

Chapter 22

Reservoir Computing as a Model for *In-Materio* Computing

Matthew Dale, Julian F. Miller and Susan Stepney

Abstract Research in substrate-based computing has shown that materials contain rich properties that can be exploited to solve computational problems. One such technique known as *Evolution-in-materio* uses evolutionary algorithms to manipulate material substrates for computation. However, in general, modelling the computational processes occurring in such systems is a difficult task and understanding what part of the embodied system is doing the computation is still fairly ill-defined. This chapter discusses the prospects of using *Reservoir Computing* as a model for *in-materio* computing, introducing new training techniques (taken from Reservoir Computing) that could overcome training difficulties found in the current *Evolution-in-Materio* technique.

22.1 Introduction

Biological organisms vastly outperform classical/conventional computing paradigms in many respects, from possessing inherent fault-tolerance to constructing highly parallel machines. Much of this is achieved by exploiting physicality and by sharing and distributing computational effort throughout the spatial system. As such, they exploit physical interactions through feedback with the real-world, utilising features such as their own morphology.

Many of these systems comprise relatively simple elements that emerge and coalesce into more complex, but robust, structural layers across different scales. Such

M. Dale (✉) · S. Stepney
Department of Computer Science, University of York, York, UK
e-mail: md596@york.ac.uk

J.F. Miller
Department of Electronics, University of York, York, UK
e-mail: julian.miller@york.ac.uk

S. Stepney
York Centre for Complex Systems Analysis, University of York, York, UK
e-mail: susan.stepney@york.ac.uk

grounding properties (and many more) have enabled these complex systems to thrive and evolve, adapting and co-evolving in real-time with their local ecosystem.

The conventional von Neumann computing architecture, although expertly refined over time, poses some fundamental inefficiencies. For example, classical computers require the transformation between high-level languages to low-level machine code, a process that requires layers of conversion through a compiler stack, making it computationally costly, slow, and highly susceptible to faults and errors. These systems typically succumb to many issues in speed, from both an inability to deal with concurrent computations and the bottleneck created by the transfer of data between separate memory and processing entities. Because of these architectural weaknesses other intertwined discrepancies arise, such as an increase in power consumption, system size, and top-down design complexity.

Unconventional computing tries to address some of these limitations by attempting to provide alternative architectures and systems that typically exploit the underlying physics and many-scale interactions of the real-world. Many forms of unconventional systems have been explored in recent years. For example, quantum computing is one such system, where two-state quantum bits, typically described by electron or photon spin/polarisation, can be exploited to perform large numbers of parallel computations (i.e. $2^{no.of\ qubits}$). This is achieved through the principles of superposition, enabling a qubit to be in both states simultaneously (see [104] for a review). Another example is reaction-diffusion computing [2], which performs computation through local chemical reactions and diffusion. By using chemical processes, the system can execute highly parallel computations, performed by the complex interactions of propagated waves of information caused by local disturbances. *Physarum polycephalum* (slime mould) is currently under investigation as an excitable, reaction-diffusion medium that could form the basis of a programmable amorphous biological computer [3, 123].

For the purpose of this chapter we discuss research into configurable *in-materio* systems (typically on the nanoscale), with the intention of identifying a computational model that could suit and conform to many *in-materio* systems.

This chapter begins with an introduction to the concepts of material computation with a brief mention of the criticality hypothesis (Sect. 22.2). Next, the field of Evolution-*in-materio* as a current investigation into the theory of material computation is reviewed, discussing its methodology and current research (Sect. 22.3). Reservoir Computing and its potential as an *in-materio* model is examined in Sects. 22.4 and 22.5. To further aid in identifying “good” reservoirs some evaluation techniques are highlighted in Sect. 22.6. Lastly, an argument is made in Sect. 22.7 to the possible benefits reservoir computing could have on the design of computational substrates.

22.2 Material Computation

This section discusses some of the principal concepts of material computation, such as some informal definitions, the role critical dynamics has in complexity and where it may fit within other computing paradigms.

Material computation is defined to be a computational process that occurs in a sufficiently complex dynamical system realised in the form of a physical material substrate.

This definition relies on the idea that matter contains physical information and that systems of collective matter can dynamically modulate and redistribute information to change state; therefore, implied here, to perform some form of computation. A simple example of this would be the physical process in which the state or phase of matter changes in relation to energy, which could (loosely) in some sense be viewed as a computational mechanism.

Stepney et al. [111] provide a more contextual definition of material computation as: “computation directly by physical and chemical processes of a complex substrate, with little or no abstraction to a virtual machine”. From this definition we can categorise material computation as an analogue process that utilises physical constructs, the tangible medium itself, and meta-processing to do computation. This “physicality” may be defined and observed as the structural topology, characteristic behaviours and information processing associated with the many-scale interactions occurring in that system.

As a concept, material computation is still in its developmental or conceptual stage, with early experiments supporting, or working in tandem with, current digital technology to form hybrid, and potentially very powerful, machines. A configurable substrate can be used to transfer some of the computational burden from the digital system to the material [111]. As such, the material can endow the system with many of the properties and advantages of analogue systems, such as speed and concurrency, in a device where memory and processing are not separable. To achieve this, engineers attempt to exploit processes and behavioural phenomena that naturally occur, properties that are governed directly by the underlying physics and chemistries of the substrate.

Exploiting computation directly from materials offers many potential advantages over classical systems where the computation performed does not depend so much on the details of the materials used. As a paradigm, substrate computing potentially offers vast amounts of computational power by capitalising on the massive parallelism and natural constraints of these rich systems. Such properties are suggested to have the potential to provide solutions “for free”, or at least computationally cheaper, and provide a rich explorable state space, aligning computation to particular trajectories [108].

Much of the current interest in material computation is to abstract a model (or models) of computation from what the substrate does naturally. It has been proposed

that the first step to producing a potential unified theory of material computation, i.e. a theory that better understands what computation is and how it occurs in materials, should take place in “primitive” (un-evolved) substrates, where the general principles are in plain sight [108]. After this, material computation could emerge with some supportive reasoning to a better understanding of computation in biological substrates.

22.2.1 *The Effects of Criticality on Complexity*

It is hypothesised that there is a critical state in which a system can exhibit maximal computational power, and therefore where maximal complexity can be acquired, labelled a region near (or at) the *edge of chaos* [60]. This concept may have a direct relationship to material computation, whereby a material can exhibit “richness” [75], and therefore be exploitable, by operating close to or within this region.

This *edge of chaos* represents the transitional border between ordered dynamics, where perturbations to the system quickly fade and the system settles, and chaotic behaviour, where perturbations significantly affect long-term stability and the system becomes unpredictable. This critical landscape can be observed by looking at the systems trajectory in the phase space by monitoring the convergence towards or away from a steady state, and thus highlighting a system’s sensitivity to initial conditions. Both behaviours are thought to be necessary to gain maximal complexity, using some ordered behaviour to maintain memory, and some chaotic behaviour to enable processing.

Langton [60] observed the effects and advantages of systems working in this transitional region, using cellular automata. At a critical point, Langton observed that a cellular automaton could optimally perform computations, imitating complex life-like behaviour. Earlier, Packard [86] observed another unique property: that genetic algorithms tend to evolve populations in these critical regions, suggesting that adaptability was therefore optimised close to the edge of chaos. Similar conclusions are proposed and demonstrated in neural networks, where vast computational power and capability in this region is described through network connectivity. Bertschinger and Natschläger [11] demonstrate these relationships in input-driven networks, by accurately determining the position of the critical line with respect to structural parameters.

It has also been suggested that *living* neural networks support the “criticality hypothesis”. Beggs [9] discusses this notion by looking at how the power-law distribution of neuronal avalanche sizes (a cascade of bursts of activity) suggests operation near a critical point. Beggs further explains, that the implications of these avalanche size distributions implies that information transmission, information storage, computational power and stability could all simultaneously optimise at the edge of chaos.

22.2.2 Configuration and Structure

Conventional programs and algorithms represent idealised mathematical objects, irrespective of their underlying hardware. In a physical system (say a biological system) computation is embodied, behaviour may not be completely captured by a closed mathematical model. As such, trying to program these embodied systems requires different techniques. The programming and manipulation of materials requires (to some extent) a complex understanding of the properties and interactions within that system. We therefore require either some convoluted top-down “programming” approach (in the traditional sense) or an alternative mechanism (e.g. through training, learning or evolution). Whichever method is applied, this “program” would alter the details of system in a controlled manner, for example, using controlling fields that affect structure and dynamics.

22.3 Evolution-In-Materio

Evolution-*in-materio* (EIM) is a term coined by Miller and Downing [75] to refer to the means by which a physical system, a complex material, could be manipulated by computer controlled evolution (CCE) to perform useful computation.

The idea of using unconstrained evolution as a search method in physical media is deep rooted in the field of Evolvable Hardware (EH) [34, 36]. Most evolved configurations in EH lead to digital products or components later embedded into physical artefacts. For example, evolving simulated models and optimisation programs, or, designing a physical system that can be manufactured after evolution; like antennas [65], robots [64, 90] or chemical systems made of oil droplets [35].

Miller argues that Evolution-*in-materio* sits between full embodiment and the realised evolved EH solutions described above [76]. In this form, physical artefacts are configured (or conceptually created) during the process and assessed/controlled by simulated Darwinian evolution. This can also be seen in some EH systems, but, typically such systems are limited to constrained silicon hardware, e.g. electronic circuits evolved on Field Programmable Gate Arrays (FPGA) [43]. EIM relies on a hybrid analogue/digital architecture where the evolutionary process (encapsulated on a digital computer) controls the writing/reading of physical signals to/from an analogue material. As such, the directed search tries to exploit the dynamics of the material by evaluating the performance of individual test configurations. Physical realisations are therefore embodied in the search process but evaluated externally. A system that operates in this manner is theoretically very powerful, allowing the manipulation of physical properties which are hitherto unknown.

An early example of EIM can be found in Thompson’s work with FPGAs [114]. Thompson attempted to evolve a frequency discriminator by allowing evolution to reconfigure circuit elements on the FPGA. In the process, he discovered that evolution had in fact used subtle electrical variations in the underlying material to form a

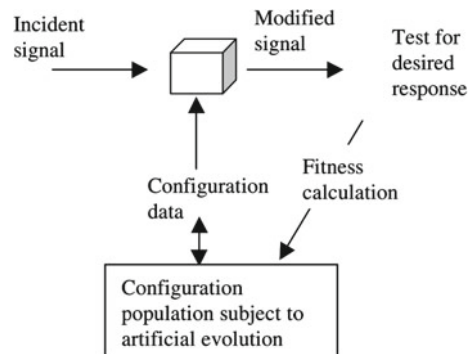
solution. It was only made evident when evolved configurations no longer solved the problem when moved around the FPGA and when areas not directly involved somehow contributed to the overall operation. Thompson's work led to an explosion of interest [36, 76]. Miller and Downing [75] label this work as a "starting point" for the exploration of intrinsic evolution as a means by which to exploit the natural properties of materials for computation.

22.3.1 Computation in Liquid Crystal

Miller and Downing [75] discuss several interesting materials that could possess the desired characteristics needed for both computation and evolution. These include, liquid crystal, conducting and electro-active polymers, voltage controlled colloids, nanoparticle suspensions and irradiated or damaged semiconductors [75]. These materials are exploited using a device that can alter the material's function through configuration parameters (discrete signals), using a configurable analogue processor (CAP) (see Fig. 22.1).

Liquid Crystal (LC) has a number of advantages for readily applying the theory. LC contains several key features including, wide availability, addressable using digital voltages, exhibits emergent behaviour, has a unique mesomorphic structure between ordered and disordered, and can relax to an initial base state. Harding and Miller [39] adopt liquid crystal as a basis material and construct a bespoke platform to solve multiple computational problems. The hardware houses a liquid crystal display (LCD) and an array of dynamically selectable input/output connections to both the LCD and external measurement devices. They demonstrate liquid crystal as an efficient evolvable material where relatively small numbers of generations can produce effective solutions. Over the course of their investigation the LC system has been applied to three separate tasks; tone discrimination [39], creating logic gates [41] and a real-time robot controller [40].

