

Chapter 11

Guided Self-Organization of Input-Driven Recurrent Neural Networks*

Oliver Obst and Joschka Boedecker

11.1 Introduction

To understand the world around us, our brains solve a variety of tasks. One of the crucial functions of a brain is to make predictions of what will happen next, or in the near future. This ability helps us to anticipate upcoming events and plan our reactions to them in advance. To make these predictions, past information needs to be stored, transformed or used otherwise. How exactly the brain achieves this information processing is far from clear and under heavy investigation. To guide this extraordinary research effort, neuroscientists increasingly look for theoretical frameworks that could help explain the data recorded from the brain, and to make the enormous task more manageable. This is evident, for instance, through the funding of the billion-dollar "Human Brain Project", of the European Union, amongst others. Mathematical techniques from graph and information theory, control theory, dynamical and complex systems (Sporns 2011), statistical mechanics (Rolls and Deco 2010), as well as machine learning and computer vision (Seung 2012; Hawkins and Blakeslee 2004), have provided new insights into brain structure and possible function, and continue to generate new hypotheses for future research.

A marked feature of brain networks is the massive amount of recurrent connections between cortical areas, especially on a local scale (Douglas et al. 2004). Since information in these recurrent connections, or loops, can circulate between many neurons in a given circuit, they are ideally suited to provide a time-context for computations leading to predictions about future events. One particular mathematical

Oliver Obst

CSIRO Computational Informatics PO Box 76, Epping, NSW 1710, Australia
e-mail: oliver.obst@csiro.au

Joschka Boedecker

Machine Learning Lab, University of Freiburg, Germany

* Both authors contributed equally to this work.

model that is used to investigate the consequences of loops for computation and optimization in neuronal circuits are *recurrent neural networks* (RNNs).

In RNNs, many detailed properties of real neurons are abstracted for the sake of tractability, but important general concepts are kept. Elements of these networks are simple nodes that combine inputs from other nodes in the network, usually in a non-linear fashion, to form their outputs. They are connected in a directed graph, which may contain cycles. In *input-driven* RNNs, a constant stream of input data drives the dynamics of the network. Dedicated output units can then use this dynamics to compute desired functions, for instance for prediction or classification tasks. Since they can make use of the temporal context provided by the recurrent connections, RNNs are very well suited for time-series processing, and are in principle able to approximate any dynamical system (Maass et al. 2007).

While the recurrent connections of RNNs enable them to deal with time-dependencies in the input data, they also complicate training procedures compared to algorithms for networks without loops (e.g., Backpropagation (Rumelhart et al. 1986) or R-Prop (Riedmiller and Braun 1993)). Notably, training RNNs with traditional training methods suffer from problems like slow convergence and vanishing gradients. This slow convergence is due to the computational complexity of the algorithms training all of the connections in a network (such as BPTT (Werbos 1990) or RTRL (Williams and Zipser 1989)), as well as to bifurcations of network dynamics during training, which can render gradient information unusable (Doya 1992). Also, derivatives of loss functions need to be propagated over many time steps, which leads to a vanishing error signal (Bengio et al. 1994; Hochreiter 1998).

The realization of these fundamental issues led to alternative ways of using and training RNNs, some of which can be summarized in the field of *Reservoir Computing* methods (see, e.g., a recent overview by Lukovsevicius and Jaeger 2009), specialized architectures like the Long Short Term Memory (LSTM) networks (Hochreiter and Schmidhuber 1997) or training by evolutionary algorithms as in the *Evolino* approach (Schmidhuber et al. 2007). The Reservoir Computing field has been an active part of RNN research over the last decade, while there was less activity in gradient-descent-like methods which appear to generate renewed interest only recently (Bengio et al. 2012), partially due to the development of more efficient training techniques as in (Martens and Sutskever 2011).

Reservoir methods implement a fixed high-dimensional reservoir of neurons, using random connection weights between the hidden units, chosen small enough to guarantee dynamic stability. Input weights into this reservoir are also selected randomly, and reservoir learning procedures train only the output weights of the network to generate target outputs. A particular appeal of reservoir methods is their simplicity, and that the computation required for training is relatively low.

Taking the echo state network approach as a specific example of a typical reservoir network (see Fig. 11.1), it will consist of the following components: A random input-matrix \mathbf{W}_{in} combines input values \mathbf{u} linearly and sends them to the units in the high-dimensional hidden layer, referred to as the *reservoir*. The units in the reservoir also have recurrent connections amongst each other, collected in the matrix \mathbf{W}_{res} . These loops implement a fading memory, so information can remain in the

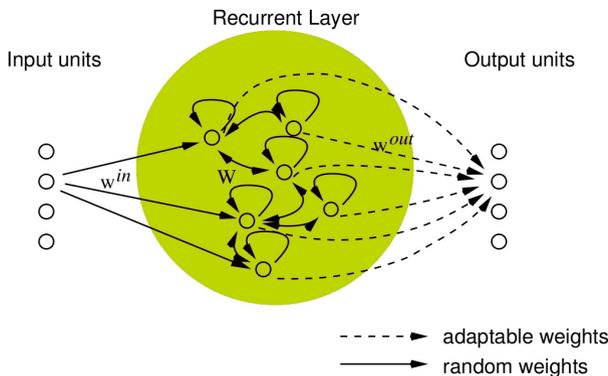


Fig. 11.1 The architecture of a typical Echo State Network (ESN), which belongs to the class of reservoir computing networks. In ESNs, the input and recurrent hidden layer (reservoir) connections are fixed randomly, and only output weights are trained. The reservoir projects the input stream nonlinearly into a high-dimensional representation, which can then be used by a linear readout layer. An important precondition for the approach to work is that the reservoir implements a fading memory, i.e. that reservoir states do not amplify, but fade out over time if no input is presented to the network.

system for some time. In this context, the metaphor of a reservoir is often used since the hidden layer can be seen as a water reservoir that gets disturbed by a drop, but slowly returns to its initial state after the ripples from the input have decayed. This reservoir state \mathbf{x} is mapped at time step $t + 1$ by an activation function $f(\cdot)$ such as a hyperbolic tangent, in the following way:

$$\mathbf{x}_{t+1} = f(\mathbf{W}_{\text{res}} * \mathbf{x}_t + \mathbf{W}_{\text{in}} * \mathbf{u}_{t+1}) \tag{11.1}$$

The input and hidden layer connections, \mathbf{W}_{in} and \mathbf{W}_{res} , are not trained in reservoir computing approaches. It is also possible to introduce feedback connections from outputs back into the reservoir (Lukovsevicius and Jaeger 2009). To approximate a specific target function, only the output weights \mathbf{w}_{out} are trained with a simple linear regression. This drastically simplifies the training procedure compared to previous approaches, while leading to excellent performance on time-series processing tasks (Jaeger and Haas 2004). It also avoids the problems of vanishing gradient information and disruptive bifurcations in the dynamics since no error gradients have to be propagated into the fixed, random parts of the network.

