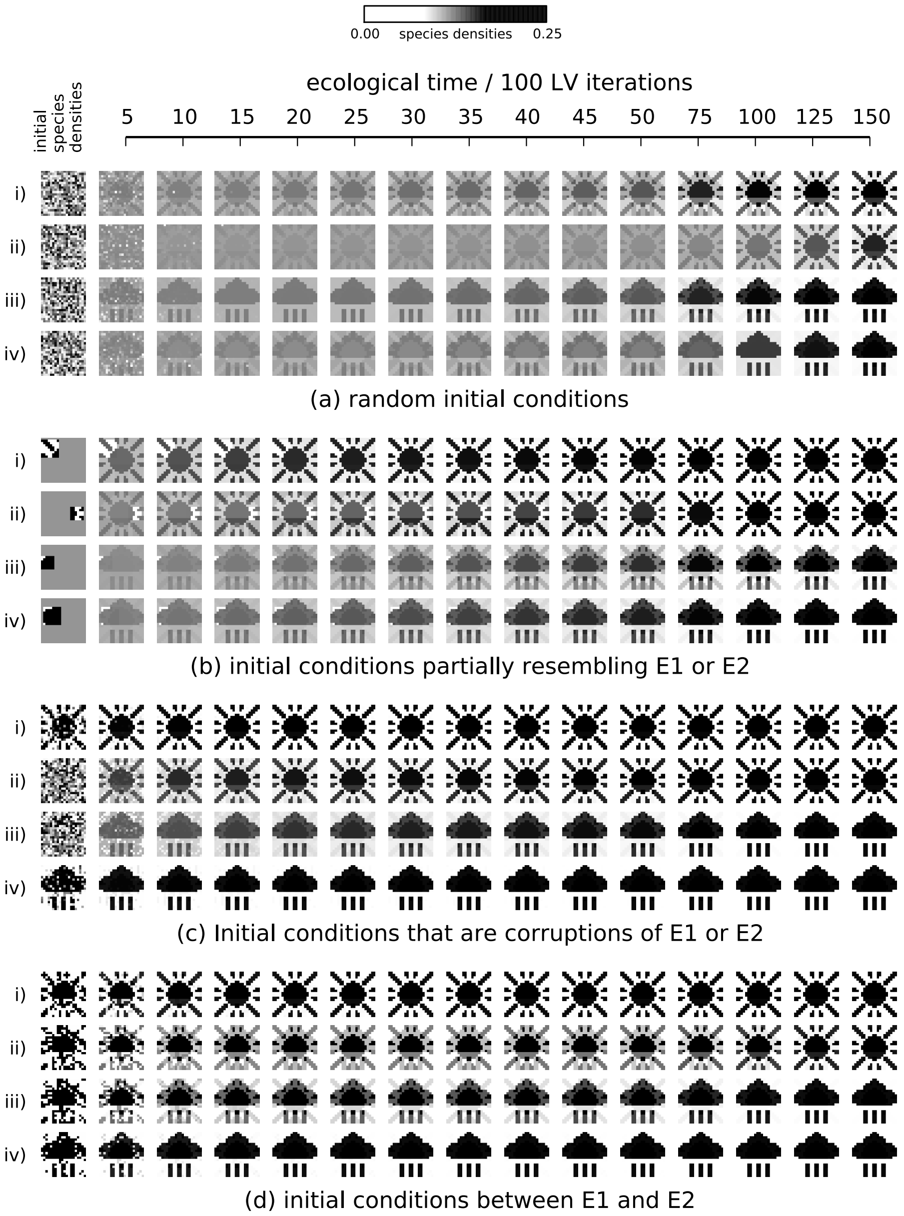

FIGURE 3.5: Ecological dynamics after evolution in varying environmental conditions, E_1 and E_2. Each row shows change in species densities due to ecological dynamics from a single demonstration, starting from random initial species densities (left hand side). All demonstrations show convergence to one of two states, attractors corresponding to either E_1 or E_2 (right hand side). a) Assembly rules cause random initial species densities to develop into one of the two 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.

ecosystem is proportional to the environmental forcing applied (Figure 3.6.a). In contrast, after the evolution of interactions (Experiment 2), the response of the system is discrete or switch-like and exhibits substantial hysteresis (Figure 3.6.b). 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_2 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 (Figure 3.7). 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; Scheffer et al., 2009; Dakos et al., 2008). Figure 3.7 also 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 (Scheffer et al., 2001). 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 B.3).

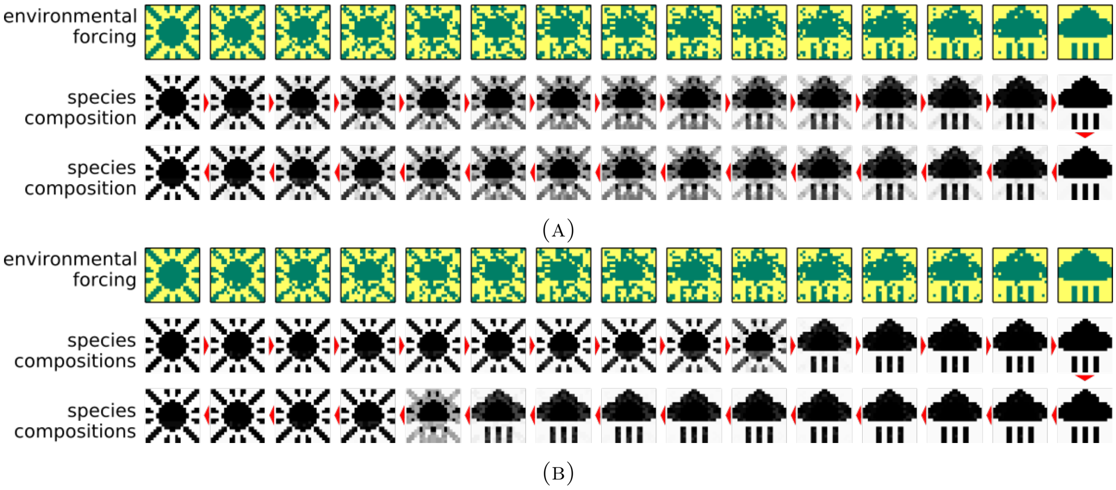

Figure 3.6: Response to environmental forcing before and after evolution in varying environmental conditions. In each test of community dynamics, we first set environmental forcing to E_1 and allow the community to assemble to a stable attractor. We then incrementally change the environmental forcing towards the E_2 pattern, allowing ecological dynamics to stabilise after each adjustment (middle row, left-to-right). Once we have reached patter E_2, we reverse the process, from E_2 to E_1 (bottom row, right-to-left). a) The system's response to forcing prior to any evolution of interactions. Changes in species densities are proportional to forcing. b) The system's response to forcing at the end of simulation; after 545 iterations of ecological and evolutionary dynamics in environments E_1 and E_2, as in Experiment 2. Species densities show an abrupt switch between attractors with hysteresis. This change in dynamical behaviour is also shown in Figure 3.7, showing the evolution of the two attractor state.

These observations demonstrate a conversion of one type of ecological alternate stable state into another. Beisner et al. (2003) 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 (discussed in Chapter 1). Figures 3.6 and 3.7 show a system that converts alternate 'ecosystem states' into alternate 'community states'; thus converting patterns from past environmental states into ecological memories.

Figure 3.8 shows a 'bestiary' of ecological attractors changing over evolutionary time, showing some diversity before settling down to the two-attractors shown in Figures 3.5, 3.6 and 3.7. 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 B.4. This effect can also be observed in Figures 3.8 and 3.7).

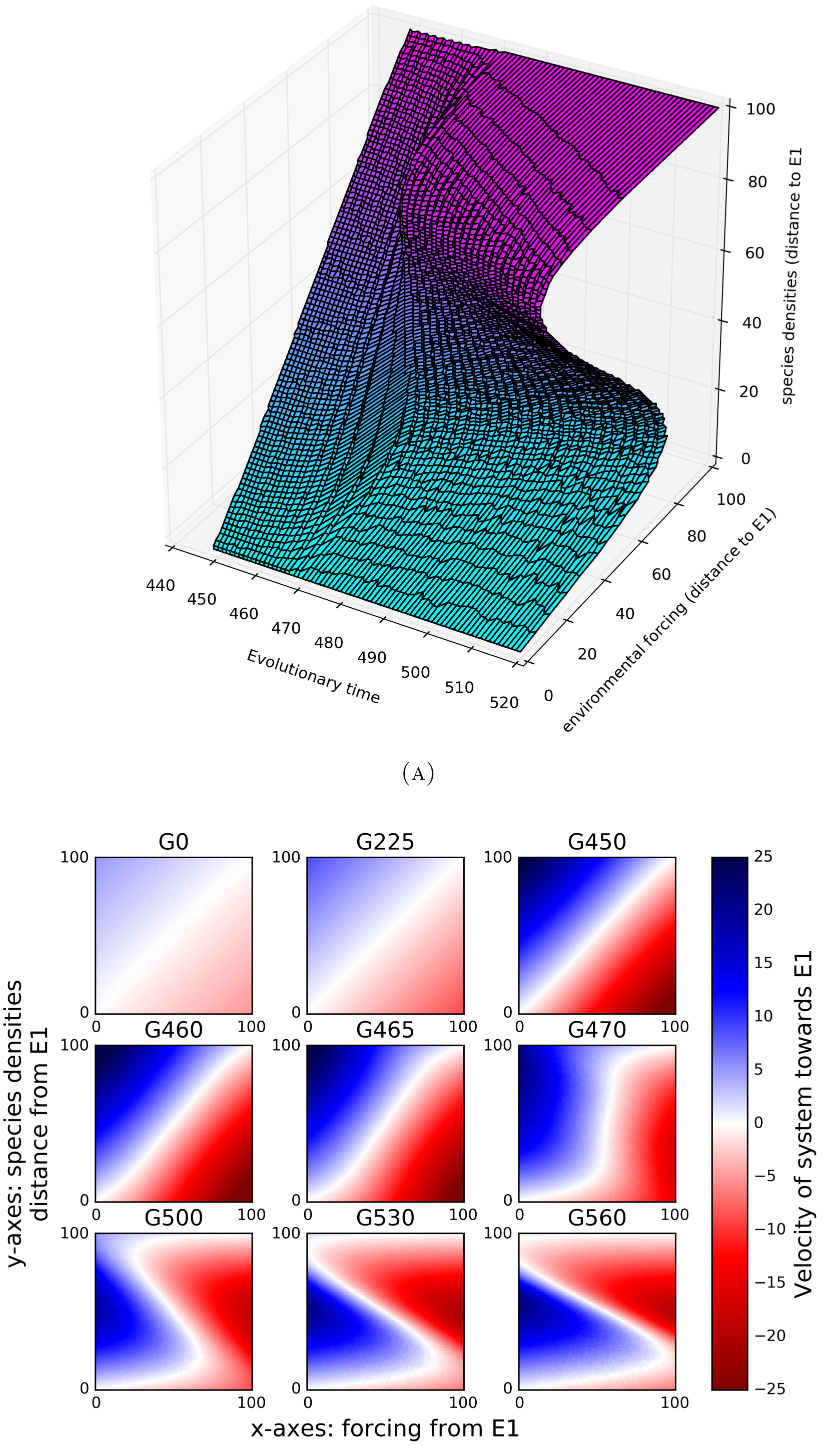

(A)

(B)

FIGURE 3.7: Evolution of the two attractor state and vector fields for population dynamics. a) Surface shows the fixed points for population dynamics. b) Panels show velocity of community composition in relation to the two environmental states, E_1 and E_2 for snapshots during simulation. Blue indicates trajectory towards E_1 (downwards), red indicates towards E_2 (upwards). Fixed points are revealed at the white boundary between these areas. Initially, change in species densities is proportional to environmental forcing. By around generation 470 non-linear but non-catastrophic transitions are observable. By generation 500 there are two stable attractors with a catastrophic transition and hysteresis. Points near the critical transition have slower population dynamics (lower velocity) than points far from critical transition. Eventually this state also breaks down, and the system remains trapped at one of the two attractors (generation 560).

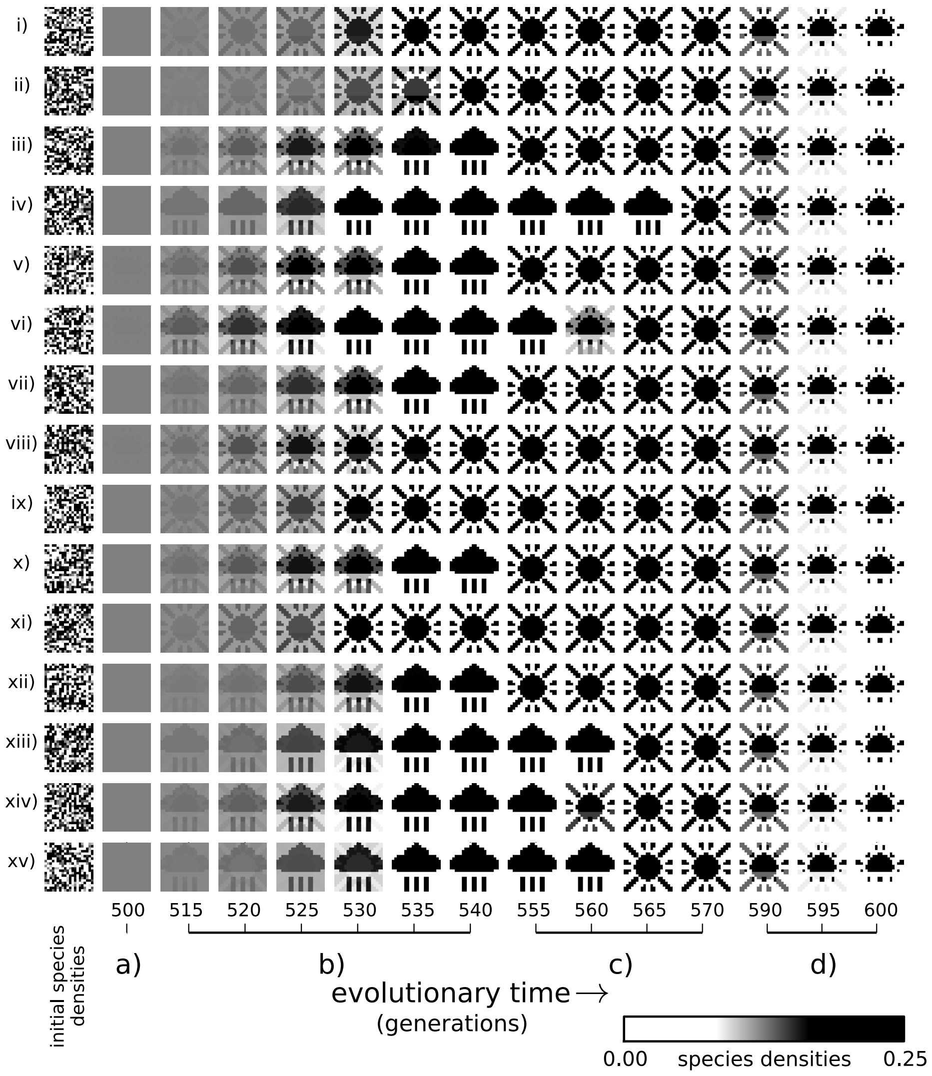

FIGURE 3.8: 'Bestiary' of ecological attractors changing over evolutionary time. From an array of different random initial species densities (left-most column), the ecological states reached in the population dynamics changes over evolutionary time. a) Initially, and up until generation 500, all initial conditions lead to the same ecological attractor (with species at similar densities). 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 (Poderoso and Fontanari, 2007).

3.4 Discussion

3.4.1 Key findings

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 B.5). 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 (Finn, 1976) rather than just becoming *less competitive* with them (see 3.4.2 below).

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 (Angeler et al., 2010).

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, as in invasional meltdown, a small number of founders (Figure 3.5.b) or, similarly, the system can recover each specific state from corruptions of that state (Figure 3.5.c). 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 (Figure 3.5.d). 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) (Figure 3.3) and adaptive capacity (robustness to environmental forcing) (Figures 3.3 and 3.6) (Holling, 1973; Gallopín, 2006; Folke, 2006; Cropp and Gabric, 2002; Angeler et al., 2010). However, if the evolutionary history of a community has included a multi-modal distribution of environmental conditions, then this can result in alternative stable states (rather than universal stability) and the community may exhibit critical transitions in changing from one state to another (Holling, 1973). This switch-like change in the community (Figure 3.6) is only exhibited when the forcing that is applied is similar to past forcing. When arbitrary forcing is applied the response may remain linear (Figure B.2). 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 (Higgins et al., 2002; Scheffer et al., 2001). 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.

3.4.2 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. Algorithmically, this is achieved via Sinkhorn-Knopp normalisation (Knight, 2008, Appendix B.1) after adjustment to interaction rates due to selective pressure (Equation 3.2). The effect of this is that row sums, and, through preservation of symmetry, column sums, are kept constant throughout simulation.

These normalising 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 increases the niche overlap with others. However, the extent to which species evolve dependencies rather than just (context-sensitive) individual advantages in natural ecosystems is an empirical matter. 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.

Our experimental work here has investigated only competitive interactions and has not investigated mutualistic interactions or asymmetric interactions such as characteristic of trophic, e.g. predator-prey, relationships. However, the observation that selected changes to interactions are Hebbian (Chapter 2) does not depend on them being symmetric (or competitive). That is, Equation 3.2 is not sensitive to any assumptions about the initial values of interaction coefficients, e.g., whether ω_{ij} and ω_{ji} are equal or even have the same sign, and therefore applies to predator-prey relationships as well as symmetric competitive interactions. Furthermore, as Equation 3.2 is insensitive to the current state or sign of ω_{ij} competitive interactions can, in theory, eventually become mutualistic (we have not not modelled this here).

Equation 3.2 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 chapter 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, as above, 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 (Watson et al., 2014) 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, 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. That is, an individual cannot evolve in a manner that avoids competition with one species, x, without also making their growth more dependent on the presence of some other species, y. 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 will therefore be 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) (Roughgarden, 1979), most evolution occurs whilst ecological dynamics are at or near equilibrium, as modelled here. Our finding here, that natural selection can incorporate correlational information into interspecific interactions and form an associative memory of specific states, does not depend upon this condition. The Hebbian behaviours we observe will act upon correlational information between species whether or not the system is at an attractor. As long as, when evolutionary processes are applied, there is some consistent correlational information in community composition, we would expect to see the dynamics similar to those in our model; the formation of associative memories related to experienced states.

That said, the behaviour of evo-eco dynamics when these processes have more similar timescales (Turcotte et al., 2011) 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 Levin, 2011). Moreover, any model assuming a single ecological attractor will overlook the interesting behaviours modelled here, regardless of whether the timescales are separated or similar.

Finally, our work here demonstrates that ecosystems are capable of exhibiting collective adaptive behaviours without being Darwinian units of selection. This result, along with recognition of the isomorphism between these systems and appropriately trained neural networks prompts inquiry as to whether natural communities are capable of more complex computational tasks. For example, natural ecosystems are under very many constraints that limit species abundance (e.g. NCP availability), but it is not known whether selection on individuals improves a system's ability to resolve these constraints. Hopfield networks are known to be able to solve complex constraint satisfaction problems (Hopfield et al., 1986). Are there conditions in which ecological community networks exhibit similar constraint satisfaction behaviours? We address these issues in the following two chapters.

Chapter 4

Evolved constraint satisfaction and implications for community functions

4.1 Introduction

As natural selection operates within individual species and not on community structure as a whole, it is not clear if, or how much, ecological community networks are holistically organised by evolutionary processes. It is not clear, for example, whether these systems are organised by natural selection so as to better resolve internal constraints, or whether such behaviour has any consequences for community function.

In Chapter 2 we identify that selective pressure on interspecific interactions is a product of species cooccurrence, and that this is isomorphic to particular models of unsupervised associative learning. In Chapter 3 we build from this isomorphism, and examine whether ecological community networks can be 'trained' by distributed associative learning mechanisms, driven by natural selection, to encode memories of community compositions relating to specific environmental states (very much like how a neural network can be trained to recall specific compositions). We find that community networks can achieve this in same manner as neural networks; by encoding a distributed associative memory in the network of interspecific interactions.

However, as well as the ability to be trained to recall specific attractor states, connectionist models of simple interacting components are known to exhibit other interesting system-level functions, without system-level direction. These include the ability of neural networks to form dynamical attractors for states that minimise internal inter-unit constraints. This behaviour is the result of *self-modelling*, whereby a dynamical system augments its behaviour with correlational information from its own attractors (Watson et al., 2011b), and results from the same distributed learning process we have identified as occurring in community networks (i.e. Hebbian changes network interactions). The presence of Hebbian dynamics in ecological networks thus leads to the intriguing possibility that conditions may exist in which natural selection on interspecific interactions can produce similar self-modelling behaviour in ecological community networks, driven by natural selection lower levels of organisation, and leading to community-level constraint satisfaction.

Here we build upon the work in previous chapters and explore additional implications arising from the isomorphism identified between evolving community networks and learning in neural networks. We address the following research questions related to constraint satisfaction behaviour in ecological community networks:

1. How does natural selection alter the dynamical attractors of community assembly rules in the absence of external forcing?
2. Are there mechanisms by which the evolution of interspecific interactions causes communities to better resolve interspecific constraints?
3. What conditions are required for the evolution of constraint satisfaction behaviour?
4. What are the consequences of collective constraint satisfaction for community behaviour and function?