Fig. 22.1 Configurable Analogue Processor (CAP): using an evolutionary guided search the material is reconfigured to solve some computational task through applied input signals [75]



Harding and Miller found that a *rich* substrate of liquid crystal supplied many more exploitable properties compared to conventional silicon hardware (i.e. Thompson's FPGA). This in turn increased the diversity of solutions, thus increasing its evolvability. Harding and Miller's work demonstrated the advantages of emergent design by configuring intractable characteristic properties with no knowledge of their existence. But, the experiments also raised other fundamental questions on applying techniques from the *Evolution-in-materio* paradigm. For example, the length of time needed to "program" materials, i.e. how long does the search need to be and what constraints are there when transitioning into the physical world? what are the difficulties in defining the boundaries of the system under evolution, i.e. what is actually doing the computation? and what effects come with non-isolated systems embedded in a physical environment? is the system/evolution utilising environmental conditions and sources of noise? and what are the affects on system/solution reproducibility? what are the consequences of varying conditions on replicating the solution? Many of these questions and more are discussed in [76].

22.3.2 *NASCENCE Project: Carbon Nanotube Substrates*

As part of the European FP7-ICT research project *NASCENCE*¹ (NAnoScale Engineering for Novel Computation using Evolution) further materials were considered, along with a new bespoke hardware platform [15]. The present hardware iteration, known as the Mecobo board [70], forms another hybrid hardware architecture to integrate digital computers with experimental materials. The system interfaces with materials placed on micro-electrode arrays (MEA) (Fig. 22.2) using a similar premise to Harding and Miller's liquid crystal system, the CAP.

The substrates currently under analysis consist of Single-Walled Carbon Nanotubes (SWCNT) mixed with either a polymer or liquid crystal. Another substrate consisting of randomly-dispersed gold nanoparticles is also under investigation, but requires temperatures of less than a few Kelvin to function. The polymer/LC mixtures disperse the nanotubes into random static networks, forming varying connection topologies and conductive pathways, possibly forming something akin to a random electrical circuit. Carbon nanotubes are used as they can exhibit either metallic or semi-conducting behaviour and contain other unique properties (e.g. ballistic conduction, thermal conductivity, self-assembly via van der Waals forces), whilst the mixing material is believed to create isolating regions, forming an insulator between elements and creating network structure.

A number of recent investigations have demonstrated the capabilities of SWCNT/polymer mixtures, in particular Poly-methyl-methacrylate (PMMA) and Poly-butyl-methacrylate (PBMA), as a potentially rich, evolvable and ubiquitous substrate. These investigations include: solving classification and optimisation problems such as frequency classification [78]; classifying various data instances [23, 80]; solving

¹*NASCENCE* homepage: nascence.no.




Fig. 22.2 Micro-electrode array used in the NASCENCE project to contain, stimulate and record activity in a carbon nanotube-based substrate. Computer controlled evolution is used to select active electrodes and mode (i.e. record or stimulate)

small numbers of cities instances of the travelling salesman problem (TSP) [22]; and applied to the (NP-hard) bin packing problem [79]. Early evidence highlights and supports the plausibility and potential of the methodology, but, in some respects it still lacks competitive results and still exhibits some of the issues raised at the end of Sect. 22.3.1.

PBMA appears to show greater stability than PMMA. The electrical percolation threshold of PBMA is said to occur around a concentration of 1% (w.r.t. polymer weight), forming a useful mixture of short and long-conductive pathways. After this, adding more nanotubes to the mixture is said to provide a negligible computational advantage as a suitable network appears to already exist. Although, interestingly, higher concentrations do demonstrate more non-linear current-voltage (I-V) behaviour in comparison. At less than 1%, the nanotube networks become fairly sparse and are argued to have reduced computational performance (and potentially more linear properties) [74]. The PMMA polymer was investigated towards the beginning of the NASCENCE project (see [56]) and is reported to have a percolation threshold between 0.17 and 0.70% in the literature, as explained in [74], but more importantly, a direct comparison is difficult to make between the two polymers as they come under different polymer groups, i.e. contain different chain lengths.

Investigations into SWCNT/liquid crystal mixtures has shown some promise, for example, non-linear I-V behaviour appears to be more prominent. It has also been demonstrated that conductivity and orientation can be changed by an in-plane electric field. But, LC has been shown to experience a longer configuration time, in-terms of LC molecule and SWCNT alignment, due to LC molecules being smaller than SWCNT ribbons, and associated relaxation times. Other issues include; Long-term stability and reconfigurability, and the exact role of the liquid crystal in nanotube alignment [121].

New engineering technologies, adaptations in the search method (or fitness criteria), changes in hardware (number of electrodes, different pitch sizes, etc.) and new materials could all have a significant effect on the field. The key components to the success of EIM lay in improvements to the fabrication of materials and the interfacing system used for stimulation and observation. This early work has also only highlighted one of many means of “programming” a material via evolution (i.e. through discrete voltage inputs). Other controlling fields, or even a combination of fields, could be utilised to manipulate/configure different materials—hopefully further separating the distinction between configuration and input signals. The ideal scenario for this field would be to pave the way for cheap, small, easily reconfigured and manufactured, multi- or single-purpose standalone computational devices.

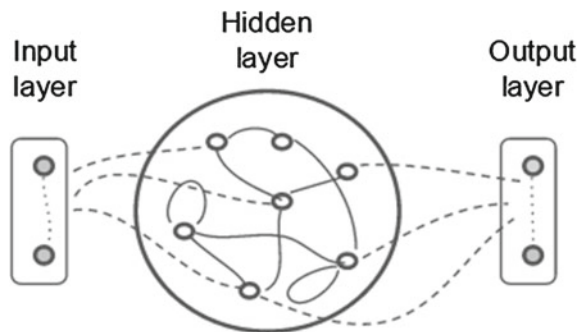
22.4 Reservoir Computing

22.4.1 What Is Reservoir Computing?

Reservoir Computing is the unification of three individually conceived methods for creating and training artificial recurrent neural networks (RNN): Echo State Networks (ESN) [44], Liquid State Machines (LSM) [72] and the Backpropagation-Decorrelation (BPDC) on-line learning rule [105].

A typical RNN model consists of a system of three layers: an input layer, a hidden layer (the core network), and an output layer (see Fig. 22.3). The hidden layer contains processing elements (neurons) that are interconnected through weighted synapses (connection weights). The input and output layers are connected to the hidden layer, again through weighted synapses. Variations on the types of connectivity, e.g. feedback from the output to hidden layer or input layer to output layer, depends on the task and method. For simplicity, we examine a simple input-to-hidden-to-output system encompassing a recurrent network in the hidden layer.

Fig. 22.3 A typical three-layer recurrent neural network. The input and output layer are connected to the hidden layer via weighted connections. The connections between neurons in the hidden layer are also weighted



When driven by an input, neurons activate, propagating information through the network to other neurons through varying connection strengths. The presence of recurrent connections can produce self-sustained activations, preserving a dynamic memory in the network's internal state. Networks such as this have been shown to be theoretically very powerful and can be both Turing equivalent [53] and good universal approximators of dynamical systems [30]. However, making the most of RNNs comes at a price, as they suffer from many training difficulties, such as the computational expense of updating large networks, bifurcation points, and sometimes falling into inescapable local optima when using gradient-descent.

Reservoir Computing offers an alternative training technique. It reduces the computational cost and removes the problem of degenerative gradient information that leads to poor convergence. Also, the concepts of reservoirs go beyond traditional neural networks and encompass (to some extent) more general dynamical systems. Evidence of this can be seen in the following implementations: in electronic circuits [95, 98], a bucket of water [28], Gene Regulation Networks (GRN) of *E. coli* bacteria [24, 52], deoxyribozyme oscillators (referred to as “DNA reservoir computing”) [33] and a cat's primary visual cortex [81].

22.4.2 Reservoir Types

There are many “flavours” of reservoir, originating from two separate research fields of machine learning and computational neuroscience. The first focuses on training dynamical systems for temporal learning tasks using artificial recurrent neural networks. The second aims at realistically modelling the computational properties of neural microcircuits. We give a summary of the two main branches of RC, and reservoirs in unconventional hardware. For more types and variations see [68].

22.4.2.1 Echo State Network

The *Echo State Network* (ESN) is a discrete-time neural network constructed from a sparse, random collection of analogue neurons (Fig. 22.4). The typical neuron model employed uses the sum of its weighted inputs, applied to a sigmoid function (generally a hyperbolic-tangent), to give the neuron state $x(n)$ at time n . The state activations $x(n)$ of these neurons are termed as echo states [44], i.e. echoes of the input history. To propagate and hold this history the network requires the *echo state property*, or more generally speaking, a fading memory. The property itself is provided by the characteristic dynamics of the system.

In ESNs, different scaling parameters, and in particular the *spectral radius* $\rho(\cdot)$, influence these dynamics. These parameters fundamentally alter and control the amount of memory and non-linearity present in the system. The $\rho(\cdot)$ parameter is used to scale the weight matrix W so that the largest absolute eigenvalue satisfies




Fig. 22.4 Echo State Network: A random, static, recurrent network of sigmoidal neurons. The input driven system projects $u(n)$ into the reservoir network. Each node possesses a one-to-one weighted connection to all inputs (via W^{in}), outputs (via W^{out}) and other nodes (via W). The extra input (labelled “1” in the figure) is provided as a bias W_{bias} . Training occurs on the output weights W^{out} (in relation to reservoir states $x(n)$) by reducing the error between $y(n)$ and $y^{target}(n)$ [66]

$\rho(W) < 1$ (typically, but not always; see [68]). Within this region the echo state property is said to be assured.

Another variant of the ESN model is the leaky-integrator neuron model, using a neuron that possesses some form of memory of previous activations. These neurons contain a leaking rate, or decay parameter α , which can control the speed of the reservoirs update dynamics (Eq. 22.1).

As Jaeger [44] describes, each neuron acts like a digital low-pass filter enabling a discrete network to approximate the dynamics of a continuous network (variations and uses can be seen in [44, 66, 68]). Dynamical systems have a natural time-scale; understanding the time-scale on which the input is changing compared to the time-scale of the system dynamics can be difficult. The leaky parameter α helps control and mediate any differences in input time-scales.

Putting together these components gives the neuron state update equation:

$$x(n) = (1 - \alpha)x(n - 1) + \alpha f(W_{in}u(n) + Wx(n - 1)) \quad (22.1)$$

22.4.2.2 Liquid State Machine

The *Liquid State Machine* (LSM) model arose as a method for defining the computational properties and power of neural microcircuits, “an alternative to Turing and attractor-based models in dynamical systems” [71]. The LSM model represents a competitive model for describing computations in biological networks of neurons. The LSM attempts to model cortical micro-columns in the neocortex, structured in cortical layers of randomly created spiking neurons based on a spatial embedding. Among other things, it has been described as a possible process used by mammalian

brains in speech recognition [27] and has been verified for cortical microcircuits in the primary visual cortex [81] and the primary auditory cortex [54].

Networks based on LSM use continuous streams of data (spike trains) to achieve real-time computations. Maass [71] has argued that classical models cannot handle real-time data streams based on spike trains. Unlike ESNs, they are generally more adaptive systems, supporting additional advanced readout features such as parallel perceptrons [72], although, in many cases, a linear readout is preferred.

Investigations using Liquid State Machines has highlighted the potential for more abstract unconventional applications, for example, pattern recognition using a physical medium (water) [28], and imitating LSMs in *E. Coli* [52].

22.4.2.3 Unconventional Hardware: Single Non-linear Dynamical Node with Delayed Feedback

Recent experimental applications of reservoir computing in optoelectronics and photonics [6, 61, 87] have demonstrated a new way of constructing a pseudo-reservoir system. Using *delay systems* theory, a system can imitate the characteristics of a recurrent network without being one. Delay systems represent a class of dynamical systems which incorporate non-linear systems with some form of delayed feedback and/or delayed coupling.

The key feature of this new RC flavour is to replace a physical network of nodes (often large in size) with a single non-linear node and a delay line. The delay system mimics a large interconnected network by creating a topology of *virtual* nodes in the delay line. This is achieved by applying time-multiplexing techniques on the input, i.e. through a combination of sample-and-hold operations mixed with an additional input mask. The sample-and-hold operation creates a stream $I(t)$ which defines the state update determined by the delay τ in the feedback loop. An additional mask M is created to represent coupled weights between stream $I(t)$ and the virtual nodes. The matrix M was initially randomly created at first, motivated by the random connectivity in reservoirs; but later optimised masks were proposed [7].

The number of virtual nodes in the system is defined by N equidistant points separated in time ($\theta = \tau/N$) along the delay interval τ . The resulting time-multiplexed input sequence becomes $J(t) = M \times I(t)$ which is then fed into the non-linear node (Fig. 22.5). Once the system has updated after time τ , the output nodes access the states in the delay line using $\sum_{i=1}^N w_i x(t - \frac{N}{\tau}(N - i))$. For more detail on the masking process see the supplementary information provided for [6, 7].