This approach works very well in practice. However, results will depend on the particular random set of weights that is drawn. In fact, there is considerable variability in performance between runs of networks with equal parameter settings, but different reservoir weights drawn each time (Ozturk et al. 2007). Striking a balance between the two extremes of fully trained RNNs and reservoir methods, it is interesting to retain some of the simplicity and efficiency of reservoir methods, but at the same time avoid some of the variability that comes with randomly created reservoirs. Self-organized methods are of interest here, because the initial random

configuration of the reservoir is in general already useful to perform the task. Each unit or connection then could, by way of local updates, contribute to an improved version of the reservoir, dependent on the data that each unit or weight processes over time. Advantages of self-organized methods are their potential for scalability, since they usually rely mainly on locally available information, making them good candidates for distributed processing.

Driving self-organization into a desired direction requires an understanding what properties a good RNN or reservoir network has. The mathematical tools to understand computation in these networks (which are instances of the larger class of input-driven dynamical systems) are still under active development (Manjunath et al. 2012). However, different perspectives, e.g., from functional analysis, dynamical systems, information theory, or statistical learning theory already offer insights towards this goal. They can provide answers to questions such as: how well can a given RNN approximate a given class of functions? How does it implement a certain function within the collective of its distributed nodes? How much memory does an RNN provide about past inputs and where is this information stored? How does information flow through the system at different time points in time and space? How well can it separate different inputs into different network states, and how well will it generalize on data that has not been seen during training? All of these aspects contribute to the successful performance of a network on a given task (or class of tasks). Understanding how to improve them will provide possible target signals to enrich and guide the self-organized optimization process of an RNN.

In this chapter, we review attempts that have been made towards understanding the computational properties and mechanisms of input-driven dynamical systems like RNNs, and reservoir computing networks in particular. We provide details on methods that have been developed to give quantitative answers to the questions above. Following this, we show how self-organization may be used to improve reservoirs for better performance, in some cases guided by the measures presented before. We also present a possible way to quantify task performance using an information-theoretic approach, and finally discuss promising future directions aimed at a better understanding of how these systems perform their computations and how to best guide self-organized processes for their optimization.

11.2 Assessing the Computational Power and Mechanisms of Information Processing of Reservoirs

In many cases, artificial neural networks are created with a specific goal in mind, for example to approximate a particular function or system. Training success and computational capability of the network with respect to this task are usually assessed on data that have not been used for training. Similarly to the training data, these are expected to match properties of the (yet unknown) application data well enough. A loss functional like the mean square error (MSE) or the cross-entropy is used to assess the quality of the trained system. For specific applications of the network, this is a standard approach that usually delivers meaningful results. When a neural

network is trained for a single purpose, it is not necessary to determine its general computational power, and the loss on the validation data or during its application may be the only relevant property.

The loss on a specific class of problems does not express the general computational power of the network, though. This property becomes more interesting when a part of the system is used for more than one task: relevant cases would be dynamical reservoirs that are used for multiple applications, networks that are trained “online” when the task changes, or to set up or to compare generic microcircuits. One of our motivations to evaluate mechanisms of information processing is to compare self-organized approaches within reservoirs. Ideally, self-organization leads to measurable effects in the reservoir which positively affect the performance of the system. In this section, we present a number of measures for different qualities of dynamical systems that are useful in this evaluation. These measures can be roughly divided into approaches that are based on or related to information theory, approaches that relate to learning theory, and dynamical systems theory.

11.2.1 Information-Theory Related Measures

Information theory and (Shannon) entropy have been used in a number of ways in neural network and complex systems research. One particular heuristic to measure (and eventually improve) RNN is to estimate and influence the entropy distribution of firing rates of a neuron. In individual biological neurons, for example, an approximate exponential distribution of the firing rate has been observed in visual cortical neurons (Baddeley et al. 1997). Under the constraint of a fixed energy budget, i.e., a fixed mean firing rate, an exponential distribution maximizes the potentially available information: it is the maximum entropy distribution for positive random values with a given mean. Triesch (2005) uses this idea to adapt the intrinsic excitability of a neuron with an online adaption. In this approach, the Kullback-Leibler divergence is used to measure the difference between the sample distribution of an individual neuron’s output, and the exponential distribution. Target distributions different from the exponential distribution are plausible dependent on specific circumstances. For example, in reservoir networks with real-valued units, normal distributions have been used (Schrauwen et al. 2008) to reflect the use of both negative and positive values. In both cases, the mechanism attempts to maximize the information per available energy unit locally at each neuron. Since energy constraints in reservoirs of artificial neural networks are typically not an issue, the maximum entropy distribution for these would in fact be the uniform. Without an energy constraint, the approach resembles the Infomax principle (Linsker 1987), where the average mutual information between input and output is maximized. As Bell and Sejnowski point out in their approach to maximize the mutual information for non-linear units (Bell and Sejnowski 1995), for invertible continuous deterministic mappings this mutual information is maximized by maximizing the output entropy alone. Due to the limited degrees of freedom of the approach, the desired target distribution cannot be approximated for every kind of input (Boedeker et al. 2009). Intrinsic plasticity

as well as its particular limitation can be related to Ashby's law of requisite variety (Ashby 1956) in that by increasing variety available in the reservoir a larger variety of outputs can be successfully approximated on one hand. On the other hand, the lack of variety in the mechanism adapting the individual neurons is also responsible for the difficulty in increasing the entropy for a variety of inputs.

The field of information dynamics (Lizier et al. 2007, 2012) provides information-theoretic measures that explicitly deal with processes or time-series. *Information storage*, as one of the tools, quantifies how much of the stored information is actually in use at the next time step when the next process value is computed. $A(X)$ is expressed as the mutual information between the semi-infinite past of the process X and its next state X' , with $X^{(k)}$ denoting the last k states of that process:

$$A(X) = \lim_{k \rightarrow \infty} A^{(k)}(X) \quad (11.2)$$

$$A(X, k) = I(X^{(k)}; X') \quad (11.3)$$

Information transfer, expressed as transfer entropy (Schreiber 2000), quantifies the influence of another process on the next state (for a formal definition, see Sect. 11.3.3 below). Boedecker et al. (2011) use these measures, to gain a better understanding of computational mechanisms inside the reservoir, and how they increase task performance for certain tasks at the phase transition between ordered and unstable dynamics. For a rote-memory task, a sharp peak can be observed in both information storage and information transfer near this phase transition, and suggests that a maximized capacity for information storage and information transfer correlates well with task performance in this dynamics regime, and this particular task. Prokopenko et al. (2011) suggest the Fisher information matrix as a way to detect phase transitions, but this has, to our knowledge, not been applied to RNN yet.