4.1.1 Constraint satisfaction and self-modelling behaviour

In connectionist models, constraint satisfaction problems are puzzles in which the objective is to learn how to find a configuration that minimises inter-unit constraints. For example, the travelling salesman problem and the eight queens problem can each be cast as constraint satisfaction problems (CSPs) in which the objective is to learn how to find state-compositions that minimise inter-unit constraints. The emphasis here is on ‘learning to solve’ rather than merely ‘solving’ a problem (the usual objective of human game-players). Neural networks used to solve CSPs must develop an attractor for (i.e. learn how to find) the state composition with lowest competition between units (the solution), such that the network is able to reliably assemble that state.

In ecological models, species pairs with high levels of interspecific competition can be regarded as the constraints to be ‘satisfied’. Mature community compositions with low levels of interspecific competition, can be said to satisfy or solve the constraints of the network better than community compositions with high numbers of competing species. Should ecological community networks develop attractors for specific states that minimise interspecific competition, this is analogous to the ‘learning to solve’ constraint satisfaction behaviour in connectionist models of learning.

Simple distributed mechanisms can incrementally alter a dynamical system such that it can find lower energy configurations more reliably and more quickly in a process known as self-modelling (Watson et al., 2011a). Specifically, when Hebbian learning is applied to the connections of a simple dynamical system undergoing repeated relaxation, the system will develop an associative memory that amplifies a subset of its own attractor states. This modifies the dynamics of the system such that it enhances its ability to find configurations that minimise total system energy, and globally resolve conflicts between interdependent variables.

4.1.2 Conditions for constraint satisfaction behaviour

Intuitively, if evolution by natural selection can alter the fitness interactions of a community in an unconstrained manner, then it would eventually result in reduction or removal of all competition with other species, since this allows the highest growth rate for any species. However, the previous chapter shows that, assuming ecological tradeoffs prevent such a non-competitive utopia, the effect of natural selection acting on ecological interactions is to change them in a Hebbian manner.

In neural networks trained with Hebb's rule, the ability of self-modelling dynamics to occur and improve network constraint satisfaction behaviour is dependent upon two conditions: 1) learning takes place whilst the system is in a variety of states, thus preventing over-learning of one state and allowing the system to observe correlations between features; and 2) the correlation between units when learning takes place tends to be consistent with correlations between units in the low-competition solution state(s). Although not a requirement for the underlying dynamical behaviour, a third condition is typically applied when demonstrating constraint satisfaction behaviour: that the constraint problem being tested should be non-trivial to solve due to very many locally optimal configurations that fail to satisfy constraints as effectively as possible. (A usual demonstration of constraint satisfaction is for a system with multiple attractors to change its connections through learning and thus form a single attractor for the configuration that minimises inter-unit competition).

We can translate the above conditions for neural networks to equivalent ecological requirements and make predictions on conditions that are necessary for constraint satisfaction behaviour in ecological networks. Thus, we predict community assembly rules will evolve so as to increase the frequency at which they assemble low competition states (i.e. good quality solutions) if: 1) there is regular disturbance, such that evolution occurs whilst the system is in a variety of community compositions, and 2) the mean correlations between species when selective pressures act are consistent with correlations between species in low competition configurations. If these two conditions are met in communities where there are multiple 'sub-optimal' stable states, then the mechanisms of distributed associative learning (driven by Hebbian evolutionary processes) have potential to enlarge the basin of attraction for low-competition states, increasing the frequency at which the community will assemble these low-competition states.

4.2 Methods

We use the framework developed in Chapter 2 and expanded in Chapter 3 to model population dynamics; evolved changes in interactions; and consequent evolutionary tradeoffs. We model these dynamics in an ecological scenario which present the community with a complex constraint satisfaction problem, and examine the consequences for eco-evo feedbacks.

4.2.1 Ecological dynamics

We use the ecological model described in Chapter 2 with the Lotka Volterra equation:

$$\frac{dx_i}{dt} = \frac{m_i x_i}{k_i}\left(k_i + \sum_{j=1}^{N} \omega_{ij} x_j\right) \tag{4.1}$$

where x_i, m_i and k_i are the density, intrinsic net growth rate of species, and carrying capacity of species i respectively. N is the number of species in the network.

4.2.2 Evolutionary dynamics

We use the rate of adaptation in interactions developed in Chapter 2:

$$v_{ij} = x_i \mu s = \frac{m_i}{k_i} g \mu x_i x_j \tag{4.2}$$

where μ is the beneficial mutation rate, and g is the size of the conferred change in ω_{ij}.

4.2.3 Ecological tradeoffs and normalising constraints

We repeat the framework for normalising ecological constraints developed in Chapter 3 (Sinkhorn-Knopp normalisation; Knight, 2008), such that for all species i and j $(j \neq i)$, $\sum_{j=1}^{N} \omega_{ij}(t) = Q_i$ and $\sum_{j=1}^{N} \omega_{ji}(t) = Q_j$, where $Q_i = Q_j < 0$ is a constant (the sum of interaction terms in row/column i at time $t = 0$).

4.2.4 Evolvable and unevolvable characteristics

Species traits contribute to interspecific competition in ecological communities in a multitude of different ways (Kunstler et al., 2016). These effects may be indirect, as in traits influencing environmental filtering, that determine which species appear in the same environment (e.g. traits governing tolerance of local conditions), or direct, such as a traits governing dependency on a shared resource, or preference for a particular habitat (although see more complex scenarios, e.g. interaction norms; Thompson, 1999). We are interested in how competition between cooccurring species shapes community dynamics, and model on traits that affect competition directly.

We assume that there are some aspects of interspecific competition that can be modified by heritable variation in the characteristics of individuals, and some that are not. Components of fitness interactions that do not evolve (or evolve slowly) result from traits in each species that are resistant to selective pressure (such as a shared bauplan, habitat preference or fundamental resource dependency). In computational terms, this network of unevolving internal constraints represents 'the problem' faced by the community (specifically, a constraint satisfaction problem; Rumelhart and McClelland, 1987). Other characteristics are more responsive to selective pressure, such as traits that alter growth rate (Zuppinger-Dingley et al., 2014) or common resource dependencies (Hutchinson, 1965), and these do evolve during simulation. Changes in these evolvable fitness-interactions affect the community assembly dynamics and, ultimately, the community composition (or stable state) that the system converges to. If the stable state reached minimises competition due to unevolving constraints in the community, this is, in computational terms, a 'solution' to the constraint satisfaction problem.

We model this framework by separating community interactions into two matrices, a matrix of unevolving constraints 'U', and a matrix of evolving constraints 'E', such that: $\Omega = U{+}E$ (and $\omega_{ij} = u_{ij}{+}e_{ij}$). This separation enables us to examine the evolutionary response of the system to differing topologies of unevolving constraints, and test whether the system satisfies unevolving constraints whilst controlling for changes in species densities due to evolved changes.

4.2.5 Unevolving characteristics

We are interested in the effects of evolution on community dynamics and, specifically, whether evolutionary changes in the network of interactions alters the frequency at which a community converges to attractors that have low interspecific competition. To examine these effects we construct a community constraint network, U, with multiple attractors, resulting from competitive exclusion between certain species due to similar traits or niche profile (Diamond, 1975).

We model a community with complex interdependencies, where F independent traits $\{T_1, T_2, ..., T_F\}$ determine competition between species, and total competition between each pair of species is the sum of their competition across all traits, such that entries in the matrix U is given by:

$$u_{ij} = \sum_{f=1}^{F} c_{f(ij)} \tag{4.3}$$

where C_F is the competition matrix for that trait. Traits values indicate whether species compete due to temporal or habitat overlap, or nutrient dependency (as in Figure 4.1 where species of the same bauplan or resource dependency compete). For each trait T_f we can define values in a competition matrix C_F as:

$$c_{f(ij)} = \begin{cases} \gamma, & \text{if } T_f(x_i) = T_f(x_j) \\ 0, & \text{otherwise} \end{cases} \tag{4.4}$$

4.2.5.1 Ecological scenario

We consider a model of a community with complex interdependencies, where related species have multiple conserved traits T for similar bauplan (T_1) and species from different phyla have convergent traits (T_2), such as for a common resource dependency. If we allow P phyla and R resource dependencies, where $P = R = N$, we can choose to classify species traits for T_1 and T_2 as functions of species index:

$$T_1(x_i) = \left\lfloor \frac{i}{N} \right\rfloor \tag{4.5}$$

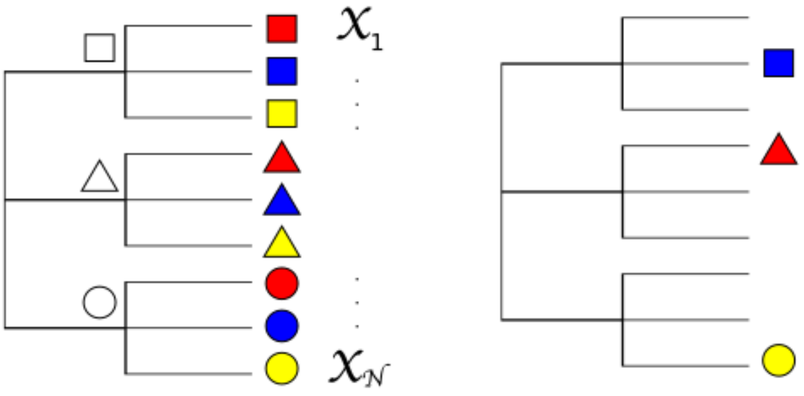

(A) Regional species pool (B) Realised community

FIGURE 4.1: Regional species pool and a realised community showing phylogenetically random structure arising from both conserved and convergent traits. a) A species pool of 9 species $\{x_i, ..., x_9\}$, where species with a common phylogeny share a similar bauplan (represented by shape) and species from different phylogenies compete for the same resource (represented by colour). b) If competition between species with a similar bauplan prevents coexistence, and competition between species utilising the same resource prevents coexistence, realised communities can have a maximum of three species, in which all pairs of species i and j must have different bauplans: $\lfloor i/N \rfloor \neq \lfloor j/N \rfloor$, and different resource dependency: $\mod (i/N) \neq \mathrm{mod}(j/N)$.

$$T_2(x_i) = i \bmod N \tag{4.6}$$

Where for each trait two species are competitive if they share the same value for that trait.

If competition (γ) is sufficiently high that it results in competitive exclusion, we obtain dynamics whereby stable configurations can have at most $\sqrt{N}$ species. The constraint satisfaction problem posed by T_1 and T_2 so far is the same scenario shown in Figure 4.1 (with $N = 9$ species). In this simple scenario, it is very straightforward for ecological dynamics to find compositions that maximally satisfy the constraints in U (all combinations of $\sqrt{N} - 1$ compatible species allow for the addition of an additional compatible species). As we wish explore evolutionary dynamics in more complex scenarios where many of the potential attractors do not maximally satisfy community constraints, we expand the above scenario so that the network topology of interactions matches a constraint problem of known difficulty (Appendix C.1.1), where coexistence of large numbers of species is possible, and there exist multiple stable but not globally efficient attractors consisting of subsets of species of size $< \sqrt{N}$, where additional species cannot be added without violating constraints. To achieve this we add additional constraints T_3 and T_4, orthogonal to T_1 and T_2, representing additional dimensions affecting compatibility of species:

$$T_3(x_i) = T_1(x_i) + T_2(x_i) \tag{4.7}$$

$$T_4(x_i) = T_1(x_i) - T_2(x_i) \tag{4.8}$$

4.2.6 Simulation

Without limitations on Ω, Equation 4.1 may exhibit unstable or even chaotic behaviour. In the following experiments 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 (May, 1975), in which case the dynamics have only fixed point equilibria (Hughes and Roughgarden, 1998).

We are interested in how disturbance (and subsequent ecological dynamics) impacts evolutionary dynamics and conduct two experiments: In our 'primary' experiments we consider conditions where disturbance has the potential to cause a system to visit multiple attractors, but where frequency of disturbance is sufficiently infrequent such that most evolutionary change happens after ecological dynamics have stabilised at, or near, a stable attractor. In our 'secondary' control experiments we randomise community composition at the beginning of simulation, but do not disturb species densities during simulation. This has the effect that the system remains at the same stable attractor throughout the experiment (Appendix C.1.2).

Parameters are described in Appendix C.1.3.

4.3 Results

We use our experiments to investigate conditions necessary for ecological networks to evolve collective constraint satisfaction behaviour; the implications of this for community function, and the relationship between the observed evolutionary dynamics and Hebbian processes in neural networks.

4.3.1 Prior to evolution, community constraint satisfaction is correlated with increased community biomass

We run the ecological scenario with $N = 64$ species and where competition (γ) is sufficiently high that it results in competitive exclusion, with the effect that the maximum number of species that can coexist without one or more species pairs violating an innate constraint in the matrix U between them is $\sqrt{N} = 8$. Prior to evolution, only 37% of random community compositions converge to a stable state with the maximum, eight, peacefully coexisting species; in all other configurations one or more species pairs violate a constraint in U. Competition avoidance in communities is correlated with higher biomass and those stable states that minimise competition have the maximum observed biomass (Figure 4.2).

4.3.2 In conditions without disturbance, evolutionary processes do not improve community constraint satisfaction

During simulation we assume that ecological dynamics are fast compared with evolutionary changes in interaction rates and hence allow sufficient ecological time for species densities to

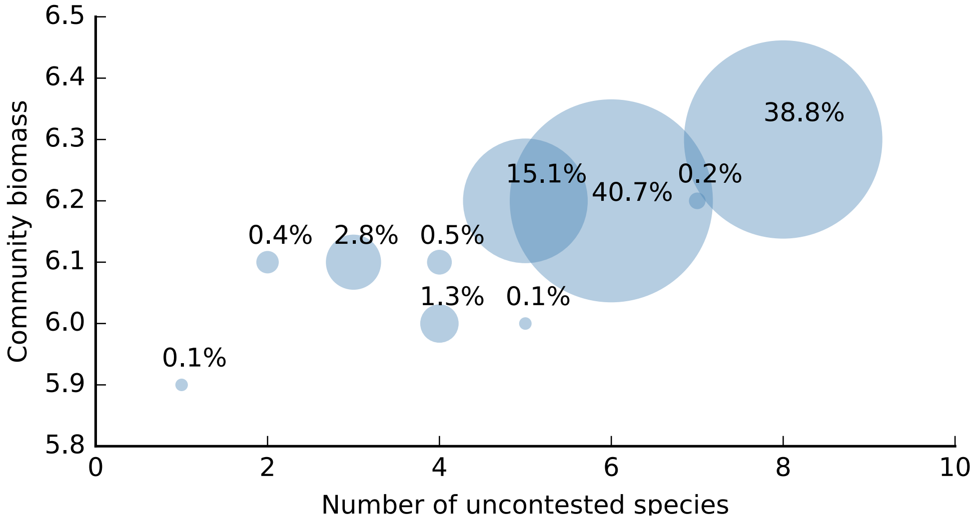

Figure 4.2: Correlation between number of uncontested species and biomass in stable states reached after 10^4 iterations of LV dynamics at t(0), before any evolution. Parameters shown in Table C.1. 'Uncontested' species are those that are compatible with all other species present according to the unevolving constraints in the matrix U. From 1000 random initials tested, 10 clusters are identified, where members of a cluster are state compositions with the same number of uncontested species and the same biomass at 1 d.p. The percentage of random community compositions that reach a stable state in each cluster (x = number of uncontested species, y = biomass) are: 0.1% at (1, 5.9), 0.4% at (2, 6.1), 2.8% at (3, 6.1), 0.5% at (4, 6.1), 1.3% at (4, 6.0), 0.1% at (5, 6.0), 15.1% at (5, 6.2), 40.7% at (6, 6.2), 0.2% at (7, 6.2) and 38.8% at (8, 6.3).

stabilise near an attractor before applying evolutionary processes. Without disturbance, community composition is relatively unchanged throughout simulation (Appendix Figure C.1). The rate of adaptation of interactions (Equation 4.2) is a function of species co-occurrence with the effect that interactions between those species present become less competitive over time. These changes to community assembly rules result in an increase in the basin of attraction of the experienced stable state (Figure 4.4b).

After simulation is complete we retroactively test how community assembly rules have changed over time. For each time step, t, we examine the community compositions assembled from random initial conditions. With no disturbance, we observe only slight changes in the frequency at which community dynamics assemble 'peaceful' community compositions - whereby all species present avoid competition with one another (Figure 4.3c). Similarly, there is little effect on biomass when controlling for evolved changes (see below; Figure 4.3d).

4.3.3 Disturbance results in evolutionary dynamics that facilitate community constraint satisfaction

If disturbance is frequent enough, but ecological dynamics still have time to stabilise between disturbances, the system visits (and evolutionary processes occur at) very many different stable states. (During the first 100 cycles of the simulation, communities reach a mean of 84.5 (std 6.5) different stable states.) Again, selective pressures, acting as a function of species cooccurrence (Equation 4.2) are strongest on interactions between species cooccurring in high density, but

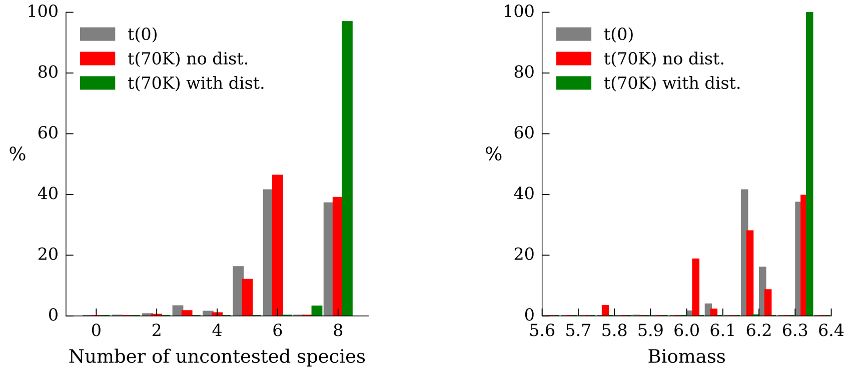

(A) Community competition (B) Community biomass

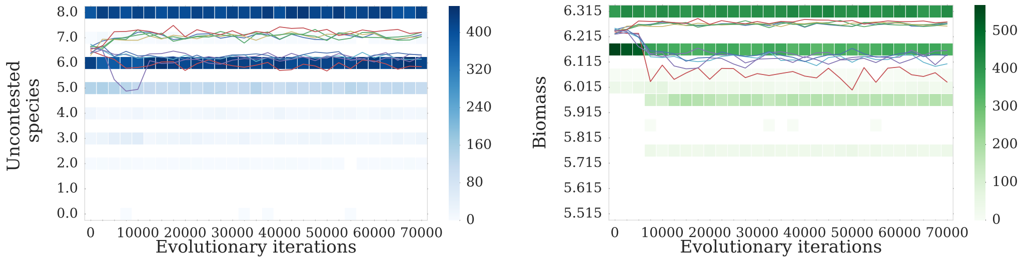

(C) Uncontested species, no disturbance (D) Biomass, no disturbance

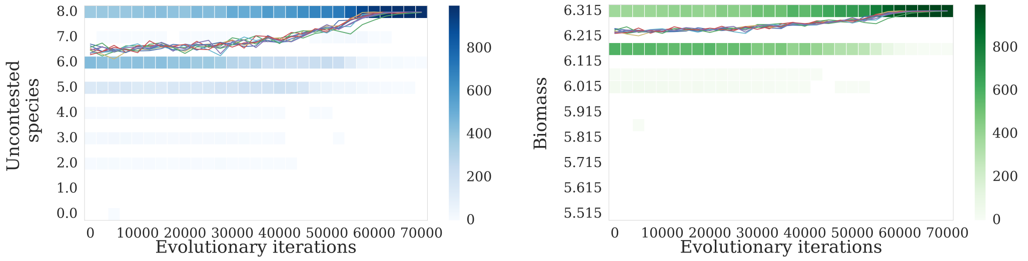

(E) Uncontested species, with disturbance (F) Biomass, with disturbance

FIGURE 4.3: Effects of evolving assembly rules on stable state composition (a, c, e) and biomass (b, d, f). a) Prior to evolution (t(0), grey bars), only 37% of 1000 random community compositions tested converge to alternate stable states in which eight species peacefully coexist (eight is the maximum number possible). At the end of simulation t(70,000), in simulations without disturbance (red bars), this behaviour is relatively unchanged by evolutionary processes but, in simulations with disturbance (green bars), almost all random community compositions reach stable states with eight peacefully coexisting species. Plots c) and e) show this change in behaviour for 10 simulations without, and with disturbance respectively. Lines show, for each single simulation, the mean number of uncontested species in stable states reached from 100 random community compositions. Heat maps show the distribution of the same metric across all simulations. b) The biomass of stable states reached from random community compositions re-stabilised in the original interaction matrix $\Omega(0)$. In simulations without disturbance there is a small increase in the frequency at which assembly rules form stable states with the highest biomass, whereas in simulations with disturbance all random community compositions converge to stable states with the maximum possible biomass. As above, plots d) and f) show how this increase in stable state biomass develops over simulation. The effect of running the simulation with disturbance is clear; by the end of simulation the system converges to compositions with lower competition and higher biomass. For the sake of clarity legends for the lower four panels are included in the appendix (Figure C.2).