The single non-linear node scheme has many interesting implications for designing hardware reservoir computers. A clear advantage is that the overall architecture needed is very simple. But using a delay loop implies a serial process (in contrast to the parallel feeding of nodes in RC). Therefore, the speed of information processing depends on the state update given by time τ . A few suggestions have been given to compensate for such discrepancies and to increase computational capability, e.g. adding additional delay lines (increasing memory capacity) and finding an optimal number of nodes that can reasonably be implemented physically: adding more nodes

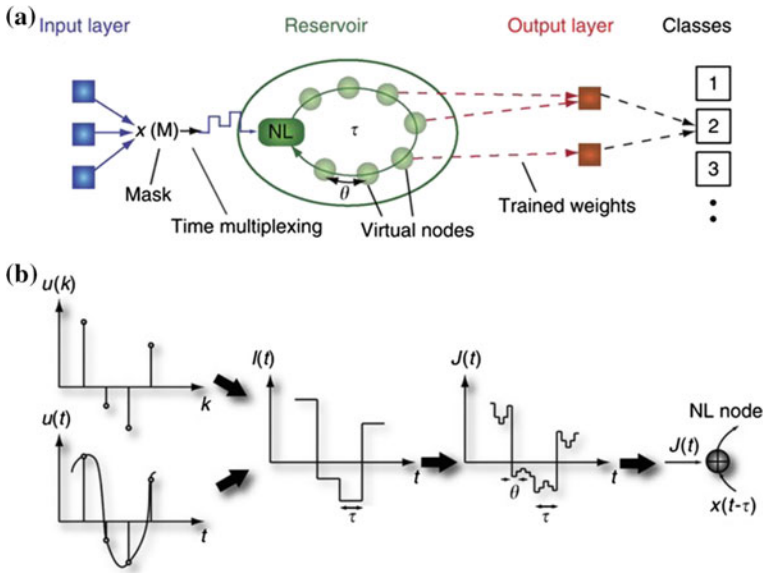


Fig. 22.5 **a** Single non-linear node with delayed feedback. A new system derived from delay systems theory where virtual nodes are created in the delay loop τ to represent the hidden layer. The number of virtual nodes N is separated by time θ given by $\theta = \tau/N$, **b** time-multiplexing of the input stream $I(t)$ (sampled from discrete $u(k)$ or continuous time $u(t)$) and a randomly created mask M that creates the input sequence $J(t)$. $J(t)$ is then added to the delayed state of the reservoir $x(t - \tau)$ before being fed back into the non-linear node [6]

will decrease τ , but having too many nodes is physically impractical given the hardware involved. Future work using this technique looks promising and could produce some ultra-high-speed computing solutions for specific tasks.

22.4.3 Reservoirs and Kernels

A reservoir can be interpreted as possessing kernel-like properties. A *kernel* acts as a pre-processor, embedding input data into a vector space known as a *feature* space. It is understood in many statistical machine learning methods that combining this feature space with a simple linear discriminant algorithm can enable the learning of complex non-linear functions.

In kernel methods this is achieved by projecting the input space $u(n)$ into a high-dimensional (possibly infinite-dimensional) feature space $x(n)$ *without* paying the price of its explicit computation, referred to as the *kernel trick* [92]. A kernel can therefore be expressed as the expansion function $x(u(n))$. However, there are two significant differences between reservoirs and kernels: reservoirs **do** explicitly compute the input transformation, i.e. do not possess the kernel trick; and kernels are typically

ill-equipped to handle temporal signals. To tackle temporal tasks the learned function $x(\cdot)$ requires some form of memory of previous inputs. Reservoirs solve this by utilising the network’s recurrent topology, which creates memory by retaining previous state activations. The final expansion function of the reservoir can therefore be represented as $x(n) = x(x(n-1), u(n))$.

Reservoirs use this pre-processing technique to map temporal features of the input into a spatially defined feature map (the network). The desired features can then be extracted, or combined, in a linear fashion to create the output $y(n)$:

$$y(n) = W^{out} x[u(n)] \quad (22.2)$$

where $W^{out} \in \mathbb{R}^{N_y \times N_x}$, and N_y, N_x are the number of output nodes and internal nodes respectively.

This enables reservoirs to tackle many temporal and dynamic real-world tasks not possible using simple non-temporal kernels. Equation (22.2) also implies the system contains a clear separation between the reservoir and the linear readout, separating the training procedure from the hidden layer, i.e. only W^{out} is trained. As such, the kernel representation offers a much faster and better converging mechanism compared to other RNN models, as it does not suffer from vanishing gradients. This representation also classes reservoirs as powerful adaptive filters. For more information on kernels see [100].

22.4.4 Reservoirs and Criticality

To design an optimal reservoir one should find a good trade-off between: (i) the transformation of the input that optimally boosts the linear classifiers capability, later referred to as the “quality” of the kernel; and (ii) a sufficiently-long (fading) memory based on the input history. These two properties often conflict; to obtain a useful memory requires ordered behaviour and a rich transformation requires dynamic behaviour. Legenstein and Maass [62] have shown that optimal reservoirs tend to experience the best trade-offs at a critical point near the “edge of chaos”.

Dynamic networks are said to exhibit emergent criticality and self-organising properties [110]. A novel example of this can be seen [10], where a self-organising structure of carbon nanotubes evolves to produce maximum entropy given a strong applied electric field. Moreover, it has been observed that dynamic networks can self-organise into critical regions where they can perform interesting computations.

These systems are characterised by motion in the phase space described as *trajectories* or state transitions. A trajectory may converge towards (i.e. be attracted to) a stable or unstable steady state; an attractor. Attractors vary from *point*—a point in the state space that attracts trajectories into its basin—to *strange* and chaotic—attracting trajectories, but inside diverge exponentially. These systems are very robust, small perturbations in the trajectory will tend to converge towards the same attractor. But, both external inputs and parameter changes in the system can drastically alter the

shape of the phase space. Such changes may perturb a trajectory, moving or “clamping” the system between different attractor basins. As a result, clamping may even create, remove or change existing attractors and thus alter the initial phase space created by the natural dynamics of the system [109]. The overall dynamics of the system can therefore, to some extent, be controlled.

As discussed in Sect. 22.2, certain critical regions have been identified to display interesting properties for computation (e.g. maximum complexity and performance). Although the reasons for maximised performance is not entirely understood, a series of quantitative measures have been used to observe the effects of phase transition on computational capability.

Early measurements of dynamics in reservoir-like systems can be found in [11] where, using a similar technique to Derrida and Pomeau [26], dynamic behaviour is measured using the Hamming distances between two output states. By observing a growth in the distance between states it can be determined that chaotic behaviour is present; a decrease in distance would indicate more ordered behaviour. This concept is similar to computing the characteristic Lyapunov Exponent for a dynamical system (see Sect. 22.6.2.2), with both measures analysing the sensitivity to differences in initial conditions.

Bertschinger and Natschläger [11] highlight two fundamental properties required for these networks: a *fading memory* and a “network mediated” separation. A fading memory is indicative of an ordered phase (memory) with some dynamics (fading). The same property is found in both Liquid State Machines and Echo State Networks, referred to as the “echo state property” in the latter. This allows the readout function to use information from recent inputs and derive functions of those inputs from the network state. Network mediated separation is deemed fundamentally important for input time-series networks, with similarities to the separation property in LSMs (see [72]). The property requires (ideally a large) diversity in network states that is the result of differences in inputs alone, allowing characteristic features to be identified in the input, and that any changes in state should not directly be a result of chaotic dynamics which could produce the same phenomena. As such, this property enables a readout function to respond effectively to any variation in the inputs.

Legenstein and Maass [63] propose two critical elements that characterise the computational capabilities of a complex dynamical system (cortical neural microcircuits in this case). These new measures, or properties, are proposed because Lyapunov exponents are only useful for analysing a systems dynamics, and are not necessarily helpful in predicting good parameter regions that create high computational performance. These measures are the *kernel-quality* and the *generalisation-capability*. The kernel-quality refers to the linear separation property found in kernels (Sect. 22.4.3). An empirical measure of this property is achieved by examining the complexity of functions that can be carried out on the inputs that boost the classification power of a subsequent linear layer. The generalisation-capability quantifies a system’s capability to generalise any learned behaviour to a new input.

Boedecker et al. [13] extend these ideas to ESNs, and create a general framework for direct and localised measurements for each neuron. Boedecker et al. give measurements indicating the memory of each neuron and the transfer of informa-

tion between each neuron. This work highlights some interesting and relevant points for all systems: a network does not necessarily need to be at the edge of chaos to do computation, the measured region where a system is at the edge of chaos is not universal for all tasks, and a critical state may maximise computational capabilities, but, such criticality may also be unnecessary or detrimental to certain tasks.

22.4.5 Training Reservoirs

This section describes the training procedure and some of the available techniques one can use to train both the linear outputs and the reservoir itself. The two methods for training linear readouts described here are: “off-line” *batch-mode* training, using simple linear or ridge regression techniques (done once all reservoir states are collected into matrix X for training length T), and “on-line” training, often gradient descent-based, Recursive Least Squares algorithm (a useful extensive investigation of RLS-type training is shown in [57]). We then describe some *pre*-processing training techniques that can be used on the reservoir itself, methods that have been identified as useful in creating tailored/optimal reservoirs.

There are many training techniques available in this diverse field. Here we just discuss methods perceived as “classical” training methods. For more examples of training techniques see [47, 66, 68].

22.4.5.1 Off-Line Training

Reservoirs are traditionally trained in a supervised manner where the temporal input $u(n)$ and coupled target output $y^{target}(n)$ are provided. Given a desired output the system can learn input-output behaviour by minimising the error (Eq. 22.3) between system output and desired output.

$$E(y, y^{target}) = \sqrt{\frac{1}{T} \frac{\sum_{t=1}^T (y(n) - y^{target}(n))^2}{\sigma^2(y^{target}(n))}} \quad (22.3)$$

To evaluate if the learned behaviour generalises accordingly, new input data is tested and the error between the two are again compared.

Reservoir computers are conceptually viewed as recurrent neural networks incorporating the three-layered topology of N_x hidden nodes (neurons), N_u inputs and N_y outputs. As discussed in Sect. 22.4.2, variations on how these are implemented are also possible, but, essentially the system still adheres to the same structural layers.

The general update state equations for most systems are defined in Eqs. (22.4) and (22.5) for, discrete time $n = 1, \dots, T$, internal state $x(n)$ and output $y(n)$:

$$x(n) = f(W^{in}u(n) + Wx(n-1) + W^{fb}y(n-1) + W^{bias}) \quad (22.4)$$

$$y(n) = f^{out}(W^{out}[x(n); u(n)]) \quad (22.5)$$

The function f is commonly represented by a sigmoid, typically a hyperbolic tangent. In echo state networks this represents a basic *tanh* neuron, but varies depending on the application. In other networks, f can be designed to form linear nodes, threshold logic gates or spiking neurons. In regards to the output $y(n)$, f^{out} may also be a non-linear function (sigmoid) but tends to be identity in most cases.

The weight matrices; input weights $W^{in} \in \mathbb{R}^{N_u \times N_x}$, reservoir/hidden layer weights $W \in \mathbb{R}^{N_x \times N_x}$ and feedback weights (from the output to the reservoir, if needed) $W^{fb} \in \mathbb{R}^{N_x \times N_y}$ are all drawn randomly from a uniform distribution at creation and remain static. The output weight matrix $W^{out} \in \mathbb{R}^{N_y \times (N_x + N_u)}$ includes weights for the inputs as they act as additional states (hence the concatenation of $[x(n); u(n)]$). Typically, the W matrix forms a sparse network with many of the weights set to zero, the other W^{in} and W^{fb} matrices can either be dense or sparse. Additional scaling parameters might also be applied to the matrices to govern properties such as non-linearity, stability and global dynamics. Techniques that can optimise/adapt each matrix (on-line or pre-processing) will be discussed later in the section.

The bias W^{bias} can be used to counteract training issues and large weights by adding noise, acting as a *regularisation* parameter, or to push the *tanh* neuron to a particular state, creating a smoothing effect.

Applying feedback W^{fb} can be useful, or detrimental, to certain tasks. Some tasks might not be learnt to a reasonable degree without feedback, or, certain systems may require dynamics beyond what is supplied by the driven input to construct a suitable output. Adding feedback comes with its own risks, feedback will ultimately change the stability of the system and requires adaptations in the training procedure. It is often advised only to use feedback when necessary. For more information see [66].