Fisher information also plays a role in quantifying the memory stored in a dynamical reservoir: Information about the recent input in reservoir networks is stored in the transient dynamics, rather than in attractor states. To quantify this information storage, Ganguli et al. (2008) use Fisher information as basis for a measure of memory traces in networks. The measure is applicable for systems with linear activations $\mathbf{f}(\mathbf{x}) = \mathbf{x}$, subject to Gaussian noise \mathbf{z} , and input $\mathbf{v}(t)$:

$$\mathbf{x}(t) = \mathbf{f}(\mathbf{W}^{\text{in}}\mathbf{v}(t) + \mathbf{W}\mathbf{x}(t-1) + \mathbf{z}(t)). \quad (11.4)$$

The Fisher Memory Matrix (FMM) between the present state of the system \mathbf{x} and the past signal is defined as

$$J_{k,l}(\mathbf{v}) = \left\langle -\frac{\delta^2}{\delta \mathbf{v}_k \delta \mathbf{v}_l} \log P(\mathbf{x}(t)|\mathbf{v}) \right\rangle_{P(\mathbf{x}(t)|\mathbf{v})}. \quad (11.5)$$

Diagonal elements $J(k) \equiv \mathbf{J}_{k,k}$ are the Fisher information that the system keeps in $\mathbf{x}(t)$ about a pulse at k steps back in the past, i.e., the decay of the memory trace of a past input. $J(k)$ is called the Fisher memory curve (FMC). Tino and Rodan (2013) investigate the relation between $J(k)$ and the short term memory capacity

MC (Jaeger 2001) (details on MC in the following subsection), and show that the two are closely related in linear systems. For these, $J(k)$ is independent of the actual input used. In the general, nonlinear case that is interesting for us, however, the FMC depends on the input signal, as the memory capacity MC does, and is hard to analyze.

A measure for Active Information Storage in input-driven systems has been proposed to quantify storage capabilities of a nonlinear system independent of particular inputs (Obst et al. 2013). The measure is an Active Information Storage (Lizier et al. 2012) where the current input u_{n+1} is conditioned out:

$$A_X^U(k) = \langle a_X^U(n+1, k) \rangle_n, \text{ with} \quad (11.6)$$

$$a_X^U(n+1, k) = \log \frac{p(x_n^{(k)}, x_{n+1} | u_{n+1})}{p(x_n^{(k)}) p(x_{n+1} | u_{n+1})} \quad (11.7)$$

$$= \log \frac{p(x_{n+1} | x_n^{(k)}, u_{n+1})}{p(x_{n+1} | u_{n+1})}. \quad (11.8)$$

The idea for this measure is to remove influences of structure in input data, and to only characterize the system itself, rather than a combination of system and input data. In theory, this influence would be removed by having the history sizes in computing the information storage converge to infinity. Large history sizes, however, require large amounts of data to estimate the involved joint probabilities, and this data, and the time required for the estimation is often not available. Active information Storage for input-driven systems assesses one aspect of the computational capabilities of a dynamical system, others, like the information transfer, would need to be defined for input-driven systems in a similar way.

11.2.2 Measures Related to Learning Theory

Legenstein and Maass (2007) propose two measures to quantify the computational capabilities of reservoirs in the context of liquid state machines, one of the two main flavors of reservoir computing networks: the linear separation property and the generalization capability. The linear separation property quantifies the ability of a computational system to map different input streams to significantly different internal states. This is useful because only then will the system be able to (linearly) map internal states to different outputs. The measure is based on the rank of an $n \times m$ matrix M whose columns are state vectors $\mathbf{x}_{u_i}(t_0)$ of circuit C after having been driven by input stream u_i up to a fixed time t_0 . These state vectors are collected for m different input streams, i.e., u_1, \dots, u_m . If the rank of M is m , then C , together with a linear readout, is able to implement any assignment of output units $y_i \in \mathbb{R}$ at time t_0 for inputs u_i .

For the generalization ability, they propose to approximate the VC-dimension of class \mathcal{H}_C of the reservoir, which includes all maps from a set S_{univ} of inputs u into $\{0, 1\}$ which can be implemented by a reservoir C . They present a theorem

(and corresponding proof sketch) stating that under the assumption that S_{univ} is finite and contains s inputs, the rank r of a $n \times s$ matrix whose columns are the state vectors $\mathbf{x}_u(t_0)$ for all inputs u in S_{univ} approximates the VC-dimension(\mathcal{H}_C), specifically $r \leq \text{VC-dimension}(\mathcal{H}_C) \leq r + 1$.

According to (Legenstein and Maass 2007), a simple difference of both (normalized) measures leads to good predictions about which reservoirs perform well on a range of tasks.

The loss or the success on a set of test functions is another possibility to characterize the systems from a learning point of view. One such measure is the short term memory capacity MC (Jaeger 2001) that we briefly mentioned above. To compute the MC, a network is trained to generate delayed versions $v(t - k)$ of a single channel input $v(k)$. The measure then is the sum of the precisions for all possible delays, expressed as a correlation coefficient:

$$\text{MC} = \sum_{k=1}^{\infty} \text{MC}_k \tag{11.9}$$

$$\text{MC}_k = \max_{\mathbf{w}_k^{\text{out}}} \frac{\text{cov}^2(v(t-k), y_k(t))}{\sigma^2(v(t)) \sigma^2(y_k(t))}, \text{ with} \tag{11.10}$$

$$y_k(t) = \mathbf{w}_k^{\text{out}} \begin{pmatrix} v(t) \\ \mathbf{x}(t) \end{pmatrix}, \text{ and } \mathbf{x}(t) = \mathbf{f}(\mathbf{W}^{\text{in}}v(t) + \mathbf{W}\mathbf{x}(t-1)).$$

The symbols cov and σ^2 denote covariance and variance, respectively. Each coefficient takes values between 0 and 1, and expresses how much of the variance in one signal is explainable by the other. As shown in (Jaeger 2001), for i.i.d. input and linear output units, the MC of N -unit RNN is bounded by N . The measure is related to the Fisher memory matrix approach above.

Another approach in this area is the *information processing capacity* of a dynamical system (Dambre et al. 2012). It is a measure based on the mean square error MSE in reconstructing a set of functions $z(t)$. The idea is to distinguish from approaches that view dynamical systems merely providing some form of memory for a – possibly nonlinear – readout. In (Dambre et al. 2012), systems are regarded as both providing memory and performing nonlinear computation. The readouts then only combine states of the system linearly, attempting to minimize the MSE for a function $z(t)$, so that all essential aspects of computation have to be covered by the dynamical system. The capacity of the system for approximating the desired output is computed using the (normalized) MSE of the optimal linear readout,

$$C_T[X, z] = 1 - \frac{\min_w \text{MSE}_T[\hat{z}]}{\langle z^2 \rangle_T} \tag{11.11}$$

This computed capacity is dependent on the input. In order to avoid an influence of structure in the input on the results, i.i.d. input is required for the purpose of measuring the capacity. To measure information processing capacity, several functions z have to be evaluated. The idea is that if z and z' are orthogonal, $C_T[X, z]$ and $C_T[X, z']$ measure independent properties of the system. The total capacity, on the other hand,

is limited by the number of variables x_i , so that a finite number of output functions is sufficient. A possible choice of output functions are Legendre polynomials, which are orthogonal over $(-1, 1)$.

