

Biochemical Connectionism

Michael A. Lones^{1,2}, Alexander P. Turner², Luis A. Fuente²

Susan Stepney³, Leo S. D. Caves⁴ and Andy M. Tyrrell²

¹School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh, UK

Departments of ²Electronics, ³Computer Science, and ⁴Biology, University of York, York, UK

Email: {michael.lones, apt503, laf509, susan.stepney, leo.caves, andy.tyrrell}@york.ac.uk

Abstract

In this paper, we discuss computational architectures that are motivated by connectionist patterns that occur in biochemical networks, and speculate about how this biochemical approach to connectionism might complement conventional neural approaches. In particular, we focus on three features of biochemical networks that make them distinct from neural networks: their diverse, complex nodal processes, their emergent organisation, and the dynamical behaviours that result from higher-order, self-modifying processes. We also consider the growing use of evolutionary algorithms in the design of connectionist systems, noting how this enables us to explore a wider range of connectionist architectures, and how the close relationship between biochemical networks and biological evolution can guide us in this endeavour.

1 Introduction

A connectionist model is one in which computation, information processing, or intelligence emerges from the activity of a network of simple, non-linear elements [95, 78, 100]. The study of connectionist models came about through understanding of how the brain functions at the physical level, and is closely associated with the development of the field of neural networks. Because of this, in modern computer science the terms are often used synonymously. However, the brain is not the only naturally-occurring network that carries out functions analogous to computation. Less well known, but arguably more ubiquitous, are the biochemical networks that underlie the behaviour of biological cells.

In this paper, we argue that biochemical networks have the potential to add to connectionist knowledge. We do not argue that biochemical networks and their computational models are somehow better than neural networks, or *vice versa*. In biological organisms, neural networks and biochemical networks are different things, and are used for different purposes. Instead, we note that brains and biochemical networks display different patterns of organisation,

and these may inform our understanding of connectionism in different ways. Up to now, models of connectionism have favoured those patterns of organisation that are highly evident in the brain; in this paper, we investigate the benefits of introducing patterns of organisation that are highly evident in biochemical networks.

In fitting with this special issue on the Frontiers of Natural Computing, this paper is partly a review of recent results, partly speculation about future avenues of research, and partly a position paper, outlining where we think the field should move in the future. We begin, in Section 2, with an introduction to connectionism, followed by a review of neural connectionist architectures in Section 3. This is followed, in Section 4, by an introduction to biochemical networks and a general discussion of the concept of biochemical connectionism in Section 5. We then focus on three general features of biochemical networks that distinguish them from neural networks, and discuss how these features might inform connectionism: the complexity and diversity of their nodal processes (Section 6), their emergent organisation (Section 7) and the existence of higher-order and self-modifying processes (Section 8). Section 9 concludes.

2 Connectionism

From a computational perspective, connectionism can be understood as the emergence of complex behaviour from networks of behaviourally simple components. The field developed from cognitive theories about how the brain computes, particularly the idea that cognition comes about through parallel distributed processing, rather than traditional sequential models of computation [78]. The early development of connectionism is closely associated with the work of McCulloch and Pitts, particularly their seminal paper from 1943 [77], where they developed the idea of an artificial neurone, describing how these could be connected together in both feedforward and recurrent networks in order to carry out computation, i.e. neural networks. However, as discussed in Teuscher’s [100] detailed historical account, early ideas of computational connectionism can also be traced back to Turing’s 1948 [103] report on *unorganized networks*, in which he outlined the potential for a randomly-connected group of logic gates to carry out computation. Remarkably, he also suggested the use of an evolutionary algorithm to find networks with specific computational properties, considerably pre-empting the current wave of activity in this area.

Most existing ideas, discussion and literature on connectionism follow on from the *neural* models of computation and cognition developed by McCulloch and Pitts [77], Rosenblatt [93] and others. Despite this, there is plenty of scope for a wider view of connectionism based around *generic* networks of computational components. Farmer [27], for instance, provides a revealing case study of four network-based computational systems—neural networks, classifier systems, immune networks, and autocatalytic networks—describing how these can be viewed within a generic mathematical framework, and thereby exposing some of the fundamental connections between these seemingly diverse computational architectures. A generic connectionist framework is an admirable goal; however, in this article, we restrict our attention to connectionist architectures motivated by the organisation of biochemical networks. This is

not merely for reasons of space. Rather, elucidating the function and structure of biochemical networks has been an important part of modern biological enquiry, giving us a timely opportunity to use detailed knowledge of a naturally occurring connectionist system to inform computational connectionism.

Biochemical connectionism may be a new term, but the idea of modelling biochemical networks for computational purposes can trace its roots to diverse early work. Perhaps best known is Kauffman's [47] use of random Boolean networks to model genetic networks, although the idea of randomly interconnecting logic elements to achieve complex behaviour was also explored in earlier work by Walker and Ashby [110] and, as already mentioned, by Turing [103]. Also of relevance to the idea of biochemical connectionism is the computational modelling of metabolic networks, whether or not this involves an explicit network model. This is most evident in the field of membrane computing [86], particularly the use of P Systems for computation [117], but also relates to other work in modelling chemical processes, such as Turing's models of morphogenesis [104], Ulam and von Neumann's work on cellular automata [107] and artificial chemistries [21].

The majority of connectionist models date from a time when understanding of biochemical networks was still in its infancy. In recent years, there has been considerable progress in mapping the structure of actual biochemical networks, driven by data produced by various genomics, proteomics, metabolomics, and other -omics projects [46]. There has also been an increasing appreciation of the complexity and diversity of mechanisms used by biochemical networks, many of which have only been elaborated in recent years, such as RNA interference [79], the central role of the cytoskeleton [51], and epigenetic processes such as chromatin remodelling [16]. Hence, it seems like a good time to revisit the field of biochemical connectionism, and consider how contemporary biological understanding may guide the field, and the wider field of connectionism, in the future.

3 Neural Connectionism

Before we discuss biochemical networks and their potential contribution to connectionist models, it is first insightful to consider how existing connectionist architectures have been influenced by understanding of the brain.

In general, neural networks model three components of the brain: neurones, neural pathways, and synaptic plasticity. Neurones are modelled as an activation function that sums or integrates signals received from other neurones. In earlier neural networks, the activation function was approximated as a step-function, but nowadays it is most often implemented as a sigmoid, capturing the soft switching behaviour seen in biological neurones. A prominent exception is spiking neural networks [70], which more accurately model inter-neurone communication as a spike train rather than a series of fixed activation levels. In this case, activation functions are derived from biological neurone models, such as the Hodgkin-Huxley model [36].

Neural pathways are the patterns of connectivity that determine signal flow through a neural network. Commonly used neural network architectures such as the multilayer perceptron (MLP) [80] model neural pathways as a feed-

forward structure. Whilst feed-forward neural pathways do occur in the brain (notably in the cerebellum [25]), the use of feed-forward patterns of connectivity in neural networks arguably has more to do with ease of training, rather than biologically-motivated modelling. More realistic, from a biological perspective, are recurrent neural networks (RNNs), e.g. the Jordan network [45], which model the presence of feedback within neural pathways. Some neural network architectures go further than this, modelling actual neural pathways within the human brain, such as the thalamocortical system [49] and patterns of connectivity within the hippocampus [81]. There are also neural networks which model non-synaptic communication within the brain: examples include GasNets [43], which capture the action of diffusive neurotransmitters; and artificial neurone-glia networks [88], which capture the role of glial cells.

Synaptic plasticity is the capacity for synapses to modulate the strength of signals being passed between neurones, and is considered the primary mechanism for learning within the brain. Within a neural network, synaptic plasticity is modelled using variable weights on the connections between neurones. These are then progressively updated using a learning algorithm, until the neural network performs its desired function. Commonly used learning algorithms include gradient methods such as backpropagation [94], metaheuristics such as evolutionary algorithms [116], and also mechanisms motivated by learning in the brain, such as Hebbian learning [30]. There are also a number of neural network architectures that dispense with this conventional view of neural plasticity. Continuous-time RNNs (CTRNNs) [35], for instance, model the brain as a dynamical system, capturing the idea that learning can occur through transitions between attractor states. This does not require synaptic plasticity, and a number of authors have demonstrated how CTRNNs can learn through dynamical transitions alone [115, 102]. Reservoir computers (e.g. echo state networks [44], liquid state machines [71]) also reflect this dynamical view of neural activity, and only require training of their linear read-out nodes.

In addition to these low-level models of brain function, there exist a number of neural network architectures that are based on higher level models of how the brain works. Prominent amongst these are various kinds of associative memories (e.g. Hopfield networks [41]), and self-organising maps (e.g. Kohonen networks [54]), with the latter modelling the unsupervised nature of adaptation in the brain.

4 Biochemical Networks

Biochemical networks differ from neural networks in a number of prominent ways. Before we go on to discuss these differences in detail, we begin with a brief overview of biochemical networks, whose structure and function come about through interactions between the tens of thousands of different biochemical species present within a biological cell [52]. These interactions are highly organised both spatially and temporally, and are orchestrated by proteins, the cell's active molecular machines, of which there are thousands of different kinds in a human cell [18]. The cell's biochemical network, in essence, is the pattern of interactions between these protein-mediated reactions. This, in turn, can be divided into three interacting components, each with distinct structures, temporal scales, and dynamical

properties: a genetic network, a metabolic network, and a signalling network.

A *genetic network* describes the regulatory interactions between genes [10]. A gene is a contiguous region of DNA which (with some exceptions) encodes the amino acid sequence of a protein. A protein is constructed from this genetic description through a regulated process of gene expression [55]; in effect, gene regulation is the function carried out within the nodes of a genetic network. The initial, and most highly regulated, step of this process is the assembling of a transcription complex at the gene's transcription start site. This involves the coordinated binding of a number of different proteins, most importantly RNA polymerase, which actually carries out the transcription. Since RNA polymerase rarely binds to DNA by itself, it must be assisted by a range of general transcription factors, resulting in a baseline level of transcription. To further increase the number of protein products (a process called up-regulation) it must be stabilised by a range of special-purpose transcription factors (activators) which bind to particular sequences (promoters and enhancers) in the surrounding regions of DNA. It is also possible for special-purpose transcription factors (repressors) to disrupt the transcription complex, leading to down-regulation of the gene. Since they are proteins, both activators and repressors must themselves be the products of gene expression. Consequently, a particular gene will only be expressed in significant quantities if the genes describing its activators are also expressed, and the genes describing its repressors are not expressed. This pattern of regulatory interactions, when extended to all genes, forms the cell's genetic network.

A *metabolic network* describes the interactions between metabolites, the small biochemicals that are involved in diverse cellular processes, including energy supply, internal structuring, growth, development and reproduction [59]. In the relatively low temperature environment of the cell, reactions between metabolites are unlikely to occur. Instead, a reaction must be catalysed by an enzyme, a type of protein which binds together other chemical species and guides them through the intermediate stages of their reaction; in effect, these enzyme-mediated reactions are the functions carried out within the nodes of a metabolic network. When the product of one enzyme-mediated reaction becomes the substrate of another, a metabolic pathway is formed. This pathway then becomes elongated as its products or substrates become involved in other enzyme-mediated reactions. Finally, product-substrate sharing between pathways results in the formation of the cell's metabolic network. Since enzymes are proteins, their expression level within a particular cell is determined by the state of that cell's genetic network. Hence, different cells in the same organism can contain different metabolic networks, which in turn is the basis of cell specialisation.