The *off-line* technique is completed in one training cycle T after the system has computed all states for the given inputs. It provides a very fast training mechanism as it essentially computes a linear model given by the known output Y , collected reservoir states X and desired output Y^{target} :

$$Y = W^{out} X \quad (22.6)$$

The collected state matrix $X \in \mathbb{R}^{N_x \times T}$ is created when the input $u(n)$ is run through the reservoir states $x(n)$. To avoid initial transients created by an initial zero state $x(0)$, a section at the beginning of the training sequence is discarded in the state matrix X . Essentially, the system goes through a “warming-up” process where states bounce around rather than returning to the equilibrium output, i.e. the system is too chaotic to retrieve any useful information about the input.

Given Eq. (22.6), we can find the optimal weights that minimise the error between $y(n)$ and $y^{target}(n)$ by solving the overdetermined system:

$$Y^{target} = W^{out} X \quad (22.7)$$

Equation(22.7) can be solved for W^{out} using linear regression. The simplest method is to use Ordinary Least-Squares (Eq.22.8), but typically this method succumbs to stability issues when inverting (XX^T) .

$$W^{out} = Y^{target} X^T (XX^T)^{-1} \quad (22.8)$$

Lukoševičius [66] recommends using either ridge regression (regression with Tikhonov regularisation β) (Eq.22.9) or the Moore–Penrose pseudo-inverse (Eq.22.10). Ridge regression is a stable and effective solution and is generally advised. Adding a regularisation parameter counteracts the problems of producing very large output weights, which often indicates very sensitive and unstable solutions.

$$W^{out} = Y^{target} X^T (XX^T + \beta I)^{-1} \quad (22.9)$$

where I is the identity matrix and β the regularisation parameter.

Setting $\beta = 0$ gives the same method for solving linear regression in Eq.(22.8). It is therefore recommended to use a logarithmic scale for selecting β where it never reaches zero [68].

The pseudo-inverse is applied in some cases typically because it is straightforward to implement in certain programming environments (e.g. MATLAB). However this comes at a price. The pseudo-inverse method is computationally expensive for large matrices of X and typically overdetermined. However, in most cases the network is made up of relatively small matrices and over-fitting depends on the difficulty of the task.

$$W^{out} = Y^{target} X^+ \quad (22.10)$$

22.4.5.2 On-Line Training

Some tasks require an on-line training method that adapts with time, minimising the error at each time step. This implicitly turns W^{out} into an adaptive linear combiner. The Recursive Least Squares (RLS) algorithm (Eq.22.11) is more commonly applied as it overcomes the severely impaired convergence performance of the Least Means Square (LMS) algorithm [46].

$$E(y, y^{target}, n) = \frac{1}{N_y} \sum_{i=1}^{N_y} \sum_{j=1}^n \lambda^{n-j} (y_i(j) - y_i^{target}(j))^2 \quad (22.11)$$

Using RLS comes at a cost: the number of weights is quadratic rather than linear, and it can still be numerically unstable in some cases. Other powerful on-line methods may be useful to a practitioner, particularly in the presence of feedback connections, such as Backpropagation-Decorrelation (BPDC) [105].

The RLS training procedure is described here, derived from [25]. First, set the error forgetting parameter λ close to but less than one; the forgetting parameter controls the contribution of previous samples. Next, initialise the autocorrelation matrix $\rho(0) = I/\delta$, with δ being a small constant and I the identity matrix. At each time step compute the output weights using the following steps:

- Step 1: Calculate the reservoir state $x(n)$ and output signal $y(n)$ for input $u(n)$.
 Step 2: Calculate the error between target output $y^{target}(n)$ and system output given previous output weights:

$$e(n) = y^{target}(n) - W^{out}(n-1)x(n) \quad (22.12)$$

- Step 3: Update the gain vector $K(n)$:

$$K(n) = \frac{\rho(n-1)x(n)}{\lambda + x^T(n)\rho(n-1)x(n)} \quad (22.13)$$

- Step 4: Update the autocorrelation matrix $\rho(n)$:

$$\rho(n) = \frac{1}{\lambda}[\rho(n-1) - K(n)x^T(n)\rho(n-1)] \quad (22.14)$$

- Step 5: Assign new output weights $W^{out}(n)$ using (22.12) and (22.13):

$$W^{out}(n) = W^{out}(n-1) + K(n)e(n) \quad (22.15)$$

For more readout training methods including feedback training (such as FORCE training), supervised, unsupervised, reinforcement learning, etc. consult the excellent review [68] and practical aid [66].

22.4.5.3 Adaptation and Pre-training

Adaptive reservoirs, ones that change weights or configuration, are inspired by natural adaptation in biological neurons. These adaptive processes are the result of persistent changes in a neuron's electrical properties, governed by unsupervised *local* adaptation rules often referred to as *Intrinsic Plasticity* (IP). These rules represent a homeostatic mechanism in which neurons self-modify their intrinsic activity (i.e. excitability). Using such learning rules has shown to increase robustness and performance when pre-training reservoirs [96, 106]. For an overview of recent investigations including both local and global adaptation schemes, see [67, 68].

In classical RC, reservoirs are generated randomly, hence the performance of each reservoir varies on creation. Reservoir computing boasts its training performance on the separation between the reservoir and readout. The readout training, at its core, is quite inexpensive, allowing the possibility of other forms of reservoir *pre-training*, i.e.

generating reservoirs deterministically for each task. Even a crude experiment such as selecting a reservoir which produces the smallest error from a pool of randomly-created reservoirs highlights the advantages of pre-training. Evolutionary algorithms are a potential search strategy for pre-training.

Investigations using evolutionary optimisation for pre-training are well documented. Many strategies have been attempted, including evolving topologies (i.e. network size), weight matrices (such as W^{in} , W , W^{fb}), global parameters (e.g. spectral radius), connection density, adapting slopes of the activation function $f(\cdot)$ and even training w^{out} when no target signal is available. Other interesting methods include *EvoLino*—evolving hidden connections to gradient-based long-short term memory (LSTM) RNNs [91] and using *Neuro-Evolution of Augmented Topologies* (NEAT) as a meta-search algorithm for evolving ESNs [20] (the “related work” section also discusses other neuro-evolution methods for constructing ESNs). All of these methods have shown great potential, highlighting the performance increases and optimisations reservoir pre-training can create for specific tasks. Pre-training and adaptation appears to be one of many key branches under investigation in the field of reservoir computing.

22.5 Modelling Materials with Reservoir Computing

Any high-dimensional dynamical system with an observable state $x(n)$ that is a result of input $u(n)$ could form the basis of a reservoir according to [68]. This implies that any material that can exhibit sufficient dynamics and a fading memory could therefore, theoretically, be adopted as a reservoir.

Given that reservoir computing is based on artificial recurrent neural networks, one implementation route would be to design hardware substrates modelled on simplified neural network-like structures, e.g. large coupled networks of non-linear elements with varying densities of connectivity. Using this structural model, semi-isolated regions of activity may form exploitable meta-states for the trainable readout. Various implementations of hardware-based artificial neural networks have existed for some time; see [77] for a review of HNNs. More recently, there has been increased popularity towards applying memristive components to neuromorphic circuits, see [58] (a review CMOS/memristor hybrids) and [101].

Kulkarni and Teuscher [59], among others, have examined and demonstrated reservoir networks built from memristor devices. Memristors appear to be ideal components for building reservoirs: they exhibit non-linear properties and maintain a memory of previous inputs. In the Kulkarni and Teuscher experiment, simulated circuits are randomly-created from networks of memristors, then a subsequent readout layer is trained using a genetic algorithm to solve some computational task. Other simulated memristor reservoirs include: simple-cycle reservoirs [16] i.e. memristive networks that form nodes instead of analogue neurons (see [88]), training more realistic volatile memristor models [19], and variation-tolerant reservoirs [17].

Fabricating random, highly-interconnected networks and devices from nanoscale switching elements is a challenging task. At these scales, features of self-assembly and self-organisation are essential: characteristics that might only be achievable through unconventional methods and materials. Konkoli and Wendin [55] provide a brief review of some non-CMOS devices and discuss the suitability of the RC model on such unconventional devices. These devices include, molecular electronic networks, memristor networks, and other substrates outlined as part of the SYMONE project [122]. At present, the SYMONE project is investigating networks of organic transistors (NOMFETs) and self-assembled networks of gold nanoparticles that could feature functionalised memristive junctions [55].

Two examples that demonstrate physical, *in-materio* reservoir computing systems are silicon-based photonic chips (based on nanophotonic crystal cavities) [29, 117], and Atomic Switch Networks (ASNs) [103, 112, 113].

The *photonics chip* primarily exploits the propagation of light through silicon. Inside these chips are photonic crystals that remove the propagation of certain frequencies of light, known as the band gap. Adding a line defect to a crystal produces a photonic crystal waveguide, effectively a process by which light is forced between the defect. Cavities are then created along the line defect to create an optical “resonator” which traps light, increasing the power inside the resonator. These resonators then form a optical reservoir which can be trained and manipulated using different types of readouts, e.g. [117] creates a linear system, and the non-linearity is added at the readout through the inherent non-linearity of the measuring equipment. Methods such as this propose an interesting optical alternative to hardware-based reservoir computing.

The *atomic switch network* approach focuses on the electrical and chemical properties of a material. These networks attempt to mimic the vast complexity, emergent dynamics, and connectivity of the brain. Highly-interconnected networks are constructed by bottom-up self-assembly of silver nanowires. Through a triggered electro-chemical reaction, coated copper seed nucleation sites spawn large quantities of silver nanowires of various lengths, from nano- to millimetre scale. Large random networks are formed, creating crossbar-like junctions between nanowires, and when exposed to gaseous sulphur create $Ag|Ag_2S|Ag$ nanoscale metal-insulator-metal (MIM) junctions. Applying external activation (a bias voltage) to these junctions creates a temporary resistance drop, leading to functional memristive junctions called Atomic Switches. Applying this construction and functionalisation process the ASN method offers some unique properties, such as scalability and practicality in creating highly-complex nanoscale substrates.

The emergent behaviour and dynamics of ASNs can be observed through fluctuations in network conductivity, a result of spontaneous switching between discrete metastable resistance states, where locally excited regions cause cascading changes in resistance throughout the system. As such, the non-linear responses to resistive switching are reported to result in higher harmonic generation (HHG), also suggested as a useful measure for quantitatively evaluating reservoir dynamics [103].

A clear advantage of the ASN technique is that it allows some degree of regulation in fabrication and further control through “resistance control” training [113].

For example, varying the parameters of the nucleation sites (copper seeds) can alter the network structure and therefore the substrates dynamics. The relative size, numbers and pitch of copper seeds can determine the length of wires, forming distant connections or confining spatial regions to dendrite-like structures, thus ultimately defining the density of connections [103].

ASNs appear to be natural candidates for reservoir computing, producing a high-dimensional recurrent network that does not require the manipulation of individual elements. ASNs have been applied to one derived RC task; the waveform generation task [103, 113]. This is a simple analogue task which measures a reservoir's ability to construct a desired output waveform from an independent input waveform using network generated harmonics. For example, given an input sine-wave, can a trained reservoir construct a sawtooth, square-wave, or any other periodic function at the output (essentially a Fourier series task using harmonic analysis). These initial experiments have proven ASNs to be capable reservoirs and has also highlighted HHG as a potential metric for evaluating reservoir properties.

22.6 Evaluating the Characteristics of Reservoirs

Creating a random reservoir (in simulation at least) is fairly straightforward, but designing one with the right properties, using the large parameter space available, is a challenging task. In many cases parameter selection is done by hand, through trial and error, and with expert knowledge of the desired characteristics. So how can we better understand and evaluate reservoirs?

One approach is to simplify its construction. In doing so, one could provide a more theoretical understanding of what makes reservoirs useful/successful. Reference [88] explores this idea by addressing three issues: (i) what is the minimal complexity of topology and parameters that produce comparable performance to standard models? (ii) what degree of randomness is needed to construct competitive reservoirs? and (iii) how do completely deterministic reservoirs compare? These are good questions for understanding underlying RC principles, but may be impractical to investigate given an already created (maybe static) physical substrate. Instead, we desire more experimental quantitative measures that individually describe the reservoir and its qualities as an efficient kernel.