The proposed approach has been used to compare different implementations of dynamical systems, like reaction-diffusion systems and reservoirs. An interesting idea that is also mentioned in (Dambre et al. 2012) would be to extend the approach so that the underlying system adapts to provide specific mappings. One possibility might be to adjust the number of internal units in an online-learning setting, e.g., when the task changes. The requirement for i.i.d. input is a limitation of the current approach, though it appears that even in the non-i.i.d. input case useful information about the system can be gathered. It might also be interesting to compare how the approach relates to the information-dynamics framework (Lizier et al. 2007, 2012) to quantify computation in non-linear systems.

11.2.3 Measures Related to Dynamical Systems Theory

To gain understanding of the internal operations that enable high-dimensional RNNs solving a given task, a recent effort by Sussillo and Barak (2013) draws on tools from dynamical systems theory. Using numerical optimization techniques, the authors identify fixed points and points of only gradual change (also called *slow points*) in the dynamics of the networks. Linearization around these points then reveals computational mechanisms like fixed points that implement memories, saddle points (fixed points with one unstable dimension) that enable switching between memory states, and approximate plane attractors that organize the storage of two consecutive input values to be memorized. For the tasks that were looked at in this work, the computational mechanisms could be inferred from the linearized dynamics around the set of fixed and slow points, and task performance of the trained networks was well explained.

In (Williams and Beer 2010), the authors argue for a complementary role of dynamical analysis, which involves, e.g., looking at attractors and switching between attractor states, and also an information-theoretic analysis when trying to understand computation in dynamical systems (including input-driven ones – even though the input might simply be considered as part of the environment and is assumed to be distributed uniformly). They evolve agents that are controlled by small continuous-time recurrent neural networks (CTRNNs) and evaluate their behavior in a relational categorization task. This involves keeping a memory about different objects the agent can sense, and reacting with avoid or catch behaviors based on the relation of both objects. Dynamical analysis shows that the state of a specific neuron in the CTRNN is correlated with the size of the first object, and switching on or off a different neuron determines whether the agent catches or avoids the second object. Both features are found to be connected through a saddle-node bifurcation in the CTRNN dynamics whose timing and location depends on properties of the second object. The desire to understand the flow of information through the

brain-body-environment system between these events leads the authors to information-theoretic measures unrolled over time (similar to the motivation and approach in (Lizier et al. 2007)). By considering the temporal evolution of measures like conditional mutual information, they are able to measure *information gain* or *information loss* of a state variable at specific time points. Similarly, they can quantify the *specific information* that a state variable carries about a particular stimulus at each time step. The behavior of the agent can then be explained by a sudden gain and then loss of information about object sizes in the first neuron, and then a rapid gain of information about relative size of the objects. In summary, the authors state that the two different ways to look at the computational mechanisms of the RNN differ, but provide coherent and even complementary information on how the agent solves the task that would be difficult to get with either approach alone.

Another approach from dynamical systems theory to understand and predict computational capabilities in RNNs builds on the concept of Lyapunov exponents. Although these concepts are only defined for autonomous dynamical systems, an analogous idea is to introduce a small perturbation into the state of one of two identical networks but not the other, and observe the time evolution of the state difference while the networks are driven with identical input. In case the perturbation fades out, the network is assumed to be in the stable phase of the dynamics. If it amplifies, the network is in the unstable, and possibly, the chaotic dynamics regime. If it approximately persists, the network is arguably at the phase-transition between stable and unstable regimes. Example applications of this approach can be found in (Bertschinger and Natschläger 2004; Boedecker et al. 2011). In (Bertschinger and Natschläger 2004), it was observed that the performance of binary threshold unit RNNs is maximized at this phase-transition for a task that requires memory and nonlinear processing to be solved successfully. This result was later replicated for analogue Echo State Networks in (Boedecker et al. 2011) for a rote-memory task; however, it was also found that some tasks do not benefit from reservoirs at the phase-transition, as observed before in the complex systems literature (e.g., Mitchell et al. 1993).

11.3 Improving Reservoir Information Processing Capabilities through Self-Organized Adaptation

A pragmatic way to evaluate the quality of a reservoir is to train the output, and evaluate it on a training or validation set (Lukovsevicius 2012a). In most circumstances, training is fast so that a number of hyper-parameter settings can be tested. Lukovsevicius (2012a) proposes a number of invaluable recipes to reservoir production. The recipes are very helpful for creating a good enough reservoir before output weights are trained. They show up promising directions for exploration, but are intended to be used as a guide rather than hard and fast rules, as some of them are mutually exclusive. The approaches selected for this section are intended to improve the reservoir itself in a self-organized way after it was created or selected. Possibly, this might happen simultaneously in combination with online learning of

output weights, or, alternatively, as a self-organized pre-training approach followed by the standard offline output weight training.

11.3.1 SORN: Self-Organized Optimization Based on 3 Local Plasticity Mechanisms

One approach that has demonstrated how self-organization can be leveraged to optimize a reservoir network can be found in (Lazar et al. 2009). SORN is a self-organizing recurrent network architecture using discrete-time binary units. The three plasticity mechanisms are: a variant of spike-time dependent plasticity (STDP), adjusting certain weights in the reservoir, a synaptic normalization rule (SN) responsible to keep the sum of afferent weights of a neuron constant, and intrinsic plasticity (IP) learning to adapt the unit firing threshold. The network state evolves using the following update functions:

$$R_i(t+1) = \sum_{j=1}^{N^E} W_{ij}^{EE}(t)x_j(t) - \sum_{k=1}^{N^I} W_{ik}^{EI}y_k(t) - T_i^E(t) \quad (11.12)$$

$$x_i(t+1) = \Theta(R_i(t+1) + v_i^U(t)) \quad (11.13)$$

$$y_i(t+1) = \Theta\left(\sum_{j=1}^{N^E} W_{ij}^{IE}(t)x_j(t) - T_j^I\right) \quad (11.14)$$

T^E and T^I are threshold values, drawn randomly from positive intervals for excitatory units and inhibitory units, respectively. Θ is the Heaviside step function, and $v_i^U(t)$ the network input drive. Matrices W^{IE} and W^{EI} are fully connected, and represent connections between inhibitory and excitatory units, and vice versa. W^{EE} holds connections between excitatory units. These are random, sparse, and without self-recurrence. Inhibitory units are not directly connected to each other. All weights are drawn from the interval $[0, 1]$, and the three matrices W^{IE} , W^{EI} , and W^{EE} are normalized, i.e., $\sum_j W_{ij} = 1$. The network state at time t is given by the two binary vectors $x(t) \in \{0, 1\}^{N^E}$, and $y(t) \in \{0, 1\}^{N^I}$, representing activity of the N^E excitatory and the N^I inhibitory units, respectively.