A *signalling network* describes the protein-mediated molecular interaction pathways through which external signals are delivered to the cell's internal environment [39]. For the most part, a signal is received by a cell when a chemical messenger binds to a protein complex present on the cell's surface. The extracellular receptor propagates the signal through the cell wall and stimulates the release of small molecules called secondary messengers. These secondary messengers then diffuse within the cell, where they are bound by spatially-localised effector proteins. Two common types of effector protein are protein kinases and protein phosphatases, both of which modify other proteins, typically leading to their activation or deactivation. When the modified protein is an enzyme, this may result in

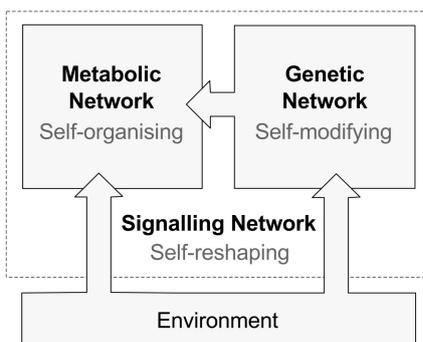


Figure 1: The main pathways of interaction between biochemical networks.

the production of new secondary messengers: amplifying the original response, causing a signalling cascade, and elongating the signalling pathway. Alternatively, it may cause a change to the cell’s metabolic network; or, when the modified protein is a transcription factor, a change in the cell’s genetic network. Signalling pathways also interact with each other, either through shared proteins or—when pathways share secondary messengers—through crosstalk, resulting in cooperative responses with complex spatial and temporal dynamics [50, 61].

The metabolic, genetic and signalling networks have been described, respectively, as the self-organising, self-reshaping, and self-modifying components of a cell’s biochemical network [74]. It is important to realise that they do not work in isolation, but are coupled (see Fig. 1). By regulating protein production, the genetic network modifies the behaviour of both the metabolic and signalling networks. By delivering chemical signals to different subcellular locations, the signalling network modulates the behaviour of both genetic and metabolic networks. In single-celled organisms, these interactions allow the cell’s metabolism to be reconfigured for different nutrient environments; in multicellular organisms, they are the basis of cellular differentiation and morphogenesis.

5 Biochemical Connectionism

There has been a fruitful relationship between the development of connectionist computational architectures and developments in the understanding of brain structure and function. Certainly there is no reason to believe that this relationship will not continue to be fruitful as further knowledge about the brain becomes available. What, then, is the value of exploring the relationship between biochemical networks and connectionism? There are several answers to this question. First, there is a benefit to having more information about the organisation of complex networks, and the various -omics projects have led to detailed information about the structure and function of biochemical networks. Second, we might argue that much of the low-lying fruit in neural understanding has already been picked, and there is more immediate benefit to exploiting the relatively unexplored synergisms between connectionism and biochemical networks. Third, and following on from this, we do not have to look far to see patterns of organisation that occur widely in biochemical networks yet not in the brain, and which appear to underlie a number of important cellular activities. Fourth, we could argue that, from the perspective of biochemical networks at least, the brain

Table 1: Characteristics of neural and biochemical networks

Level	Characteristic	Neural networks	Biochemical networks
Node	Function types (§6)	Mostly homogenous	Heterogenous
	Interactions (§6.1)	Additive	Combinatorial
	Function complexity (§6.2)	Simple	Complex
Network	Connections (§7)	Explicit	Implicit
	Spatiality (§7.2)	Minor role	Major role
	Coupling (§8.2)	First order	First and higher order
Structure	Evolution (§7.3)	Indirect	Direct
	Self-modification (§8.3)	Only through growth	Various mechanisms

appears to be a relatively conservative connectionist architecture; biochemical networks arguably provide a wider pool of inspiration for exploring connectionist architectures.

The idea of deriving computational systems from models of biochemical processes is certainly not new. The resulting computational systems include *in silico* models, such as cellular automata [107], membrane systems [86], Boolean networks [47] and artificial chemistries [6], *in vitro* processes such as chemical [1] and DNA [2] computers, and even *in vivo* approaches, such as the use of a slime mould to control a robot [101]. The application of concepts from biochemical networks to connectionist architectures is also not new [27, 73], and many of the approaches listed above can be readily mapped to a connectionist perspective. However, the aim of this paper is not to show how this can be done, or even to describe these different approaches in detail. Rather, our focus is on general principles of biochemical networks and how they might contribute to connectionist models as a whole.

Table 1 summarises some of the most evident differences between biochemical networks and neural networks. At the nodal level, the neurones in the brain are relatively homogenous, interact additively and are individually quite simple. Biochemical networks, on the other hand, comprise thousands of diverse biochemicals, which interact in a generally combinatorial fashion, and the resulting nodal processes can be very complex; in Section 6, we explore how understanding of these kind of nodal processes might inform connectionism. At the network level, the most prominent difference is the lack of explicit physical interconnections between nodes in a biochemical network, i.e. there is no analogue of neural axons and synapses. Instead, connectivity emerges from the interactions of individual biochemicals, each of which is subject to different spatial constraints on its movement. The patterns that result from this emergent organisation, and the benefits they might have for connectionist architectures, are discussed in Section 7. A further consequence of the relatively unconstrained interactions within biochemical networks is the capacity for higher-order interactions to occur between both nodes and networks. From a connectionist perspective, these represent a distinct source of novelty. Section 8 addresses the potential benefits of using higher-order interactions, and self-modifying processes in general, within connectionist architectures.

6 Functional Components

In comparison to neural networks, the nodes of biochemical networks are very diverse in terms of their operation, structure and biological role. From a connectionist perspective, this raises the question of whether there is a benefit to having heterogeneous processes occurring at the network’s nodes, or from using nodal processes of greater complexity than those seen in neural networks. Whilst this question has not been a focus of research in neural networks, a number of non-sigmoidal activation functions have been used, and have shown some benefits in terms of learning speed and generality [24]. Another potential benefit, discussed in [23], is that heterogeneous neural networks would allow certain computational behaviours to be expressed using relatively small networks. However, a significant barrier to using heterogeneous or alternative functions in neural computing is the widespread use of gradient-based learning methods, where there is a dependency between the training algorithm and the network’s transfer functions. In this respect, the growing use of evolutionary algorithms for training is opening up new avenues in the space of connectionist architectures. Examples include functionally-heterogeneous generalisations of the MLP, such as compositional pattern-producing networks [97], and process-oriented architectures such as genetic network programming [72].

6.1 Biochemical interactions

The functional activities occurring at a biochemical network’s nodes are a consequence of physical interactions between biochemicals. These interactions can be modelled in a number of different ways, and in general these models are different to those used for neural networks (although exceptions include the use of sigmoid kinetics to model cooperative binding in enzyme-mediated reactions [92]). It is therefore interesting to consider whether models of biochemical interactions can be usefully applied to connectionist architectures, and what these models might look like.

There already exists a relatively well-known example of a connectionist architecture modelled upon a biochemical network: the Boolean network (or Random Boolean Network, RBN) [47]. A Boolean network models a genetic network as a directed graph, with the vertices representing genes and the edges representing regulatory interactions. Notably, it models the process of gene regulation using a Boolean function. Given the complexity of the process of gene expression, this may seem rather simplistic. However, Boolean networks have been shown capable of modelling the dynamics of biological gene regulatory circuits [3], so clearly Boolean models are sufficient—in some cases, at least. From a computational perspective, Boolean functions are quite different to the threshold functions commonly used in neural networks, since they are based on combinatorial rather than merely additive relationships between variables [118]. In this sense, Boolean functions are individually more expressive than weighted sigmoids. This suggests that Boolean networks might have advantages for certain computational tasks; the few examples where they have been applied to computational tasks tend to support this notion [56, 22, 14].

However, as a general-purpose connectionist architecture, Boolean networks do have limitations. For example,

inputs and outputs must be binary encoded. Assuming these are transferred to/from the system via the expression states of individual genes, this entails that the network must be at least as large as the bit length of the largest binary-encoded input or output. Large networks, in turn, are more difficult to induce than small networks. Conversely, small Boolean networks are limited in their ability to express complex dynamics, since they can only exist in 2^N states, where N is the size of the network. A potential solution to both of these problems is to use continuous-valued generalisations of Boolean functions. Such models have the added advantage that they can be used to model quantitative (in addition to qualitative) aspects of real biochemical circuits, an idea that has been explored in a number of recent papers [113, 32]. In [113], for example, the authors describe a method of transforming Boolean functions into equivalent ODEs, such that their behaviour is equivalent for concentrations of 0 and 1, but differ for intermediate values. This approach would, of course, have an impact on the efficiency of a connectionist architecture, since it would be necessary to simulate ODEs. A more efficient alternative could be to use multi-valued logics rather than fully-continuous models. In [32], for example, the authors present a method of constructing equivalent multi-valued logics from Boolean functions, and observe that a 3-valued logic offers important benefits over a Boolean model in terms of correctly capturing the dynamics of a biological regulatory circuit.

6.2 Dynamical processes

In Section 4, we discussed the formation of a gene transcription complex, noting how its correct placement and stabilisation depends upon the activities of many different transcription factors. In a complex multicellular organism, such as a human, a transcription complex can comprise tens of different transcription factors, each an intricate molecular machine assembled from hundreds or thousands of interacting amino acids [57]. Moreover, the formation of the transcription complex is just the first stage of the gene expression process. Following this, the gene is copied into an RNA transcript, edited to remove introns, translated into an amino acid chain, and then folded into a protein; whilst being subject to further regulation at (and between) all these stages [55]. When considering this level of complexity, and the ongoing interactions between its component parts, it makes more sense to think of transcription not as a functional mapping, but as a dynamical process. The same is true of other ‘nodal processes’ within biochemical networks, such as the enzyme-guided transformation of biochemicals, or the assembling of a signalling module around the tail of a cell surface receptor.

Given the complexity (and, in some cases, our limited understanding) of these processes, it would be challenging (and perhaps not very useful) to model them accurately. An alternative and more abstract approach, which we have considered in our work [68, 63], is to place general models of complex, dynamical processes within the nodes of a connectionist architecture. In particular, we have investigated the utility of non-linear discrete maps for this purpose. For the most part, these are numerical models of diverse dynamical processes that occur in biological and other naturally-occurring complex systems; the archetypal example being the logistic map [75], which models ordered and chaotic regimes that occur in certain models of population growth. A desirable property of discrete maps is

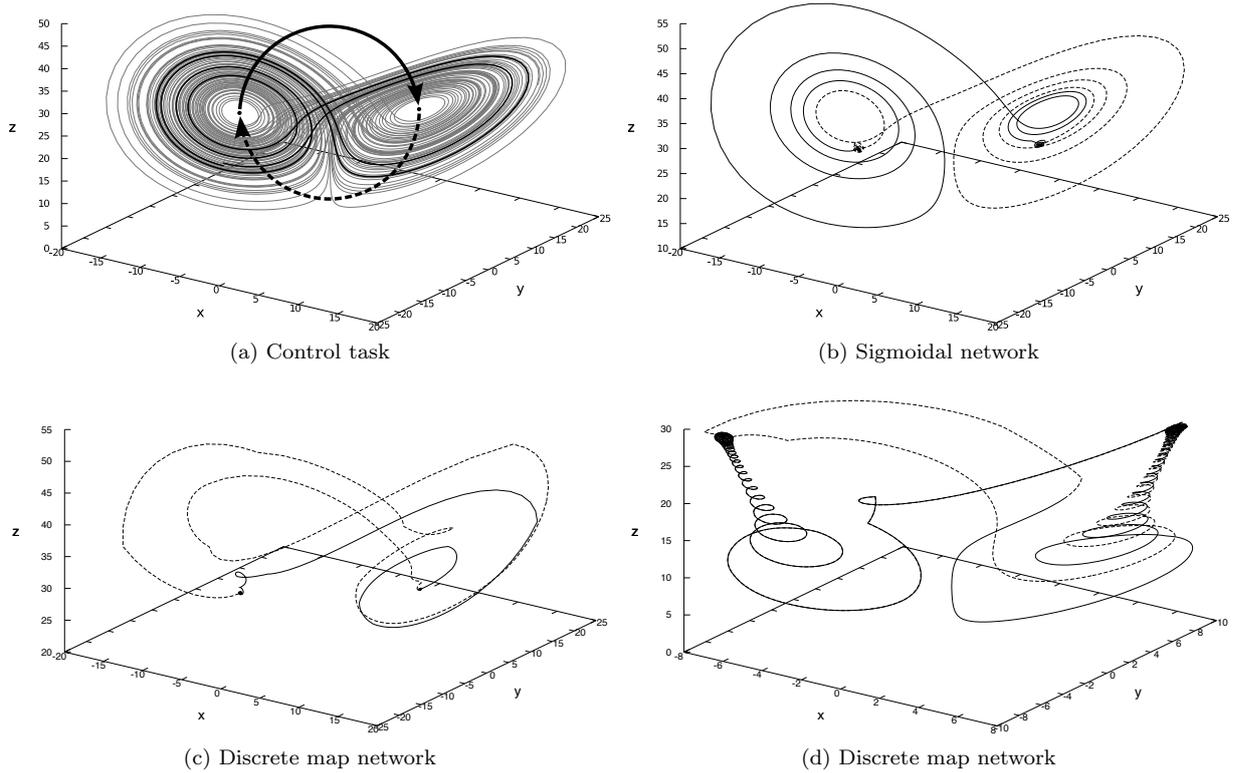


Figure 2: Controlling trajectories in the Lorenz system using evolved recurrent networks. (a) The objective is to control a trajectory so that it moves as fast as possible from the unstable point in the leftmost lobe to the unstable point in the rightmost lobe (unbroken lines in b-d), where it must then stabilise the trajectory, and then repeat the task in reverse (broken lines). (b) The best control strategy found using a sigmoidal network. (c) A near-optimal control strategy used by an evolved discrete map network (d) An example of a more exotic control strategy used by a discrete map network. Adapted from [63].