Determining, or evaluating, reservoir quality and performance can be achieved in two ways, either through direct measurement of performance on a given task, or by cumulatively assessing individual properties of the reservoir. Using the latter method provides a mechanism in which performance could be partly predicted for any task. As [82] explains, a good reservoir that scores well in all properties may be able to facilitate the process of “learning transfer”, where the reservoir can be trained to some objective function that will increase its capability without seeing any output task. The objective function, in the mentioned case, measures how well a reservoir (an LSM) separates different classes of inputs into different reservoir states. As such, it was shown that improving separability and instilling an adaptive balance between

chaotic and ordered behaviour (through changes in structure) a reservoir can increase its performance across different artificial problems.

In this section we discuss some of the measures proposed in the reservoir computing literature, accompanied by a variety of different benchmark tasks used to assess performance. Such measures and benchmarks may be an effective method for evaluating the potential of substrate-based reservoirs.

22.6.1 Kernel Quality and Separation

22.6.1.1 Kernel Quality

Kernel quality evaluates a reservoir's ability to produce diverse and complex mappings (functions) of the input stream u that can consequently increase the classification performance of a linear decision-hyperplane [62].

Kernel quality, also known as the *linear separation property*, was first introduced by Legenstein and Maass [62], along with an accompanying metric referred to as the *generalisation capability* of a reservoir. The two properties are closely coupled and can be measured using a similar ranking mechanism. The first, kernel quality, measures a reservoir's ability to produce diverse reservoir states given significantly different input streams. The second measures the reservoir's capability to generalise given similar input streams with respect to a target function. Both measurements can be carried out using the method in [18], by computing the rank r of an $n \times m$ matrix M , with the two methods differing only in the selection of m input streams u_i, \dots, u_m , i.e. input streams being largely different or of similar type/class.

The rank is assessed as follows; Given the input stream u_i the reservoir state vector x_{u_i} of length n is collected to form each column in the matrix M . The rank r of each matrix is then estimated by Singular Value Decomposition (SVD). Büsing et al. [18] explains that a good reservoir should possess a high kernel quality and a low generalisation rank, and also identify a correlation in the measurement to the reservoir's dynamics. For example, reservoirs in ordered regimes produce low ranking values in both measures, and in chaotic regimes produces high values in both measures. A similar connection is also observed in [21], where the Lyapunov exponent and the kernel quality strongly correlate.

22.6.1.2 Class Separation

Class separation is a metric that corresponds directly to different classes of input stimuli. Demonstrations of class separation can be found in [21, 31, 82]. Separation is measured as the average distance between resulting states, once again, given the assumption that significantly different inputs should generate significantly different reservoir states. To calculate separation requires the division of the input and state vectors into discrete classes; [31] provides an alternative measure characterised on

the original assumption. For example, given two different input vectors $u_j(n)$ and $u_k(n)$ the euclidean distance between inputs should be large and positive, as described by D :

$$D := \|u_j(n) - u_k(n)\| \tag{22.16}$$

If the reservoir exhibits a good separation property the reservoir states $(x_j(n)$ and $x_k(n))$ should increase in distance, or be equal to the original distance:

$$D \leq \|x_j(n) - x_k(n)\| \tag{22.17}$$

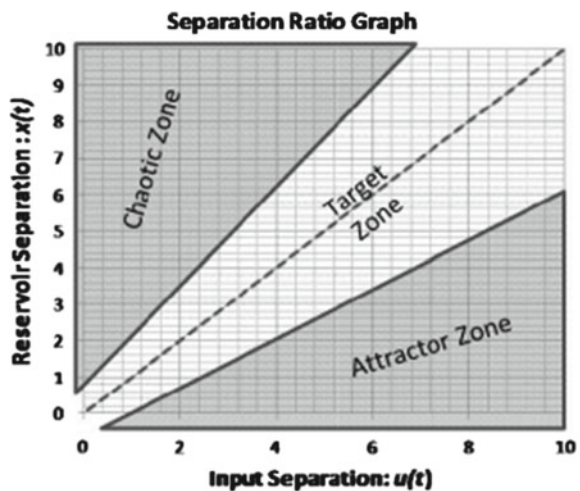
which can be represented as the ratio:

$$\frac{\|x_j(n) - x_k(n)\|}{\|u_j(n) - u_k(n)\|} \geq 1 \tag{22.18}$$

This simplified measure has been extended [31] into *Separation Ratio Graphs* to produce a visual representation of separation and the phase transition of correlated dynamic behaviour (see Fig. 22.6).

Konkoli and Wendin [55] offer another comparable method for identifying reservoir quality in memristor networks. This metric is again based on the assumption that quality can be measured by observing the reservoir’s ability to generate different dynamic states at the output. In this case, it is observed by measuring the dissimilarity between output states and a linear combination of the inputs, i.e. determining if the non-linear frequency response of a network cannot be approximated by a linear mixture of delayed inputs. Dissimilarity is measured in the Fourier space (ω) between outputs $o(n)$ and a linear combination of the time shifted inputs $z(n)$, given by:

Fig. 22.6 Separation Ratio Graph [31]. Graphical representation of the phase transition between chaos and order. Systems in the target zone are said to possess both a good separation property and ideal dynamic behaviour to produce optimal reservoirs



$$\delta = \frac{\|o(\omega) - z(\omega)\|}{\|o(\omega)\|} \quad (22.19)$$

A large dissimilarity (large δ) is ideal in a reservoir as it describes a complex projection of the input. A small δ on the other hand may simply describe a linear propagation of the input, highlighting the absence of richness in the reservoir.

22.6.2 Quantifying Dynamics

Producing a reservoir with rich dynamics is evidently desirable, but how can we quantify dynamical behaviour? how can we pin-point the critical region at the edge of chaos? In the literature, there are several prominent themes in classifying and measuring dynamics, including observation of trajectories and attractor behaviour, internal scaling of input-driven activity, memory and retention, and higher-harmonic generation, particularly in reservoirs comprising resistive switches.

22.6.2.1 Spectral Radius

In Echo State Networks, the spectral radius $\rho(\cdot)$ (the largest absolute eigenvalue) of the internal weight matrix W is used to determine and control reservoir dynamics [44]. The parameter ρ globally scales the internal weights, moving the system between different regimes, altering the non-linearity and impulse response of the reservoir. Effectively, the scaling parameter alters the internal time-scales of the system, providing the *echo state property*.

Parameter selection of the spectral radius is typically centred around a value of one; smaller than one is attributed to a stable regime (a fading response to input stimuli), if larger than one, a system will typically be unstable, causing undesirable interference to new inputs.

22.6.2.2 Lyapunov Exponents

A popular measurement for criticality, or chaotic behaviour, in a reservoir is the empirical estimation of the Lyapunov exponents (LE) for a dynamical system [11, 21, 31, 63, 118]. To calculate the Lyapunov Exponents, and quantify a system's criticality, we measure the divergence between two close trajectories due to some small perturbation. For example, if an orbit is close to an attractor, in an ordered phase, small changes should dissipate over time. In a chaotic phase, an applied change will diverge exponentially from an orbit, persisting or increasing over time.

Gibbons [31] provides a simple approximation of the largest (maximal) Lyapunov exponent (derived from [89]), as the largest tends to dominate. Various formats and interpretations exist in calculating different Lyapunov Exponents due to different

approximation methods (see [51, 63, 118]). For example, Verstraeten et al. [118] examines the local Lyapunov spectrum and the Jacobian of the reservoir, suggesting performance can be better predicted by the maximum of the minimal Lyapunov Exponent. Boedecker et al. [13] takes a similar approach to Gibbons, estimating the mean exponential rate of divergence of the trajectories for an n -dimensional phase space.

A system with a maximal Lyapunov exponent $\lambda \approx 0$ is somewhere near the *edge of chaos*. A chaotic system is present at $\lambda > 0$ and an ordered, or sub-critical, system at $\lambda < 0$.

Another suggested measure of criticality is the measure of instantaneous entropy of reservoir states, defined as the *average state entropy* (ASE) in [85]. Applying Rényi quadratic entropy, one can measure the distribution of instantaneous amplitudes in reservoir states, providing some measure of the “richness” of dynamics. The entropy measurement is associated with an expectation that increased diversity of amplitudes, i.e. increased spread of amplitudes away from some concentrated point, will increase the readout’s ability construct the desired response.

22.6.2.3 Memory Capacity

Measuring the short-term memory capacity of a reservoir was first outlined by Jaeger [45] as a quantitative measurement to observe the echo state property (fading memory). To determine the memory capacity of a reservoir we simply measure how many delayed versions of the input $u(n - k)$ the outputs can recall or recover with precision. As Jaeger describes, using the equation in (22.20), we can measure memory capacity by how much variance of the delayed input can be recovered, summed over all delays [45]. This is carried out by training individual output units to recall the input at time k with the maximum capacity MC of an N node reservoir typically bounded by its size, i.e. $MC \leq N$.

$$MC = \sum_{k=1}^{\infty} MC_k = \sum_{k=1}^{\infty} \frac{cov^2(u(n - k), y(n))}{\sigma^2(u(n))\sigma^2(y(n))} \quad (22.20)$$

This measurement has direct connection to the dynamic behaviour of a system. It can be helpful in identifying the boundaries between static structure that provides memory, and the point of complex dynamics that gives us processing. As such, one might expect a chaotic system to lose information regarding previous inputs at a faster rate and a more ordered regime to increase (to some extent) input retention.

22.6.2.4 Harmonic Generation

The generation of higher harmonics in Atomic Switch Networks has been identified as a technique for examining emergent behaviour and network connectivity

[103, 113]. In ASNs, higher harmonic generation (HHG) is attributed to the non-linear frequency response of the system due to both input amplitude and memristive “hard” switching behaviour. To examine HHG, one can plot the frequency response of the system, which can be used to identify connectivity and system dynamics, i.e. changes in network response (an on-set of HHG) is typically due to an increase in “hard” switching memristive connections past a percolation threshold [103]).

22.6.3 Evaluation Through Benchmark Tasks

As a means of direct assessment, we can measure performance of reservoirs and their subsequent readouts by applying them to specific tasks. Reservoir computing (and neural networks as a whole) possesses an abundance of benchmark tasks, from simple classification and time-series prediction to robot navigation [5, 25] and non-linear channel equalisation [87, 88].

In this section we discuss benchmark tasks that are the most prevalent in the reservoir computing literature.

22.6.3.1 Waveform Generation

This task requires a rich transformation of a temporal input waveform (a periodic signal) to create an entirely new waveform. It is based on Sillin et al. [103]’s physical adaptation of Jaeger’s [47] sine-wave generator task and is linked directly to Fourier series/analysis. The task is to train the system to produce three different waveforms, given an input sine wave. In [103, 113], this is achieved by applying a 10Hz input sine-wave (at one electrode in the ASN) to produce a 10Hz square-wave, and sawtooth, and a 20Hz sine-wave at the output ($y(n)$), via the combination of other weighted electrode readings (e.g. recorded states $x(n)$). The task is said to be an excellent precursor to testing potential reservoir substrates on more difficult temporal problems [103], as it highlights an abundance of temporal features (phase shifts, delays, harmonic generation, recurrence etc.).

A similar task is the continuous-time multiple superimposed oscillator (MSO) task. In this benchmark, the reservoir’s role is to predict the evolution of, and generate a superposition of, two or more sinusoidal waves with different frequencies.

$$s(t) = \sin(\omega_1 t) + \sin(\omega_2 t) \quad (22.21)$$

where $\omega_1 = 0.2$ and $\omega_2 = 0.311$. The task has been demonstrated in photonics experiments [29] and other non-traditional reservoirs [107].

22.6.3.2 Time-Series Prediction and Generation

Two prominent benchmark tasks in reservoir computing are Non-linear Auto-regressive Moving Average (NARMA) dynamical system modelling and the Mackey–Glass chaotic time-series prediction task.

The NARMA task originates from Atiya and Parlos [8]’s work on training recurrent network; its goal is to evaluate a reservoir’s ability to model an n th order, highly non-linear dynamical system where the system state depends on the incoming input as well as its own history. The challenging aspect of the NARMA task is that it contains long-term dependencies created by the n th order time-lag. Typically, the benchmark is carried out on 10th and 30th ordered systems [6, 46, 47, 106].

A description of the 10th ordered task is as follows; Given white noise $u(n)$ from a uniform distribution of interval $[0, 0.5]$, the reservoir should predict an output $y(n)$ close to the target $y(n + 1)$, calculated by:

$$y(n + 1) = 0.3y(n) + 0.05y(n) \left(\sum_{i=0}^9 y(n - i) \right) + 1.5u(n - 9)u(n) + 0.1 \quad (22.22)$$

Mackey–Glass chaotic time-series prediction is another common benchmark, where the system is trained to predict one time-step into the future (for examples see [44, 49, 118]).