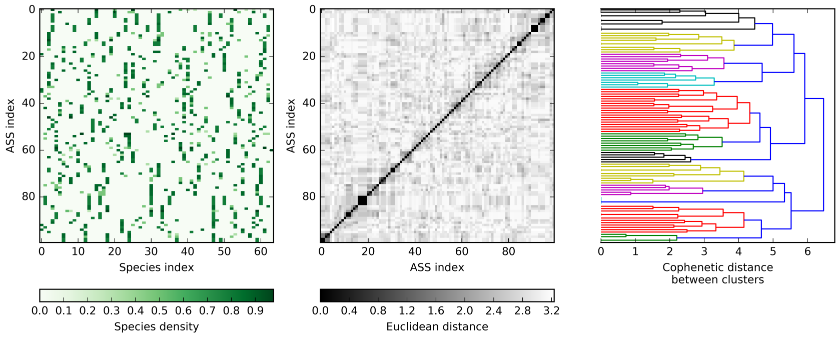

(A) Distribution of 100 SS at t(0)

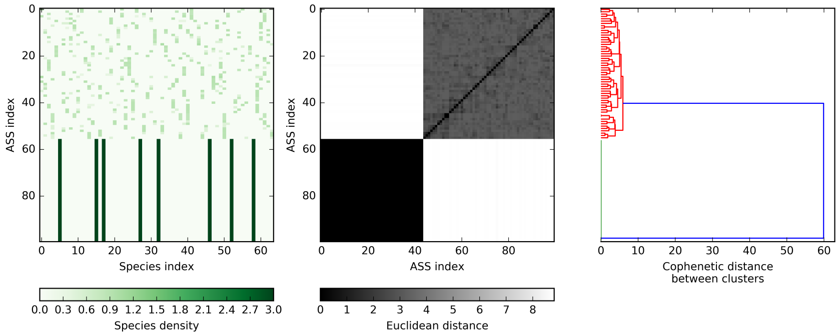

(B) Distribution of 100 SS at t(70,000), simulation without disturbance

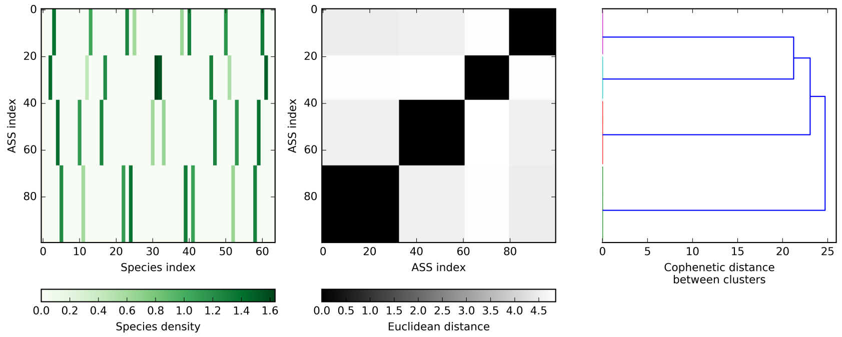

(C) Distribution of 100 SS at t(70,000), simulation with disturbance

FIGURE 4.4: Distribution of stable states (SS) at beginning and end of simulation for selected sample simulations. Left panels show species densities in 100 stable states reached from random community compositions, grouped by Euclidean distance. Centre panels shows Euclidean distance between these stable states. This data used to cluster stable states, with Cophenetic distance between clusters shown in right hand panels. a) Prior to simulation assembly rules converge to very many, different stable states. b) With no disturbance the system develops an associative memory for the configuration it evolves at. c) With disturbance the system develops associative memories for four distinct stable states, corresponding to states that minimise competition.

now the effects of disturbance cause the system to visit many different stable states, and many different combinations of cooccurring species. Thus the net change in interactions is a function of species cooccurrence across many stable states (rather than just one). If the evolutionary changes are slow enough, and the system visits sufficiently many different stable states, this pattern of changes causes the dynamical landscape to alter such that the community only has attractors for stable states in which no species pairs are in high competition and, in all of 10 repeats, by $t(70,000)$, all initial conditions converge to community compositions with eight peacefully coexisting species (Figure 4.3e; Figure 4.4c).

This is an important result, demonstrating that, given sufficient disturbance, low-level evolutionary processes can result in ecological communities that consistently converge to community compositions that satisfy all innate constraints in the network U.

4.3.4 Community constraint satisfaction results in increased community biomass, stability and resilience

It is not surprising that if species are able to evolve their interspecific interactions with trophic competitors, that there is an increase in the total biomass of the system (Appendix Figure C.3). What we are interested in however, is not just the effect that evolution has on each individual species, but whether the evolution of interactions has any effects on assembly rules that alter the community's fundamental composition (which species are present) and the consequences of this for biomass. To do this, for each time step we re-examine the stable state reached by random community compositions, and then additionally re-stabilise the resultant community compositions in ecosystems where values in E are set to their initial value. By removing the effects that evolving interactions have on biomass we observe that the system is not only finding configurations of species that maximally resolve intrinsic constraints, but that these compositions also have highest community biomass in the original interaction matrix (Figure 4.3f).

These coevolutionary mature communities exhibit greater ecological stability than stable states at the coevolutionary naive communities at the start of simulation, with faster recovery from disturbance, and increased ability to recover composition following disturbance (Figure 4.5).

Coevolutionarily mature communities also have lower diversity (Figure C.4), as a small number of species tend to dominate.

4.3.5 Changes in interspecific interactions correlate with Hebbian patterns of change

As in previous chapters, we observe isomorphic behaviours between evolutionary dynamics and Hebbian processes. In our experiment here, this results in a pattern of change in interspecific interactions in our model that is highly consistent with the pattern of change that would be observed in a neural network trained on the same compositions using Hebbian learning. Throughout the simulation, interactions evolve in a manner consistent with Hebbian patterns of change that increase the basin of attraction for network configurations that maximally satisfy constraints in U (Figures 4.6, Appendix Figure C.5).

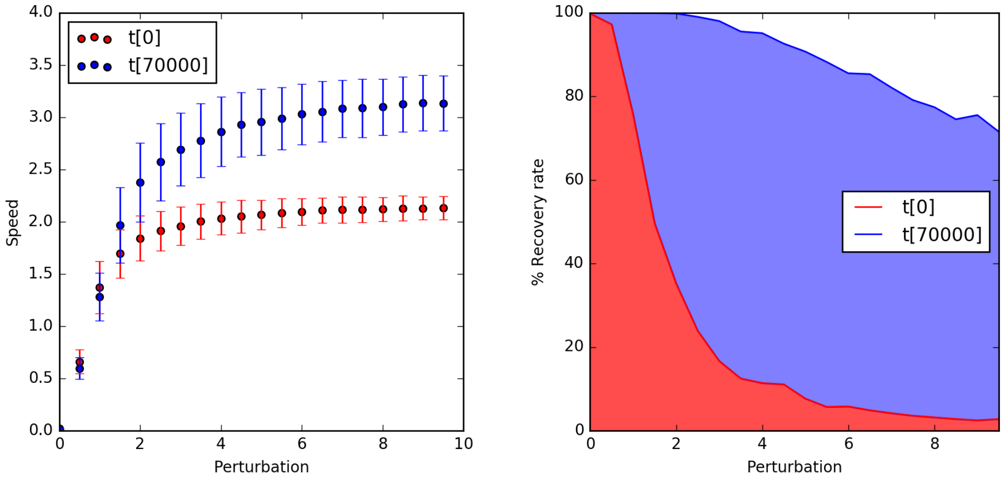

(A) (B)

Figure 4.5: Ecological stability at beginning and end of simulation. Data from a randomly selected simulation in experiment one, with ecological disturbance. At t[0], 1000 initial community compositions assemble to 604 different stable states. At t[70,000] 1000 random compositions assemble only 4 distinct stable states. Once each community has reached a stable state, we test stability by examining response to disturbance. For each species we reduce species densities by d, where d is random number between 0 and p. Minimum densities are capped at 0.001 to prevent disturbance eliminating species. a) The coevolved community (blue) exhibits greater speed of recovery from disturbance (i.e. greater elasticity) than the coevolutionary naive community (red). Speed is measured by Euclidean distance between community compositions after one iteration of Equation 4.1. b) The coevolved community is also more likely to recover from disturbance (it has higher amplitude). Recovery rate indicates number of communities that return to the same stable state following disturbance.

At the start of simulation the system visits many different stable states, including stable states in the set of the 92 configurations that satisfy all the constraints in U. Although the system visits many different stable states, over multiple iterations the Hebbian patterns of change that occurs in the network is consistent with those in a neural network trained to learn these 92 'maximally satisfying' configurations (Figure 4.6 a-c, Appendix Figure C.6a). Similarly, over the whole simulation, the evolved pattern of change is consistent with Hebbian patterns of change required to train a neural network to learn the four compositions that the system converges to at the end of simulation (Figure 4.6 d-f, Appendix Figure C.6b). Consistency between evolved and Hebbian changes occurs through the decreased competition between units that cooccur and the increasing competition between other units under the assumption of niche packing.

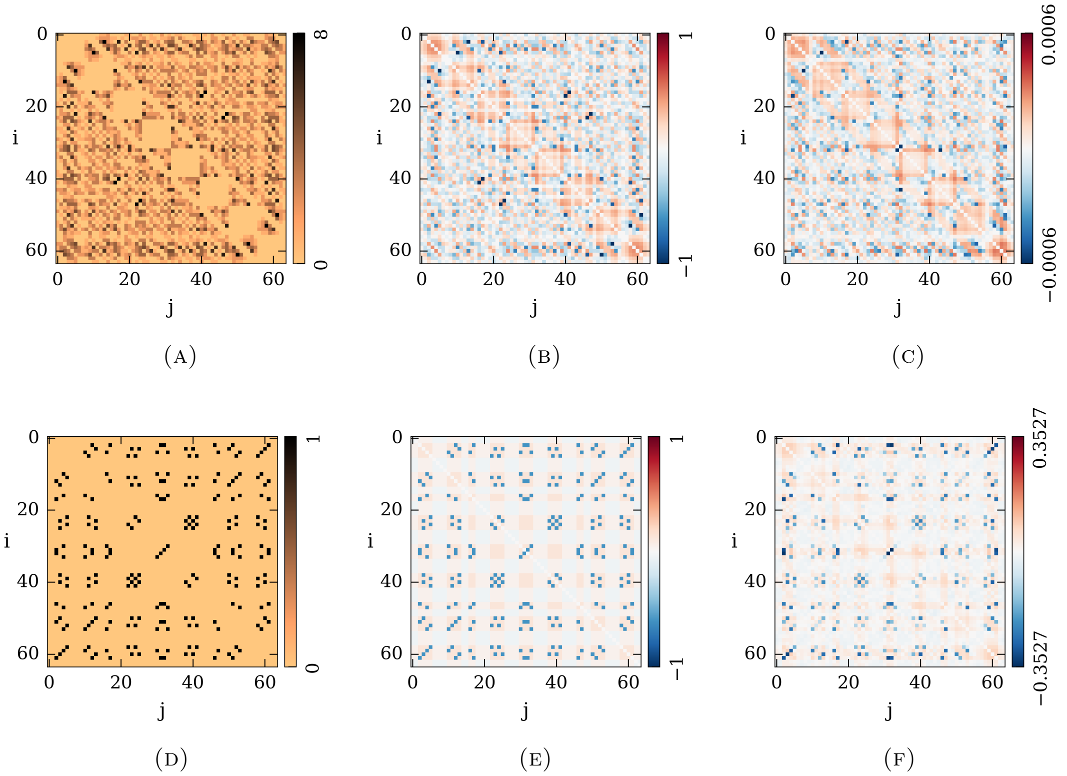

(A) (B) (C)

(D) (E) (F)

FIGURE 4.6: Hebbian patterns of change at the beginning of (top row), and over an entire (bottom row) simulation for a selected repeat from an experiment with regular disturbance. With the constraint topology set by U, there are 92 SS, that maximally satisfy all constraints in U. Entries in a) indicate co-occurrence between species i and j at these 92 attractors. b) The pattern of change in a neural network with 64 nodes trained using Hebb's rule to recall the same 92 compositions. c), the pattern of change in the community network during simulation, using the first 1000 iterations of scenario A. The pattern of change in b) correlates well with that in c): 82.81% of entries have the same sign; $R^2 = 67.6\%$; see Figure C.6a. At the end of simulation the system converges to dynamics where all random community compositions assemble to one of four community compositions. d) The pattern of co-occurrence of species in these four attractors. e) Hebbian changes to a neural network trained to recall the same four compositions. f), the pattern of change in the network over a complete run of scenario A. The pattern of change in e) correlates very well with that in f): 88.09% of entries have the same sign; $R^2 = 83.8\%$; see Figure C.6b.

4.4 Discussion

As in the previous chapter, we observe that distributed associative learning mechanisms operate in ecological community networks, just as they do in connectionist models of learning. By recognising the equivalence between models of ecological community networks and artificial neural networks we are able to make predictions about the types of behaviours exhibited by community networks, making several contributions to our understanding of community assembly dynamics.

4.4.1 Ecological community networks can be organised to assemble constraint-minimising compositions.

Our experimental work shows how, given certain conditions are met, distributed associative learning mechanisms in eco-evo dynamics modify community network structure so as to enable these systems to increase the frequency at which they assemble mature community compositions that minimise interspecific constraints (Figure 4.3). Just as a neural network under appropriate training can learn how to solve a complex constraint satisfaction problem, so ecological community networks can 'learn how to find' community compositions that 'solve' interspecific constraints. This behaviour in ecological community networks is possible despite the lack of any community-level direction; it is the result of individual selection acting on component species.

4.4.2 Constraint satisfaction is enabled through *self-modelling* behaviour in evolutionary dynamics.

As the system visits different attractors (community compositions) associative learning mechanisms, driven by Hebbian evolutionary dynamics, update the connections of the network with correlational information from its own attractors. In this way the system 'self-models' (Watson et al., 2011a) as it learns from its own constraints (contrasted with the experiment in Chapter 3, where the ecological network modelled external environmental constraints).

The effect of this behaviour is that the system develops an enlarged basin of attraction for compositions that are consistent with the mean correlation of species across all attractors. The end result in the simulation is an ecological network with four attractors, each of which maximally resolves internal network constraints.

We note two ways in which interspecific interactions are affected: 1) interactions between species that are incompatible due to fundamental constraints, which become more competitive. 2) Interactions between species that are compatible / incompatible due to secondary effects of (1), which become more cooperative / competitive as a result of correlation reinforcement at attractor compositions. Both of these effects are necessary for the system to learn the correlation structure and in order to drive the observed constraint satisfaction dynamics.

4.4.3 Constraint satisfaction behaviour in ecological community networks occurs through the same processes as in neural networks

The emergence of this system-level behaviour occurs without any group selection mechanisms. Instead, this behaviour arises from correlation learning behaviour driven by the Hebbian nature of evolutionary changes in species interactions. These processes are isomorphic to those found in learning neural networks, in which computational processes also develop at a system level without system-level direction, (i.e, without 'supervision', Watson and Szathmáry, 2016).

4.4.4 We identify conditions necessary for the evolution of enhanced constraint satisfaction behaviour.

In neural networks, training data must meet two conditions for associative learning processes to produce constraint satisfaction behaviour. Firstly, data must be sufficiently diverse for Hebbian learning (Hebb, 1949) in neural networks to develop a reasonable representation of the problem. Insufficiently diverse training data results in a neural network 'over-learning' features of the specific (sub-optimal) solutions it experiences, rather than the general features of the problem (Watson and Szathmáry, 2016). Secondly, as Hebbian learning in a network operates on the mean correlation of units across the history of the network's experience, the mean correlation of units the network experiences (i.e., across training data) must be consistent with the mean correlation of units in good solutions. These conditions are often met in the distribution of local optima (stable states) found by the convergent dynamical processes of constraint networks built from pairwise (or low order) interactions (Watson et al., 2011a; Watson and Szathmáry, 2016; Watson et al., 2011b).

We identify equivalent conditions required for ecological community networks to learn / evolve constraint satisfaction behaviour below, contributing to the parallels in mechanisms of distributed associative learning between these two dynamical systems (Table 4.1).

4.4.4.1 Regular disturbance enables the system to experience a variety of different states

Our control experiment illustrates how, should natural communities experience limited disturbance, ecological processes are not able to acquire correlational information from multiple different attractors. Assembly rules evolve so as to increase the frequency of assembling the experienced stable state(s), forming an ecological memory of the historically experienced configurations. (In machine learning terms, there is insufficiently diverse training data, resulting in 'over-learning'.) Regular ecological disturbance counteracts this effect. With disturbance, and given sufficient time, evolutionary processes act across multiple different compositions, (providing diversity in the 'training data') with the result that the system forms an ecological memory for a subset of its own attractor states.

	Unsupervised correlation learning	**Coevolution in ecological networks**
4.1.0	Conditions required for neural networks to learn to solve constraint satisfaction problems.	Conditions required for ecological community networks to evolve assembly rules with enhanced constraint satisfaction behaviour, resulting in lower interspecific competition, increased biomass and improved stability and resilience.
4.1.1	Learning takes place whilst the system is in a variety of states, preventing overlearning of one state and allowing the system to observe correlations between features.	Regular disturbance, such that evolution occurs whilst the system is in a variety community compositions. This prevents the system forming a distributed associative memory for a single experienced state and enables the system to generalise over correlations across experienced states.
4.1.2	Correlations between units in training data tend to be consistent with correlations between units in the solution state.	Mean correlations between species during evolution are consistent with correlations between species in low competition configurations.

Table 4.1: Equivalent conditions are required for distributed associative learning to result in improved constraint satisfaction behaviour in neural networks and community networks.

4.4.4.2 Hebbian dynamics and the accumulation of correlational information

We observe that the compositions that are amplified in the distribution of attractors are low-competition and high biomass compositions, but that is not the reason why they are amplified. These compositions are amplified because the mean correlation of species at these specific attractors is strongly correlated with the mean correlation of species across all attractors. In other words, the attractors that are amplified have species compositions that are most like the 'average' of species compositions across all attractors.

This effect is illustrated in Figure 4.6 where, as in Chapter 3, we observe Hebbian dynamics in the changes in interspecific interactions. We highlight two important observations: Firstly, that at the start of simulation, over multiple iterations, the pattern of change in interactions is consistent with Hebbian changes required to learn the 92 maximally satisfying conditions - despite the fact that most random community compositions do not converge to one of the 92 'maximally satisfying' configurations. One might expect that the evolution that occurred at other stable states would obscure the effect of evolution that occurred whilst the system was at one of these 92 configurations. This does not happen because the concurrence between species across all stable states is highly correlated with cooccurrence at these 92 stable states ($R^2 = 75.6\%$, Appendix Figure C.5).

Secondly, we note how the system increases the basin of attraction for a specific subset of its own attractor states (the final four stable states). This behaviour occurs as the system accumulates correlational information from all local attractor states visited. The system is not simply being exposed to low energy states and remembering them, but incrementally building an associative memory for these four low energy states from correlational information across all states.

4.4.4.3 Implications for 'predictive' self-modelling

The experiment in this chapter shows how eco-evo feedbacks can enlarge the basin of attraction for states that are low-competition and high biomass. At the start of simulation only 38.8% of random community compositions assemble a mature composition with lowest-possible competition. During simulation, the community enlarges the basin of attraction for certain states using information from other state compositions, such that, at the end of simulation, 100% of random community compositions assemble lowest-possible competition states.

The observation, that associative learning mechanisms can form attractors for particular states from correlational information at other state compositions leads to the intriguing suggestion that, if the mean correlations between units at stable states is consistent with correlation of units at a configuration that maximally satisfies constraints (or has any other interesting property), then there is potential for the system to develop a generalised model of its past experience from stable states visited and develop an attractor for the maximally satisfying stable state, even if it has not yet visited that composition. We explore this possibility in the following chapter.

4.4.5 Consequences for community functions

4.4.5.1 Biomass

The community compositions that correspond to increased constraint satisfaction and competition reduction also correspond to high biomass communities (Figure 4.2). Thus, as conditions enable constraint satisfaction, so do they enable communities to increase collective biomass. This happens even when controlling for evolved changes in interactions that reduce competition (and would increase biomass). That is, the community learns to assemble compositions that have higher inherent biomass (Figure 4.3f).

It is important to distinguish between the two (related) effects that lead to reduced competition and high biomass in mature coevolved communities. Firstly, it is not surprising that, since the effect of natural selection is to reduce competition between species that co-occur, that the total competition in the community is reduced by evolution. What we additionally demonstrate in the experiments is something more subtle (and more interesting), that low interspecific competition results not only from evolved reductions in competition, but also due to the community increasingly assembling stable states where the species present have low levels of intrinsic competition with each other. (And biomass increases are due, not only to gains in individual growth rate, but also due to community increasing ability to assemble stable states with low levels of intrinsic interspecific competition.) This change in assembly dynamics (which causes the community to assemble low-competition, high-biomass stable states) results from species modifying their interactions with one-another, and while this is driven by (very many) individual-level selective processes, the dynamics of the modified assembly rules are best regarded as a collective community behaviour. This effect is more than the sum of its parts in a formal sense: the adaptation that occurs within each species cannot be separated into adaptations that improve the fitness of that species (Figure 4.6); rather it depends on community members evolving interactions with one another (such that, e.g., the growth of species i is enhanced by the presence of species j and vice versa). This is exactly the same sense in which neural computation occurs at the network level and is not meaningfully decomposed into the sum of computations carried out by individual neurons (Watson et al., 2016).