computational efficiency; although capable of expressing complex dynamical behaviours, they are implemented using simple iterative equations. In our work, we have found that, on average, networks containing discrete maps either perform better or learn more quickly than the same network architectures containing conventional sigmoidal functions when applied to a range of computational tasks [65, 63]. Figure 2 shows an example of this, comparing the ability of sigmoidal and discrete map recurrent networks to solve a complex control task in the Lorenz system. Notably, we were able to evolve discrete map networks that could solve this task better, and more consistently, than sigmoidal networks. In particular, we observed that discrete map networks generate much more diverse behaviours, particularly in the early stages of learning when diversity is most important. In the context of the Lorenz system, this meant that they were able to generate relatively unusual, but near-optimal, trajectories such as the one shown in Figure 2c, rather than the more common, slower, trajectories generated by evolved sigmoidal networks (Figure 2b). In effect, discrete maps appear to provide a ready means of expressing complex dynamics, without requiring the evolution of a dedicated sub-network of sigmoidal nodes.

6.3 Constructive functions

In general, it is difficult—and perhaps even impossible—to say how a particular nodal model affects the overall dynamics of a connectionist system. In networks of discrete maps, for example, it has been observed that there is no obvious relationship between the overall dynamics of the network and the dynamical behaviour of the individual maps [15]. Furthermore, even if this relationship could be predicted, it is difficult to know what kinds of network dynamics are appropriate for different computational tasks, and therefore which nodal models would be useful. A potential solution to this dilemma can be seen throughout biochemical networks: the use of constructive functions. By constructive, we mean the construction of a functional behaviour through the mutual assembly of complementary-shaped biomolecules. Examples include the construction of a regulatory function through the mutual assembly of transcription factors, and the construction of new biochemical species as a product of enzyme-mediated reactions. A significant consequence of this is that the function of the whole is not determined by the functions of the individual parts, but rather through their interactions. This, in turn, has the appealing property that novel functions can be created by the recruitment of new biomolecules to the complex. In computational terms, functions can be induced by the system as appropriate, rather than having to be predefined.

The idea of using constructive functions in computational systems has been explored by a number of authors [6, 26, 28]. The key to this approach is to use primitive elements that are sufficiently expressive when assembled into different composite forms. One source of such primitives is mathematics. In [6], for example, the authors developed a constructive system based upon interactions between numerical strings and matrices. Algebra, in particular, seems like a promising source of ideas for constructive functions [26]; group theory, for example, describes many interesting mathematical spaces, including those, such as Lie groups, which model fundamental aspects of the physical world. Ideas can also be taken from existing biochemical connectionist models: in our own work [28], for example, we have

looked at a constructive system based upon interactions between Boolean network fragments.

However, most of this work has been done in the context of well-stirred artificial chemistries, systems in which the components interact more-or-less freely with one another. As such, this still leaves the question of how to integrate constructive functions into a more conventional connectionist architecture. One possibility would be to associate a primitive function with each edge of the network. Each node would then assemble the functions of its incoming edges in order to construct the node's function. If the resulting function is conventional, in the connectionist sense, then it can be applied to a vector or scalar sum of its inputs, generating a new scalar quantity to deliver to downstream nodes. However, for many of the existing constructive functions, there is no distinction made between function and data—in effect, functional structures are applied to other functional structures in order to generate new functional structures. Hence, we might imagine these functional structures being the elements that flow around the network, rather than scalar quantities. This does present the issue of how to encode external inputs and outputs within functional structures, although we can envision that in certain circumstances (e.g. matrices) this might be more natural than splitting them into their individual numeric components.

7 Emergent Organisation

In the previous section, we discussed how biochemical and neural networks differ in terms of the functional operations carried out at the nodes of the network. In this section, we move up a level, and consider the overall structure of biochemical networks. From this perspective, one of the most evident differences in biochemical networks is the lack of a physical network structure, i.e. there is no analogue of neural axons and synapses. Rather, connections between the network's nodes emerge as a consequence of implicit properties encoded in each biochemical's shape and chemical make-up. In this sense, the structure of a biochemical network is emergent from the properties of its component parts, as we saw with constructive functions in the previous section.

7.1 Decoupled state

One consequence of this emergent organisation is that function and state need not be tightly coupled, as they are in a neural network where each neurone is associated with a single activation level. For example, in a metabolic network, multiple enzymes ('functions') may competitively modulate the concentration of a single biochemical species ('state'). Likewise, in a genetic network, multiple genes (the 'functions' in this sense) can express the same transcription factor, whose overall concentration level ('state') then affects the expression of other genes.

However, these decoupled 'state' components are not only under the influence of directly competing protein-mediated reactions; they are also under the indirect influence of mass conservation. Unlike in the brain, where state is encoded in relatively abundant electrical signals, biochemical state requires chemical mass. The supply of this is strictly limited, and must be divided amongst the different processes occurring within the cell. The consequence of

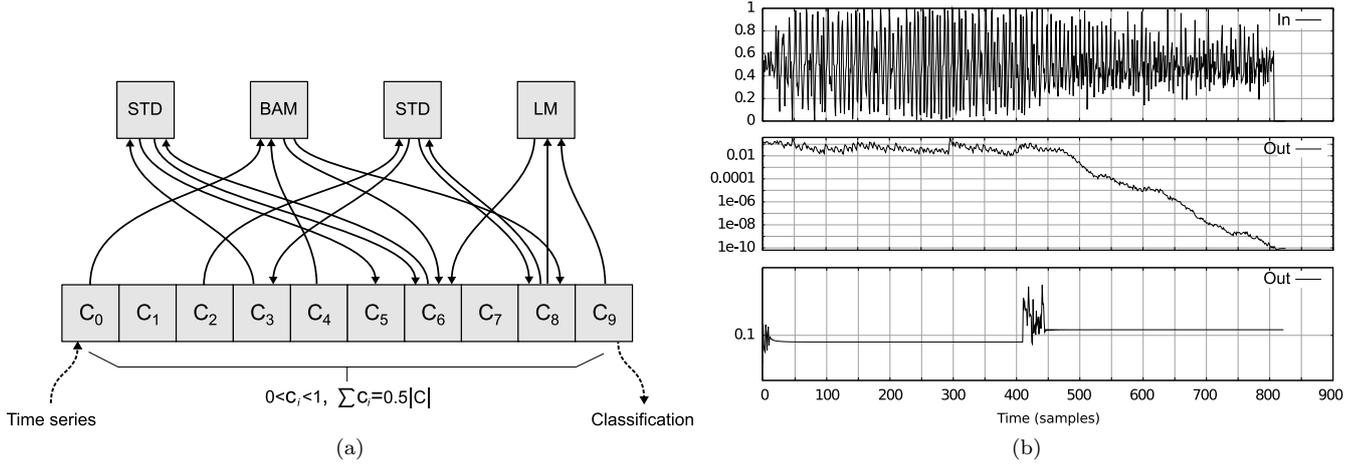


Figure 3: (a) A conserved decoupled connectionist architecture evolved to diagnose Parkinson’s disease. The sum of the state components, $C = \{c_0, \dots, c_9\}$, is conserved, according to the equation shown, via a linear scaling process. Time series data is input a value at a time by setting c_0 , synchronously updating the network after each input. The classification is read from the final value of c_9 . The network uses four function nodes, each a discrete map: STD=standard map, BAM=baker’s map, LM=logistic map. (b) An example of the network responding [middle] to a decay in amplitude of the movements of a Parkinson’s patient [top], and [bottom] the network’s response when the input is removed and the network is later perturbed with an impulse to c_0 . See [66] for more details.

this is that mass balance leads to coupling between the state components of different biochemical networks, since an increased concentration of a chemical in one part of the cell must be counterbalanced by a decreased concentration in another part, or in another chemical with overlapping constituents.

In our work, we have looked at how a conservation law influences the dynamics of a decoupled connectionist architecture [68, 63]. Figure 3 illustrates this architecture: showing how state and function are represented by separate nodes, and how the sum of the concentrations of the state components is conserved after each iteration of the network. In general, we have found that mass conservation makes it easier to induce networks with a desired input-output mapping, making the inductive process less sensitive to design factors such as choice of nodal processes and the mechanism used to introduce external inputs to the network. We have speculated that, at least in part, this is because it is easier to recruit signals originating from an external source or from other parts of the network, since these are propagated through the network by the conservation process. Interestingly, we found these conserved decoupled networks to be particularly effective at a signal processing task: diagnosing Parkinson’s disease from movement time series data. In general, they achieved higher accuracies and higher success rates at this task than both genetic programming [64] and a range of other connectionist architectures, including recurrent neural networks [65, 66]. Analysis of evolved classifiers suggested that the conservation process was acting as a dampening mechanism, preventing long-lived chaotic signals from dominating the dynamics of the network. This can be seen in the lower plot in Figure 3b, which shows that the network’s dynamics soon become dampened after input is removed and after the network is hit with an impulse. In this respect, it resembles the dampening constraints used in the construction of reservoir computers (e.g. the echo state property [44]).

7.2 Spatiality

A cell is a physical space. Each biochemical occurs at a particular location within this space, and it may only interact with other biochemicals that occur in its immediate neighbourhood. Because of this, spatiality is a fundamental concept for biochemical networks. As biological knowledge has grown, it has become increasingly evident that cells are highly organised at the spatial level [51, 112]. They are not the spatially unstructured ‘bag of chemicals’ they were once viewed to be; most biochemicals are localised within a particular part of the cell, and in many cases are physically tied to the cell’s membranes or cytoskeleton. This physical localisation even extends to certain biochemical pathways, such as signalling complexes, in which the components of a signalling pathway are physically assembled and held together by scaffold proteins [8].

Nevertheless, biochemical pathways do contain motile components, entailing both that a component may be missing from a particular pathway at a particular time, or that an extra component may join a particular pathway at a particular time. These kind of processes are particularly significant in biochemical pathways that function at low concentrations of their constituent elements, where it has been suggested that this ‘noisiness’ may be beneficial [76]. From a connectionist perspective, this behaviour implies that the network structure has a degree of dynamism, with edges appearing and disappearing over the course of time. One way of modelling this within a connectionist architecture is to have some kind of dynamical process that adds and removes edges to the network (an idea which is explored further in Section 8). Another approach is to switch between different nodal functions. This idea has been explored in [96], where the authors describe a Boolean network whose nodes switch between different Boolean functions according to a probability distribution. The nodes of a GasNet [43] (see Section 3) also switch their behaviour over time, though in this case based on the value of a time-varying state component. Notably, they use a spatial coordinate system, in which individual neurones are embedded, and across which neurotransmitters can diffuse, causing neurones to switch between different sigmoidal activation functions.