STDP and synaptic scaling update connections of excitatory units of the reservoir, while IP changes their thresholds. Inhibitory neurons and their connections remain unchanged. In the SORN the STDP for some small learning constant η_{stdp} is formalized as:

$$\Delta W_{ij}^{EE}(t) = \eta_{stdp}(x_i(t)x_j(t-1) - x_i(t-1)x_j(t)). \quad (11.15)$$

Synaptic scaling normalizes the values to sum up to one:

$$\Delta W_{ij}^{EE}(t) = W_{ij}^{EE}(t) / \sum_j W_{ij}^{EE}(t). \quad (11.16)$$

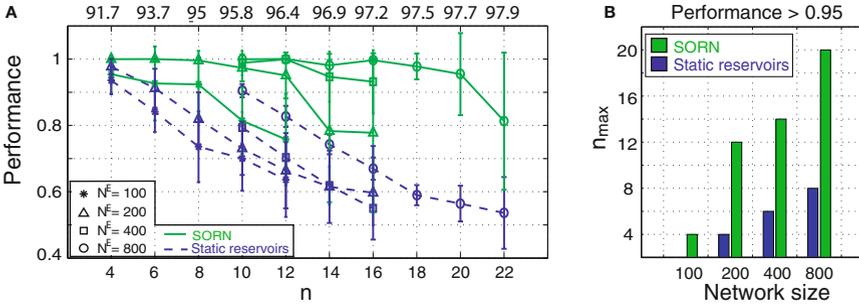


Fig. 11.2 (left) Normalized performance versus task difficulty as indicated by n , the number of repeated characters of a word which the network should predict. Different network sizes were tested. The numbers on top indicate the maximum possible performance – which is limited by the inherent randomness of the first character of a word within the sequence. Standard deviation among trials is indicated by the error bars. (right) Highest value of n for which a network achieved more than 95% of maximum performance as function of network size. The plastic SORN networks are able to deal with significantly harder tasks than the static reservoirs at this performance level. Graphs reproduced from (Lazar et al. 2009).

IP learning is responsible for spreading activations more evenly, using a learning rate η_{ip} , and a target firing rate of H_{IP} :

$$T_i^E(t + 1) = T_i^E(t) + \eta_{ip}(x_i(t) - H_{IP}) \tag{11.17}$$

Lazar et al. (2009) show that the SORN outperforms static reservoir networks using a letter prediction task. The network has to predict the next letter in a sequence of two different artificial words of length $n + 2$. These words are made up of three different characters, with the second character repeated n times. The first character of a word is random and the network cannot do better than randomly guessing which one will come up. If the reservoir is able to efficiently separate the repeated character in the middle part of the word, though, the network can learn to count these characters and predict the rest of the sequence correctly. Figure 11.2 compares the normalized performance of SORNs and static reservoir networks of different sizes on instances of the task with increasing n (increased difficulty). SORNs are able to outperform static reservoirs clearly on this task. A PCA analysis in (Lazar et al. 2009) reveals that the SORN indeed shows a much better separation property and maps repeated inputs to distinct network states, while the states of static reservoirs are much more clustered together and thus harder to distinguish by the linear readout.

The combination of the three mechanisms appears to be a key to successful self-organization in an RNN. Figure 11.3 illustrates that the dynamics of SORN reservoir become sub-optimal if only two of the three plasticity mechanisms are active. Without synaptic normalization, the network units become highly synchronized. This severely restricts the representational power of the reservoirs. If IP learning is switched off, the activity of neurons in the network becomes unbalanced. Some

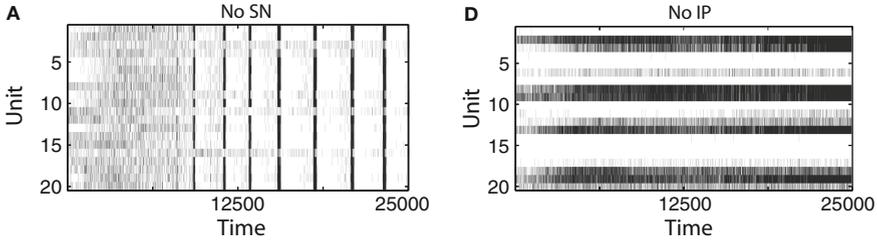


Fig. 11.3 Activity snapshots for 50 randomly selected reservoir neurons. (left) A reservoir without synaptic scaling develops highly synchronized, seizure-like firing patterns. (right) Without the IP mechanism, neuron activity is unevenly distributed with some neurons firing almost constantly while others are nearly silent. Graphs reproduced from (Lazar et al. 2009).

neurons fall nearly silent while others are active almost all the time. This is in contrast to the case with IP where activity is more evenly distributed, enabling a richer representation of information in the reservoir.

Though some of the self-organizing mechanisms like STDP are biologically plausible, there are not too many examples of successful applications for training RNNs, or, as Lazar et al. (2009) states, “Understanding and controlling the ensuing self-organization of network structure and dynamics as a function of the network’s inputs is a formidable challenge”. For time-series prediction and system identification tasks, an extension of the approach to analog units would be required. Also, an investigation of the information dynamics during and after adaptation may provide insights, for example into the relation between reservoir configuration and information transfer.

11.3.2 Hierarchical Self-Organizing Reservoirs

A different approach based on self-organized optimization of reservoirs is presented in (Lukovsevičius 2012b). The author compares classical ESNs and recurrent RBF-unit based reservoir networks (called Self-Organizing Reservoirs, SORs) which resemble Recurrent Self-Organizing Maps (RSOMs) (Voegtlin 2002). The input and reservoir weights of the SOR are adapted by learning rules traditionally used for Self-Organizing Maps (SOMs) (Kohonen 2001) and NeuralGas networks (Martinetz and Schulten 1991).

The update equations for the SOR are:

$$\tilde{x}_i(n) = \exp(-\alpha \|\mathbf{v}_i^{in} - \mathbf{u}(n)\|^2 - \beta \|\mathbf{v}_i - \mathbf{x}(n-1)\|^2), \quad i = 1, \dots, N_x, \quad (11.18)$$

$$\mathbf{x}(n) = (1 - \gamma)\mathbf{x}(n-1) + \gamma\tilde{\mathbf{x}}(n). \quad (11.19)$$

Here, the internal reservoir neuron states at time n are collected in vector $\mathbf{x} \in \mathbb{R}^{N_x}$ and their update in vector $\tilde{\mathbf{x}} \in \mathbb{R}^{N_x}$. The factor $\gamma \in (0, 1]$ is the leak-rate. The vector $\mathbf{u} \in \mathbb{R}^{N_u}$ contains the input-signal, while matrices \mathbf{V}^{in} and \mathbf{V} are the input and recurrent weight matrix, respectively, whose i th column vectors are denoted by

$\mathbf{v}_i^{in} \in \mathbb{R}^{N_x}$ and $\mathbf{v}_i \in \mathbb{R}^{N_x}$. Parameters α and β scale the input and recurrent distances, and $\|\cdot\|$ denotes the Euclidean norm.