4.4.5.2 Diversity

During simulations with disturbance, community dynamics change such that mature communities assemble attractor compositions that minimise interspecific competition. Prior to evolution, the ecological community assembles a broad range of different community compositions, many of which are comproised of more than $\sqrt{N}$ species, but where species are at lower densities than those at attractors of exactly $\sqrt{N}$ species (Figure 4.4a). At the end of simulation a small number of species dominate each mature community (Figure 4.4c). This reduction in diversity and increased dominance by a subset of species can be observed in reductions of Shannon diversity throughout the simulation (Figure C.4).

4.4.5.3 Stability

The evolution of enhanced constraint satisfaction behaviour and associated dynamics has implications for studies of community stability and tipping points between stable states (e.g. Dakos et al., 2015). We find that the correlation reinforcement that occurs in coevolutionary mature communities causes these communities to have fewer stable states, and that the stable states that persist are those which resolve high numbers of internal constraints. These communities recover faster following disturbance (they have higher elasticity; Westman, 1978), and can return to the same stable state from greater perturbations than coevolutionary naive populations (they have higher amplitude) (Figure 4.5). These dynamical changes are also relevant for invasional ecology, as coevolved communities that avoid competition with each other and (through niche packing) exclude mutual competitors are likely to be less susceptible to invasion than random assemblages. Dependencies between species in the invaded community also contribute to invasional meltdown - as the invaded community is repeatedly weakened after each subsequent invasion. Similarly coevolutionary experienced species-groups that have modified interactions with one-another so as to be more compatible, may also be more able to invade other communities (Simberloff and Von Holle, 1999). Finally, we note the potential of pre-existing constraints between species (e.g. due to relatedness) to be reinforced by Hebbian reinforcement and amplify phylogenetic effects on species distribution at the community level.

4.4.5.4 Choice of parameters

Our goal in this chapter is to demonstrate that the consequences of ecological organisation at the community level (without selective pressure at this level) has non-trivial consequences for community function (i.e. it enables constraint satisfaction behaviours). Previously the lack of selective pressure at the community level has meant that community-level behaviours such as these have not been widely considered. It is therefore valuable to demonstrate how community organisation may occur, and its potential consequences. With this focus in mind, we have chosen a parameter set that causes the model to clearly exhibit the constraint satisfaction behaviours we are interested in. We hope that doing so will encourage further investigation, including investigation of how these results transfer to natural ecological scenarios.

4.4.6 Summary

We identify conditions necessary for ecological community networks to evolve enhanced constraint satisfaction behaviour, causing them to increase the frequency at which they assemble low-competition, high-biomass community compositions. This behaviour is not driven by community level effects, but arises due to distributed associative learning mechanisms, driven by natural selection. These behaviours occur in the same way (and depend upon equivalent conditions) as distributed associative learning processes in neural networks, and can be predicted by recognising the natural isomorphism that exists between these systems.

Chapter 5

Eco-evo feedbacks drive self-modelling behaviour in community assembly rules

5.1 Introduction

For a pool of potential community members there are many possible assembly sequences and endpoints, with sensitivity to initial conditions influencing the trajectory of assembly dynamics (Thompson et al., 2001). When released from an arbitrary initial configuration, an ecological community network will assemble a mature community composition that locally resolves constraints between units. However, in situations where there are very many interdependencies between variables, the likelihood of a community assembling a configuration that globally minimises interspecific constraints from random initial conditions can be very low.

Previously, the lack of natural selection at the community-level has generally precluded thinking about collective adaptive behaviours at the community level, and it has not been clear whether coevolved communities are better able to resolve constraints than co-evolutionarily naive populations. However, our finding, that the structures of ecological networks evolve as a function of their cooccurrence, creating dynamics that are isomorphic with connectionist models of learning, invites new possibilities for our understanding of how these systems are organised, and their potential to optimise for constraint satisfaction. In this chapter we consider whether distributed associative learning mechanisms enable community networks to perform 'predictive' self-modelling behaviour, and improve their ability to resolve collective constraints in computationally challenging ecological scenarios.

5.1.1 'Predictive' self-modelling in neural networks

Simple distributed mechanisms can incrementally alter a dynamical system such that it finds lower energy configurations, more reliably and more quickly. This can occur as feedback from correlational information from its own attractors augments the system's ability to resolve internal constraints in a process known as self-modelling (Watson et al., 2011a). Specifically, when Hebbian learning is applied to the connections of a simple dynamical system such as a neural network undergoing repeated relaxation, the system can develop an attractor for composition that reflects the mean correlations between units at experienced states. This modifies the dynamics of the network such that it increases the frequency at which it finds configurations that 1) minimise total system energy, and 2) globally resolve conflicts between interdependent variables. Unlike the training process for storing a specific associative memory, the system is not merely 'remembering' low energy states that have been previously visited, but is able to 'predict' their location by generalising over local attractor states that have already been visited.

5.1.2 Predictive self-modelling in community networks

The model in the previous chapter demonstrates the ability of eco-evo feedbacks to enlarge the basin of attraction for global optimal configurations for constraint satisfaction. These behaviours are not directed, e.g. by selective effects at the community level, but are simply a consequent effect of eco-evo feedbacks enhancing the network of interactions with correlational information from community compositions. However, as the basin of attraction for those states was large enough that the system regularly visited those configurations, that experiment was not conclusive in demonstrating the ability of a self-modelling system to develop an attractor for a globally optimum configuration without having visited that state (it did not demonstrate the 'predictive' behaviour described above).

To explore this issue we chose a community network structure matching a challenging constraint satisfaction problem of known difficulty, the puzzle game Sudoku. When converted into a network of pairwise constraints, this puzzle has the property that there are an enormous number of local optima, and a unique, globally optimal solution for constraint satisfaction. Using the evolving community model we have developed in earlier chapters, we can test whether improvements in constraint satisfaction abilities in community assembly rules enable community networks to 'solve' these puzzles, by learning how to find a unique configuration of species that minimises interspecific constraints. We find that community networks can, in fact, solve many puzzles easily and can also solve some the hardest known examples of this puzzle. Whilst the resource allocation constraints faced by natural populations are unlikely to be this difficult, such puzzles serve to illustrate the potential of simple evolutionary processes to drive system level behaviours and constraint satisfaction in a recognisably challenging problem.

5.2 Methods

We employ a model identical to that in Chapter 4, using the same processes for ecological dynamics (Equation 4.1), rate of change in interactions (Equation 4.2), and ecological constraints (Section 4.2.3). We also maintain the separation of the weight matrix Ω into evolvable (E) and unevolvable (U) components (Section 4.2.4).

We examine ecological dynamics in a new ecological scenario, derived from the puzzle game Sudoku, where constraints relating to the rules of the game are captured in unevolvable interspecific constraints, and constraints relating to a specific puzzle instance (the given cells) are captured by elevating the carrying capacity for particular species.

Simulation parameters are described in Appendix D.1.

5.2.1 Ecological scenario

The aim in Sudoku is to complete the entries of a 9×9 matrix. Each cell must be occupied by a number $1 - 9$ and each number must appear in any cell, row, column or 3×3 square exactly once. We cast these constraints in an ecological setting by considering a community of species competing for limited resources. As in forbidden species combinations (Diamond, 1975), species with similar resource profiles may exclude each other from a habitat. The 81 cells in the puzzle matrix and the nine possible entries for each cell define the size of the ecological community needed to capture the rules of the game in pairwise constraints: $81 \times 9 = 729$ species, whereby each species represents a single digit $(1 - 9)$ in a single location (row and column).

As in Chapter 3, we define traits $\{T_1, ..., T_4\}$ to represent competition between species pairs representing two species either: in the same grid cell (T_1); for the same digit in the same row (T_2); for the same digit in the same column (T_3); or for the same digit in the same box (T_4). Again, for each trait T_f, two species are competitive if they share the same value for that trait, and the unevolvable matrix, U, is given by the sum of competition across all traits (Equation 4.3).

5.2.2 Ecological representation of a specific Sudoku puzzle

Community composition is shaped by both biotic interactions between species and the abiotic constraints of the environment. Abiotic constraints such as the availability of light, space and water shape the carrying capacity of the environment for species in the regional pool. These constraints are incorporated in our model through elevated carrying capacities for the particular subset of species that define the configuration of that particular puzzle (i.e. the numbers that are provided to set the puzzle). Carrying capacities of those species are raised by α, and the carrying capacities of the remaining species are reduced by the same amount. This ensures that the species representing the given entries appear in mature community compositions. We use a suite of three set puzzles to test whether evolution by natural selection can enable ecological communities to resolve these constraints and produce community compositions consistent with solved Sudoku solutions (Appendix D.1).

5.2.3 Converting a community vector to a puzzle 'answer'

To translate a vector of species densities to a 9×9 Sudoku matrix (e.g. to see whether a system has solved a puzzle), for each cell in a Sudoku grid we select the species in highest density from all species representing digits in that cell. Puzzle answers are scored by the number of cells that do not violate a constraint of the game. Vectors that translate to solutions of the game, thus receive the maximum score, 81.

5.3 Results

5.3.1 Community dynamics prior to evolution

Fixing specific species (a puzzle instance) creates a network topology for which there is only one community composition corresponding to a Sudoku solution. With network size much greater than in Chapter 3, for all puzzles in the test suite, prior to evolution none of 1000 random community compositions reach the unique stable state corresponding to the Sudoku solution (Figure 5.1, left hand panels).

5.3.2 Evolution of constraint satisfaction

As in Chapter 3, during simulation we run ecological dynamics until the system is at or near a stable state before applying evolutionary processes. We disturb species densities and repeat the process, with the effect that the system visits very many stable states. The effect of raising carrying capacities for certain species is that these species consistently appear in the stable state reached. This, in turn, has the effect that other high density species in community compositions tend to be those that are compatible with these ‘fixed’ species. Our analytical result in Chapter 2 identifies that the selective pressures on interspecific interactions are a function of species’ cooccurrence (Equation 4.2). As evolutionary processes drive the greatest reductions in competition (or increased correlation) between species in highest densities, competitive interactions between compatible species (across $\{T_1, ..., T_4\}$) are those most reduced / made less competitive. The effect of this correlation learning is that by the end of simulation, for all simulations, community dynamics converge to stable states that resolve a higher number of interspecific constraints than at the start of simulation (Figure 5.1). Furthermore, for puzzles s15a and s15c, we observe instances where community dynamics evolve to a point where all random community compositions converge to community compositions that maximally satisfy interspecific constraints, community compositions that correspond to Sudoku solutions (Figure 5.1 a-d). Six repeats of the scenario with puzzle s15a and two repeats with puzzle s15c converge to a stable state matching the unique Sudoku configuration in these puzzles. The change in ecological dynamics for a selected ‘solution-finding’ simulation is shown in Figure D.2. In this instance ecological dynamics prior to evolution assemble from random community compositions (Sudoku score 9) to a stable state with Sudoku score 62. At the end of simulation, from the same initial composition, assembly rules converge to a stable state with Sudoku score 81, corresponding to a solved Sudoku puzzle.

5.3.3 The importance of correlational information

The ability of the system to find low-competition configurations is only possible because natural selection is able to act upon correlational information from species pairs. That is, natural selection is only able to shape assembly rules such that the system converges to low-competition states because the selective pressure on interactions is a function of species cooccurrence, rather than just their occurrence. This enables evolutionary processes to acquire information from the correlational relationship between species; which species ‘go together’ rather than just which species ‘should be present’.

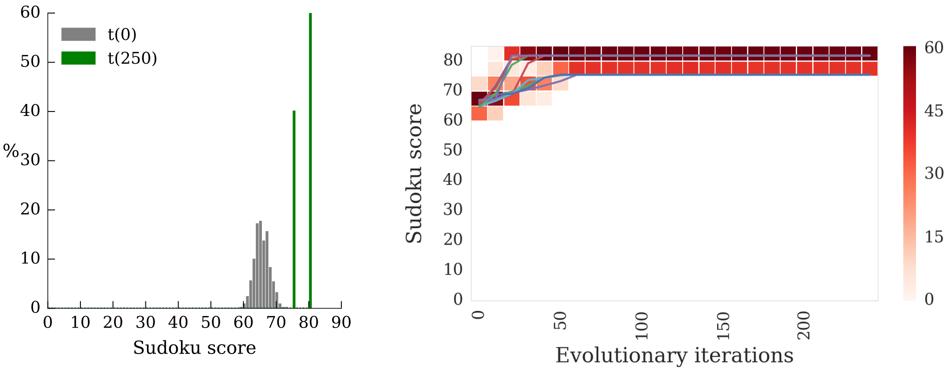

(A) puzzle s15a

(B)

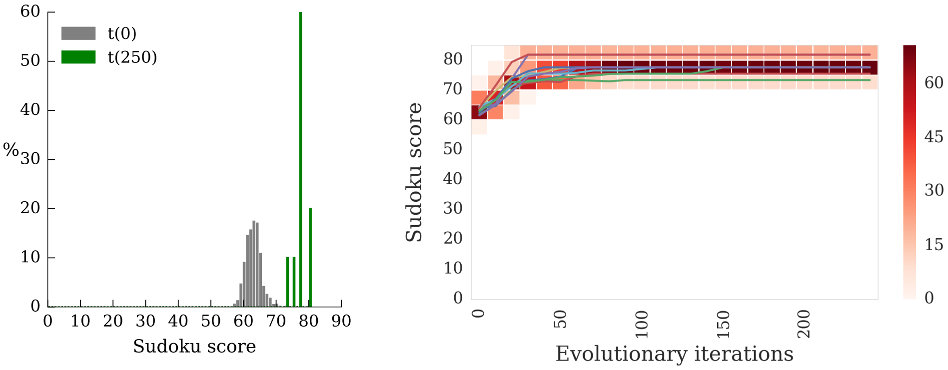

(C) puzzle s15c

(D)

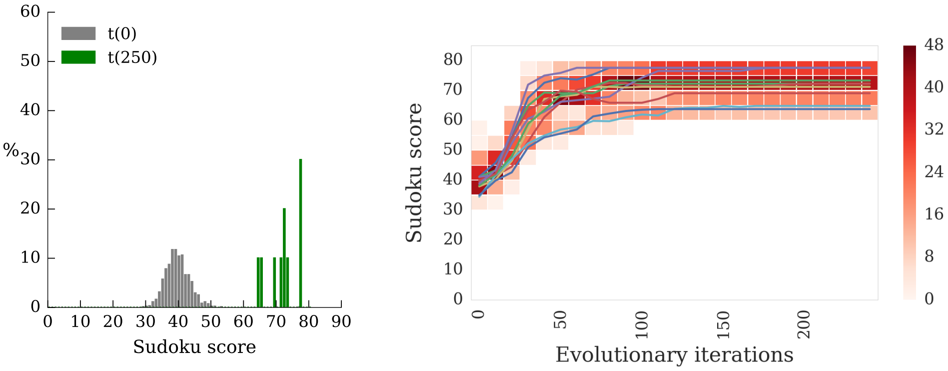

(E) AI Escargot

(F)

Figure 5.1: Effects of evolutionary processes on community dynamics for simulations run with puzzles s15a (a, b), s15c (c, d) and AI Escargot (e, f). Left hand panels show the distribution of Sudoku scores of stable states reached by assembly rules at t(0) (grey bars); and t(250) (green bars). Data aggregated from 100 random community compositions tested in each of 10 simulation runs. Right hand panels show the change in Sudoku score throughout 10 simulations for each member of the test suite. Heat maps show the binned data (bin size = 5) aggregated from all simulations. From these results we find that six repeats with puzzle s15a and two repeats with puzzle s15c converge to a stable state matching the unique Sudoku configuration, with the maximum Sudoku score possible, 81 (equivalent to solving these Sudoku puzzles) and, that for all puzzles and simulations, the ability of the system to find configurations that resolve internal constraints improves over evolutionary time. For the sake of clarity legends for right hand panels are included in the appendix (Figure C.2).

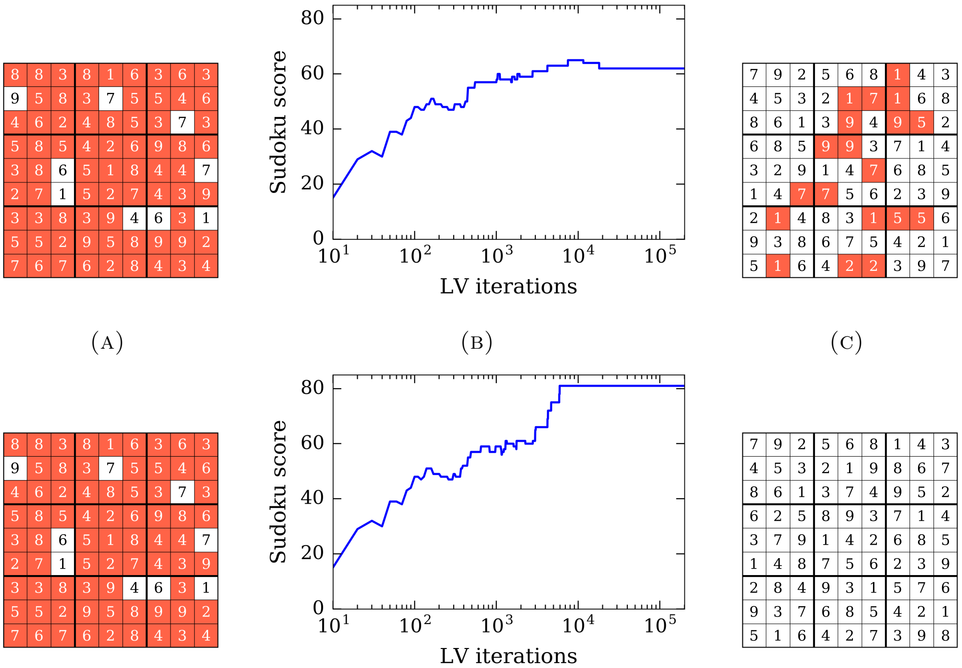

(A)

8	8	3	8	1	6	3	6	3
9	5	8	3	7	5	5	4	6
4	6	2	4	8	5	3	7	3
5	8	5	4	2	6	9	8	6
3	8	6	5	1	8	4	4	7
2	7	1	5	2	7	4	3	9
3	3	8	3	9	4	6	3	1
5	5	2	9	5	8	9	9	2
7	6	7	6	2	8	4	3	4

(B)

(C)

7	9	2	5	6	8	1	4	3
4	5	3	2	1	7	1	6	8
8	6	1	3	9	4	9	5	2
6	8	5	9	9	3	7	1	4
3	2	9	1	4	7	6	8	5
1	4	7	7	5	6	2	3	9
2	1	4	8	3	1	5	5	6
9	3	8	6	7	5	4	2	1
5	1	6	4	2	2	3	9	7

(D)

8	8	3	8	1	6	3	6	3
9	5	8	3	7	5	5	4	6
4	6	2	4	8	5	3	7	3
5	8	5	4	2	6	9	8	6
3	8	6	5	1	8	4	4	7
2	7	1	5	2	7	4	3	9
3	3	8	3	9	4	6	3	1
5	5	2	9	5	8	9	9	2
7	6	7	6	2	8	4	3	4

(E)

(F)

7	9	2	5	6	8	1	4	3
4	5	3	2	1	9	8	6	7
8	6	1	3	7	4	9	5	2
6	2	5	8	9	3	7	1	4
3	7	9	1	4	2	6	8	5
1	4	8	7	5	6	2	3	9
2	8	4	9	3	1	5	7	6
9	3	7	6	8	5	4	2	1
5	1	6	4	2	7	3	9	8

FIGURE 5.2: Ecological dynamics, at t(0), before simulation (a-c); and t(250), after simulation (d-f), for a selected solution-finding simulation using puzzle s15c. a) A random community composition corresponds to a Sudoku grid with very many conflicting cells (in red). b) Assembly rules in the unevolved community form a stable state that corresponds to c) a Sudoku grid with fewer, but not zero, conflicts. d) - f) At the end of simulation evolutionary changes to assembly rules mean the community now converges to a stable state with zero conflicts corresponding to the unique Sudoku solution.

We demonstrate the importance of this effect by testing the (in)ability of a non-evolving system to find good Sudoku scores through the aggregation of information from species' densities only, rather than from correlations between species (Figure 5.3). We define two functions $A1$ and $A2$, each of which aggregates information from species densities to generate Sudoku grids and explore whether aggregating information from species densities improves Sudoku score as we increase the number of inputs (M), where inputs are stable states reached from random community compositions in the unevolved system. Specifically, we define $A1$ as the modal Sudoku answer from Sudoku answers derived from M stable states, and $A2$ as the Sudoku answer derived from the mean species densities from M stable states. Neither method of aggregating information in species densities improves the Sudoku score of the collective answer as M is increased. Of course, if we consider the mechanics of the game this is not surprising. In an empty Sudoku grid, any digit can be placed in any cell, but the addition of each digit further constrains which other digits can be added; it is the interactions between digits that are fundamental in defining the score of the system. The evolving system learns, not because the distribution of species at each stable state is consistent with the Sudoku solution, but because the correlations between species is consistent with correlations in the solution configuration. It is only through the acquisition of correlational information that the system is able to augment its dynamical behaviour, generalising over the attractors it has experienced, and amplify the basin of attraction for the global optimum, without having ever visited that configuration.