The use of explicit notions of space in neural networks has also been explored in [29] and [108], where the authors note that whilst this approach can promote useful behaviours such as modularity, excessive spatial constraints can limit a learning algorithm’s ability to explore behavioural space more generally. Explicit notions of space are also found in various other computational models of biochemical processes. Perhaps the best known of these are cellular automata (CA) [107] which, in their elementary form, can be considered a kind of Boolean network in which interconnections only occur between nodes that are immediate spatial neighbours (and all nodes use the same Boolean function). At the other end of the spatial spectrum are artificial chemistries [21], most of which model biochemical systems as ‘well-stirred’ reactors in which all function and state elements are able to interact over the course of time. Between these two extremes are models such as P Systems [86] in which components are organised into compartments, within which they are free to interact, but between which interactions can not directly occur. This models the manner in which cells are organised into membrane-bound compartments. However, none of these is arguably a good model of real biological systems, in which different components have different degrees of containment

and motility in order to facilitate the requirements of different biological processes. For example, those that require isolation from other processes tend to be physically co-located and protected from interference by scaffolding [37] or cellular membranes. By comparison, those that integrate diverse signals tend to be spatially distributed and contain motile components.

Signalling networks provide a good example of these two extremes. As we mentioned above, integrated components of signalling pathways are often physically co-located in signalling complexes in order to provide cohesion and prevent interference. However, at the other extreme is the widespread incidence of crosstalk, which occurs when a particular biochemical is involved in more than one unrelated biochemical pathway whose spatial or temporal domains overlap. This can result in a signal generated internally by one pathway being delivered to another pathway. Crosstalk is particularly commonplace in signalling networks, where a relatively small number of secondary messenger molecules are used to propagate a broad spectrum of extra-cellular signals to locations within a cell. In human-engineered systems, crosstalk is generally seen as a form of interference that is disruptive to correct functioning. However, its prevalence in signalling networks has led to an understanding that crosstalk can also have beneficial roles, particularly as a mechanism for providing complex non-linear responses to combinations of stimuli [82] [61].

In our own work, we have looked at the benefits of using a simple model of crosstalk within a connectionist architecture. Mimicking the structure of a signalling network, different types of external input are initially delivered to different sub-networks. These sub-networks (which are comparable to pathways) have no explicit interconnections. However, they do have crosstalk nodes, in which signals from one sub-network are allowed to leak into another using a simple mixing function. This allows external inputs to be treated in a semi-independent manner, providing a mechanism to introduce new inputs with minimal effect upon existing sub-networks. As a demonstration of this approach, we have shown that a control problem with strong dependence between its inputs can be solved using a connectionist architecture comprising separate networks for each input linked by crosstalk [34]. However, we would expect this kind of architecture to be particularly useful for problem domains in which a connectionist system has to react to many potential inputs: for example, in a robotic system, where inputs concerning both the environment and the robot's internal conformation are received from various sensors. In [33], we considered this approach as part of a layered architecture for controlling a hexapedal robot on rough terrain (see Figure 4). In particular, we evolved a coupled artificial signalling network, in which each of the component networks uses local sensory inputs to control the movement of a single leg. The behaviour of the whole robot then emerges as a result of evolved crosstalk pathways between the networks, providing a conduit for global exchange of information between the component networks. This is in contrast to the normally predetermined architectures associated with neural network-based approaches to legged robot control, and proved a successful method for generating adaptive gaits in a challenging environment. The utility of crosstalk has also been investigated in [20], where the authors developed a learning classifier system in which crosstalk is used to transfer information within the system.

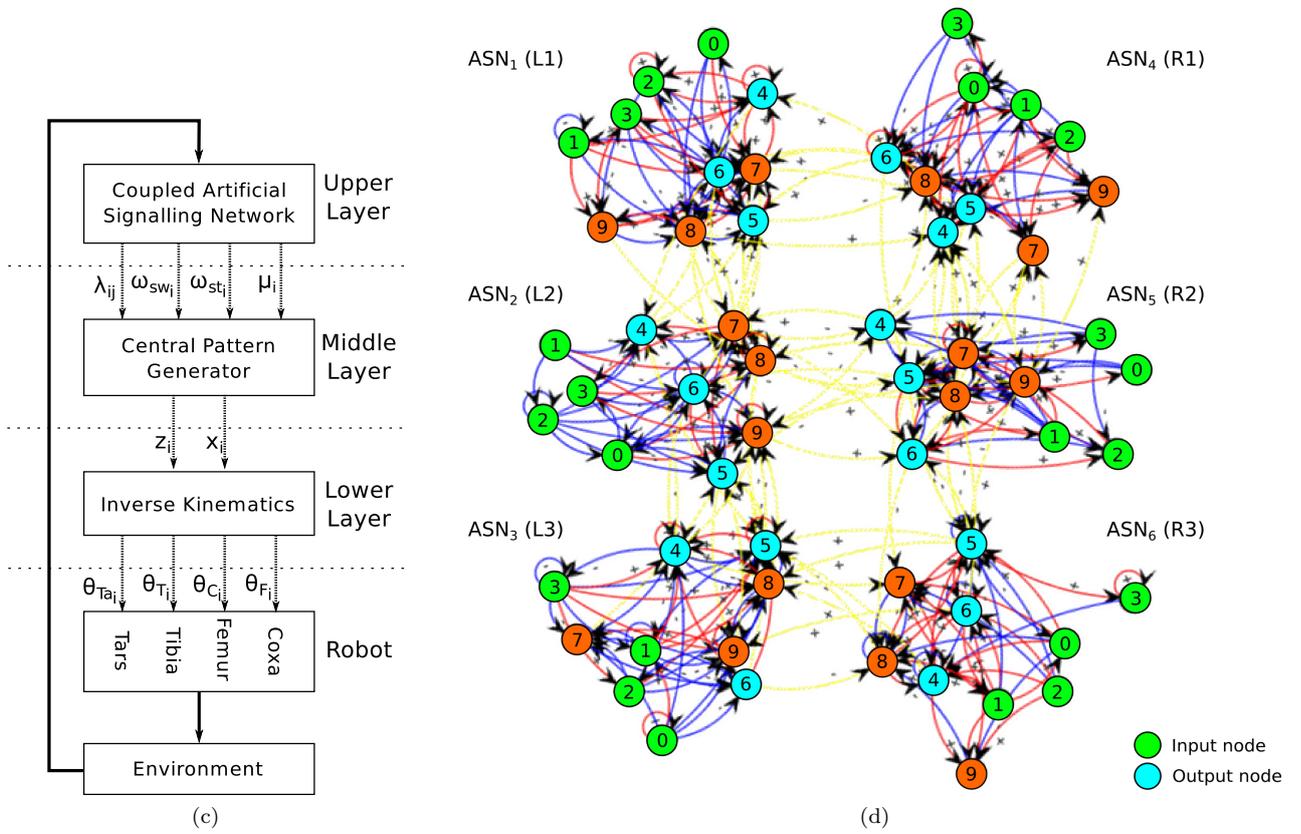
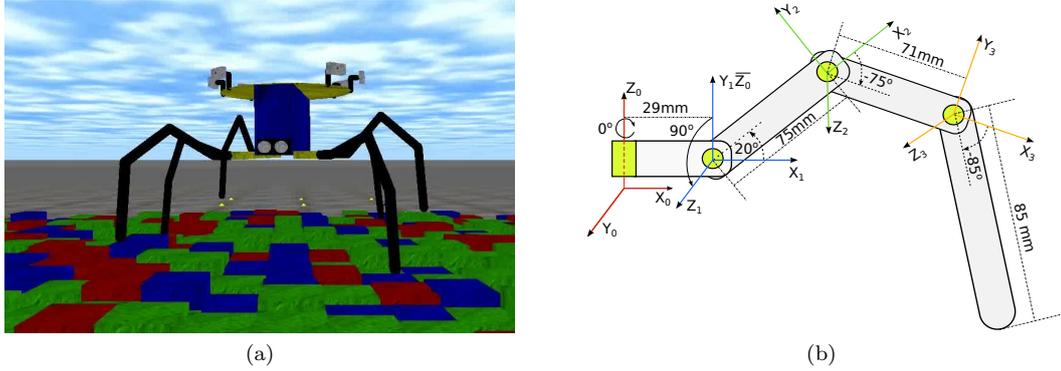


Figure 4: Using a connectionist architecture motivated by crosstalk-coupled signalling networks to control a legged robot. (a) A simulation of the hexapedal robot used in this study navigating a rugged terrain. (b) Each leg has 4 degrees of freedom. (c) The layered controller architecture, in which a coupled artificial signalling network controls the parameters of a central pattern generator (CPG), whose output drives the inverse kinematics that control the movements of individual joints. (d) A coupled artificial signalling network evolved to provide an adaptive gait within rugged environments. The inputs, for each component network, are from the contact sensor on the corresponding leg, and the three infrared sensors on the corresponding side of the robot. The outputs are CPG control parameters. Crosstalk edges are shown in yellow (light grey). For more details, see [33].

7.3 Evolvability

Evolutionary algorithms are a useful means of learning connectionist architectures. Unlike gradient-based methods, such as back-propagation, they are not limited to optimising weights; they can also be used to design a network's topology, and they can handle heterogeneous architectures. In principle they can be used with any representation, so long as appropriate variation operators (i.e. mutation, recombination) are available. However, in practice it is beneficial to use a representation that is *evolvable*, so that when it is mutated or recombined there is a tendency for the evolutionary algorithm to explore fitter variants—or, at least, variants that are not significantly worse on average [62, 42]. In computer science in general, there are many examples of representations that are not evolvable, i.e. that when randomly mutated, lead to dysfunctional variants. By comparison, evolvability is commonplace in biological systems [17, 53, 87]. This is particularly so for biochemical networks: since proteins are encoded in DNA, they are directly exposed to the machinery of natural evolution, and must be able to react in an appropriate manner. Hence, we might expect biochemical networks to be a particularly useful source of information regarding evolvable ways of representing connectionist structures.

One prominent pattern of organisation known to promote evolvability is the use of multiple weak interactions [17, 109]. We have already discussed a major example of this in biochemical networks: the interactions between the initiation site, transcription factors, enhancers, and other control signals that lead to the formation of a transcription complex in higher organisms. It is the cumulative effect of all of these interactions that controls the activation level, meaning that the strength of activation can be varied gradually through the addition and removal of components and interactions, and can respond flexibly to new sources of regulation [53]. These new sources of regulation can, in turn, be any biochemical whose shape is complementary to an existing part of the transcription complex, and can potentially come about from a huge variety of source processes—including crosstalk, which is another notable source of weak linkage in biochemical systems. An important source of transcriptional regulation, and one that has been hypothesised to account for much of the complexity seen in higher organisms, is gene duplication [84]. This is an example of another prominent pattern of organisation known to promote evolvability, redundancy [17]. Redundancy allows a system to maintain its function whilst exploring novel behaviour through a process of duplication and divergence of the system's existing components. In genetic networks, this process occurs at multiple levels: the level of whole genes, at the level of transcription factor binding sites, and even at the level of chromosomes and whole networks. Given its hypothesised importance within biological evolution, a number of authors have looked at the roles that duplication can play within the evolution of artificial genetic network models [4, 60, 11].

Underlying the effectiveness of duplication events in biological systems is a fundamental, yet often overlooked, property of biochemical systems: function-structure duality. In essence, the physical structure of a biochemical determines how it interacts with other biochemicals, and therefore determines its function within the system. This in turn leads to the emergence of higher level biological processes such as the binding of antigens by the immune system, the binding of substrates to enzymes, and the assembly of transcription factor complexes. Moreover, from

an evolutionary perspective, it means that previously unseen biochemicals can be introduced to the system yet still have a defined role within the system; when new biochemicals are introduced via duplication events, the duplicates are able to compete against the original to fulfil the same functional role, or potentially to fulfil new roles. In either case, there is a tendency to preserve the existing behaviour whilst providing scope to explore new adaptations and functions.