The unsupervised training of the SOR updates the input and recurrent weights as:

$$\mathbf{v}_i^{all}(n+1) = \mathbf{v}_i^{all}(n) + \eta(n)h(i,n)([\mathbf{u}(n); \mathbf{x}(n)] - \mathbf{v}_i^{all}(n)), \quad (11.20)$$

where $\mathbf{v}_i^{all}(n) \equiv [\mathbf{v}_i^{in}; \mathbf{v}]$ and $\eta(n)$ is a time-dependent learning rate. The learning-gradient distribution function h is defined either as:

$$h(i,n) = \exp(-d_h(i, \text{bmu}(n))^2 / b_h(n)^2), \quad (11.21)$$

where $d_h(i, j)$ is the distance between reservoir units i and j on a specific topology, $\text{bmu}(n) = \arg \max_i (x_i(n))$ is a function returning the index of a *best matching unit* (BMU), and $b_h(n)$ is the time-dependent of the learning gradient distribution. With this definition of h , the learning proceeds according to the SOM algorithm. To implement NeuralGas-like learning, it suffices to change this definition to:

$$h_{ng}(i, n) = \exp(-d_{ng}(i, n) / b_h(n)), \quad (11.22)$$

where $d_{ng}(i, n)$ denotes the index of node i in the descending ordering of activities $x_i(n)$ (see (Lukovsevicius 2012b) for additional details). Both algorithms were found to be similarly effective to improve the pattern separation capability of reservoirs compared to standard ESNs when tested on detection of certain signal components on a synthetic temporal pattern benchmark, and on classification of handwritten digits from a stream of these characters. Further improvements are reported if these SORs are stacked on top of each other in a hierarchy, trained in a layer-by-layer fashion. However, results only improve if enough time is given for the self-organization process to find suitable representations. If layers are stacked with very little training time for each of them, performance actually worsens.

11.3.3 Guided Self-Organization of Reservoir Information Transfer

In (Obst et al. 2010), the information transfer between input data and desired outputs is used to guide the adaptation of the self-recurrence in the hidden layer of a reservoir computing network. The idea behind this step is to change the memory within the system with respect to the inherent memory in input and output data (see Section 11.4 below for a discussion which develops these ideas further).

The network dynamics is updated as:

$$\mathbf{x}(k+1) = \text{diag}(\mathbf{a})\mathbf{W}\mathbf{y}(k) + (\mathbf{I} - \text{diag}(\mathbf{a}))\mathbf{y}(k) + \mathbf{w}^{in}\mathbf{u}(k) \quad (11.23)$$

$$\mathbf{y}(k+1) = \mathbf{f}(\mathbf{x}(k+1)), \quad (11.24)$$

where $x_i, i = 1, \dots, N$ are the unit activations, \mathbf{W} is the $N \times N$ reservoir weight matrix, \mathbf{w}^{in} the input weight vector, $\mathbf{a} = [a_1, \dots, a_N]^T$ a vector of local decay factors,

\mathbf{I} is the identity matrix, and k denotes the discrete time step. As a nonlinearity, $f(x) = \tanh(x)$ is used. The a_i represent the coupling of a unit's previous state with the current state, and are computed as:

$$a_i = \frac{2}{1 + m_i},$$

where m_i represents the memory length of unit i ($m_i \in \{1, 2, 3, \dots\}$), initialized to $m_i = 1$. Increasing individual m_i through adaptation increases the influence of a unit's past states on its current state. The information transfer is quantified as a conditional mutual information or transfer entropy (Schreiber 2000):

$$T_{X \rightarrow Y} = \lim_{k, l \rightarrow \infty} T_{X \rightarrow Y}^{(k, l)}, \text{ with} \quad (11.25)$$

$$T_{X \rightarrow Y}^{(k, l)} = I(X^{(l)}; Y^l | Y^{(k)}). \quad (11.26)$$

Parameters k and l are history sizes, which lead to finite-sized approximations of the transfer entropy for finite values.

In a first step, the required history size l is determined which maximizes the information transfer $T_{u \rightarrow v}$ from input u to output v . This value will increase for successively larger history sizes, but the increases are likely to level off for large values of l . Therefore, l is determined as the smallest value which is still able to increase $T_{u \rightarrow v}$ by more than a threshold ε :

$$T_{\mathbf{u} \rightarrow \mathbf{v}}(1, \hat{l} + 1) \leq T_{\mathbf{u} \rightarrow \mathbf{v}}(1, \hat{l}) + \varepsilon \quad \text{and} \quad (11.27)$$

$$T_{\mathbf{u} \rightarrow \mathbf{v}}(1, l) > T_{\mathbf{u} \rightarrow \mathbf{v}}(1, l - 1) + \varepsilon \quad \text{for all } l < \hat{l}. \quad (11.28)$$

In a second step, the local couplings of the reservoir units are adapted so that the transfer entropy from the input of each unit to its respective output is optimized for the particular input history length \hat{l} , as determined in step one. Over each epoch θ of length ℓ , we compute the transfer entropy from activations $x_i^{(\ell)}$ to output $y_i^{(\ell)}$ for each unit i :

$$te_i^\theta = T_{x_i^{(\ell)} \rightarrow y_i^{(\ell)}}(1, \hat{l}).$$

If the information transfer during the current epoch θ exceeds the information transfer during the past epoch by a threshold (i.e., $te_i^\theta > te_i^{\theta-1} + \varepsilon$), the local memory length m_i is increased by one. In case $te_i^\theta < te_i^{\theta-1} - \varepsilon$, the local memory length is decreased by one, down to a minimum of 1. The decay factors a_i are fixed once they stabilize, which ends the pre-training phase.

In (Obst et al. 2010), the method is tested on a one-step ahead prediction of unidirectionally coupled maps and of the Mackey-Glass time series benchmark. Showing results for the former task as an example, Figure 11.4 (left) displays the mean square errors of the prediction over the test data for different coupling strengths e and fixed order parameter ω for both echo state learning with and without adaptation of information transfer in the reservoir (averages over 50 trials). For each individual trial the same reservoir and time series have been used once with and without adaptation.