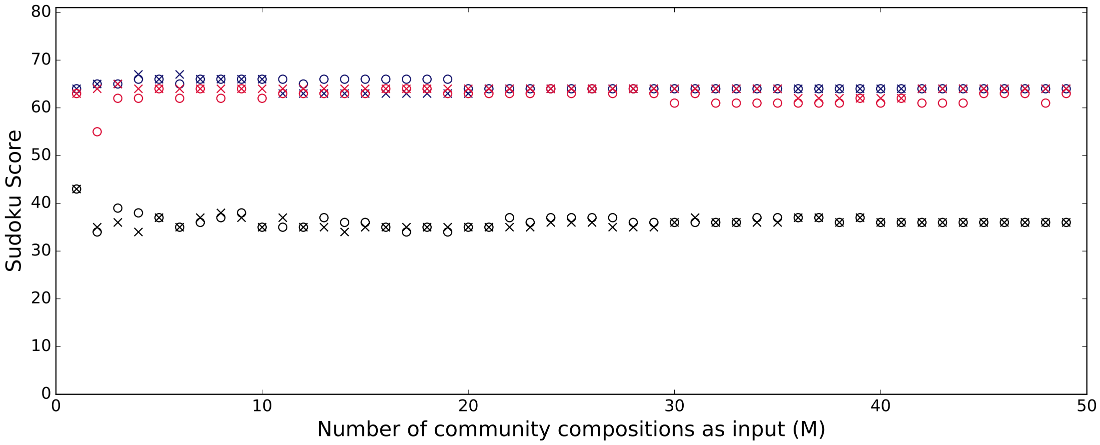

FIGURE 5.3: Rather than applying any evolutionary processes, here we examine the Sudoku scores of answers generated by aggregating information from species densities from multiple (M) stable states. We define two functions $A1$ and $A2$ where $A1$ is the modal Sudoku answer from Sudoku answers derived from M stable states, and $A2$ is the Sudoku answer derived from the mean species densities from M stable states. Neither function, $A1$ (circles), nor $A2$, (crosses) improves the Sudoku score of the collective answer as M is increased. The lack of success here highlights the importance of correlational information in creating the observed effect in our experiments. Colour represent tests conducted using each puzzle in the test suite: AI Escargot (black), s15a (blue), and s15c (red).

5.4 Discussion

Our result here demonstrates the capacity of an ecological community network to develop an associative memory for the globally optimal, lowest-competition state (corresponding to a Sudoku solution) within timescales in which a non-evolving system would not yet have visited that state. For those simulations where the system successfully converges to a solution configuration, the system enlarges the basin of attraction for that configuration before it is visited for the first time. The system is, in effect, generalising over the configurations it experiences in a random sample of higher-competition attractor compositions, enabling the basin of attraction for the solution configuration to be enlarged even though it has not yet been visited (Watson et al., 2011a).

The general process in the ecological network is as follows. We configure constraints in our model so that assembly dynamics corresponds to a constraint optimisation problem of known difficulty. As the system is repeatedly perturbed and allowed to re-assemble from many different initial conditions, it visits very many different mature community compositions (dynamical attractors). If the frequency of disturbance is low compared with the time allowed for ecological assembly, this causes the system to spend most of its time at these mature compositions. But, as there are so many local optima, most are never visited, and those that are visited, tend to be visited only once. Evolutionary processes occur at the end of each ecological cycle, when the community is at one of these mature compositions. Evolutionary dynamics act to reduce competition between species in high density, causing these species pairs to be more likely to cooccur in future, and creating feedback between ecological and evolutionary processes. These eco-evo feedbacks, in proportion to species cooccurrence, have the effect of enlarging the dynamical basin of attraction for the system's current state configuration, and therefore would make the system more likely to visit state configurations that have already been visited (simple positive feedback). But, for the timescales of interest these very small changes in basin size for local optima have very little effect on likelihood of the system re-forming that state. However, changes to the interaction structure of the network also enlarge the basin of attraction for other state configurations, in proportion to their similarity to state configurations that have already been visited. In constraint satisfaction problems built out of low-order constraints, correlations between units in the global optima are highly correlated with correlations between units across local optima (discussed below). Whereas the increases in basin size for individual local optima are not cumulative (because local optima are only visited once) the changes to the global optima *are* cumulative. The effect of these cumulative changes is that the system develops an associative memory for the global optimum, without having ever visited that state, increasing the likelihood that this configuration is assembled following future disturbance.

As such, the ecological network exhibits the same 'predictive' self-modelling known to occur in neural networks. This finding marks the final contribution of this thesis towards identifying equivalent processes between neural networks and ecological community networks (Table 5.1).

The dynamics we demonstrate show how correlational information from local attractors can accumulate to cause the ecological community network to form new attractors (or enlarge attractors) for compositions where correlations between species are consistent with those at local attractors. If these compositions are global optima for any property, such as low competition, the community consequently develops attractors for that property. In the case we examine the community develops attractors for low competition states. It is not necessarily the case these

	Unsupervised correlation learning	Coevolution in ecological networks
5.1.0	Ability of a neural network to learn to solve complex constraint satisfaction problems by augmenting its own dynamical behaviour with correlational information from its own attractors.	Ability of an ecological community network to evolve its assembly rules so as to solve complex resource-allocation problems.
5.1.1	Learning processes take place whilst the system is at a variety of states.	Evolution occurs whilst the community is in a variety of compositions (Table 4.1.1).
5.1.2	Hebbian learning modifies intercations between units in proportion to correlations across trained states.	Species evolve their interactions in proportion to their cooccurrence, amplifying trends in their correlations.
5.1.3	Changes to interactions alter the network's dynamical behaviour, increasing the basin of attraction for compositions in proportion to their similarity to the mean correlations between units.	Evolving interspecific interactions alter assembly dynamics, increasing the frequency of assembling community compositions in proportion to their similarity to experienced interspecific correlations.
5.1.4	Increased basin of attraction leads to increased frequency of assembling this state. Potentially assembling this state without having experienced it.	If the mean correlation between species is consistent with the correlations between species in low-competition states (Table 4.1.2), this leads to increased frequency of assembling low-competition states. Potential to assemble the globally optimal state for constraint satisfaction without having experienced (and evolved at) that composition.
5.1.5	The system *learn how to find* solutions to complex combinatorial problems (*self-modelling*; Watson et al., 2011a).	The system *evolves to assemble* resource optimising compositions.

TABLE 5.1: Self-modelling in behaviour in neural networks and ecological community networks arises due to equivalent dynamics in each system.

behaviours are replicated in natural community networks. However, we suggest that for community networks where interspecific constraints are low-order (e.g. the sum of pairwise constraints, as in a Lotka-Volterra model), this relationship between local and global optima is a reliable heuristic (Watson et al., 2011a). Exploring this will be an important area of future work.

By showing how self-modelling constraint satisfaction behaviour enables our simulation model to, in effect, solve Sudoku puzzles, we aim to drive curiosity about the consequences of these dynamics in natural scenarios. Whilst the Sudoku-solving behaviour we demonstrate in our computer model will not literally be replicated in tangled ecological conditions, the general dynamics we describe here, where evolutionary processes drive self-modelling behaviour and constraint satisfaction, will be testable in natural scenarios.

Chapter 6

Discussion

This thesis introduces distributed associative learning as a framework to expand our understanding of the effects of eco-evo dynamics on collective ecological behaviours. Working within this framework we find that, despite not being an evolutionary unit, an ecological community can behave like an unsupervised learning system. This behaviour can be understood by recognising the natural isomorphism between the selective pressure on interspecific interactions in a community network and the change in interactions caused by Hebbian learning in a neural network. Building from this isomorphism, this work finds that evolving ecological community networks are capable of performing several adaptive behaviours also known to occur in neural networks: creating internal organisations that collectively habituate to past environmental conditions; actively recalling past responses to those conditions; and resolving complex internal constraints through self-modelling behaviour.

6.1 Main claims

6.1.1 Selective pressure on interspecific interactions produces positive feedback on species cooccurrence

Our first key result is to identify that the selective pressure on interspecific interactions produces positive feedback on species' cooccurrence (Chapter 2). That is, interactions between species that tend to cooccur, become less competitive, with the result that these species are more likely to cooccur in future.

Formalising the relationship between species' cooccurrence and the rate of change of their interspecific interactions bridges a gap between microevolutionary processes and community-level consequences, providing new insight into how community networks are shaped by natural selection at lower levels of selection. Previously there have only been two choices in how to interpret the relationship between selective processes and collective behaviours in ecosystems: either collective behaviours have no system-level organisation or some mechanism of group selection must be involved. Our findings demonstrate that there is a third possibility. Collective behaviours in ecological communities can arise from the positive feedback of individual natural selection and ecological population dynamics without the requirement of group selection.

6.1.2 A natural isomorphism exists between selective pressure on interspecific interactions and Hebbian learning in neural networks

Our second claim, developed in Chapters 2 and 3, is to identify that the relationship between species cooccurrence and selective pressure on their interactions has a natural isomorphism with the effects of unsupervised correlation learning in neural networks. Given the presence of evolutionary trade-offs, the effect of individual-level natural selection acting on interspecific relationships is dynamically equivalent to Hebbian learning on interactions in a neural network and, thus, ecosystems can exhibit organised collective behaviours via the same principles of connectionist learning that apply to neural networks.

This isomorphism results from analogous relationships between components in each system. In neural networks trained with Hebb's rule, connections between one node and another are adjusted in proportion to the product of the output from both nodes. The Hebbian pattern of change in network connections is often paraphrased as 'neurons that fire together, wire together' (Lowel and Singer, 1992). In community networks the selective pressure on each species to reduce its interaction with competitor species is determined by both its own population density (providing potential for mutations) and the density of the competitor species (providing incentive to avoid interaction), with the result that the rate of adaptation on interactions is proportional to the product of the species' densities (i.e. species co-occurrence). Interactions between species that co-occur (fire together) evolve to become less competitive, with the result that they are more likely to co-occur in future (wire together).

These dynamics drive equivalent system-level behaviours in each system (distributed associative learning, discussed below). There is a key difference between the drivers of the isomorphic behaviour: in neural networks, distributed associative learning is the result of a directional learning mechanism; whereas in ecological systems distributed associative learning is the result of natural selection, acting on random heritable variation in interspecific fitness dependencies.

The recognition of this isomorphism allows the useful transfer of knowledge between domains. We begin this transfer by translating known consequences of these organisational behaviours in neural networks to make predictions for the organisation of system-level behaviour in ecological communities. We explore the consequences of these mechanisms in our simulation models, making contributions to our understanding of eco-evo feedbacks for memory and constraint satisfaction behaviours.

6.1.3 Coevolution in community networks can create organised collective behaviours driven by natural selection at lower levels of organisation

Distributed associative learning mechanisms, isomorphic to Hebbian processes in neural networks drive collective organisation at the community level, without community-level selection. The result of this is that the mature ecological communities are non-random assemblages of species, coevolved such that they are better able to coexist.

This finding, that coevolved communities assemble mutually compatible compositions of species is reminiscent of Clements' 'organismal' metaphor for plant communities (Clements, 1936) - that proposed that assembly processes occur along discrete pathways, culminating in a well defined ecological community. Clement's view was famously at odds with that of Gleason (1926). Gleason viewed community compositions as being determined more by random events, such as climate, soil properties or physiography, with interspecific interactions playing a relatively minor role in determining community assembly.

Our work has focussed on the role of interspecific interactions on community assembly, and we do not compare the effect of these with other processes affecting assembly dynamics. Yet the results in this thesis do highlight the potential of ecological communities to evolve assembly processes in a consistent, organised manner so as to encode information relating to particular community states (thus alternate community states, become alternate ecological states; Beisner et al., 2003). The mature community compositions that arise in our models are coevolved, non-random, dynamically stable compositions, much like Clement's super-organismal plant communities.

6.1.4 Coevolution in community networks can produce a distributed associative memory of multiple past environmental states

In Chapter 3 we show how distributed associative learning in ecological communities enables these networks to form ecological memories of historically experienced states, encoded in the network of interactions between community members, in the same way that multiple memories of specific states can be encoded in the connections of a neural network. Specifically, we demonstrate that community networks can form distributed associative memories of multiple stable states and recall each of these specific and distinct states with high fidelity. This behaviour is encoded by individual natural selection in the network of interactions between community members. This result demonstrates how, despite not being Darwinian units, communities can display collective behaviours, creating internal conditions that habituate to past environmental conditions and actively recalling those states.

In Chapters 3 and 4 we also note how the dynamical landscape is shaped by the formation of ecological memories, and consequent implications for the stability of coevolved communities, including increased elasticity (speed of recovery from disturbance) and amplitude (increased likelihood of assembling the same state after recovery from distrurbance). Stability in these metrics tends to increase as community members form better relationships, but if the evolutionary history of a community has created a multimodal dynamical landscape this can result in the formation of alternative stable states, rather than universal stability. We observe that a system may respond very differently to forcing towards an alternatively stable state, with the possibility of critical transition in the trajectory of the system towards a different attractor.

6.1.5 Coevolution in community networks can produce enhanced constraint satisfaction behaviour

In Chapters 4 and 5 we show how ecological communities can evolve enhanced constraint satisfaction behaviours by augmenting their own dynamical behaviour with correlational information acquired whilst evolving at locally-optimal configurations. This behaviour is equivalent to the way in which a neural network 'self-models', augmenting its behaviour with an associative memory of its own attractor states (Watson et al., 2011a).

In Chapter 4 we identify that the same conditions are required for a community network to *evolve* enhanced constraint satisfaction behaviour, as are required for a neural network to *learn via training* to exhibit the same behaviour. In each case this behaviour is facilitated through the encoding of correlational information in the interaction network.

Collective constraint satisfaction behaviours have implications for community biomass. As low-competition states have higher biomass, the ability of distributed associative learning mechanisms to evolve attractors for low competition states means that ecological communities also evolve attractors for high biomass states, even when controlling for the effects of natural selection on individual species' biomass. This effect occurs through distributed associative learning mechanisms driven by individual-level natural selection at lower levels of organisation.

6.2 Adaptive behaviours in ecological communities

In each of the scenarios modelled, distributed associative learning mechanisms, dynamically equivalent to unsupervised correlation learning in neural networks, result in ecosystems exhibiting adaptive behaviours via the same principles of connectionist learning that apply to neural networks.

Unpicking the analogy further, we can ask 'what is it that ecosystems learn?'. Our results show that coevolution in ecological networks has the capacity to cause the network to form ecological attractors for particular community states (influenced by either external or internal conditions) by forming a distributed associative memory of corresponding community states. The effect of this learning is that communities can either recreate experienced states (Chapter 3), or assemble specific states that resolve internal constraints (Chapters 4 and 5). In each of these experiments communities learn *where to go* (where an attractor is a climax community resulting from a successional process (Law and Morton, 1996) with one of the properties we study in this thesis - memory or constraint satisfaction behaviours), *how to get there* (i.e. assembly dynamics or successional process evolved in a manner isomorphic to Hebbian learning), and *how to stay there* (i.e. the coevolved relationships that increase the resilience and stability of those mature ecological states).

We identify the mechanisms of this learning, demonstrating *how communities learn* (by augmenting their dynamical behaviour with correlational information from experienced states), and we understand the *requirements for learning* (in particular, the presence of normalising conditions that prevent selection from eliminating competition, and the importance of ecological disturbance to prevent over-learning a single state).

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 thesis. 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 adaptive behaviours, as demonstrated by the claims above.

6.3 Correlational learning

Our fundamental result (and each of our consequent findings) stems from the observation that species evolve their interactions in proportion to their cooccurrence. Interactions between species that co-occur evolve to become less competitive, with the result that they are more likely to co-occur in future. However, whilst the mechanism of distributed associative learning does depend on this general result, it does not depend on the exact formulation of the above argument that we present in Chapter 2 (Equations 2.3 and 2.4). In Chapter 2 we demonstrate that our conclusions are robust to the choice of underlying model, and in Chapter 3 we demonstrate that alternative models generate the same patterns of change in community interactions that are necessary for associative learning to take place (see also Appendix B.2).

The essential requirement in the dynamics of the eco-evo feedbacks is that interactions are modified in response to interspecific *correlations*. We demonstrate the importance of this in Chapter 5 (Figure 5.3) by showing how a system that aggregates information only from species densities, and not the correlations between species, is unable to show any improvement in constraint satisfaction behaviour. It is, in fact, impossible for networks of interactions to demonstrate the behaviour we have shown (e.g. forming associative models of multiple memories and recall these with high fidelity) without the acquisition of correlational information.

We find that evolution drives correlational learning in interspecific interactions leading to higher-level organisation. In doing so we introduce a mechanistic, process-based framework bridging the gap from micro-evolutionary processes to macro-evolutionary changes in system dynamics, expanding understanding regarding the impact that processes such as character displacement can have when replicated across a network. This finding places emphasis on the importance of considering the potential of species to evolve their interactions with one another - especially when seeking to understand the impact of low-level selective processes on collective behaviours at higher levels of organisation. Failure to recognise the ability of natural selection to shape correlations between units, and the potential impact that this can have, leaves researchers reliant upon sum-of-the-parts explanations. With this limitation all community functions must be explained as either the results of natural selection on individual units in the system, or chance / outside processes. Incorporating the effects of eco-evo feedbacks facilitates the view of communities as complex adaptive systems (Levin, 1998) with non-linear dynamics, historical dependencies and multiple possible outcomes from assembly dynamics.

6.4 Ecological communities as complex adaptive systems

Complex adaptive systems are systems comprised of large numbers of entities, among which local interactions create multiple levels of collective structure and organisation. As well as ecological communities, examples include many other biological systems, ranging from gene regulation networks to human social systems (Farmer and Lo, 1999), as well as artificial systems such as the Internet (McCabe et al., 2011). A special feature of complex systems is the emergence of non-trivial structures which can dominate the system's behaviour but which cannot be easily traced back to the properties of the constituent entities. Not only do higher emergent features of complex systems arise from lower-level interactions, but the system-level patterns that they create also affect these lower levels, creating feedback loops (Bourgine et al., 2009). As complex adaptive systems require analysis at many different spatial and temporal scales, researchers face radically new challenges when trying to observe complex systems, in learning how to describe them effectively, and in developing theories of their behaviour and control.

Our work has utilised transfer of theory between disciplines of complex systems study, from connectionism to evolutionary ecology. Doing so has enabled us to make testable predictions regarding the behaviour of evolving ecological networks, as well as the conditions required for these behaviours. Of course, it would be possible to understand emergent behaviours in ecological networks without any reference to connectionism - but by transferring information between disciplines we have been able target our experimental work to test the limits of these similarities, creating efficiencies in our research efforts.

Just as our work has benefited from the transfer of knowledge from other complex systems, so too are our findings likely to be relevant to studies of other domains. We have modelled a complex system in which the simple reinforcement of correlations creates learning processes that drive adaptive behaviours at the system level, despite that system not being an evolutionary unit. This result may be transferable to other domains where selfish components self-organise into structures that enhance global adaptation (Watson et al., 2011b)

As an example, in developmental biology, gene regulation networks (GRNs) interact so as to govern the expression of genes (Szathmáry et al., 2001; Wagner et al., 2007; Watson et al., 2014, 2016). These networks are evolved so as to produce fit phenotypes, with the general assumption being that evolution of these networks happens in a conventional manner - via selection acting on random variation amongst the molecular regulators that make up these networks. However, developmental processes involve complex interactions that can introduce correlations between phenotypic traits, causing some traits to co-vary, and creating patterns of expression by molecular regulators that are partially non-random. This covariation in expression means that there is potential for selection to strengthen correlations in expression between molecular regulators. Kouvaris et al. (2017) explored this possibility and found Hebbian patterns of change in interactions, just as we have identified in ecological networks. These changes to interactions alter the variability of regulated expression at the system level, which in turn shapes the selective pressures on the GRN. These findings, showing that both ecological community networks and GRNs may exhibit correlational learning and distributed associative learning processes, imply that there may be useful transfer of knowledge between these dynamically similar systems, and (perhaps more interestingly) that the principles we identify may be useful in other biological systems and at multiple levels of biological organisation - both in cases where the network as a whole are within a single evolutionary units (as in Kouvaris et al., 2017) and in cases where it is not (as here, see also Watson and Szathmáry, 2016).

6.5 Functional integration of species and evolutionary transitions

Our finding, that collective organisation emerges from lower levels of selection in ecological networks, contributes to work on the origins of biological organisation and evolutionary transitions in levels of Darwinian individuality (Watson et al., 2016; Watson and Szathmáry, 2016; Watson and Thies, 2019). One of the problematic aspects of the transitions that make them difficult to accommodate in conventional thinking is that the establishment of a new evolutionary unit requires many complex and coordinated changes to rescale the variability, heritability and selectability of the collective - e.g. particle plasticity, niche construction and new modes of reproduction that solve the 'homogeneous fitness heterogeneous function' problem (Watson and Thies, 2019). Once a new evolutionary unit has been established, selection at this new level can explain all sorts of adaptations to its component subunits that support or stabilise the unit in this way - but before that new evolutionary unit exists, selection at this new level cannot be a causal factor in creating these changes. This seems to leave the only available conclusion being that the changes necessary for a transition to occur can only occur by chance (and maybe thereafter be adapted for the purpose of stabilising the transition).