In part, the notion of function-structure duality is a more fundamental perspective on the constructive functions we talked about in Section 6.3. Indeed, constructive functions are one way in which we could implement this notion within a connectionist architecture. However, this approach is dependent upon the expressiveness of the underlying primitive elements, since a function can only be expressed through composition of these elements. Another approach, often used in connectionist models of genetic networks that are used for computation [7], is to introduce an indirect reference space in which functional components can identify one another. An example of this is the template-matching scheme described in [91]. Motivated by the organisation of genetic regions in biological chromosomes, this approach [91] uses evolved patterns to identify the nodes of the network. Over the course of evolution, there is an expectation that a relationship will emerge between these patterns and the function of the nodes they reference; in effect, this amounts to a loose form of function-structure duality in which the evolved patterns play the role of structural recognition. However, there is no guarantee that such a relationship will emerge, especially over the relatively short periods of evolution found in evolutionary algorithms. A potential solution to this is to encode details of the node's functional behaviour within its pattern, a method we have previously looked in our work on implicit context representation [67]. Another issue with indirect reference spaces is the need to carry out pattern matching in order to identify the edges of the network. In a spatially unstructured connectionist architecture, this could involve comparing every pair of nodes in order to identify the strongest matches, which may be prohibitive for large networks. Hence, there may be some benefit to combining an indirect reference space with a spatially-structured architecture (see Section 7.2) so that pattern matching only needs to be done in the local neighbourhood of each node.

8 Higher-Order Processes

As well as its contribution to evolvability, function-structure duality has a number of other implications for the organisation of biochemical networks. One of these is the widespread occurrence of higher-order, self-modifying, processes. These include proteins that modify the activity of other proteins, networks that modify the structure of other networks, and various processes that occur at the chromosomal level, such as retro-transposition [90] and chromatin restructuring [114]. These are possible because biochemicals interact through recognition of structure, and this structural reference space includes all biochemicals: not just those with informational roles (e.g. metabolites), but also those with functional roles (i.e. proteins) or information storage roles (e.g. chromosomes) [40]. In general, there is no analogue of these kind of higher-order processes in the brain, where neurones can not physically alter

each other's activity, and where physical alterations at the network level only occur over long time-scales as a result of growth. Consequently, they are of particular interest as a potential source of novelty within connectionist architectures.

8.1 Higher-order functions

Perhaps the most fundamental higher-order process in biochemical systems is the modification of a protein by another protein. We have already mentioned the most common example of this kind of protein-protein interaction: protein kinases in signalling pathways. These change the function of other proteins by phosphorylating their binding sites, causing changes in their structural or chemical interactivity. A single protein kinase can modify the activity of many target proteins, leading to significant downstream changes in a cell's biochemical network. For this reason, their activity is tightly regulated in biological systems, often via phosphorylation, and hence through the activity of other protein kinases [83].

From a computational perspective, the modification of a protein by another protein is analogous to the kind of higher-order function application seen in the λ -calculus and functional programming languages. Indeed, in one of the earliest artificial chemistry models, Fontana [31] used λ -expressions as the basic elements of the system, arguing that higher-order interactions such as these are good candidates for producing complex, adaptive behaviour. In effect, the λ -expressions are used as another kind of constructive function. From a more conventional connectionist perspective, protein modification can also be seen as a network rewriting operation, causing one or more interactions to be added or removed from the network. This idea could be introduced to existing architectures by adding new edges, which act as switches within their target nodes, up or down-regulating a node's overall activity according to the activation level of a source node. In a multi-layer perceptron, for example, this might be done by scaling the output of a node's activation function. In a Boolean network, it would be similar to the use of canalising functions (such as AND) in which one or more inputs strongly determine a node's output [48], and which have been hypothesised to improve robustness by stabilising a network's attractors [99]. It also resembles multiplicative coupling in autocatalytic networks, a mechanism that allows nodes to exercise control over other nodes, and which is discussed in [27] in the context of connectionist architectures.

8.2 Higher-order networks

Protein modification often leads to changes at the network level. For instance, a protein kinase in a signalling pathway may activate a dormant enzyme, leading to the production of new metabolites: in effect, the signalling network has modified the metabolic network. Likewise, when a protein kinase activates a transcription factor, the signalling network modifies the genetic network. These network level interactions are not only due to protein interactions: by modulating the production of signalling and transcription components and the availability of energy, the metabolic network can modify both the signalling and genetic networks, a mechanism which plays a significant role in a wide

range of cellular processes [111]. However, arguably the most important network-level interaction is the modification of the metabolic and signalling networks by the genetic network. This is possible because genes encode the proteins that form the functional elements of the metabolic and signalling networks, and the regulatory interactions between genes can modulate the levels at which these proteins are expressed. Given that almost all the proteins present in the cell are a product of gene expression, this means that the genetic network can make very large scale changes to biochemical networks. Since each protein is individually regulated, modulation of gene expression can also result in very specific patterns of change. A prominent example of this large scale, yet specific, modification of a cell's biochemical networks is cellular differentiation, whereby a cell's internal processes become increasingly specialised in order to carry out a particular biological function.

This relationship between the genetic network and the cell's other biochemical networks is interesting from a connectionist perspective, since it amounts to a higher-order relationship in which one network controls the structure of another network. In effect, one network is responsible for selecting a behaviour, and the other network is responsible for implementing it. In particular, this seems like a potentially useful pattern of organisation for systems that need to switch between different behaviours over the course of time. To investigate this idea, we implemented a connectionist architecture in which one network controls the expression of functional components in another network [69, 63]. In mimicry of a biological cell, the controlling network is modelled upon a genetic network (although it also closely resembles a recurrent neural network) and the controlled network is modelled upon a metabolic network (similar to the one shown in Fig. 3). As in a biological cell, the two networks are coupled together through gene expression, such that the expression level of a 'gene' in the controlling network determines the activation level of the corresponding 'enzyme' in the controlled network. The enzyme's activation level, in turn, determines its relative influence upon the concentrations of a set of metabolites, which are also used to encode external inputs and outputs.

In general, we have found this architecture to be beneficial in situations where there is a need to move quickly between different kinds of behaviour [63]. An example of this is the Lorenz control task, described in Section 6.2, where there is a need to move between different behaviours depending upon whether the controller is transitioning between the two unstable points of the system, stabilising at these points, or moving backwards or forwards between the points. We found higher-order networks to be particularly good at solving this task [63], outperforming standalone genetic and metabolic networks in terms of both optimality and consistency of solutions. In particular, we found that in evolved networks, the genetic network reconfigures the metabolic network at points where the behaviour of the controller is required to change quickly. A more tangible illustration of this behaviour is shown in Fig. 5, which shows a higher-order network controlling the locomotion of an unstable two-legged robot. The objective was to move the robot in a way that maximises the distance covered during the evaluation period: which, in this case, is achieved using an intricate backwards somersaulting manoeuvre, showing how the extra degrees of freedom present in the higher-order network can promote solutions with surprisingly complex dynamics. This example shows a phenomenon we often see in evolved controllers: a separation of time-scales, in which the genetic network responds to relatively

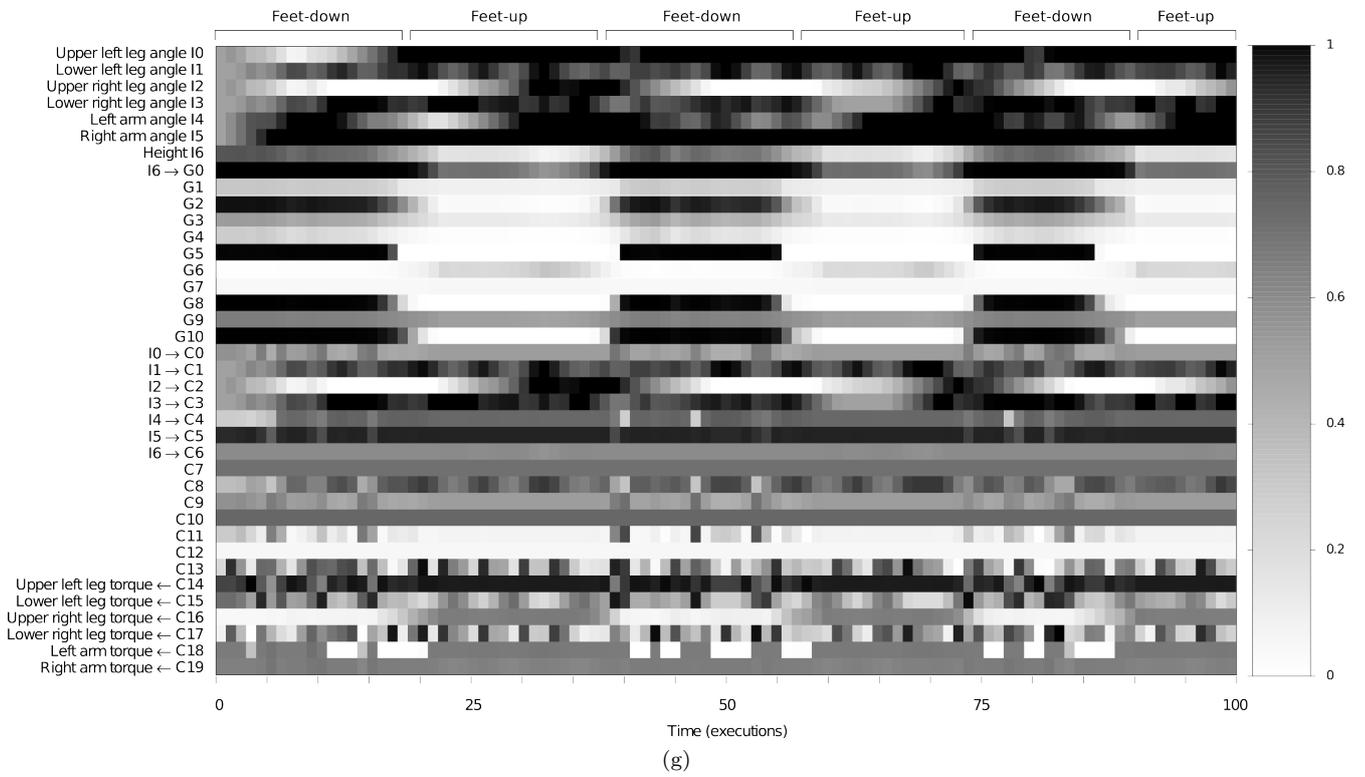
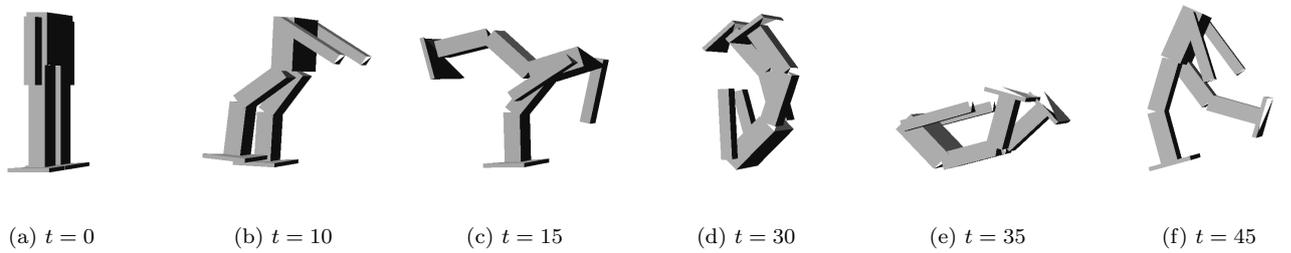


Figure 5: (a–f) Higher-order network controlling a two-legged robot in order to carry out a backwards somersault, showing timestamped snapshots from a single cycle of movement. (g) Time series view showing how an artificial metabolic network modulates the concentrations of metabolites (C0–C19) in order to map joint angles (I0–I5) into actuator torques (C14–C19) during a series of somersaults. Expression of functional components in the metabolic network is modulated by an artificial genetic network (G0–G10), which changes its expression state based on the height of the robot’s torso (I6). Expression levels and metabolite concentrations can vary between 0 (white) and 1 (black). For more details, see [69].

slow-changing dynamics (roughly the feet-down and feet-up phases of movement, as labelled, and visible as dark and light vertical bands in gene expression) and the metabolic network responds to fast-changing dynamics (the actuator kinematics). In effect, the genetic network reconfigures itself twice per rotation, causing the metabolic network to switch between two different behaviours according to whether the robot is in a feet-down or feet-up phase. In most problems we have looked at, we have also found that there is an advantage to choosing different function sets for the genetic network and metabolic network. This functional heterogeneity reflects findings in another recent paper that considered a model of interacting genetic and metabolic networks [12]. In that work, the author explored coupled Boolean networks, and found that topological heterogeneity was advantageous. Hence, there seems to be an indication that heterogeneity is a useful property in higher-order networks: perhaps reflecting the fact that many problems can be broken down into sub-tasks with different dynamical characteristics that, in turn, benefit from different dynamical solutions.