The system is described by the Mackey–Glass delay differential equation:

$$\dot{y}(n) = \alpha y(n - \tau) / (1 + y(n - \tau)^\beta) - \gamma y(n) \quad (22.23)$$

As [44] explains, parameters for the MG task are typically set to $\alpha = 0.2$, $\beta = 10$ and $\gamma = 0.1$ with the parameter τ set to 17 to produce a mildly chaotic attractor: the system has a chaotic attractor for $\tau > 16.8$.

Other time-series prediction benchmarks are summarised in [88], including predicting laser activations in the Santa Fe Laser dataset (originally used in [50]), predicted next output in the Hénon Map dataset, and IPIX Radar and Sunspot series datasets.

22.6.3.3 Classification Tasks

Simple classification problems are wide and varied in the field of machine learning, some of which can be seen in both the RC and EIM literature. For example, typical tasks for EIM are tone and frequency discriminators, and Iris and Lenses dataset classification [38, 78, 80]. Examples in reservoir computing include, signal classification (discriminating between two waveforms) [87], various n-bit parity problems [11, 25, 94, 97] and other time-independent classification tasks [4].

Possibly the most adopted classification task in RC involves the recognition of isolated digits from multiple speakers [61, 87, 88, 93, 96, 119, 120]. Taken from a

subset of the T146 speech corpus dataset, the task uses a total of 500 speech fragments collected from five different participants listing the digits zero to nine (repeated ten times). The reservoir interprets these speech fragments through a preprocessing filter in the form of a digital model of the human cochlea. Linear classifiers are then trained to be sensitive to individual digits, with a final classification being made on the temporal mean of the output. Performance on the task is measured using cross-validation by calculating the number of misclassified digits using the Word Error Rate (WER).

Reservoir computing is said to be very competitive to, or outperform, many of the state-of-the-art approaches on these tasks.

22.6.4 Material Properties and Considerations

In order to perform material computation we require some ability to manipulate and control certain aspects of physical structure and behaviour. To observe such effects requires a means of observation and measurement. In the two examples presented (ASNs and EIM), we have discussed one method by which this can be achieved, through the application and recording of electrical voltages to a micro-electrode array. There are other possible methods, for example: optical stimulus/measurement and other stimuli across the electromagnetic spectrum (e.g. optoelectronic and photonic reservoirs); image recognition for observation (e.g. a method also used in Fernando and Sojakka [28]’s bucket of water); control and observation through magnetic fields (e.g. manipulating ionised gases or observing nuclear magnetic resonance (NMR) [108]); chemical excitation and reaction (e.g. reaction-diffusion computers and slime mould [1, 2]).

There are many physical properties and considerations that require discussion when talking about using any novel material for computation. Here we focus on four key factors that possess some relevance to substrate-based reservoirs: (i) a means by which to observe network connectivity and activity, (ii) assuring non-linearity is present in the system, (iii) methods for modelling activity and behaviour, and (iv) the impact of environmental factors.

22.6.4.1 Network Connectivity and Activity

The computational capability of a material is often directly related to conductivity and the density of connections inside it. Variations in these concentrations can have an adverse or favourable effect on conductivity and task performance. For example, to optimise an ASN one can control the densities of silver nanowires [103]. Massey et al. [74] identify a similar relationship, where nanotube concentration directly alters the conductivity and computational performance of SWCNT/PBMA composites.

How can we measure connectivity in materials and analyse distributed activity? One possible method is demonstrated in [112], where ultra-sensitive infra-red (IR)




Fig. 22.7 Infrared image of carbon nanotube/polymer mixture. Infrared has been identified as a useful method for observing current flow and local regions of activity in both EIM and ASNs [73]

imaging is used to help identify network conductivity and behaviour in ASNs. IR imaging was used to observe and measure possible dominant conductive pathways and identify the activity of local regions to stimulus by thermal emission. Similar experiments have also been carried out on carbon nanotube/polymer substrates within the NASCENCE project (see Fig. 22.7).

Other forms of network observation include: observation of power-law scaling in the power spectral density [103], optical microscopy such as examining structure through an optical microscope (an example can be seen in Fig. 22.8 with a SWCNT/LC substrate), electron microscopy and scanning-probe microscopy such as scanning electron microscopy (SEM) [32] and scanning-force microscopy (SFM) [14], observing current flow in polymer substrates by electron-beam-induced-current (EBIC) [37], investigating structural properties using Fourier transform infrared and Raman spectroscopy (FTIR) (a method used to observe the interaction of molecules in liquid crystal/SWCNT composites in [121]), and other absorption and emission spectroscopy techniques [99].

22.6.4.2 Non-linearity

Non-linearity within a material can be measured through current-voltage (I-V) characteristics. In ASNs, non-linearity is observed by the presence of pinched hysteresis curves as a function of input amplitude (produced by applying slow voltage sweeps). This non-linear I-V characteristic is said to be the result of changing switch behaviour (in this case towards a “hard-switching” regime) and increased harmonic generation [103].




Fig. 22.8 Image of SWCNT/LC behaviour to voltage stimulus, observed through a microscope. Observation at this scale shows detail in the complexity of local interactions and carbon nanotube alignment to electric fields [73]

Similar I-V measurements have also been used in both SWCNT/PBMA and SWCNT/LC composites proposed in the NASCENCE project [74, 121].

22.6.4.3 Modelling

Simulating the internal properties of specific materials would be desirable, but it is typically impractical to create an exact representation of one individual material. Instead, we could use an abstraction, a model, of what the material is doing. Some methods that model specific features of the system, and others that are potential accompaniments to the reservoir model, include: the modelling of electron transport using a Monte Carlo approach, percolation visualisation for electrode pathways, Voronoi diagrams for visualising electrode activity [102], Volterra series and Wiener series models, using Random Boolean Networks as a model, abstract neural networks for modelling global activity, NARMAX (non-linear autoregressive moving average model with exogenous inputs) modelling, and using Cellular Automata as a reservoir/substrate model [83, 84].

22.6.4.4 Environment and Noise

The local environment, thermal noise, other noise related fluctuations and quantum effects are tangible concerns in the physical domain. Sensitive systems require good isolation and compensation techniques to reduce the effects of both internal and external noise. Effectively, the systems susceptibility to noise will determine its robustness and adaptability. In many cases, noise and unwanted variability could be filtered using conventional techniques, but requires some care in implementation (i.e.

does not impede on the boundaries of the system or unduly affect what is doing the computation). Evolution-*in-materio* on the other hand has been found to exploit such external influences. Using an evolutionary search in physical systems has shown an uncanny ability to utilise strange characteristics that are clearly attributed to external influences. For example, Bird and Layzell [12]’s evolvable motherboard was directly influenced by local laboratory equipment in producing evolved oscillators.

22.7 Feedback Design: Designing Substrates to Be Good Reservoirs

Beyond simply applying the reservoir model, could we utilise some of its highlighted features to aid in the design of *in materio* systems?

Often the difficulty in designing and engineering computational substrates is knowing what range of features are necessary, or even exploitable. Some aspects of the physical system are unknown. There are clearly some attributes that are desirable (for evolution at least): reconfigurable, input-driven, reproducible, some stability, etc. Can we use known computational phenomena from reservoir computing to help reduce the search space of suitable materials? Could the abstract model guide material design, and in return, serve to create a more realistic and efficient model? This could create a virtuous feedback-loop in design, where new novel materials can be evaluated for computational capability and reservoir quality.

Going further, could multiple smaller reservoirs reproduce the overall dynamics of a larger network? As stated in [66], a powerful extension of ESN consists of many small ESNs in parallel, where an averaged output has drastically improved performance, but, at a price in terms of memory capacity, i.e. fewer nodes typically equals less memory. Another approach would be to use multiple reservoirs to overcome hardware limitations, such as the number of electrode contacts available. These networked reservoirs could, as a by-product, provide increased robustness (further distributing computation) and added variability in states, i.e. promote a good overall separation property. Additionally, such systems could be constructed from various materials, each with different kernel properties and time-scales, allowing a global weighting system or training method (possibly backpropagation) to optimally choose which material to use for certain tasks. Moreover, this could lead to a more generic reservoir network suitable to multiple tasks.

Hierarchical reservoirs represent an up-and-coming, possibly highly advantageous, avenue worth pursuing [68]. A hierarchical system attempts to overcome some of difficulties found in classical reservoirs such as scalability, learning complex intelligent tasks and working with multiple time-scales simultaneously. Some early examples of these architectures can be found in [48] where high-level reservoirs extract features from low-level ones using a “feature-voting” system. Another can be seen in a system comprised of decoupled sub-reservoirs with inhibitory connections [124] where the inhibitory connections predict the activation of sub-reservoirs. Other

examples include, hierarchically clustered ESNs and the impact substructures have on stability [51], increased performance on speech recognition tasks [115] and acoustic modelling [116], and a hierarchical architecture used for hand-writing recognition to sort delivery parcels [69].

22.8 Conclusion

We believe that reservoir computing can be used as a stepping stone to modelling and training the currently somewhat static in structure *in-materio* systems. Reservoir computing is considered mostly on the basis that the systems themselves may not be able to undergo “on-line” training/adaptation (manually or autonomously). Thus, the material is perceived as an initially-configurable black-box, or a suitably rich kernel, removed from the final training process.

The reservoir method provides many advantages when dealing with unknown or intractable properties, but, it is also limited by the complexity of the task it is required to learn. For example, a good reservoir requires a high-dimensional expansion of the input that can be exploited by a subsequent readout layer. This implies that for more complex tasks the reservoir may require an exponential growth in exploitable features from an exceptionally large feature space, which itself is potentially impractical, or at least very difficult to extract useful features from. Harnessing these features then becomes a significant challenge when implemented in a physical device. Standard artificial neural networks try to overcome this problem by building the required non-linear features internally. This is achieved through internal training, which is traditionally achieved through backpropagation.

To apply backpropagation in a physical system requires a material to contain the right attributes for internal training and manipulation, however. A recent example of a physical system that does use error-backpropagation as a training method is demonstrated in [42]. The experiment and implementation is carried out on an acoustic system, using the propagation of sound waves between a speaker and microphone, and on an electro-optical system, harnessing the reciprocal transmission of light through an optical circuit. This experiment provides a unique insight into backpropagation applied to physical systems and presents a potentially competitive alternative to digital neural networks.

The research areas discussed in this chapter, if combined, provide a rich avenue to explore. Reservoir computing presents us with a convenient theoretical model. It also provides some indication of what properties are required to increase the performance: the right dynamic criticality, good kernel projection, etc.

Optimisation and pre-training can improve performance in reservoirs, by allowing a reservoir to be *pre-set* with some properties that effectively increase the performance of the main training process. Evolution as a form of pre-training in the physical domain may provide a crucial and efficient technique for manipulating/configuring matter into suitable reservoirs, e.g. through structural alignment/deformation, or creating rich local regions of varying activity. As is described in the *Evolution-in-materio*

doctrine, evolution may be the best practical technique to exploring and exploiting properties that are currently intractable or hitherto unknown, properties that could produce the most interesting and competent physical reservoirs.

Acknowledgments Matthew Dale is funded by a Defence Science and Technology Laboratory (DSTL) Ph.D. studentship.