The work presented in this thesis, in contrast, shows us that distributed associative learning can create stable functional organisation in a system that is not (yet) an evolutionary unit. The ability to find better solutions to constraint problems such as resource allocation problems enables solutions to division of labour problems and provides selective access to fitness components that belong to the collective - a key prerequisite for a transition (Watson and Thies, 2019) - before the new unit is established. This then has the potential to provide selective incentive for particles to maintain and stabilise the collective and complete the transition, and potentially participate in the next level of organisation (Watson and Szathmáry, 2016). This functional integration of units of selection (Bouchard and Huneman, 2013) resulting in modified selective pressure on evolutionary units) is a component of an egalitarian evolutionary transition (Queller, 1997) - whereby evolutionary entities differentiate functional roles (such as via speciation) before forming a new evolutionary unit (Watson et al., 2016). (This contrasts with fraternal transitions, whereby entities change the scale of the evolutionary unit first and then differentiate in their functional roles.)

6.6 Correlational consistency between local and global optima in ecological networks

Chapter 5 demonstrates the ability of coevolution in ecological networks to drive self-modelling behaviour and evolve (learn) an enlarged basin of attraction for a unique community composition that maximally satisfies interspecific constraints. The coevolving community network achieves this is by aggregating correlational information from very many local optima, and forming a new associative memory from these correlations. In the models we have tested, constraint satisfaction problems are constructed through pairwise interactions, and the global optima for constraint satisfaction has the property that it is consistent with correlations between units at local optima. This condition is common in many models of convergent dynamical processes in constraint networks built from pairwise interactions (Watson et al., 2011b,a; Watson and Szathmáry, 2016) and we suspect it may be a feature of natural communities, but this should be tested. Note that it is only the specific properties of the evolved attractors which we question here; the general finding, that ecological community networks self-model and generalise over their experiences and not sensitive to this condition.

6.7 A framework for examining eco-evo interactions

We have introduced here a simple framework for examining feedback between ecological and evolutionary dynamics, and how the interplay of these can affect system-level behaviour. Eco-evo dynamics are often treated naively as separate processes with limited interaction or possibility for feedback but, in many cases, changes in the allelic composition of a population and species sorting can happen on overlapping timescales. Of course, there are existing models for the study of eco-evo dynamics (Govaert et al., 2018), but the one we introduce is perhaps the simplest model capable of demonstrating the potential of ecological communities to perform associative learning behaviours. The use of such a simple model facilitates understanding the drivers of observed behaviour and eases parametrisation for demonstrating the specific behaviours we are interested in.

We identify several aspects of the model that should be examined in future work to increase our understanding and general applicability of our result.

6.7.1 Ecological constraints in simulation models

An assumption in our simulation models is the presence of normalising constraints that prevent selection from eliminating all competition between species. This has the effect that any decrease in competition between two high density competitors is counterbalanced by increased competition between those species and other community members. The effect of unconstrained changes to interactions is solely to make high density species more cooperative with one another, rather than making them dependent on the simultaneous high density state of specific species (and hence less fit in some conditions). Given this, it is important for future work to consider how alternative methods of modelling these constraints impact the behaviours we have illustrated in our simulation models. It would be useful to explore, for example, whether the observed behaviours are maintained in scenarios modelling the evolution of species resource utilisation profiles (e.g. in response to niche shifts away from competitors).

6.7.2 Simplifying assumptions in Lotka-Volterra dynamics in simulation models

Our simulation models have only considered interactions between trophic competitors, as might be found in plant, or certain microbial or herbivorous animal communities. We have not considered mutualistic or predator-prey dynamics, and all our simulation work has been carried out under the assumption of symmetrical interactions. However, our analysis in Chapter 2 where we show that evolved changes in interactions are Hebbian, does not rely upon this assumption of symmetry (discussed in more detail in Chapter 2). 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 up on this symmetry. Very many important collective behaviours take place in non-symmetric neural networks, many of which also incorporate Hebbian learning (including classification, clustering, and data compression; Rumelhart and McClelland, 1987; O'Reilly and Munakata, 2000). Nevertheless, relaxing this assumption during future modelling is an important direction for future research.

6.8 Other areas of future work

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 (Adams, 1998; Fernando et al., 2012).

Finally, we note the potential of pre-existing constraints between species (e.g. due to relatedness) to be reinforced by hebbian reinforcement and amplify phylogenetic effects on species distribution at the community level.

6.9 Summary

Species interactions evolve as a function of their cooccurrence. This process is isomorphic to the unsupervised learning processes in models of correlational learning. Consequently, conditions exist in which ecological community networks exhibit distributed associative learning behaviours, just as in suitably trained neural networks. These learning mechanisms impart ecological community networks with many of the same behaviours as neural networks. Note that the formal link between the disciplines does not depend on the specific scenarios we simulate when investigating ecological memory and constraint satisfaction behaviours, nor on the simplifications that are necessary for the simulation models.

Until now there has not been an established framework linking selective pressures at lower levels of organisation to distributed associative learning at the community level. Without such a framework it is possible that learning processes in the evolution of ecological community networks might be overlooked. Similarly, given the isomorphism between these processes and those in neural networks, theoretical tools may already exist to address issues of ecological organisation, but remain unused. We have constructed our experiments with the aim of drawing attention to the potential of evolution by natural selection to organise community behaviours in a consistent manner, and with powerful effects. We hope by doing so to encourage further theoretical investigation and empirical studies. Natural ecosystems are not likely to be solving actual Sudoku puzzles, but our findings here suggest that their adaptive capacity is likely to be far greater than previously imagined.

Appendix A

Additional material for Chapter 2

A.1 Additional figure

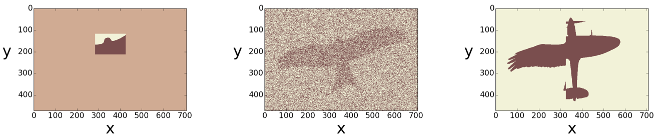

(A) Inputs to trained networks

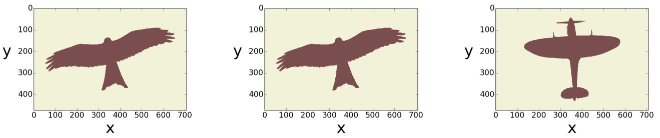

(B) Outputs from Hopfield network trained with Hebbs rule

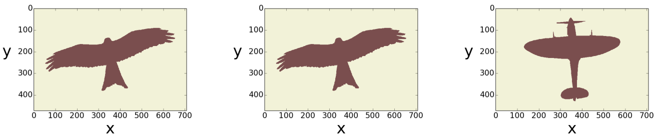

(C) Outputs from Hopfield network trained with Equation 2.4 from this thesis

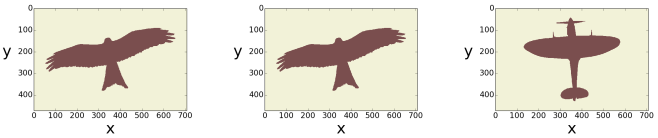

(D) Outputs from Hopfield network trained with Equation 7 from Weissman and Barton (2012)

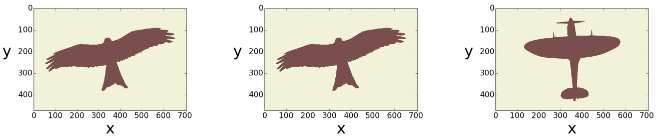

(E) Outputs from Hopfield network trained with Equation 12 from Neher et al. (2010)

Figure A.1: Equivalent dynamical behaviours resulting from training a Hopfield network using Hebb's rule and the equations discussed in this chapter (Hopfield, 1982; Hebb, 1949). Row A: Input patterns to Hopfield network. Patterns specified so as to demonstrate the ability perform pattern reconstruction, error or noise correction and classification of an ambiguous input. Row B-E: Corresponding outputs from a Hopfield network trained to recall bird and plane patterns using Hebbian learning, Equation 2.4 from this thesis, Equation 7 from Weissman and Barton (2012) and Equation 12 from (Neher et al., 2010), respectively. All training mechanisms produce the same effect - i.e., each is able to produce a distributed associative memory in the trained neural network.

Appendix B

Additional material for Chapter 3

B.1 Additional methods: Normalisation; variable environments; measuring ecological attractors; model parameters

B.1.1 Normalisation methods

In each evolutionary step, all interaction terms in $\Omega(t)$ are updated by natural selection according to Equation 3.2 to produce $\Omega'(t)$ and then renormalised to produce $\Omega(t+1)$. Sinkhorn-Knopp normalisation (Knight, 2008) 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:

$$M(k+1) = \text{column\_norm}(\text{row\_norm}(M(k))) \tag{B.1}$$

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

$$\text{row\_norm}(m_{ij(i \neq j)}) = \frac{m_{ij}(k)}{\sum_{j=1(j \neq i)}^{N} m_{ij}(k)} \tag{B.2}$$

and

$$\text{column\_norm}(m_{ij(i \neq j)}) = \frac{m_{ij}(k)}{\sum_{i=1(i \neq j)}^{N} m_{ij}(k)} \tag{B.3}$$

B.1.2 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

TABLE B.1: Simulation parameters for Chapter 3

Name	Symbol	Value
Community size (number of species)	N	400
Initial values for interspecific competition in matrix E	β	-0.2
Strength of competitive interactions in U (between incompatible species)	γ	0.5
Base carrying capacity	k	10.0
Variation in carrying capacity	α	0.1
Growth rate for all species	m	0.5
Ecological time (iterations of LV equation between resets)	τ	2 x 10^5
Beneficial mutation rate	μ	10^{-6}
Range of values for species densities during randomisation	-	[0,0.01]
Number of evolutionary changes applied in each environment before switching	T	1
Constant of proportionality in selection-limited evolution	g	0.01

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 Figure 3.2 in the main text.

B.1.3 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 (Hopfield, 1982; Watson et al., 2011a, 2014). 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. $E' = 2\bar{E} - E$).

B.1.4 Model parameters

The simulation in Chapter 3 uses the parameters shown in Table B.1.

Interaction coefficients are initialised as follows:

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

$Q_i = \sum_{j=1(j\neq i)}^{N} \omega_{ij}(t=0)$, 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 B.2.

B.2 Equivalence of Hebbian and evolved changes in more complex cases

In the main text the rate of adaptation of each interspecific interaction coefficient is modelled with Equation 3.2 corresponding to the case where there is no interference between simultaneously segregating alleles at different loci. In Chapter 2 we discuss how 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 B.1 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.

B.3 Response to environmental forcing that is *not* similar to environments experienced during evolution.

Figure B.2 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.

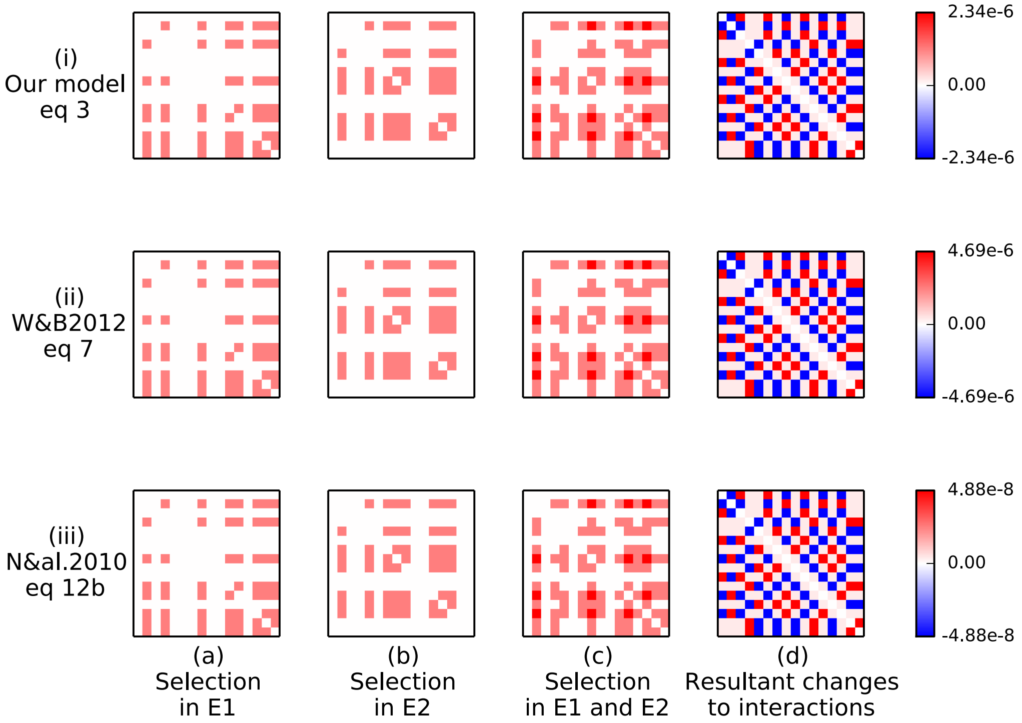

FIGURE B.1: Evolved changes to interactions are Hebbian in more complex population conditions. Change in interactions between the first 16 species are shown under evolution in a changing environment. Rate of adaptation is controlled by our equation (top row), that from Weissman and Barton (2012) (middle row) and Neher et al. (2010) (bottom row). a-b) the change in interactions due to direct selection effects (see Fig. 3.4 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 (Figure 3.4.e. 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.

B.4 Development and breakdown of multiple attractors over long evolutionary timescales.

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 (Watson et al., 2011a), and positive feedback causes one (slightly stronger) attractor to outcompete the other (Figure B.3).

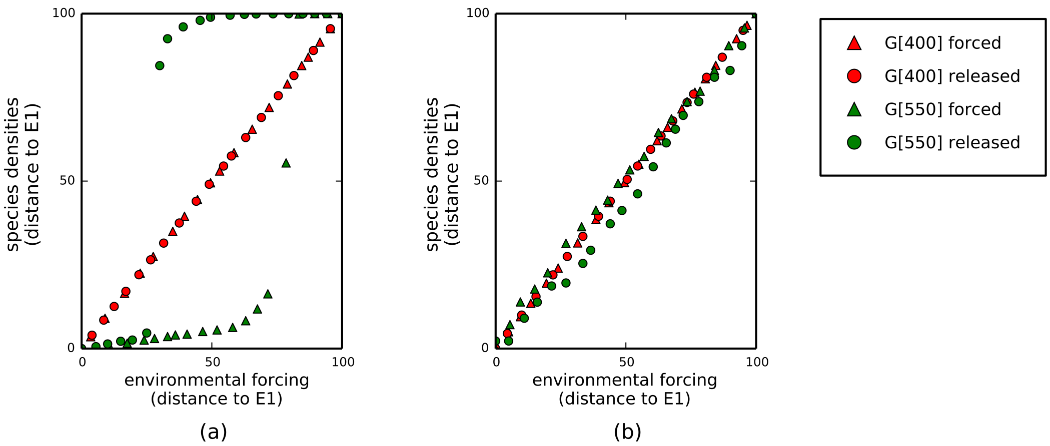

FIGURE B.2: Response to environmental forcing in different directions. a) Environmental forcing that is similar to environments experienced during evolution (i.e. toward E_2, see thumbnail pictogram), b) Environmental forcing that is not similar to environments experienced during evolution (i.e. toward an arbitrary pattern, see thumbnail pictogram).

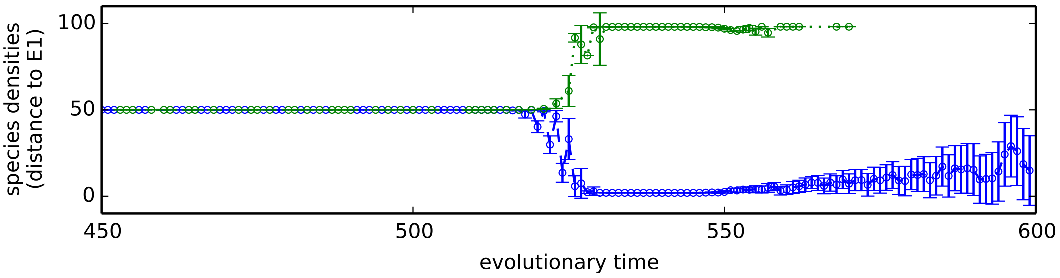

FIGURE B.3: In the long term the two-attractor system breaks down. Monte Carlo sampling of the ecological attractor states from random initial species densities during evolutionary time. Initially, all attractor states contain species densities that are only minor deviations from the default attractor (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 $\sim$575 generations), this attractor also degrades, i.e. the equilibrium magnitudes no longer match the original target closely (Figure 3.8).

B.5 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 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’; Beisner et al., 2003) - for example, 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 C

Additional material for Chapter 4

C.1 Additional methods

C.1.1 Equivalence of ecological scenario with the eight queens problem

In experiment one, the set of constraints on the community network resulting from $\{T_1, ..., T_4\}$ are mathematically equivalent to the constraints needed to specify the N queens puzzle from chess. In this puzzle N queens must be placed on a $N \times N$ chessboard such that all avoid interference with one-another. Constraints arising from T_1 and T_2 correspond to the constraints regarding rows and columns, T_3 and T_4 correspond to those regarding diagonals. A result of this equivalence is that for the experiment described in chapter 4, the indices for two species i and j for which $u_{ij} = 0$, are also indices for chess squares at which two queens can be placed without interference.

C.1.2 Dynamics in control experiment

During experiments without disturbance, species densities are randomised only at the beginning of the simulation. The result of this is that community composition is comprised of a fixed subset of species throughout simulation. This causes interactions between those species present to be very strongly affected (made much less competitive) by evolutionary dynamics. To prevent interactions becoming cooperative and destabilising the system we cap values in omega at 0. Simulations with disturbance do not experience this effect in the timescales simulated, and values in Ω remain negative throughout simulation.

C.1.3 Parameters

The simulation in Chapter 4 uses the parameters shown in Table C.1.

Table C.1: Simulation parameters for Chapter 4

Name	Symbol	Value
Community size (number of species)	N	64
Initial values for interspecific competition in matrix E	β	-0.4
Strength of competitive interactions in U (between incompatible species)	γ	-0.5
Base carrying capacity	k	3.0
Variation in carrying capacity	α	1.0
Growth rate for all species	m	0.5
Ecological time (iterations of LV equation between resets)	Et	10^4
Beneficial mutation rate	μ	2 x 10^{-3}
Range of values for species densities during randomisation	-	[0,0.01]
Constant of proportionality in selection-limited evolution	g	0.01

C.2 Additional figures

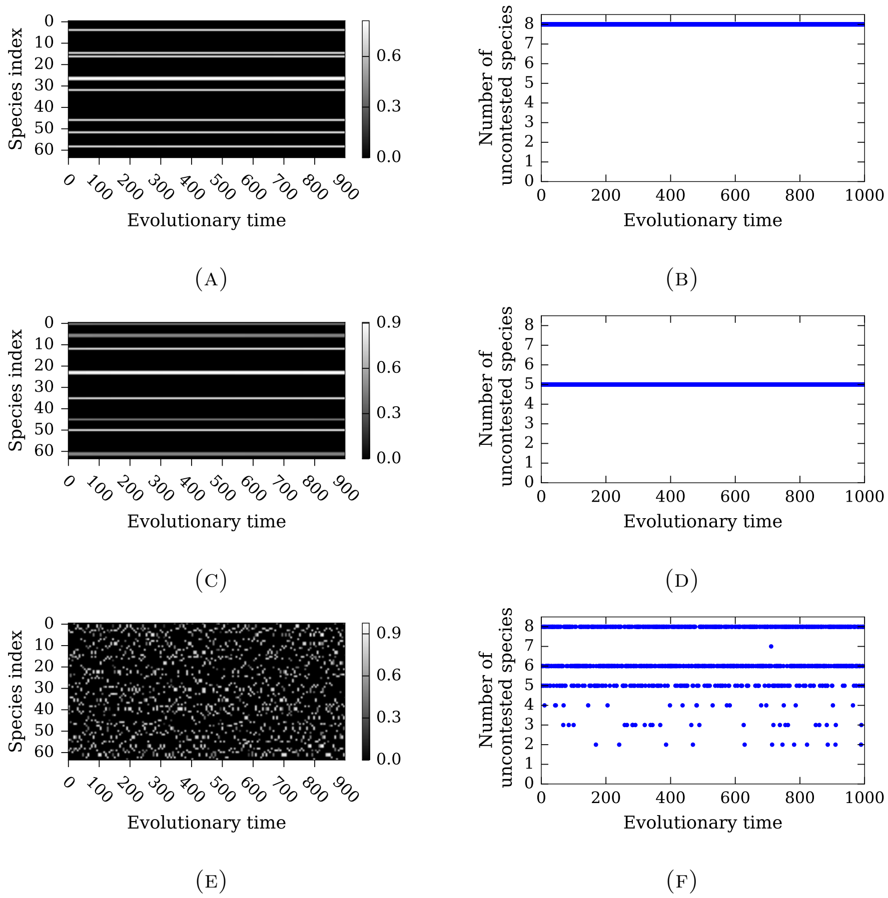

Figure C.1: State compositions at the end points of ecological dynamics over first 1000 cycles of simulation (left panels) and the corresponding number of uncontested species (right panels). The top two rows show example from simulations *without* disturbance: The top row (a & b) shows an example where the community stabilised at a composition with eight non-competing species and the middle row (c & d) shows an example where the community stabilised at a composition with five non-competing species. The bottom row (e & f) shows an example simulation *with* disturbance, where the system converges to a different stable state at each time step, with different stable states corresponding to different numbers of uncontested species.

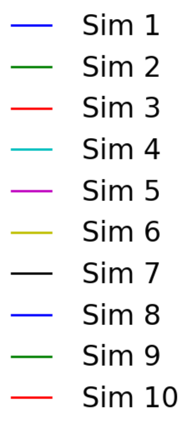

FIGURE C.2: Legend labels for figures 4.3 and 5.1.