8.3 Genomic self-modification

When thinking about genomes and genetic networks, it is common to abstract away the role of DNA and chromosomes. However, it has become increasingly apparent that both DNA and chromosomes have significant roles in gene expression beyond that of information storage. DNA is not a one-dimensional string—it has a three-dimensional structure, and the local conformation at any point is influenced by both the pattern of base pairs and the action of bound proteins [85]. This local structure, in turn, influences the binding of the transcription complex. On top of this, in higher organisms DNA is packaged into chromosomes formed of chromatin, an amalgamation of spindle-like scaffold proteins over which the DNA is first coiled, then supercoiled [114]. In order for a transcription complex to bind to a region of DNA, the chromatin must first be locally unwound. This unwinding is done by proteins, which are both the product of other genes and subject to regulation by other genes. Hence, the genetic network is regulated at both the genetic and chromatin level. Since chromatin is usually wound or unwound over a length of DNA corresponding to multiple genes, and since interacting proteins are often encoded proximally in the genome (a phenomenon known as epistatic clustering), chromatin remodelling can result in whole pathways being turned on and off. Because of this ability to make wide-scale changes, chromatin regulation plays an important role in cellular differentiation, and is thought key to the complexity of higher organisms [89]. Furthermore, chromatin structure is preserved when a cell divides: hence, any chromatin modifications that occur in germ cells will be inherited by the organism’s offspring. This, in turn, underlies a number of epigenetic processes [58].

In [105] and [106], we considered the benefits of introducing an analogue of chromatin remodelling to a connectionist architecture. This works by introducing extra nodes into the network, each of which is used to regulate the activity of a chromatin frame. Using an indirect reference space (see Section 7.3), a chromatin frame covers a defined region of the network; when turned on, it prevents this part of the network from being ‘expressed’, setting the activation levels of the covered nodes to zero. Hence, different parts of the network can be turned on and off during the course

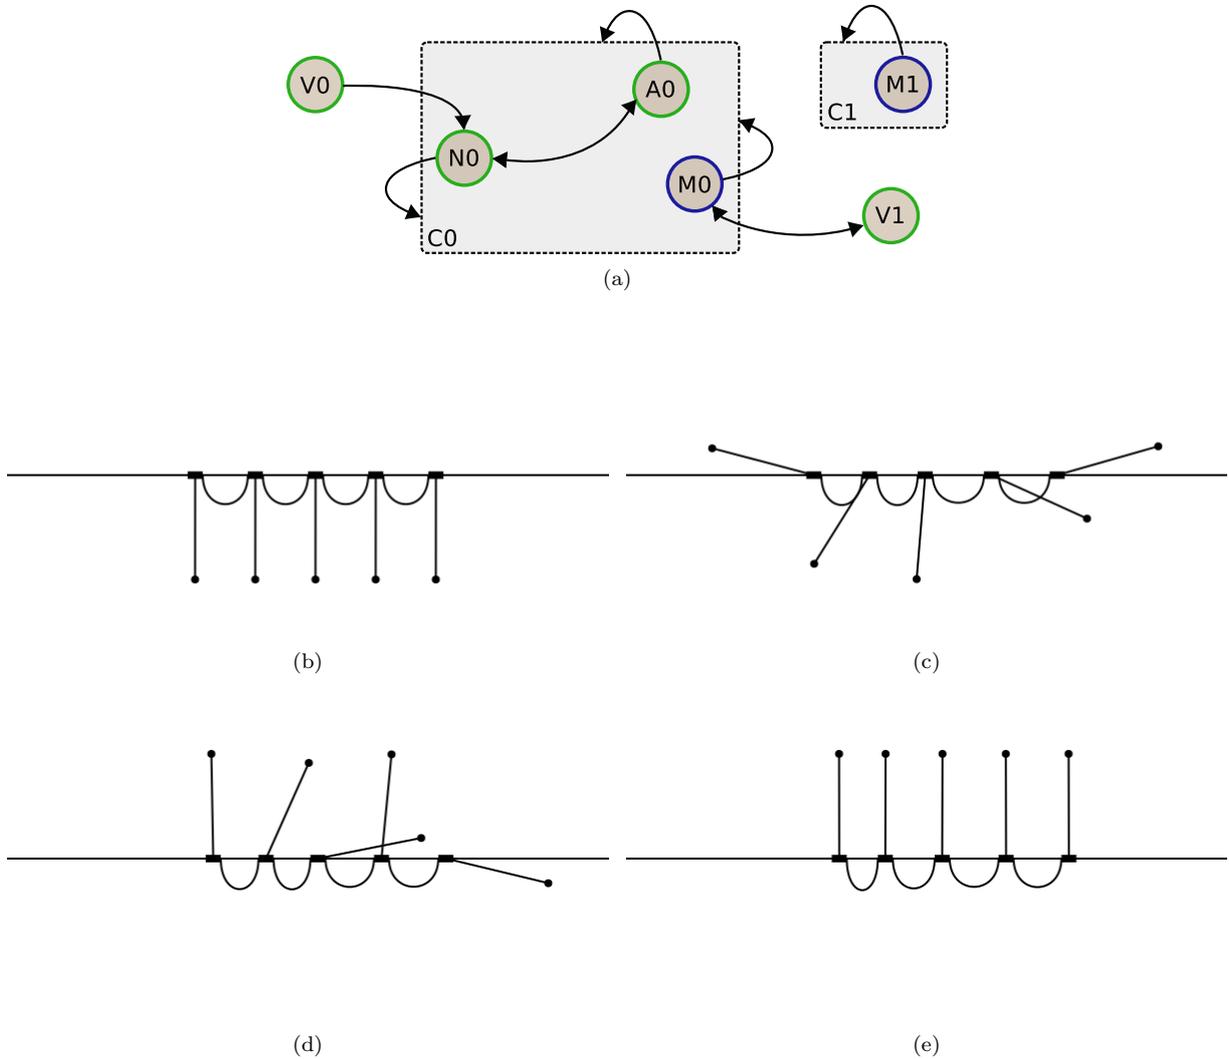


Figure 6: Example of an evolved recurrent network with an analogue of chromatin remodelling controlling the sideways movements of five coupled carts, each mounted with a pendulum. (a) The evolved controller used by each cart. C0 and C1 are chromatin analogues; when active, these deactivate the nodes within their bounds. V0 and V1 are angular velocities in the upper quadrants, and A0 is the pendulum angle. M0 and M1 are the differential motor controls for the cart. M1 and C1 are notable for forming an oscillator circuit, a common motif in the circuits evolved for this task. Starting with the pendulums hanging beneath the carts (b), the objective is to swing the pendulums to an upright position (e) and then maintain them in this position for a defined period of time. Adapted from [105].

of execution. As with the higher-order network described above, this results in the network changing structure over time. However, by allowing small groups of nodes to be added and removed, it provides a more specific mechanism for achieving this. Because of this, it is beneficial in situations when a controller must make small changes to its behaviour in order to navigate different regions of state space, and in these situations outperforms networks in which chromatin is not used. An example of this, illustrated in Figure 6, is the control of a system of coupled inverted pendulums, where the objective is to manoeuvre the pendulums to an upright position and then maintain them in this position via lateral movement of the carts to which they are mounted. Coupling between the carts makes this problem challenging [38], since collisions can readily propagate and destabilise the system. However, using an analogue of chromatin remodelling within a sigmoidal recurrent network, we were able to evolve controllers that could control five inverted pendulums at once, a behaviour that could not be achieved readily with the non-reconfigurable networks alone [105]. Analysis of the evolved networks suggests they achieve this by modulating a relatively small number of nodes. A further advantage of this approach is that the chromatin frames tend to partition the behaviours of the controller, such that manually switching a frame on or off causes transitions between the controller’s swinging and stabilising phases.

Another self-modifying process that occurs at the genomic level is caused by mobile genetic elements, such as retrotransposons and DNA transposons [90]. These are regions of DNA that are able to copy or move themselves around the genome. Whilst they sometimes encode proteins that interact with other elements of the cell’s biochemical networks, and have various evolutionary consequences, their main effect upon an organism during its lifetime is as a dynamic mutation operator. They achieve this by inserting themselves into both the coding and regulatory regions of genes, disrupting the behaviour or expression of the gene’s product within the cell. Hence, whilst both chromatin remodelling and mobile genetic elements lead to a dynamic network structure, the former achieves this by turning existing parts of the network on and off, whereas the latter achieves it by creating new nodes or by altering the connections between nodes. Some of the consequences for connectionist architectures have been explored in [13], where the author describes a simple analogue of mobile genetic elements within a random Boolean network. This is implemented using a rewiring procedure that explores different connections during the course of execution. An interesting observation made by the author is that this mechanism leads to the occurrence of more complex attractors.

9 Discussion and Conclusions

Previous research on connectionist architectures has been strongly motivated by understanding of neural function in the mammalian brain. In this paper, we have revisited, in the light of contemporary biological understanding, the idea that connectionist architectures might be similarly inspired by biochemical networks, the networks that underlie many of the complex dynamical behaviours seen within biological organisms. In particular, we have noted that these networks display a number of prominent patterns of organisation that are not seen in neural networks,

suggesting that they can guide connectionism in a way that is complementary to work on neural networks. We have structured our discussion around three ways in which biochemical networks differ from neural networks: the complexity and diversity of processes occurring at their nodes, the lack of physical wiring between these nodes, and (as a consequence of the latter) the existence of various higher-order and self-modifying processes. In each case, we have reviewed a selection of connectionist approaches motivated by these differentiating features, and noted that in general these lead to computational benefits. We have illustrated this narrative using a series of examples from our own work, showing how biochemical connectionist architectures can be successfully applied to difficult control and signal processing problems, reflecting two of the major functions of biochemical networks in cells.

Neural networks are closely associated with gradient-based learning algorithms such as back-propagation. However, in recent years, there has been a trend towards using evolutionary algorithms (e.g. NEAT [98]) to learn connectionist models. Unlike gradient-based methods, these place relatively few constraints on the space of solutions that can be explored, opening up the potential for new kinds of connectionist architectures. An underlying theme of this paper is that these new architectures need not be based on neural systems. Indeed, the close relationship between biochemical networks and biological evolution suggests that connectionist architectures motivated by biochemical organisation may be particularly well suited for use with evolutionary algorithms. Even so, this does not mean that it is necessary to follow a wholesale approach to biochemical connectionism in order to capture the evolvability of biochemical systems, since techniques such as template matching (see Section 7.3) could be readily applied to existing neural architectures.

This is also true more generally, since many of the methods we have discussed could in principle be applied to existing neural network models. For example, discrete maps could be used as activation functions within neurones, augmenting a sigmoidal neurone population—and our results suggest that this kind of heterogeneity is generally beneficial [63]. Likewise, patterns of coupling seen between biochemical networks could be used to couple together neural networks. Certainly, the coupling together of heterogeneous neural networks is a more accurate model of the brain, and may provide a better way of realising heterogeneous processes occurring over different timescales. There is also no reason why self-modifying processes could not be applied to neural networks. Self-modification appears to be an important means through which biochemical systems achieve adaptive transitions in behaviour, and we envisage that it could carry out similar roles in other connectionist systems.