References

1. Adamatzky, A.: *Physarum Machines: Computers from Slime Mould*, vol. 74. World Scientific, Singapore (2010)
2. Adamatzky, A., Costello, B., Asai, T.: *Reaction-Diffusion Computers*. Elsevier, Amsterdam (2005)
3. Adamatzky, A., Erokhin, V., Grube, M., Schubert, T., Schumann, A.: Physarum chip project: growing computers from slime mould. *IJUC* **8**(4), 319–323 (2012)
4. Alexandre, L.A., Embrechts, M.J., Linton, J.: Benchmarking reservoir computing on time-independent classification tasks. In: *International Joint Conference on Neural Networks IJCNN 2009*, pp. 89–93. IEEE (2009)
5. Antonelo, E.A., Schrauwen, B., Van Campenhout, J.: Generative modeling of autonomous robots and their environments using reservoir computing. *Neural Process. Lett.* **26**(3), 233–249 (2007)
6. Appeltant, L., Soriano, M.C., Van der Sande, G., Danckaert, J., Massar, S., Dambre, J., Schrauwen, B., Mirasso, C.R., Fischer, I.: Information processing using a single dynamical node as complex system. *Nature Commun.* **2**, 468 (2011)
7. Appeltant, L., Van der Sande, G., Danckaert, J., Fischer, I.: Constructing optimized binary masks for reservoir computing with delay systems. *Sci. Rep.* **4**(3629) (2014)
8. Atiya, A.F., Parlos, A.G.: New results on recurrent network training: unifying the algorithms and accelerating convergence. *IEEE Trans. Neural Netw.* **11**(3), 697–709 (2000)
9. Beggs, J.M.: The criticality hypothesis: how local cortical networks might optimize information processing. *Philos. Trans. R. Soc. Lond. A: Math., Phys. Eng. Sci.* **366**(1864), 329–343 (2008)
10. Belkin, A., Hubler, A., Bezryadin, A.: Self-assembled wiggling nano-structures and the principle of maximum entropy production. *Sci. Rep.* **5**(8323) (2015)
11. Bertschinger, N., Natschläger, T.: Real-time computation at the edge of chaos in recurrent neural networks. *Neural Comput.* **16**(7), 1413–1436 (2004)
12. Bird, J., Layzell, P.: The evolved radio and its implications for modelling the evolution of novel sensors. In: *Proceedings of the Congress on Evolutionary Computation CEC'02*, vol. 2, pp. 1836–1841. IEEE (2002)
13. Boedecker, J., Obst, O., Lizier, J.T., Mayer, N.M., Asada, M.: Information processing in echo state networks at the edge of chaos. *Theory Biosci.* **131**(3), 205–213 (2012)
14. Bose, S.K., Lawrence, C.P., Liu, Z., Makarenko, K.S., van Damme, R.M.J., Broersma, H.J., van der Wiel, W.G.: Evolution of a designless nanoparticle network into reconfigurable boolean logic. *Nature Nanotechnol.* (2015). doi:[10.1038/nano.2015.207](https://doi.org/10.1038/nano.2015.207)
15. Broersma, H., Gomez, F., Miller, J., Petty, M., Tufte, G.: Nascence project: nanoscale engineering for novel computation using evolution. *Int. J. Unconv. Comput.* **8**(4), 313–317 (2012)
16. Bürger, J., Goudarzi, A., Stefanovic, D., Teuscher, C.: Composing a reservoir of memristive networks for real-time computing. [arXiv:1504.02833](https://arxiv.org/abs/1504.02833) (2015)
17. Burger, J., Teuscher, C.: Variation-tolerant computing with memristive reservoirs. In: *2013 IEEE/ACM International Symposium on Nanoscale Architectures (NANOARCH)*, pp. 1–6. IEEE (2013)

18. BÜsing, L., Schrauwen, B., Legenstein, R.: Connectivity, dynamics, and memory in reservoir computing with binary and analog neurons. *Neural Comput.* **22**(5), 1272–1311 (2010)
19. Carbajal, J.P., Dambre, J., Hermans, M., Schrauwen, B.: Memristor models for machine learning. *Neural Comput.* **27**(3), 725–747 (2015)
20. Chatzidimitriou, K.C., Mitkas, P.A.: A NEAT way for evolving echo state networks. In: *ECAI 2010*, pp. 909–914. IOS Press (2010)
21. Chrol-Cannon, J., Jin, Y.: On the correlation between reservoir metrics and performance for time series classification under the influence of synaptic plasticity. *PLoS One* **9**(7), e101792 (2014)
22. Clegg, K.D., Miller, J.F., Massey, M.K., Petty, M.: Travelling salesman problem solved ‘in materio’ by evolved carbon nanotube device. In: *Parallel Problem Solving from Nature—PPSN XIII*, pp. 692–701. Springer, Heidelberg (2014)
23. Clegg, K.D., Miller, J.F., Massey, M.K., Petty, M.C.: Practical issues for configuring carbon nanotube composite materials for computation. In: *IEEE International Conference on Evolvable Systems, ICES 2014*, pp. 61–68. IEEE (2014)
24. Dai, X.: Genetic regulatory systems modeled by recurrent neural network. In: *Advances in Neural Networks-ISNN 2004*, pp. 519–524. Springer, Heidelberg (2004)
25. Dasgupta, S., Wörgötter, F., Manoonpong, P.: Information theoretic self-organised adaptation in reservoirs for temporal memory tasks. In: *Engineering Applications of Neural Networks*, pp. 31–40. Springer, Heidelberg (2012)
26. Derrida, B., Pomeau, Y.: Random networks of automata: a simple annealed approximation. *EPL (Europhys. Lett.)* **1**(2), 45 (1986)
27. Dominey, P.F.: Complex sensory-motor sequence learning based on recurrent state representation and reinforcement learning. *Biol. Cybern.* **73**(3), 265–274 (1995)
28. Fernando, C., Sojakka, S.: Pattern recognition in a bucket. In: *Advances in Artificial Life*, pp. 588–597. Springer, Heidelberg (2003)
29. Fiers, M., Van Vaerenbergh, T., Wyffels, F., Verstraeten, D., Dambre, J., Schrauwen, B., Biehnman, P.: Nanophotonic reservoir computing with photonic crystal cavities to generate periodic patterns. *IEEE Trans. Neural Netw. Learn. Syst.* **25**(2), 344–355 (2014)
30. Funahashi, K., Nakamura, Y.: Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Netw.* **6**(6), 801–806 (1993)
31. Gibbons, T.E.: Unifying quality metrics for reservoir networks. In: *IJCNN 2010, The International Joint Conference on Neural Networks*, pp. 1–7. IEEE (2010)
32. Goldstein, J., Newbury, D.E., Echlin, P., Joy, D.C., Romig, A.D Jr., Lyman, C.E., Fiori, C., Lifshin, E.: *Scanning Electron Microscopy and X-ray Microanalysis: A Text for Biologists, Materials Scientists, and Geologists*. Springer Science & Business Media, Heidelberg (2012)
33. Goudarzi, A., Lakin, M.R., Stefanovic, D.: DNA reservoir computing: a novel molecular computing approach. In: *DNA Computing and Molecular Programming*, pp. 76–89. Springer, Heidelberg (2013)
34. Greenwood, G.W., Tyrrell, A.M.: *Introduction to Evolvable Hardware: A Practical Guide for Designing Self-Adaptive Systems*, vol. 5. Wiley, New York (2006)
35. Gutierrez, J.M., Hinkley, T., Ward Taylor, J., Yanev, K., Cronin, L.: Evolution of oil droplets in a chemorobotic platform. *Nature Commun.* **5** (2014)
36. Haddow, P.C., Tyrrell, A.M.: Challenges of evolvable hardware: past, present and the path to a promising future. *Genet. Program. Evolvable Mach.* **12**(3), 183–215 (2011)
37. Hanoka, J.I., Bell, R.O.: Electron-beam-induced currents in semiconductors. *Ann. Rev. Mater. Sci.* **11**(1), 353–380 (1981)
38. Harding, S., Miller J.F.: Evolution in materio: a tone discriminator in liquid crystal. In: *CEC 2004, Congress on Evolutionary Computation*, vol. 2, pp. 1800–1807. IEEE (2004)
39. Harding, S., Miller J.F.: Evolution in materio: initial experiments with liquid crystal. In: *2004 NASA/DoD Conference on Evolvable Hardware*, pp. 298–305. IEEE (2004)
40. Harding, S., Miller J.F.: Evolution in materio: a real-time robot controller in liquid crystal. In: *2005 NASA/DoD Conference on Evolvable Hardware*, pp. 229–238. IEEE (2005)

41. Harding, S., Miller J.F.: Evolution in materio: evolving logic gates in liquid crystal. In: ECAL 2005 Workshop on Unconventional Computing: From cellular automata to wetware, pp. 133–149. Beckington, UK (2005)
42. Hermans, M., Burm, M., Van Vaerenbergh, T., Dambre, J., Bienstman, P.: Trainable hardware for dynamical computing using error backpropagation through physical media. *Nature Commun.* **6**, (2015)
43. Higuchi, T., Iwata, M., Kajitani, I., Yamada, H., Manderick, B., Hirao, Y., Murakawa, M., Yoshizawa, S., Furuya, T.: Evolvable hardware with genetic learning. In: IEEE International Symposium on Circuits and Systems, ISCAS'96, vol. 4, pp. 29–32. IEEE (1996)
44. Jaeger, H.: The “echo state” approach to analysing and training recurrent neural networks with an erratum note. Bonn, Germany: German National Research Center for Information Technology GMD Technical Report **148**, 34 (2001)
45. Jaeger, H.: Short term memory in echo state networks. GMD-Forschungszentrum Informationstechnik (2001)
46. Jaeger, H.: Adaptive nonlinear system identification with echo state networks. In: Advances in Neural Information Processing Systems, pp. 593–600 (2002)
47. Jaeger, H.: Tutorial on training recurrent neural networks, covering BPPT, RTRL, EKF and the “echo state network” approach. GMD-Forschungszentrum Informationstechnik (2002)
48. Jaeger, H.: Discovering multiscale dynamical features with hierarchical echo state networks. Technical report No. 9 (2007)
49. Jaeger, H., Haas, H.: Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science* **304**(5667), 78–80 (2004)
50. Jaeger, H., Lukoševičius, M., Popovici, D., Siewert, U.: Optimization and applications of echo state networks with leaky-integrator neurons. *Neural Netw.* **20**(3), 335–352 (2007)
51. Jarvis, S., Rotter, S., Egert, U.: Extending stability through hierarchical clusters in echo state networks. *Front. Neuroinformatics* **4**, (2010)
52. Jones, B., Stekel, D., Rowe, J., Fernando, C.: Is there a liquid state machine in the bacterium *escherichia coli*? In: IEEE Symposium on Artificial Life, 2007. ALIFE'07, pp. 187–191. IEEE (2007)
53. Kilian, J., Siegelmann, H.T.: The dynamic universality of sigmoidal neural networks. *Inf. Comput.* **128**(1), 48–56 (1996)
54. Klampfl, S., David, S.V., Yin, P., Shamma, S.A., Maass, W.: Integration of stimulus history in information conveyed by neurons in primary auditory cortex in response to tone sequences. In: 39th Annual Conference of the Society for Neuroscience, Program, vol. 163 (2009)
55. Konkoli, Z., Wendin, G.: On information processing with networks of nano-scale switching elements. *Int. J. Unconv. Comput.* **10**(5–6), 405–428 (2014)
56. Kotsialos, A., Massey, M.K., Kaisar, F., Zeze, D.A., Pearson, C., Petty, M.C.: Logic gate and circuit training on randomly dispersed carbon nanotubes. *Int. J. Unconv. Comput.* **10**(5–6), 473–497 (2014)
57. Küçükemre, A.U.: Echo state networks for adaptive filtering. Ph.D. thesis, University of Applied Sciences (2006)
58. Kudithipudi, D., Merkel, C., Soltiz, M., Garrett, S.R., Robinson, E.P.: Design of neuromorphic architectures with memristors. In: Network Science and Cybersecurity, pp. 93–103. Springer, Heidelberg (2014)
59. Kulkarni, M.S., Teuscher, C.: Memristor-based reservoir computing. In: IEEE/ACM International Symposium on Nanoscale Architectures, NANOARCH, 2012, pp. 226–232. IEEE (2012)
60. Langton, C.G.: Computation at the edge of chaos: phase transitions and emergent computation. *Phys. D: Nonlinear Phenom.* **42**(1), 12–37 (1990)
61. Larger, L., Soriano, M.C., Brunner, D., Appeltant, L., Gutiérrez, J.M., Pesquera, L., Mirasso, C.R., Fischer, I.: Photonic information processing beyond turing: an optoelectronic implementation of reservoir computing. *Opt. Express* **20**(3), 3241–3249 (2012)
62. Legenstein, R., Maass, W.: Edge of chaos and prediction of computational performance for neural circuit models. *Neural Netw.* **20**(3), 323–334 (2007)