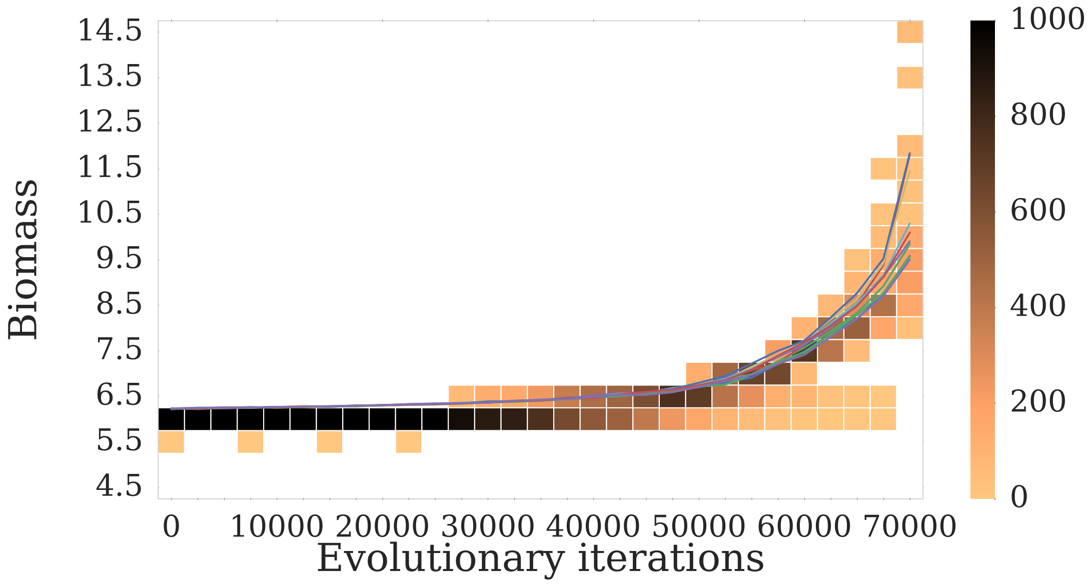

FIGURE C.3: Change in biomass over simulation. Data from 10 simulations without disturbance. Each lines shows, for a single simulation, the mean biomass of the state composition reached from 100 random community compositions. Heat maps show the distribution of the same metric across all simulations. As species are able to evolve their interspecific interactions with trophic competitors there is an increase in the total biomass of the system throughout simulation. In our main analysis we control for this effect by subsequently re-stabilising the state composition in Ω at t(0), controlling for increases in biomass due to evolved changes in interactions (Figure 4.3).

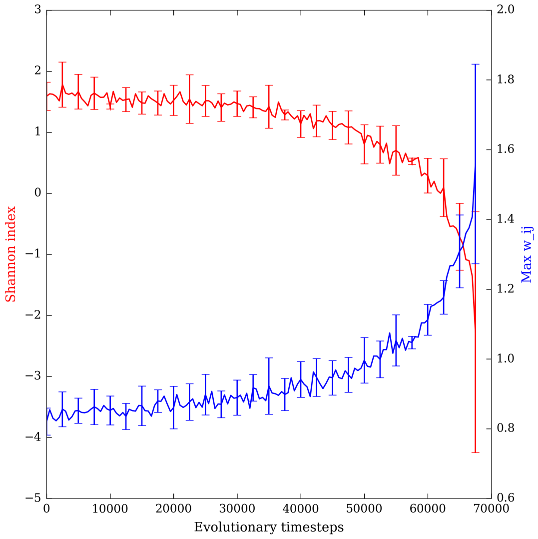

Figure C.4: Shannon diversity throughout simulation. Shannon diversity (red line, calculated as $-\sum_i^N x_i ln(x_i)$ decreases throughout the simulation as a small number of species begin to dominate the community network. This decrease in diversity coincides with decreased cooperation between those species in the stable states that the system develops attractors for, as show by changes in the maximum values in Ω (blue).

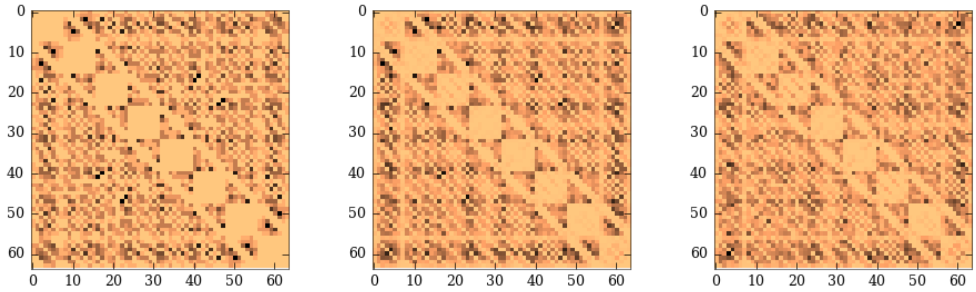

(A) Species cooccurance in the 92 optimal SS

(B) Species cooccurance in 1000 SS at t(0).

(C) As b, but with all optimal SS removed

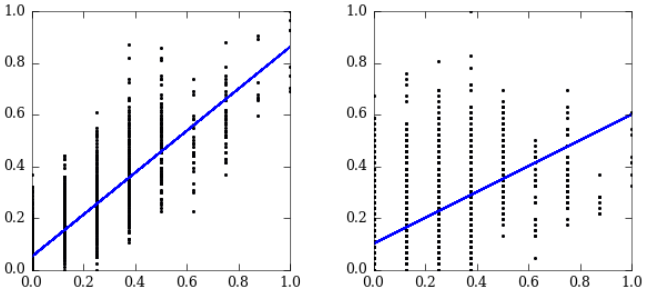

(D) Correlation between cocurrance in a) and b)

(E) Correlation between cocurrance in a) and c)

Figure C.5: Patterns of species cooccurrence between a) all 92 stable states (SS) that maximise the number of compatible species present; b) stable state compositions from 1000 random community compositions at t(0); and c) stable state compositions from random community compositions at t(0), but with those compositions that maximise compatible species present removed. d) There is good correlation between cooccurrence in a and b ($R^2 = 75.6$). e) Even with the removal of state compositions that maximise number of compatible species present, there is still correlation between species cooccurrence in a and c ($R^2 = 34.4$).

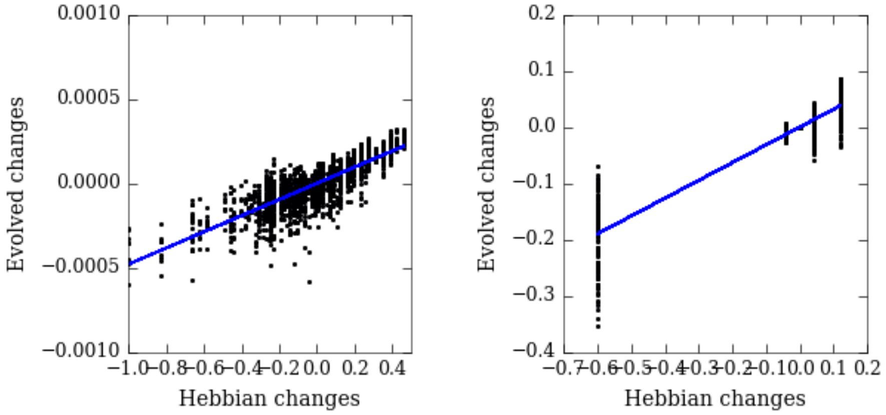

(A) First 1000 time steps (B) Whole simulation

Figure C.6: Evolved patterns of change in network interactions correlate with Hebbian patterns of change. This figure shows correlation between weight matrices shown in Figure 4.6. a) Data points show the change in an interaction in a network trained with Hebbian learning to recall the 96 attractors that maximally satisfy the constraints in scenario A versus the evolved pattern of change over the first 1000 time steps of simulation ($R^2 = 67.6\%$). b) Change in an interaction network trained with Hebbian learning to recall the 4 attractors that the evolving system eventually learns, versus the evolved pattern of change over a whole simulation ($R^2 = 83.8\%$).

Appendix D

Additional material for Chapter 5

D.1 Additional methods: parameters and Sudoku test suite

The simulation in Chapter 5 uses the parameters shown in Table D.1.

D.1.1 Sudoku test suite

We use a test suite of three Sudoku problems. We select the two hardest puzzles from the test suite published by Mantere and Koljonen (2008) (accessible online http://lipas.uwasa.fi/ timan/-sudoku/ [last accessed 12 June 2018]). We add puzzle ‘AI Escargot’, the hardest puzzle in a test suite by Inkala (2006), who claimed it as the world’s most challenging puzzle.

TABLE D.1: Simulation parameters for Chapter 5

Name	Symbol	Value
Community size (number of species)	N	729
Initial values for interspecific competition in matrix E	β	-0.2
Strength of competitive interactions in U (between incompatible species)	γ	-0.5
Base carrying capacity	k	10.0
Variation in carrying capacity	α	3.0
Growth rate for all species	m	0.5
Ecological time (iterations of LV equation between resets)	Et	2 x 10^5
Beneficial mutation rate	μ	10^{-6}
Range of values for species densities during randomisation	-	[0,0.01]
Constant of proportionality in selection-limited evolution	g	0.01

D.2 Additional figures

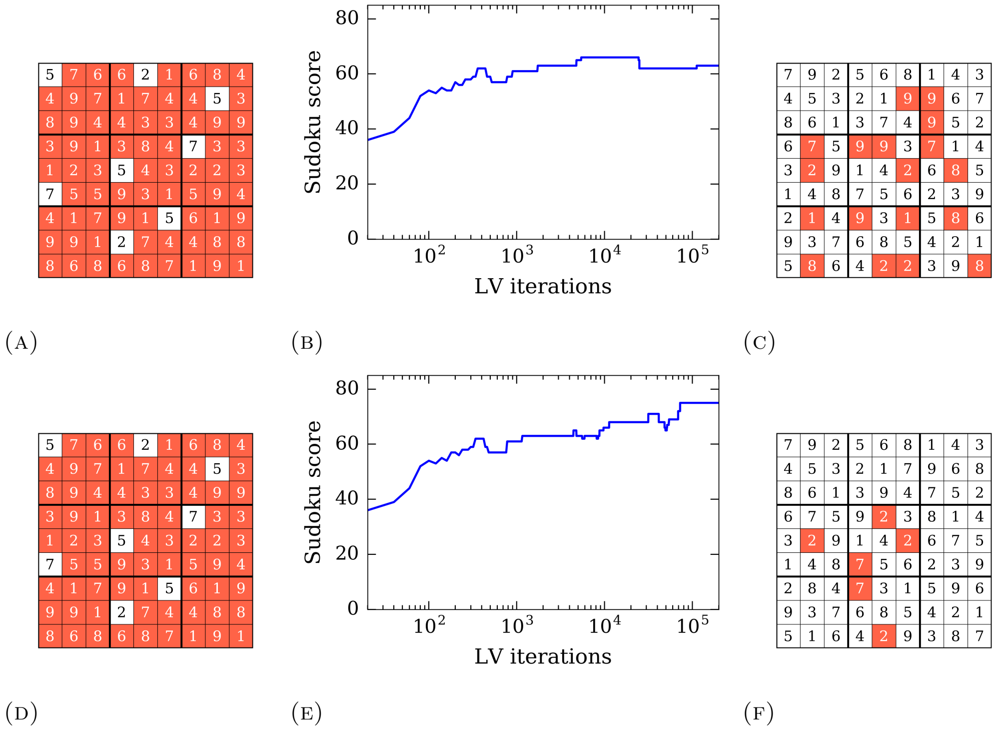

(A)

5	7	6	6	2	1	6	8	4
4	9	7	1	7	4	4	5	3
8	9	4	4	3	3	4	9	9
3	9	1	3	8	4	7	3	3
1	2	3	5	4	3	2	2	3
7	5	5	9	3	1	5	9	4
4	1	7	9	1	5	6	1	9
9	9	1	2	7	4	4	8	8
8	6	8	6	8	7	1	9	1

(C)

7	9	2	5	6	8	1	4	3
4	5	3	2	1	9	9	6	7
8	6	1	3	7	4	9	5	2
6	7	5	9	9	3	7	1	4
3	2	9	1	4	2	6	8	5
1	4	8	7	5	6	2	3	9
2	1	4	9	3	1	5	8	6
9	3	7	6	8	5	4	2	1
5	8	6	4	2	2	3	9	8

(D)

5	7	6	6	2	1	6	8	4
4	9	7	1	7	4	4	5	3
8	9	4	4	3	3	4	9	9
3	9	1	3	8	4	7	3	3
1	2	3	5	4	3	2	2	3
7	5	5	9	3	1	5	9	4
4	1	7	9	1	5	6	1	9
9	9	1	2	7	4	4	8	8
8	6	8	6	8	7	1	9	1

(F)

7	9	2	5	6	8	1	4	3
4	5	3	2	1	7	9	6	8
8	6	1	3	9	4	7	5	2
6	7	5	9	2	3	8	1	4
3	2	9	1	4	2	6	7	5
1	4	8	7	5	6	2	3	9
2	8	4	7	3	1	5	9	6
9	3	7	6	8	5	4	2	1
5	1	6	4	2	9	3	8	7

FIGURE D.1: Sample ecological dynamics using puzzle s15a, at t(0), before simulation (a-c); and t(250), after simulation (d-f). a) A random community composition corresponds to a Sudoku grid with very many conflicting cells (in red). b) Assembly rules in the unevolved community form a stable state that corresponds to c) a Sudoku grid with fewer, but not zero, conflicts. d) - f) At the end of simulation evolutionary changes to assembly rules mean the community now converges to a stable state with six conflicts. This puzzle is the second hardest in the test suite. From ten simulation experiments conducted with puzzle s15a, two runs resulted in community network with dynamical attractors for the solution state.

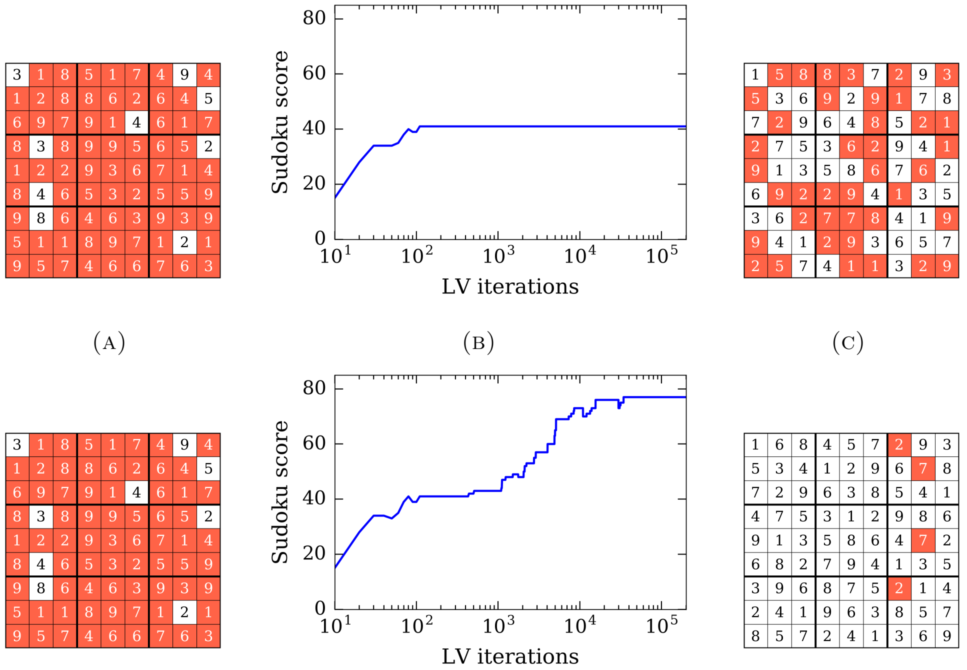

3	1	8	5	1	7	4	9	4
1	2	8	8	6	2	6	4	5
6	9	7	9	1	4	6	1	7
8	3	8	9	9	5	6	5	2
1	2	2	9	3	6	7	1	4
8	4	6	5	3	2	5	5	9
9	8	6	4	6	3	9	3	9
5	1	1	8	9	7	1	2	1
9	5	7	4	6	6	7	6	3

(A)

(B)

1	5	8	8	3	7	2	9	3
5	3	6	9	2	9	1	7	8
7	2	9	6	4	8	5	2	1
2	7	5	3	6	2	9	4	1
9	1	3	5	8	6	7	6	2
6	9	2	2	9	4	1	3	5
3	6	2	7	7	8	4	1	9
9	4	1	2	9	3	6	5	7
2	5	7	4	1	1	3	2	9

(C)

3	1	8	5	1	7	4	9	4
1	2	8	8	6	2	6	4	5
6	9	7	9	1	4	6	1	7
8	3	8	9	9	5	6	5	2
1	2	2	9	3	6	7	1	4
8	4	6	5	3	2	5	5	9
9	8	6	4	6	3	9	3	9
5	1	1	8	9	7	1	2	1
9	5	7	4	6	6	7	6	3

(D)

(E)

1	6	8	4	5	7	2	9	3
5	3	4	1	2	9	6	7	8
7	2	9	6	3	8	5	4	1
4	7	5	3	1	2	9	8	6
9	1	3	5	8	6	4	7	2
6	8	2	7	9	4	1	3	5
3	9	6	8	7	5	2	1	4
2	4	1	9	6	3	8	5	7
8	5	7	2	4	1	3	6	9

(F)

FIGURE D.2: Ecological dynamics using puzzle AI Escargot, at t(0), before simulation (a-c); and t(200), after simulation (d-f). a) A random community composition corresponds to a Sudoku grid with very many conflicting cells (in red). b) Assembly rules in the unevolved community form a stable state that corresponds to c) a Sudoku grid with fewer conflicts. d) - f) At the end of simulation evolutionary changes to assembly rules mean the community now converges to a stable state with four conflicts. This puzzle is the hardest in the test suite and, at the time it was published, considered to amongst the hardest known Sudoku puzzles (Inkala, 2006).

References

David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. A learning algorithm for Boltzmann machines. *Cognitive science*, 9(1):147–169, 1985.

Paul Adams. Hebb and Darwin. *Journal of theoretical biology*, 195(4):419–438, 1998.

Madhur Anand, Andrew Gonzalez, Frédéric Guichard, Jurek Kolasa, and Lael Parrott. Ecological systems as complex systems: challenges for an emerging science. *Diversity*, 2(3):395–410, 2010.

David G Angeler, Cristina Trigal, Stina Drakare, Richard K Johnson, and Willem Goedkoop. Identifying resilience mechanisms to recurrent ecosystem perturbations. *Oecologia*, 164(1): 231–241, 2010.

Robert Axelrod and William D Hamilton. The evolution of cooperation. *science*, 211(4489): 1390–1396, 1981.

Jan Baedke. The epigenetic landscape in the course of time: Conrad hal waddington's methodological impact on the life sciences. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 44(4):756–773, 2013.

Carolyn M Beans. The case for character displacement in plants. *Ecology and evolution*, 4(6): 852–865, 2014.

Beatrix E Beisner, Daniel T Haydon, and Kim Cuddington. Alternative stable states in ecology. *Frontiers in Ecology and the Environment*, 1(7):376–382, 2003.

Frédéric Bouchard and Philippe Huneman. *From groups to individuals: evolution and emerging individuality.* MIT Press, 2013.

Paul Bourgine, David Chavalarias, Edith Perrier, Frederic Amblard, Francois Arlabosse, Pierre Auger, Jean-Bernard Baillon, Olivier Barreteau, Pierre Baudot, Elisabeth Bouchaud, et al. French roadmap for complex systems 2008-2009. *arXiv preprint arXiv:0907.2221*, 2009.

William L Brown and Edward O Wilson. Character displacement. *Systematic zoology*, 5(2): 49–64, 1956.

Stephen R. Carpenter. Alternate states of ecosystems: evidence and some implications. In Simon Levin, editor, *Ecology: achievement and challenge*, pages 357–83, Oxford, 2001. Blackwell Science, Blackwell Science.

Stephen R Carpenter, Donald Ludwig, and William A Brock. Management of eutrophication for lakes subject to potentially irreversible change. *Ecological applications*, 9(3):751–771, 1999.

Ted J Case, Robert D Holt, Mark A McPeek, and Timothy H Keitt. The community context of species' borders: ecological and evolutionary perspectives. *Oikos*, 108(1):28–46, 2005.

Jeannine Cavender-Bares, Kenneth H Kozak, Paul VA Fine, and Steven W Kembel. The merging of community ecology and phylogenetic biology. *Ecology letters*, 12(7):693–715, 2009.

Frederic E Clements. Nature and structure of the climax. *Journal of ecology*, 24(1):252–284, 1936.

Frederic Edward Clements. *Plant succession: an analysis of the development of vegetation.* Carnegie Institution of Washington, Washington, DC, United States, 1916.

Edward F Connor and Daniel Simberloff. The assembly of species communities: chance or competition? *Ecology*, pages 1132–1140, 1979.

Roger Cropp and Albert Gabric. Ecosystem adaptation: Do ecosystems maximize resilience? *Ecology*, 83(7):2019–2026, 2002.

Vasilis Dakos, Marten Scheffer, Egbert H van Nes, Victor Brovkin, Vladimir Petoukhov, and Hermann Held. Slowing down as an early warning signal for abrupt climate change. *Proceedings of the National Academy of Sciences*, 105(38):14308–14312, 2008.