Finally, it should be noted that we have only discussed a few of the most evident patterns of organisation seen in biochemical networks, and even then we have only considered a few of the possible lessons for connectionist architectures. Factors such as the spatial organisation of pathways, the role of conserved network motifs [5], and the existence of various epigenetic processes [9], to name but a few, all seem important for the evolution and function of biochemical networks, and may have similarly beneficial lessons for connectionist architectures. As with neural networks, it seems likely that progress in biology will further refine our understanding of these systems, and will bring new general insights in the future. Equally, analysis of specific biochemical pathways [18] will likely inform us

about specific biological solutions to various problems [19], suggesting connectionist architectures suitable for solving these kind of problems in engineered systems.

Acknowledgements

This research was supported by the EPSRC under the grant “Artificial Biochemical Networks: Computational Models and Architectures” (ref. EP/F060041/1).

References

- [1] Andy Adamatzky, Ben De Lacy Costello, and Tetsuya Asai. *Reaction-Diffusion Computers*. Elsevier, 2005.
- [2] L. M. Adleman. Molecular computation of solutions to combinatorial problems. *Science*, 266(5187):1021–1024, 1994.
- [3] Réka Albert and Hans G. Othmer. The topology of the regulatory interactions predicts the expression pattern of the segment polarity genes in drosophila melanogaster. *Journal of Theoretical Biology*, 223(1):1–18, 2003.
- [4] M. Aldana, E. Balleza, S. Kauffman, O. Resendiz, et al. Robustness and evolvability in genetic regulatory networks. *Journal of theoretical biology*, 245(3):433–448, 2007.
- [5] Uri Alon. Network motifs: theory and experimental approaches. *Nat Rev Genet*, 8(6):450–461, Jun 2007.
- [6] W. Banzhaf. Artificial chemistries—towards constructive dynamical systems. *Solid State Phenomena*, 97/98:43–50, 2004.
- [7] Wolfgang Banzhaf. Artificial regulatory networks and genetic programming. In Rick L. Riolo and Bill Worzel, editors, *Genetic Programming Theory and Practice*, chapter 4, pages 43–62. Kluwer, 2003.
- [8] M. J. Berridge. *Cell Signalling Biology*, chapter Spatial and Temporal Aspects of Signalling. Portland Press Limited, 2012.
- [9] A. Bird. Perceptions of epigenetics. *Nature*, 447(7143):396–398, 2007.
- [10] H. Bolouri. *Computational Modeling of Gene Regulatory Networks—A Primer*. Imperial College Press, 2008.
- [11] L. Bull. Evolving Boolean networks on tunable fitness landscapes. *Evolutionary Computation, IEEE Transactions on*, 16(6):817–828, 2012.
- [12] Larry Bull. A simple computational cell: coupling Boolean gene and protein networks. *Artificial Life*, 18(2):223–236, 2012.

- [13] Larry Bull. Consideration of mobile DNA: new forms of artificial genetic regulatory networks. *Natural Computing*, 2013.
- [14] Larry Bull and Richard Preen. On dynamical genetic programming: Random Boolean networks in learning classifier systems. In Leonardo Vanneschi et al., editors, *Genetic Programming*, volume 5481 of *Lecture Notes in Computer Science*, pages 37–48. Springer Berlin / Heidelberg, 2009.
- [15] Hugues Chaté and Jerome Losson. Non-trivial collective behavior in coupled map lattices: A transfer operator perspective. *Physica D: Nonlinear Phenomena*, 103(1-4):51 – 72, 1997.
- [16] Cedric R Clapier and Bradley R Cairns. The biology of chromatin remodeling complexes. *Annual review of biochemistry*, 78:273–304, 2009.
- [17] M. Conrad. The geometry of evolution. *BioSystems*, 24:61–81, 1990.
- [18] David Croft, Gavin O’Kelly, Guanming Wu, et al. Reactome: a database of reactions, pathways and biological processes. *Nucleic Acids Research*, 39(suppl 1):D691–D697, 2011.
- [19] W. de Ronde, F. Tostevin, and P. R. ten Wolde. Multiplexing biochemical signals. *Phys. Rev. Lett.*, 107(4):048101, Jul 2011.
- [20] James Decraene, George G. Mitchell, and Barry McMullin. Evolving artificial cell signaling networks: Perspectives and methods. In Falko Dressler and Iacopo Carreras, editors, *Advances in Biologically Inspired Information Systems*, pages 167–186. Springer, 2007.
- [21] P Dittrich, J Ziegler, and W Banzhaf. Artificial chemistries—a review. *Artificial Life*, 7:225–275, 2001.
- [22] Elena Dubrova, Maxim Teslenko, and Hannu Tenhunen. A computational scheme based on random Boolean networks. In Corrado Priami et al., editors, *Transactions on Computational Systems Biology X*, volume 5410 of *Lecture Notes in Computer Science*, pages 41–58. Springer Berlin / Heidelberg, 2008.
- [23] W. Duch and K. Grabczewski. Heterogeneous adaptive systems. In *Neural Networks, 2002. IJCNN’02. Proceedings of the 2002 International Joint Conference on*, volume 1, pages 524–529. IEEE, 2002.
- [24] W. Duch and N. Jankowski. Survey of neural transfer functions. *Neural Computing Surveys*, 2(1):163–212, 1999.
- [25] J.C. Eccles, M. Ito, and J. Szentágothai. *The cerebellum as a neuronal machine*. Springer-Verlag, Berlin, 1967.
- [26] Attila Egri-Nagy, Paolo Dini, Chrystopher L. Nehaniv, and Maria J. Schilstra. Transformation semigroups as constructive dynamical spaces. In Fernando Antonio Basile Colugnati, Lia Carrari Rodrigues Lopes, and Saulo Faria Almeida Barretto, editors, *Digital Ecosystems*, volume 67 of *Lecture Notes of the Institute for Computer*

- Sciences, Social Informatics and Telecommunications Engineering*, pages 245–265. Springer Berlin Heidelberg, 2010.
- [27] J. Doyne Farmer. A rosetta stone for connectionism. *Physica D: Nonlinear Phenomena*, 42(13):153 – 187, 1990.
- [28] Adam Faulconbridge, Susan Stepney, Julian F. Miller, and Leo S. D. Caves. RBN-World: A sub-symbolic artificial chemistry. In George Kampis, István Karsai, and Eörs Szathmáry, editors, *Advances in Artificial Life. Darwin Meets von Neumann, Part 1*, volume 5777 of *Lecture Notes in Computer Science*, pages 377–384. Springer, 2012.
- [29] Peter Fine, Ezequiel Paolo, and Andrew Philippides. Spatially constrained networks and the evolution of modular control systems. In Stefano Nolfi et al., editors, *From Animals to Animats 9*, volume 4095 of *Lecture Notes in Computer Science*, pages 546–557. Springer Berlin Heidelberg, 2006.
- [30] D. Floreano, F. Mondada, et al. Evolution of plastic neurocontrollers for situated agents. In *From Animals to Animats IV: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, volume 4. MIT Press, Cambridge, MA, USA, 1996.
- [31] W. Fontana. Algorithmic chemistry. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, pages 159–210. Addison-Wesley, 1992.
- [32] Raimo Franke, Fabian J Theis, and Steffen Klamt. From binary to multivalued to continuous models: The lac operon as a case study. *Journal of Integrative Bioinformatics*, 7(1):151, 2010.
- [33] Luis A. Fuente, Michael A. Lones, Alexander P. Turner, Leo S. Caves, Susan Stepney, and Andy M. Tyrrell. Adaptive robotic gait control using coupled artificial signalling networks, hopf oscillators and inverse kinematics. In *Proc. 2013 IEEE Congress on Evolutionary Computation (CEC 2013)*. IEEE Press, 2013.
- [34] Luis A. Fuente, Michael A. Lones, Alexander P. Turner, Susan Stepney, Leo S. Caves, and Andy M. Tyrrell. Computational models of signalling networks for non-linear control. *Biosystems*, 112(2):122 – 130, 2013.
- [35] K. Funahashi and Y. Nakamura. Approximation of dynamical systems by continuous time recurrent neural networks. *Neural networks*, 6(6):801–806, 1993.
- [36] W. Gerstner and W.M. Kistler. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge University Press, 2002.
- [37] Matthew C. Good, Jesse G. Zalatan, and Wendell A. Lim. Scaffold proteins: hubs for controlling the flow of cellular information. *Science Signaling*, 332(6030):680, 2011.

- [38] Heiko Hamann, Thomas Schmickl, and Karl Crailsheim. Coupled inverted pendulums: a benchmark for evolving decentral controllers in modular robotics. In *Proceedings of the 13th Annual Genetic and Evolutionary Computation Conference, GECCO*, pages 195–202. ACM New York, NY, USA, 2011.
- [39] J. Hancock. *Cell Signalling*. Oxford University Press, 2010.
- [40] S. Hickinbotham, S. Stepney, A. Nellis, T. Clarke, E. Clark, M. Pay, and P. Young. Embodied genomes and metaprogramming. In T. Lenaerts et al., editors, *Advances in Artificial Life, ECAL 2011: Proc. 11th European Conference on the Synthesis and Simulation of Living Systems*, pages 334–341. MIT Press, 2011.
- [41] J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.
- [42] Ting Hu and Wolfgang Banzhaf. Evolvability and speed of evolutionary algorithms in light of recent developments in biology. *Journal of Artificial Evolution and Applications*, 2010:568375, 2010.
- [43] P. Husbands, T. Smith, N. Jakobi, and M. O’Shea. Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, 10(3-4):185–210, 1998.
- [44] H. Jaeger. Adaptive nonlinear system identification with echo state networks. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 593–600. MIT Press, Cambridge, MA, 2003.
- [45] Michael I. Jordan. Attractor dynamics and parallelism in a connectionist sequential machine. In Joachim Diederich, editor, *Artificial neural networks*, pages 112–127. IEEE Press, Piscataway, NJ, USA, 1990.
- [46] Andrew R Joyce and Bernhard Ø Palsson. The model organism as a system: integrating ‘omics’ data sets. *Nature Reviews Molecular Cell Biology*, 7(3):198–210, 2006.
- [47] S. A. Kauffman. Metabolic stability and epigenesis in randomly constructed genetic nets. *J Theor Biol*, 22(3):437–467, Mar 1969.
- [48] S. A. Kauffman. *The origins of order: Self-organization and selection in evolution*. Oxford University Press, USA, 1993.
- [49] M. M. Khan, D. R. Lester, L. A. Plana, A. Rast, X. Jin, E. Painkras, and S. B. Furber. SpiNNaker: mapping neural networks onto a massively-parallel chip multiprocessor. In *Neural Networks, 2008. IJCNN 2008. (IEEE World Congress on Computational Intelligence)*. *IEEE International Joint Conference on*, pages 2849–2856. IEEE Press, 2008.
- [50] Boris N Kholodenko. Cell-signalling dynamics in time and space. *Nature reviews Molecular cell biology*, 7(3):165–176, 2006.