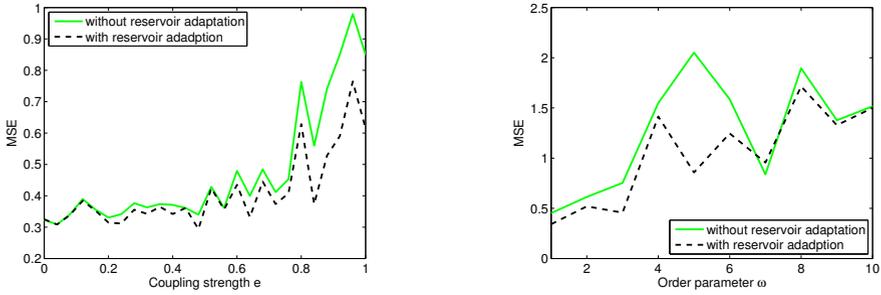


Fig. 11.4 (left) Mean squared errors of the prediction over the test data for different coupling strengths and fixed $\omega = 0$. (right) Mean squared error for different ω using a fixed coupling of $e = 0.75$.

The prediction using the reservoir adaptation is better over almost the entire range of e , with the improvement becoming more significant as the influence of the input time series becomes larger. Figure 11.4 (right) is a plot of the mean square error for different ω using a fixed coupling e . In all but one cases the reservoir adaptation improves results.

11.4 Quantifying Task Complexity

Most currently existing measures capture some of the generic computational properties of recurrent neural networks (as an important class of input-driven system), such as memory capacity or entropy at the neuron-level, but do not take task complexity into account. Optimization of the network properties based on these generic measures therefore will only do a “blind” adjustment of parameters while no optimality guarantee for the task at hand can be given. More positively put, the philosophy behind these measures is that a maximization of some of them leads to reservoirs that are capable to solve a variety of tasks. The self-organizing mechanisms in Sect. 11.3 are one way to achieve this maximization. In situations where no teaching signal is given, e.g., in clustering tasks, one can do no better than that; however, if the desired output signal is available, it can be used to quantify the task complexity as a relationship between inherent difficulty of predicting the output based on its own history, and to what extent the input data can contribute to improve these predictions. This would inform us how achievable a task is, and may also be used to trade off complexity of the system against the expected quality of the solution.

It is possible to use some of the tools that we introduced above, and take an information-theoretic approach to tackle this problem. Essentially we are interested in quantifying how difficult it is for a system to produce its next output. The systems we are interested in take a time series X as an input, and have the goal to generate output Y , another time series. To produce the next state y_{t+1} , both the output’s past $(y_1 \dots y_t)$ as well as the input up to the current step (y_1, \dots, y_{t+1}) can be considered.

The Active Information Storage A_Y (Lizier et al. 2012) can be used to capture the influence of previous outputs in producing the next output: how much information

is contained in the past of Y that can be used to compute its next state? This is expressed as the average mutual information between past $Y^{(k)}$ of the output and the next state Y' :

$$A_Y = \lim_{k \rightarrow \infty} A_Y(k), \text{ with} \quad (11.29)$$

$$A_Y(k) = I(Y^{(k)}; Y'). \quad (11.30)$$

We use $A_Y(k)$ to represent finite- k estimates. Now, A_Y and $A_Y(k)$ allow for two kind of measurements: (a) higher values for A_Y indicate better predictability of Y from its own past, i.e., A_Y is one component of the overall task difficulty. (b) With increasing values of $k_i = 1, 2, \dots, n$, estimates $A_Y(k_i)$ indicate the amount of memory that is in use. As the information that can be used to predict the next state increases with larger values of k_i , $A_Y(k_i)$ will monotonically increase with k_i , and asymptotically converge to some maximum. Finding k^* so that $A_Y(k^*) \geq A_Y(k) + \varepsilon$, for some small $\varepsilon > 0$ and $k \rightarrow \infty$, thus gives us a useful quantity for the amount of memory required, and at the same time $A_Y(k^*) \approx A_Y$ quantifies the difficulty in predicting Y' from its own past.

The other component that plays a role in the task is the input X . Its contribution to producing the next output Y' of the system, too, can be quantified, using the transfer entropy introduced above. The transfer entropy indicates how much information the input X contributes to the next state Y' , given that the past of Y is known. Increasing the input history size l increases the information available in computing Y' , for fixed output history size k^* . Large $T_{X \rightarrow Y}$ suggest that the input X helps in computing the next output, i.e., the task for the system is less difficult than for smaller transfer entropies. Finding an input history size l^* so that $T_{X \rightarrow Y}^{(k^*, l^*)} \approx \lim_{l \rightarrow \infty} T_{X \rightarrow Y}^{(k^*, l)}$ gives us another useful quantity for the amount of memory required.

Unfortunately, using these quantities to compare tasks or to design systems is not entirely straightforward, for a number of reasons. For continuous-valued time series X and Y , estimating mutual informations is cumbersome, and requires larger amounts of data in particular for larger history sizes l and k . To compare task difficulties, it would also be helpful to normalize both quantities, e.g., for the Active Information Storage by the joint entropy $H(Y; Y')$, to values between 0 and 1. The true output history Y may also be simply not available to the system, dependent on how it operates. For example, in batch mode, the only information that is available is the input X and the *estimated* output \hat{Y} . The true history of Y is usually only accessible if the system operates online. Finally, the two components A_Y and $T_{X \rightarrow Y}$ cannot be simply added to specify the overall task difficulty since input and output may redundantly share some information.

We will reserve a detailed investigation of applying both measures to a later publication. As a concept to explain contributions of input and output history, they can be an indicator for how complex the information processing system needs to be. It will also be interesting to see how other measures relate to them, and to show which aspect of the computation they measure. As an example, the memory capacity MC, as a sum of correlation coefficients can be seen as a linear measure of the potential

information transfer between input X and the desired output Y . Mutual information expresses a nonlinear relationship between two variables, and so does the transfer entropy, a conditional mutual information between X and Y . In contrast to MC, the $T_{X \rightarrow Y}$ measures actual information transfer between input and desired output, i.e., $T_{X \rightarrow Y}$ is a purely a property of the task. MC is a property of the RNN, but as it is using task specific input, it combines the properties of the RNN with properties of the input. The two quantities could be used to adjust the architecture of a neural network for better performance on a specific task.

As another example, measuring the individual distributions of unit activations in the reservoir and their divergence from a maximum entropy distribution capture properties of the input combined with properties of the network. On the other hand, Active Information Storage for input-driven systems, applied to reservoirs or individual reservoir units, expresses the amount of information in the system that is in use to predict the next state, and is meant to measure capabilities of the system only.

Related work on complexity measures includes Grassberger's forecast complexity (Grassberger 1986, 2012), which considers the difficulty of making an optimal prediction of a sequence created by a stochastic process. A sequence can be compressed up to its entropy rate, and the forecast complexity is the computational complexity of the algorithm responsible for the decompression. The probability of the next symbol is needed for this decompression. Related ideas can be found in Minimum Description Length (MDL) approaches (Rissanen 1978) and Kolmogorov complexity (Kolmogorov 1965) as measures for complex objects. Also introduced by Grassberger (1986, 2012) is the Effective Measure Complexity (EMC), the relative memory required to calculate the probability distribution of the next symbol of a sequence. The EMC is a lower bound on the forecast complexity. Both forecast complexity and EMC look at sequences, e.g., the output of an autonomous system without regard to its input, whereas we are interested in systems that produce an output based on some input. More complexity measures can be found in a special issue on "Measures of Complexity from Theory to Applications", with (Crutchfield and Machta 2011) as an introductory article.