63. Legenstein, R., Maass, W.: What makes a dynamical system computationally powerful. In: *New Directions in Statistical Signal Processing: From Systems to Brain*, pp. 127–154 (2007)
64. Lipson, H., Pollack, J.B.: Automatic design and manufacture of robotic lifeforms. *Nature* **406**(6799), 974–978 (2000)
65. Lohn, J.D., Linden, D.S., Hornby, G.S., Kraus, W.F., Rodriguez-Arroyo, A.: Evolutionary design of an X-band antenna for NASA's space technology 5 mission. In: *NASA/DoD Conference on Evolvable Hardware*, pp. 155–155. IEEE (2003)
66. Lukoševičius, M.: A practical guide to applying echo state networks. In: *Neural Networks: Tricks of the Trade*, pp. 659–686. Springer, Heidelberg (2012)
67. Lukoševičius, M., Jaeger, H.: Overview of reservoir recipes. Technical report 11, Jacobs University Bremen (2007)
68. Lukoševičius, M., Jaeger, H.: Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.* **3**(3), 127–149 (2009)
69. Lukoševičius, M., Jaeger, H., Schrauwen, B.: Reservoir computing trends. *KI-Künstliche Intelligenz* **26**(4), 365–371 (2012)
70. Lykkebø, O.R., Harding, S., Tufte, G., Miller, J.F.: Mecobo: A hardware and software platform for in materio evolution. In: *Unconventional Computation and Natural Computation*, pp. 267–279. Springer, Heidelberg (2014)
71. Maass, W.: Liquid state machines: motivation, theory, and applications. In: *Computability in Context: Computation and Logic in the Real World*, pp. 275–296 (2010)
72. Maass, W., Natschläger, T., Markram, H.: Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* **14**(11), 2531–2560 (2002)
73. Massey, M.K.: Presentation at NASCENCE Consortium Progress Meeting. Totnes, UK (2015)
74. Massey, M.K., Kotsialos, A., Qaiser, F., Zeze, D.A., Pearson, C., Volpati, D., Bowen, L., Petty, M.C.: Computing with carbon nanotubes: optimization of threshold logic gates using disordered nanotube/polymer composites. *J. Appl. Phys.* **117**(13), 134903 (2015)
75. Miller, J.F., Downing, K.: Evolution in materio: looking beyond the silicon box. In: *NASA/DoD Conference on Evolvable Hardware 2002*, pp. 167–176. IEEE (2002)
76. Miller, J.F., Harding, S., Tufte, G.: Evolution-in-materio: evolving computation in materials. *Evol. Intell.* **7**(1), 49–67 (2014)
77. Misra, J., Saha, I.: Artificial neural networks in hardware: a survey of two decades of progress. *Neurocomputing* **74**(1), 239–255 (2010)
78. Mohid, M., Miller, J.F., Harding, S., Tufte, G., Lykkebo, O.R., Massey, M.K., Petty, M.C.: Evolution-in-materio: a frequency classifier using materials. In: *International Conference on Evolvable Systems, ICES 2014*, pp. 46–53. IEEE (2014)
79. Mohid, M., Miller, J.F., Harding, S., Tufte, G., Lykkebo, O.R., Massey, M.K., Petty, M.C.: Evolution-in-materio: solving bin packing problems using materials. In: *International Conference on Evolvable Systems, ICES 2014*, pp. 38–45. IEEE (2014)
80. Mohid, M., Miller, J.F., Harding, S., Tufte, G., Lykkebø, O.R., Massey, M.K., Petty, M.C.: Evolution-in-materio: solving machine learning classification problems using materials. In: *PPSN XIII, Parallel Problem Solving from Nature*, pp. 721–730. Springer, Heidelberg (2014)
81. Nikolić, D., Haeusler, S., Singer, W., Maass, W.: Temporal dynamics of information content carried by neurons in the primary visual cortex. In: *Advances in Neural Information Processing Systems*, pp. 1041–1048 (2006)
82. Norton, D., Ventura, D.: Improving liquid state machines through iterative refinement of the reservoir. *Neurocomputing* **73**(16), 2893–2904 (2010)
83. Ozgur, Y.: Reservoir computing using cellular automata. [arXiv:1410.0162](https://arxiv.org/abs/1410.0162) [cs.NE] (2014)
84. Ozgur, Y.: Connectionist-symbolic machine intelligence using cellular automata based reservoir-hyperdimensional computing. [arXiv:1503.00851](https://arxiv.org/abs/1503.00851) [cs.ET] (2015)
85. Ozturk, M.C., Xu, D., Príncipe, J.C.: Analysis and design of echo state networks. *Neural Comput.* **19**(1), 111–138 (2007)
86. Packard, N.H.: Adaptation toward the edge of chaos. In: *Kelso, J.A.S., Mandell, A.J., Shlesinger, M.F. (eds.) Dynamic Patterns in Complex Systems*, pp. 293–301. World Scientific, Singapore (1988)

87. Paquot, Y., Dupont, F., Smerieri, A., Dambre, J., Schrauwen, B., Haelterman, M., Massar, S.: Optoelectronic reservoir computing. *Sci. Rep.* **2**, (2012)
88. Rodan, A., Tino, P.: Minimum complexity echo state network. *IEEE Trans. Neural Netw.* **22**(1), 131–144 (2011)
89. Rosenstein, M.T., Collins, J.J., De Luca, C.J.: A practical method for calculating largest Lyapunov exponents from small data sets. *Phys. D: Nonlinear Phenom.* **65**(1), 117–134 (1993)
90. Samuelson, E., Glette, K.: Real-world reproduction of evolved robot morphologies: automated categorization and evaluation. In: *Applications of Evolutionary Computation*, vol. 9028. LNCS, pp. 771–782. Springer, Heidelberg (2015)
91. Schmidhuber, J., Wierstra, D., Gagliolo, M., Gomez, F.: Training recurrent networks by evolino. *Neural Comput.* **19**(3), 757–779 (2007)
92. Scholkopf, B., Smola, A.J.: *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT press (2001)
93. Schrauwen, B., Defour, J., Verstraeten, D., Van Campenhout, J.: The introduction of time-scales in reservoir computing, applied to isolated digits recognition. In: *Artificial Neural Networks–ICANN 2007*, pp. 471–479. Springer, Heidelberg (2007)
94. Schrauwen, B., Büsing, L., Legenstein, R.A.: On computational power and the order-chaos phase transition in reservoir computing. In: *Advances in Neural Information Processing Systems*, pp. 1425–1432 (2008)
95. Schrauwen, B., D’Haene, M., Verstraeten, D., Van Campenhout, J.: Compact hardware liquid state machines on fpga for real-time speech recognition. *Neural Netw.* **21**(2), 511–523 (2008)
96. Schrauwen, B., Wardermann, M., Verstraeten, D., Steil, J.J., Stroobandt, D.: Improving reservoirs using intrinsic plasticity. *Neurocomputing* **71**(7), 1159–1171 (2008)
97. Schumacher, J., Toutounji, H., Pipa, G.: An analytical approach to single node delay-coupled reservoir computing. In: *Artificial Neural Networks and Machine Learning–ICANN 2013*, pp. 26–33. Springer, Heidelberg (2013)
98. Schürmann, F., Meier, K., Schemmel, J.: Edge of chaos computation in mixed-mode vlsi-a hard liquid. In: *Advances in Neural Information Processing Systems*, pp. 1201–1208 (2004)
99. Shah, J.: *Ultrafast Spectroscopy of Semiconductors and Semiconductor Nanostructures*, vol. 115. Springer Science & Business Media, Heidelberg (1999)
100. Shawe-Taylor, J., Cristianini, N.: *Kernel Methods for Pattern Analysis*. Cambridge University Press, Cambridge (2004)
101. Sheridan, P., Ma, W., Lu, W.: Pattern recognition with memristor networks. In: *IEEE International Symposium on Circuits and Systems, ISCAS 2014*, pp. 1078–1081. IEEE (2014)
102. Sillin, H.O.: *Neuromorphic hardware: the investigation of atomic switch networks as complex physical systems*. Ph.D. thesis, University of California, Los Angeles (2015)
103. Sillin, H.O., Aguilera, R., Shieh, H., Avizienis, A.V., Aono, M., Stieg, A.Z., Gimzewski, J.K.: A theoretical and experimental study of neuromorphic atomic switch networks for reservoir computing. *Nanotechnology* **24**(38), 384004 (2013)
104. Steane, A.: Quantum computing. *Rep. Prog. Phys.* **61**(2), 117 (1998)
105. Steil, J.J.: Backpropagation-decorrelation: online recurrent learning with $O(n)$ complexity. In: *2004 IEEE International Joint Conference on Neural Networks*, vol. 2, pp. 843–848. IEEE (2004)
106. Steil, J.J.: Online reservoir adaptation by intrinsic plasticity for backpropagation-decorrelation and echo state learning. *Neural Netw.* **20**(3), 353–364 (2007)
107. Steil, J.J.: Several ways to solve the MSO problem. In: *ESANN*, pp. 489–494 (2007)
108. Stepney, S.: The neglected pillar of material computation. *Phys. D: Nonlinear Phenom.* **237**(9), 1157–1164 (2008)
109. Stepney, S.: Nonclassical computation: a dynamical systems perspective. In: Rozenberg, G., Bäck, T., Kok, J.N. (eds) *Handbook of Natural Computing*, vol. 4, pp. 1979–2025. Springer, Heidelberg (2012)
110. Stepney, S., Braunstein, S.L., Clark, J.A., Tyrrell, A., Adamatzky, A., Smith, R.E., Addis, T., Johnson, C., Timmis, J., Welch, P.: Journeys in non-classical computation I: a grand challenge for computing research. *Int. J. Parallel, Emergent Distrib. Syst.* **20**(1), 5–19 (2005)

111. Stepney, S., Abramsky, S., Adamatzky, A., Johnson, C., Timmis, J.: Grand challenge 7: Journeys in non-classical computation. In: Visions of Computer Science, London, UK, September 2008, pp. 407–421. BCS (2008)
112. Stieg, A.Z., Avizienis, A.V., Sillin, H.O., Martin-Olmos, C., Aono, M., Gimzewski, J.K.: Emergent criticality in complex Turing B-type atomic switch networks. *Adv. Mater.* **24**(2), 286–293 (2012)
113. Stieg, A.Z., Avizienis, A.V., Sillin, H.O., Aguilera, R., Shieh, H., Martin-Olmos, C., Sandouk, E.J., Aono, M., Gimzewski, J.K.: Self-organization and emergence of dynamical structures in neuromorphic atomic switch networks. In: *Memristor Networks*, pp. 173–209. Springer, Heidelberg (2014)
114. Thompson, A.: An evolved circuit, intrinsic in silicon, entwined with physics. In: *Evolvable Systems: From Biology to Hardware*, pp. 390–405. Springer, Heidelberg (1997)
115. Triefenbach, F., Jalalvand, A., Schrauwen, B., Martens, J.: Phoneme recognition with large hierarchical reservoirs. In: *Advances in Neural Information Processing Systems*, pp. 2307–2315 (2010)
116. Triefenbach, F., Jalalvand, A., Demuynck, K., Martens, J.: Acoustic modeling with hierarchical reservoirs. *IEEE Trans. Audio, Speech, Lang. Process.* **21**(11), 2439–2450 (2013)
117. Vandoorne, K., Mechet, P., Van Vaerenbergh, T., Fiers, M., Morthier, G., Verstraeten, D., Schrauwen, B., Dambre, J., Bienstman, P.: Experimental demonstration of reservoir computing on a silicon photonics chip. *Nature Commun.* **5**, (2014)
118. Verstraeten, D., Schrauwen, B.: On the quantification of dynamics in reservoir computing. In: *Artificial Neural Networks—ICANN 2009*, pp. 985–994. Springer, Heidelberg (2009)
119. Verstraeten, D., Schrauwen, B., Stroobandt, D., Van Campenhout, J.: Isolated word recognition with the liquid state machine: a case study. *Inf. Process. Lett.* **95**(6), 521–528 (2005)
120. Verstraeten, D., Schrauwen, B., d’Haene, M., Stroobandt, D.: An experimental unification of reservoir computing methods. *Neural Netw.* **20**(3), 391–403 (2007)
121. Volpati, D., Massey, M.K., Johnson, D.W., Kotsialos, A., Qaiser, F., Pearson, C., Coleman, K.S., Tiburzi, G., Zeze, D.A., Petty, M.C.: Exploring the alignment of carbon nanotubes dispersed in a liquid crystal matrix using coplanar electrodes. *J. Appl. Phys.* **117**(12), 125303 (2015)
122. Wendin, G., Vuillaume, D., Calame, M., Yitzchaik, S., Gamrat, C., Cuniberti, G., Beiu, V.: Symone project: synaptic molecular networks for bio-inspired information processing. *Int. J. Unconv. Comput.* **8**(4), 325–332 (2012)
123. Whiting, J., de Lacy Costello, B., Adamatzky, A.: Slime mould logic gates based on frequency changes of electrical potential oscillation. *Biosystems* **124**, 21–25 (2014)
124. Xue, Y., Yang, L., Haykin, S.: Decoupled echo state networks with lateral inhibition. *Neural Netw.* **20**(3), 365–376 (2007)