Vasilis Dakos, Stephen R Carpenter, Egbert H van Nes, and Marten Scheffer. Resilience indicators: prospects and limitations for early warnings of regime shifts. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 370(1659):20130263, 2015.

Lawrence A David, Corinne F Maurice, Rachel N Carmody, David B Gootenberg, Julie E Button, Benjamin E Wolfe, Alisha V Ling, A Sloan Devlin, Yug Varma, Michael A Fischbach, et al. Diet rapidly and reproducibly alters the human gut microbiome. *Nature*, 505(7484):559–563, 2014.

Richard Dawkins. *The selfish gene.* Oxford university press, 1976.

Tamar Dayan and Daniel Simberloff. Ecological and community-wide character displacement: the next generation. *Ecology Letters*, 8(8):875–894, 2005.

Peter C de Ruiter, Anje-Margriet Neutel, and John C Moore. Energetics, patterns of interaction strengths, and stability in real ecosystems. *Science*, 269(5228):1257–60, 1995.

Jared M Diamond. Assembly of species communities. In Martin L Cody and Jared M Diamond, editors, *Ecology and evolution of communities*, page 342, Harvard, 1975. Harvard University Press.

Ulf Dieckmann, Michael Doebeli, Johan AJ Metz, and Diethard Tautz. *Adaptive speciation.* Cambridge University Press, 2004.

J Doyne Farmer. A Rosetta stone for connectionism. *Physica D: Nonlinear Phenomena*, 42(1): 153–187, 1990.

J Doyne Farmer and Andrew W Lo. Frontiers of finance: evolution and efficient markets. *Proceedings of the National Academy of Sciences*, 96(18):9991–9992, 1999.

Karoline Faust and Jeroen Raes. Microbial interactions: from networks to models. *Nature Reviews Microbiology*, 10(8):538–550, 2012.

Chrisantha Fernando, Eörs Szathmáry, and Phil Husbands. Selectionist and evolutionary approaches to brain function: a critical appraisal. *Frontiers in computational neuroscience*, 6, 2012.

Chrisantha T Fernando, Anthony ML Liekens, Lewis EH Bingle, Christian Beck, Thorsten Lenser, Dov J Stekel, and Jonathan E Rowe. Molecular circuits for associative learning in single-celled organisms. *Journal of The Royal Society Interface*, 6(34):463–469, 2009.

Deborah L Finke and William E Snyder. Niche partitioning increases resource exploitation by diverse communities. *Science*, 321(5895):1488–1490, 2008.

John T Finn. Measures of ecosystem structure and function derived from analysis of flows. *Journal of theoretical Biology*, 56(2):363–380, 1976.

Carl Folke. Resilience: the emergence of a perspective for social–ecological systems analyses. *Global environmental change*, 16(3):253–267, 2006.

Carl Folke, Steve Carpenter, Thomas Elmqvist, Lance Gunderson, Crawford S Holling, and Brian Walker. Resilience and sustainable development: building adaptive capacity in a world of transformations. *AMBIO: A journal of the human environment*, 31(5):437–440, 2002.

Kristoffer Forslund, Shinichi Sunagawa, Jens Roat Kultima, Daniel R Mende, Manimozhiyan Arumugam, Athanasios Typas, and Peer Bork. Country-specific antibiotic use practices impact the human gut resistome. *Genome research*, 23(7):1163–1169, 2013.

GF Fussmann, M Loreau, and PA Abrams. Eco-evolutionary dynamics of communities and ecosystems. *Functional Ecology*, 21(3):465–477, 2007.

Belinda Gallardo and David C Aldridge. Is Great Britain heading for a Ponto–Caspian invasional meltdown? *Journal of Applied Ecology*, 52(1):41–49, 2015.

Gilberto C Gallopín. Linkages between vulnerability, resilience, and adaptive capacity. *Global environmental change*, 16(3):293–303, 2006.

Georgii Frantsevitch Gause. Experimental analysis of Vito Volterra's mathematical theory of the struggle for existence. *Science*, 79(2036):16–17, 1934.

Henry A Gleason. The individualistic concept of the plant association. *Bulletin of the Torrey Botanical Club*, pages 7–26, 1926.

Michael Golinski, Chris Bauch, and Madhur Anand. The effects of endogenous ecological memory on population stability and resilience in a variable environment. *Ecological Modelling*, 212(3): 334–341, 2008.

Lynn Govaert, Emanuel A Fronhofer, Sébastien Lion, Christophe Eizaguirre, Dries Bonte, Martijn Egas, Andrew P Hendry, Ayana De Brito Martins, Carlos J Melián, Joost AM Raeymaekers, et al. Eco-evolutionary feedbacks—theoretical models and perspectives. *Functional Ecology*, 2018.

Gary R Graves and Nicholas J Gotelli. Assembly of avian mixed-species flocks in Amazonia. *Proceedings of the National Academy of Sciences*, 90(4):1388–1391, 1993.

Thilo Gross and Hiroki Sayama. *Adaptive networks - theory, models and applications*. Springer Berlin, Germany, 2009.

William Harcombe. Novel cooperation experimentally evolved between species. *Evolution*, 64 (7):2166–2172, 2010.

Donald O Hebb. *The organization of behavior: A neuropsychological theory*. New York: Wiley, 1949.

George E Heimpel, Lee E Frelich, Douglas A Landis, Keith R Hopper, Kim A Hoelmer, Zeynep Sezen, Mark K Asplen, and Kongming Wu. European buckthorn and Asian soybean aphid as components of an extensive invasional meltdown in North Americasian soybean aphid as components of an extensive invasional meltdown in north america. *Biological Invasions*, 12 (9):2913–2931, 2010.

RJ Hendry and JM McGlade. The role of memory in ecological systems. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 259(1355):153–159, 1995.

Paul AT Higgins, Michael D Mastrandrea, and Stephen H Schneider. Dynamics of climate and ecosystem coupling: abrupt changes and multiple equilibria. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 357:647–656, 2002.

Geoffrey E Hinton, Terrence Joseph Sejnowski, and Tomaso A Poggio. *Unsupervised learning: foundations of neural computation*. MIT press, 1999.

Crawford S Holling. Resilience and stability of ecological systems. *Annual review of ecology and systematics*, pages 1–23, 1973.

John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.

John J Hopfield. Searching for memories, Sudoku, implicit check bits, and the iterative use of not-always-correct rapid neural computation. *Neural computation*, 20(5):1119–1164, 2008.

John J Hopfield, David W Tank, et al. Computing with neural circuits - a model. *Science*, 233 (4764):625–633, 1986.

Jennifer B Hughes and Joan Roughgarden. Aggregate community properties and the strength of species' interactions. *Proceedings of the National Academy of Sciences*, 95(12):6837–6842, 1998.

George Evelyn Hutchinson. The niche: an abstractly inhabited hypervolume. *The ecological theatre and the evolutionary play*, pages 26–78, 1965.

Arto Inkala. AI Sudoku top 10. online, www.aisudoku.com [accessed 14/02/2016], November 2006.

Ferenc Jordán and István Scheuring. Network ecology: topological constraints on ecosystem dynamics. *Physics of Life reviews*, 1(3):139–172, 2004.

Philip A Knight. The Sinkhorn–Knopp algorithm: convergence and applications. *SIAM Journal on Matrix Analysis and Applications*, 30(1):261–275, 2008.

Kostas Kouvaris, Jeff Clune, Loizos Kounios, Markus Brede, and Richard A Watson. How evolution learns to generalise: Using the principles of learning theory to understand the evolution of developmental organisation. *PLoS computational biology*, 13(4):e1005358, 2017.

Georges Kunstler, Daniel Falster, David A Coomes, Francis Hui, Robert M Kooyman, Daniel C Laughlin, Lourens Poorter, Mark Vanderwel, Ghislain Vieilledent, S Joseph Wright, et al. Plant functional traits have globally consistent effects on competition. *Nature*, 529(7585):204, 2016.

Kevin Laland, John Odling-Smee, and Scott Turner. The role of internal and external constructive processes in evolution. *The Journal of physiology*, 592(11):2413–2422, 2014.

Kevin N Laland, John Odling-Smee, Marcus W Feldman, and Jeremy Kendal. Conceptual barriers to progress within evolutionary biology. *Foundations of science*, 14(3):195–216, 2009.

Richard Law and R Daniel Morton. Permanence and the assembly of ecological communities. *Ecology*, pages 762–775, 1996.

John H Lawton. Are there general laws in ecology? *Oikos*, pages 177–192, 1999.

Egbert Giles Leigh and Geerat Jacobus Vermeij. Does natural selection organize ecosystems for the maintenance of high productivity and diversity? *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1421):709–718, 2002.

Timothy M Lenton. Clarifying Gaia: regulation with or without natural selection. *Scientists Debate Gaia*, pages 15–25, 2004.

Timothy M Lenton and Marcel van Oijen. Gaia as a complex adaptive system. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1421):683–695, 2002.

Simon A Levin. Ecosystems and the biosphere as complex adaptive systems. *Ecosystems*, 1(5): 431–436, 1998.

Simon A Levin. Self-organization and the emergence of complexity in ecological systems. *Bioscience*, 55(12):1075–1079, 2005.

Simon A Levin. Evolution the ecosystem level: On the evolution of ecosystem patterns. *Contributions to science*, 7(1):11–16, 2011.

Simon A Levine. Fragile dominion, 1999.

Richard Charles Lewontin. The meaning of stability. diversity and stability in ecological systems: 13-24. In *Brookhaven Symposia in Biology. Brookhaven National Laboratory, Brookhaven, New York.[Links]*, 1969.

Siegrid Lowel and Wolf Singer. Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science*, 255(5041):209–212, 1992.

Robert MacArthur and Richard Levins. The limiting similarity, convergence, and divergence of coexisting species. *American naturalist*, pages 377–385, 1967.

Timo Mantere and Janne Koljonen. Sudoku research page http://lipas.uwasa.fi/ timan/sudoku/. Online, June 2008. URL `http://lipas.uwasa.fi/~timan/sudoku/`.

Blake Matthews, Anita Narwani, Stephen Hausch, Etsuko Nonaka, Hannes Peter, Masato Yamamichi, Karen E Sullam, Kali C Bird, Mridul K Thomas, Torrance C Hanley, et al. Toward an integration of evolutionary biology and ecosystem science. *Ecology Letters*, 14(7):690–701, 2011.

Robert M May. Some notes on estimating the competition matrix, a. *Ecology*, pages 737–741, 1975.

Margaret M Mayfield and Jonathan M Levine. Opposing effects of competitive exclusion on the phylogenetic structure of communities. *Ecology letters*, 13(9):1085–1093, 2010.

John Maynard Smith and Eörs Szathmáry. *The major transitions in evolution*. Oxford University Press, 1995.

Ernst Mayr. Wallace's line in the light of recent zoogeographic studies. *The Quarterly Review of Biology*, 19(1):1–14, 1944.

Connor McCabe, Richard A Watson, Jane Prichard, and Wendy Hall. The web as an adaptive network: coevolution of web behavior and web structure. In *Proceedings of the 3rd International Web Science Conference*, page 22. ACM, 2011.

Warren S McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4):115–133, 1943.

Robert P McIntosh. HA Gleason's 'individualistic concept' and theory of animal communities: a continuing controversy. *Biological Reviews*, 70(2):317–357, 1995.

Johan AJ Metz, Stefan AH Geritz, Géza Meszéna, Frans JA Jacobs, and Joost S Van Heerwaarden. Adaptive dynamics, a geometrical study of the consequences of nearly faithful reproduction. *Stochastic and spatial structures of dynamical systems*, 45:183–231, 1996.

AS Mikhailov, IV Mit'kov, and NA Sveshnikov. Molecular associative memory. *BioSystems*, 23 (4):291–295, 1990.

Bruce T Milne. Motivation and benefits of complex systems approaches in ecology. *Ecosystems*, 1(5):449–456, 1998.

Richard A Neher, Boris I Shraiman, and Daniel S Fisher. Rate of adaptation in large sexual populations. *Genetics*, 184(2):467–481, 2010.

William E Neill. The community matrix and interdependence of the competition coefficients. *American Naturalist*, pages 399–408, 1974.

VW Noonburg. A neural network modeled by an adaptive Lotka-Volterra system. *SIAM Journal on Applied Mathematics*, 49(6):1779–1792, 1989.

Scott L Nuismer and Luke J Harmon. Predicting rates of interspecific interaction from phylogenetic trees. *Ecology letters*, 18(1):17–27, 2015.

Samir Okasha. *Evolution and the levels of selection*. Oxford University Press, 2006.

Randall C O'Reilly and Yuko Munakata. *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. MIT press, 2000.

Greg Paperin, David G Green, and Suzanne Sadedin. Dual-phase evolution in complex adaptive systems. *Journal of The Royal Society Interface*, page rsif20100719, 2011.

Garry D Peterson. Contagious disturbance, ecological memory, and the emergence of landscape pattern. *Ecosystems*, 5(4):329–338, 2002.

Fábio C Poderoso and José F Fontanari. Model ecosystem with variable interspecies interactions. *Journal of Physics A: Mathematical and Theoretical*, 40(30):8723, 2007.

David M Post and Eric P Palkovacs. Eco-evolutionary feedbacks in community and ecosystem ecology: interactions between the ecological theatre and the evolutionary play. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1523):1629–1640, 2009.

Stephen R Proulx, Daniel EL Promislow, and Patrick C Phillips. Network thinking in ecology and evolution. *Trends in Ecology & Evolution*, 20(6):345–353, 2005.

Adam Prugel-Bennett and Jonathan L Shapiro. Statistical mechanics of unsupervised Hebbian learning. *Journal of Physics A: Mathematical and General*, 26(10):2343, 1993.

David C Queller. Cooperators since life began, 1997.

Jonathan L Richardson. The organismic community: resilience of an embattled ecological concept. *BioScience*, 30(7):465–471, 1980.

Joann Romano-Keeler and Jörn-Hendrik Weitkamp. Maternal influences on fetal microbial colonization and immune development. *Pediatric research*, 77(1-2):189–195, 2014.

J Roughgarden. *Theory of population genetics and evolutionary ecology: an introduction.* Macmillan New York NY United States 1979., 1979.

Jonathan Roughgarden. Evolution of niche width. *American Naturalist*, pages 683–718, 1972.

David E Rumelhart and James L McClelland. Parallel distributed processing, explotation in the microstructure of cognition-vol. 1: Foundations. *Computational Models of Cognition and Perception, Cambridge: MIT Press, 1987*, 1987.

Marten Scheffer, Steve Carpenter, Jonathan A Foley, Carl Folke, and Brian Walker. Catastrophic shifts in ecosystems. *Nature*, 413(6856):591–596, 2001.

Marten Scheffer, Jordi Bascompte, William A Brock, Victor Brovkin, Stephen R Carpenter, Vasilis Dakos, Hermann Held, Egbert H Van Nes, Max Rietkerk, and George Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009.

Thomas W Schoener. The newest synthesis: understanding the interplay of evolutionary and ecological dynamics. *Science*, 331(6016):426–429, 2011.

Ernst-Detlef Schulze and Harold A Mooney. *Biodiversity and ecosystem function.* Springer Science & Business Media, 2012.

Ashley Shade, Hannes Peter, Steven D Allison, Didier L Baho, Mercè Berga, Helmut Bürgmann, David H Huber, Silke Langenheder, Jay T Lennon, Jennifer BH Martiny, et al. Fundamentals of microbial community resistance and resilience. *Frontiers in microbiology*, 3, 2012.

Daniel Simberloff and Betsy Von Holle. Positive interactions of nonindigenous species: invasional meltdown? *Biological invasions*, 1(1):21–32, 1999.

J Maynard Smith, Richard Burian, Stuart Kauffman, Pere Alberch, John Campbell, Brian Goodwin, Russell Lande, David Raup, and Lewis Wolpert. Developmental constraints and evolution: a perspective from the mountain lake conference on development and evolution. *Quarterly Review of Biology*, pages 265–287, 1985.

Diane S Srivastava, Marc W Cadotte, A Andrew M MacDonald, Robin G Marushia, and Nicholas Mirotchnick. Phylogenetic diversity and the functioning of ecosystems. *Ecology letters*, 15(7): 637–648, 2012.

Phillip Staniczenko, Owen T Lewis, Nick S Jones, and Felix Reed-Tsochas. Structural dynamics and robustness of food webs. *Ecology Letters*, 13(7):891–899, 2010.

Eörs Szathmáry, Ferenc Jordán, and Csaba Pál. Can genes explain biological complexity? *Science*, 292(5520):1315–1316, 2001.

John N Thompson. The evolution of species interactions. *Science*, 284(5423):2116–2118, 1999.

John N Thompson. *The geographic mosaic of coevolution.* University of Chicago Press, Chicago, Il, United States, 2005.

John N Thompson, OJ Reichman, Peter J Morin, Gary A Polis, Mary E Power, Robert W Sterner, Carol A Couch, Laura Gough, Robert Holt, David U Hooper, et al. Frontiers of ecology: As ecological research enters a new era of collaboration, integration, and technological sophistication, four frontiers seem paramount for understanding how biological and physical processes interact over multiple spatial and temporal scales to shape the earth's biodiversity. *BioScience*, 51(1):15–24, 2001.

Martin M Turcotte, David N Reznick, and J Daniel Hare. The impact of rapid evolution on population dynamics in the wild: experimental test of eco-evolutionary dynamics. *Ecology letters*, 14(11):1084–1092, 2011.

Robert E Ulanowicz. *Growth and development: ecosystems phenomenology.* Springer-Verlag, New York, NY, United States, 1986.

Robert E Ulanowicz. Aristotelean causalities in ecosystem development. *Oikos*, pages 42–48, 1990.

Robert E Ulanowicz. Utricularia's secret: the advantage of positive feedback in oligotrophic environments. *Ecological Modelling*, 79(1):49–57, 1995.

Jiri Vohradsky. Neural model of the genetic network. *Journal of Biological Chemistry*, 276(39): 36168–36173, 2001.

Günter P Wagner, Mihaela Pavlicev, and James M Cheverud. The road to modularity. *Nature Reviews Genetics*, 8(12):921, 2007.

Alfred Russel Wallace. *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*, volume 1. Cambridge University Press, Cambridge, UK, 2011.

R.A. Watson and C. Thies. Are developmental plasticity, niche construction, and extended inheritance necessary for evolution by natural selection? The role of active phenotypes in the minimal criteria for darwinian individuality. In K Uller, Laland, editor, *Evolutionary Causation*. MIT Press, 2019.

Richard A Watson and Eörs Szathmáry. How can evolution learn? *Trends in ecology & evolution*, 31(2):147–157, 2016.

Richard A Watson, Christopher L Buckley, and Rob Mills. Optimization in self-modeling complex adaptive systems. *Complexity*, 16(5):17–26, 2011a.

Richard A Watson, Rob Mills, and Christopher L Buckley. Global adaptation in networks of selfish components: emergent associative memory at the system scale. *Artificial Life*, 17(3): 147–166, 2011b.

Richard A Watson, Günter P Wagner, Mihaela Pavlicev, Daniel M Weinreich, and Rob Mills. The evolution of phenotypic correlations and "developmental memory". *Evolution*, 68(4): 1124–1138, 2014.

Richard A Watson, Rob Mills, CL Buckley, Kostas Kouvaris, Adam Jackson, Simon T Powers, Chris Cox, Simon Tudge, Adam Davies, Loizos Kounios, and Daniel Power. Evolutionary connectionism: algorithmic principles underlying the evolution of biological organisation in evo-devo, evo-eco and evolutionary transitions. *Evolutionary biology*, 43(4):553–581, 2016.

Campbell O Webb, David D Ackerly, Mark A McPeek, and Michael J Donoghue. Phylogenies and community ecology. *Annual review of ecology and systematics*, 33(1):475–505, 2002.

Evan Weiher and Paul Keddy. *Ecological assembly rules: perspectives, advances, retreats*. Cambridge University Press, Cambridge, UK, 2001.

Daniel B Weissman and Nicholas H Barton. Limits to the rate of adaptive substitution in sexual populations. *PLoS genetics*, 8(6):e1002740, 2012.

Walter E Westman. Measuring the inertia and resilience of ecosystems. *BioScience*, 28(11): 705–710, 1978.

Thomas G Whitham, Joseph K Bailey, Jennifer A Schweitzer, Stephen M Shuster, Randy K Bangert, Carri J LeRoy, Eric V Lonsdorf, Gery J Allan, Stephen P DiFazio, Brad M Potts, et al. A framework for community and ecosystem genetics: from genes to ecosystems. *Nature Reviews Genetics*, 7(7):510–523, 2006.

GC Williams. Natural selection and adaptation, 1966.

David Sloan Wilson. *The natural selection of populations and communities*. Benjamin/Cummings Pub. Co., Menio Park, CA, United States, 1980.

David Sloan Wilson. Complex interactions in metacommunities, with implications for biodiversity and higher levels of selection. *Ecology*, pages 1984–2000, 1992.

Vero Copner Wynne-Edwards. Population control in animals. *Scientific American*, 211:68–74, 1964.

Debra Zuppinger-Dingley, Bernhard Schmid, Jana S Petermann, Varuna Yadav, Gerlinde B De Deyn, and Dan FB Flynn. Selection for niche differentiation in plant communities increases biodiversity effects. *Nature*, 515(7525):108–111, 2014.