- [51] Seyun Kim and Pierre A Coulombe. Emerging role for the cytoskeleton as an organizer and regulator of translation. *Nature Reviews Molecular Cell Biology*, 11(1):75–81, 2010.
- [52] T. Kind, M. Scholz, and O. Fiehn. How large is the metabolome? A critical analysis of data exchange practices in chemistry. *PloS one*, 4(5):e5440, 2009.
- [53] M. Kirschner and J. Gerhart. Evolvability. *Proceedings of the National Academy of Science (USA)*, 95:8420–8427, July 1998.
- [54] T. Kohonen. Self-organized formation of topologically correct feature maps. *Biological cybernetics*, 43(1):59–69, 1982.
- [55] Suzanne Komili and Pamela A Silver. Coupling and coordination in gene expression processes: a systems biology view. *Nature Reviews Genetics*, 9(1):38–48, 2008.
- [56] Eugene V. Koonin, Yuri I. Wolf, Georgy P. Karev, Pau Fernández, and Ricard V. Solé. The role of computation in complex regulatory networks. In *Power Laws, Scale-Free Networks and Genome Biology*, Molecular Biology Intelligence Unit, pages 206–225. Springer US, 2006.
- [57] Roger D. Kornberg. The molecular basis of eukaryotic transcription. *Proceedings of the National Academy of Sciences*, 104(32):12955–12961, 2007.
- [58] Satya K. Kota and Robert Feil. Epigenetic transitions in germ cell development and meiosis. *Developmental Cell*, 19(5):675 – 686, 2010.
- [59] V. Lacroix, L. Cottret, P. Thebault, and M.-F. Sagot. An introduction to metabolic networks and their structural analysis. *Computational Biology and Bioinformatics, IEEE/ACM Transactions on*, 5(4):594 –617, oct.-dec. 2008.
- [60] A. Leier, P.D. Kuo, and W. Banzhaf. Analysis of preferential network motif generation in an artificial regulatory network model created by duplication and divergence. *Advances in Complex Systems*, 10(02):155–172, 2007.
- [61] Emmanuel D Levy, Christian R Landry, and Stephen W Michnick. Signaling through cooperation. *Science Signaling*, 328(5981):983, 2010.
- [62] M. A. Lones. *Enzyme Genetic Programming: Modelling Biological Evolvability in Genetic Programming*. PhD thesis, Department of Electronics, University of York, 2003.
- [63] M. A. Lones, L. A. Fuente, A. P. Turner, L. S. D. Caves, S. Stepney, S. L. Smith, and A. M. Tyrrell. Artificial biochemical networks: Evolving dynamical systems to control dynamical systems. *Evolutionary Computation, IEEE Transactions on*, 2013. in press.

- [64] M. A. Lones, S. L. Smith, J. E. Alty, S. E. Lacy, K. L. Possin, D. R. S. Jamieson, and A. M. Tyrrell. Evolving classifiers to recognise the movement characteristics of Parkinson’s disease patients. *Evolutionary Computation, IEEE Transactions on*, 2013. in press.
- [65] M. A. Lones, S. L. Smith, A. M. Tyrrell, J. E. Alty, and D. R. S. Jamieson. Evolving computational dynamical systems to recognise abnormal human motor function. In M. A. Lones et al., editors, *Information Processing in Cells and Tissues, Proc. 9th Int. Conf.*, volume 7223 of *Lecture Notes in Computer Science*, pages 177–182. Springer, March 2012.
- [66] M. A. Lones, S. L. Smith, A. M. Tyrrell, J. E. Alty, and D. R. S. Jamieson. Characterising neurological time series data using biologically-motivated networks of coupled discrete maps. *BioSystems*, 112(2):94–101, 2013.
- [67] M. A. Lones and A. M. Tyrrell. Modelling biological evolvability: Implicit context and variation filtering in enzyme genetic programming. *BioSystems*, 76(1):229–238, 2004.
- [68] M. A. Lones, A. M. Tyrrell, S. Stepney, and L. S. Caves. Controlling complex dynamics with artificial biochemical networks. In A. I. Esparcia-Alczar et al., editors, *Proc. 2010 European Conference on Genetic Programming (EuroGP 2010)*, volume 6021 of *Lecture Notes in Computer Science*, pages 159–170. Springer Berlin / Heidelberg, 2010.
- [69] M. A. Lones, A. M. Tyrrell, S. Stepney, and L. S. D. Caves. Controlling legged robots with coupled artificial biochemical networks. In Tom Lenaerts et al., editors, *Advances in Artificial Life, ECAL 2011: Proc. 11th European Conference on the Synthesis and Simulation of Living Systems*, pages 465–472. MIT Press, August 2011.
- [70] W. Maass. Networks of spiking neurons: the third generation of neural network models. *Neural Networks*, 10(9):1659–1671, 1997.
- [71] Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, 14(11):2531–2560, 2002.
- [72] S. Mabu, K. Hirasawa, and J. Hu. A graph-based evolutionary algorithm: genetic network programming (GNP) and its extension using reinforcement learning. *Evolutionary Computation*, 15(3):369–398, 2007.
- [73] C. MacLeod and N.F. Capanni. Artificial biochemical networks: a different connectionist paradigm. *Artificial intelligence review*, 33(1):123–134, 2010.
- [74] P. C. Marijuán. Enzymes, artificial cells and the nature of biological information. *BioSystems*, 35:167–170, 1995.
- [75] R. M. May. Simple mathematical models with very complicated dynamics. *Nature*, 261:459–467, 1976.

- [76] H.H. McAdams and A. Arkin. It's a noisy business! genetic regulation at the nanomolar scale. *Trends in Genetics*, 15(2):65–69, 1999.
- [77] W. S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4):115–133, 1943.
- [78] David A Medler. A brief history of connectionism. *Neural Computing Surveys*, 1:18–72, 1998.
- [79] Gunter Meister and Thomas Tuschl. Mechanisms of gene silencing by double-stranded RNA. *Nature*, 431:343–349, 2004.
- [80] M. Minsky and P. Seymour. *Perceptrons: an introduction to computational geometry*. MIT press, 1969.
- [81] M. Mokhtar, D. M. Halliday, and A. M. Tyrrell. Hippocampus neurons and place cells/place field representation to provide path navigation. In *Neural Networks, 2007. IJCNN 2007. International Joint Conference on*, pages 795–800. IEEE Press, aug. 2007.
- [82] M. Natarajan, K.M. Lin, R.C. Hsueh, P.C. Sternweis, and R. Ranganathan. A global analysis of cross-talk in a mammalian cellular signalling network. *Nature Cell Biology*, 8(6):571–580, 2006.
- [83] A.C. Newton. Protein kinase C: structure, function, and regulation. *Journal of Biological Chemistry*, 270(48):28495–28498, 1995.
- [84] S. Ohno et al. *Evolution by gene duplication*. London: George Alien & Unwin Ltd. Berlin, Heidelberg and New York: Springer-Verlag., 1970.
- [85] T. Ohyama. *DNA conformation and transcription*. Molecular Biology Intelligence Unit. Springer US, 2005.
- [86] Gh. Păun. Computing with membranes. *Journal of Computer and System Sciences*, 61(1):108–143, 2000.
- [87] A. M. Poole, M. J. Philips, and D. Penny. Prokaryote and eukaryote evolvability. *BioSystems*, 69:163–186, 2003.
- [88] A.B. Porto-Pazos, N. Veiguela, P. Mesejo, M. Navarrete, A. Alvarellos, O. Ibáñez, A. Pazos, and A. Araque. Artificial astrocytes improve neural network performance. *PloS one*, 6(4):e19109, 2011.
- [89] T.P. Rasmussen et al. Embryonic stem cell differentiation: a chromatin perspective. *Reprod Biol Endocrinol*, 1:100, 2003.
- [90] Rita Rebollo, Mark T. Romanish, and Dixie L. Mager. Transposable elements: An abundant and natural source of regulatory sequences for host genes. *Annual Review of Genetics*, 46(1):21–42, 2012.

- [91] Torsten Reil. Dynamics of gene expression in an artificial genome—implications for biological and artificial ontogeny. In *Proceedings of the 5th European Conference on Artificial Life (ECAL'99)*, volume 1674 of *Lecture Notes in Artificial Intelligence*, pages 457–466. Springer-Verlag, 1999.
- [92] J. Ricard and A. Cornish-Bowden. Co-operative and allosteric enzymes: 20 years on. *European Journal of Biochemistry*, 166(2):255–272, 2005.
- [93] Frank Rosenblatt. The perceptron. *Psych. Rev.*, 65(6):386–408, 1958.
- [94] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.
- [95] D.E. Rumelhart and J.L. McClelland. *Parallel distributed processing: explorations in the microstructure of cognition. Volume 1. Foundations*. MIT Press, Cambridge, Ma, 1986.
- [96] I. Shmulevich, E.R. Dougherty, and W. Zhang. From Boolean to probabilistic Boolean networks as models of genetic regulatory networks. *Proceedings of the IEEE*, 90(11):1778–1792, 2002.
- [97] K. O. Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines*, 8(2):131–162, 2007.
- [98] Kenneth O. Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127, 2002.
- [99] A. Szejka and B. Drossel. Evolution of canalizing Boolean networks. *The European Physical Journal B-Condensed Matter and Complex Systems*, 56(4):373–380, 2007.
- [100] Christof Teuscher. *Turing's connectionism: an investigation of neural network architectures*. Springer Verlag, 2002.
- [101] Soichiro Tsuda, Stefan Artmann, and Klaus-Peter Zauner. The phi-bot: A robot controlled by a slime mould. In *Artificial Life Models in Hardware*, pages 213–232. Springer London, 2009.
- [102] E. Tuci, M. Quinn, and I. Harvey. Evolving fixed-weight networks for learning robots. In *Evolutionary Computation, 2002. CEC'02. Proceedings of the 2002 Congress on*, volume 2, pages 1970–1975. IEEE, 2002.
- [103] A. M. Turing. Intelligent machinery. Technical report, National Physical Laboratory, 1948.
- [104] A. M. Turing. The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 237(641):37–72, 1952.
- [105] A. P. Turner, M. A. Lones, L. A. Fuente, S. Stepney, L. S. D. Caves, and A. M. Tyrrell. The artificial epigenetic network. In *Proc. 2013 IEEE Symposium Series on Computational Intelligence (SSCI 2013)*. IEEE Press, 2013.

- [106] A. P. Turner, M. A. Lones, L. A. Fuente, A. M. Tyrrell, S. Stepney, and L. S. D. Caves. Controlling complex tasks using artificial epigenetic regulatory networks. *BioSystems*, 112(2):56–62, 2013. to appear.
- [107] S. Ulam. Random processes and transformations. In *Proc. Int. Congress of Mathematicians*, volume 2, pages 264 – 275. American Mathematical Society, 1952.
- [108] Patricia A. Vargas, Ezequiel A. Paolo, and Phil Husbands. Preliminary investigations on the evolvability of a non-spatial GasNet models. In Fernando Almeida e Costa et al., editors, *Advances in Artificial Life*, volume 4648 of *Lecture Notes in Computer Science*, pages 966–975. Springer Berlin Heidelberg, 2007.
- [109] L.G. Volkert. Enhancing evolvability with mutation buffering mediated through multiple weak interactions. *Biosystems*, 69(2-3):127–142, 2003.
- [110] C.C. Walker and W.R. Ashby. On temporal characteristics of behavior in certain complex systems. *Kybernetik*, 3(2):100–108, 1966.
- [111] K.E. Wellen and C.B. Thompson. A two-way street: reciprocal regulation of metabolism and signalling. *Nature Reviews Molecular Cell Biology*, 2012.
- [112] Maxwell Z Wilson and Zemer Gitai. Beyond the cytoskeleton: mesoscale assemblies and their function in spatial organization. *Current opinion in microbiology*, 16(2):177–183, 2013.
- [113] D.M. Wittmann, J. Krumsiek, J. Saez-Rodriguez, D.A. Lauffenburger, S. Klamt, and F.J. Theis. Transforming Boolean models to continuous models: methodology and application to T-cell receptor signaling. *BMC systems biology*, 3(1):98, 2009.
- [114] C.L. Woodcock and R.P. Ghosh. Chromatin higher-order structure and dynamics. *Cold Spring Harbor perspectives in biology*, 2(5), 2010.
- [115] Brian M. Yamauchi and Randall D. Beer. Sequential behavior and learning in evolved dynamical neural networks. *Adaptive Behavior*, 2(3):219–246, 1994.
- [116] X. Yao. Evolving artificial neural networks. *Proceedings of the IEEE*, 87(9):1423–1447, 1999.
- [117] Claudio Zandron, Claudio Ferretti, and Giancarlo Mauri. Solving NP-complete problems using P systems with active membranes. In *Unconventional Models of Computation, UMC'2K*, pages 289–301. Springer, 2001.
- [118] J.G.T. Zanudo, M. Aldana, and G. Martínez-Mekler. Boolean threshold networks: Virtues and limitations for biological modeling. *Information Processing and Biological Systems*, pages 113–151, 2011.