11.5 Conclusion

We presented methods to assess different computational properties of input-driven RNNs, and reservoir computing networks in particular, in the first part of the paper. These methods were drawn from information-theory, statistical learning theory, and dynamical systems theory, and provided different perspectives on important aspects of information processing in these systems. They help to quantify properties like the memory capacity a certain network provides, the flow of information through the system and its modification over time, the ability to separate similar inputs and generalize to new, unseen data, and others. In addition to their usefulness in their own right when trying to understand how RNNs implement the functions they are trained for, they also have the potential to be used as target signals to guide

self-organized optimization procedures aimed at improving the quality of reservoirs for a specific task over random initialization.

In the second part of the paper, we presented some recent efforts at implementing self-organized optimization for reservoir computing networks. One approach combined different plasticity mechanisms to improve coding quality and separation ability of the network, while a different approach was using methods similar to recursive self-organizing maps with SOM and NeuralGas-like learning rules. The final approach we presented proceeds in two phases: determining a learning goal in terms of information transfer between input and desired output, and using this quantity to guide local adjustments to the self-recurrence of each reservoir unit. All of these methods showed the potential of self-organized methods to improve network performance over standard, random reservoirs while avoiding problems associated with back propagation of error-gradients throughout the whole networks.

As a next step towards methods that are able to automatically generate or optimize recurrent neural networks for a specific task (or class of tasks), it seems worthwhile to combine measures for network properties and task complexity, and devise algorithms that adjust the former based on the latter. The approach taken in (Dambre et al. 2012) of using orthogonal functions to measure information processing capacity could be extended to construct suitable dynamical systems for a task when the requirements for a specific task can be measured in a similar way.

A comparison of how the measures of the information dynamics framework (Lizier et al. 2007, 2012), the information processing capacity for dynamical systems (Dambre et al. 2012), measures of criticality (Bertschinger and Natschläger 2004; Prokopenko et al. 2011) or of memory capacity (Jaeger 2001; Ganguli et al. 2008) relate to each other should reveal some interesting insights (see, e.g., Tino and Rodan 2013), since they all cover some aspects of dynamical systems. Establishing the relation between the information dynamics framework, with recent extension for input-driven systems, and information processing capacity, for example, could help to overcome requirements for i.i.d. input in the latter, to better understand dynamical systems with arbitrary input.

References

- Ashby, W.R.: *An Introduction to Cybernetics*. Chapman & Hall, London (1956)
- Baddeley, R., Abbott, L.F., Booth, M.C.A., Sengpiel, F., Freeman, T., Wakeman, E.A., Roll, E.T.: Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proc. R. Soc. Lond. B* 264, 1775–1783 (1997)
- Bell, A.J., Sejnowski, T.J.: An information-maximization approach to blind separation and blind deconvolution. *Neural Computation* 7(6), 1129–1159 (1995)
- Bengio, Y., Boulanger-Lewandowski, N., Pascanu, R.: *Advances in Optimizing Recurrent Networks*. arXiv preprint 1212.0901, arXiv.org (2012)
- Bengio, Y., Simard, P., Frasconi, P.: Learning long-term dependencies with gradient descent is difficult. *IEEE Transaction on Neural Networks* 5(2), 157–166 (1994)
- Bertschinger, N., Natschläger, T.: Real-time computation at the edge of chaos in recurrent neural networks. *Neural Computation* 16(7), 1413–1436 (2004)

- Boedecker, J., Obst, O., Lizier, J.T., Mayer, N.M., Asada, M.: Information processing in echo state networks at the edge of chaos. *Theory in biosciences Theorie in den Biowissenschaften* 131(3), 1–9 (2011)
- Boedecker, J., Obst, O., Mayer, N.M., Asada, M.: Initialization and self-organized optimization of recurrent neural network connectivity. *HFSP Journal* 3(5), 340–349 (2009)
- Crutchfield, J.P., Machta, J.: Introduction to focus issue on “Randomness, Structure, and Causality: Measures of complexity from theory to applications”. *Chaos* 21(3), 037101 (2011)
- Dambre, J., Verstraeten, D., Schrauwen, B., Massar, S.: Information processing capacity of dynamical systems. *Scientific Reports* 2, 514 (2012)
- Douglas, R., Markram, H., Martin, K.: Neocortex. In: Shepherd, G. (ed.) *Synaptic Organization In the Brain*, pp. 499–558. Oxford University Press (2004)
- Doya, K.: Bifurcations in the learning of recurrent neural networks. In: *IEEE International Symposium on Circuits and Systems*, pp. 2777–2780. IEEE (1992)
- Ganguli, S., Huh, D., Sompolinsky, H.: Memory traces in dynamical systems. *Proceedings of the National Academy of Sciences* 105(48), 18970–18975 (2008)
- Grassberger, P.: Toward a quantitative theory of self-generated complexity. *International Journal of Theoretical Physics* 25(9), 907–938 (1986)
- Grassberger, P.: Randomness, information, and complexity. Technical Report 1208.3459, arXiv.org (2012)
- Hawkins, J., Blakeslee, S.: *On Intelligence*. Times Books (2004)
- Hochreiter, S.: The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 6(2), 107–116 (1998)
- Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Computation* 9(8), 1735–1780 (1997)
- Jaeger, H.: Short term memory in echo state networks. Technical Report 152, GMD – German National Research Institute for Computer Science (2001)
- Jaeger, H., Haas, H.: Harnessing Nonlinearity: Predicting Chaotic Systems and Saving Energy in Wireless Communication. *Science* 304(5667), 78–80 (2004)
- Kohonen, T.: *Self-Organizing Maps*, 3rd, extended edn. Springer (2001)
- Kolmogorov, A.N.: Three approaches to the quantitative definition of information. *Problemy Peredachi Informatsii* 1(1), 3–11 (1965)
- Lazar, A., Pipa, G., Triesch, J.: SORN: a self-organizing recurrent neural network. *Frontiers in Computational Neuroscience* 3, 23 (2009)
- Legenstein, R., Maass, W.: What makes a dynamical system computationally powerful. In: Haykin, S., Principe, J.C., Sejnowski, T., McWhirter, J. (eds.) *New Directions in Statistical Signal Processing: From Systems to Brains*, pp. 127–154. MIT Press (2007)
- Linsker, R.: Towards an organizing principle for a layered perceptual network. In: Anderson, D.Z. (ed.) *NIPS*, pp. 485–494. American Institute of Physics (1987)
- Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Detecting non-trivial computation in complex dynamics. In: Almeida e Costa, F., Rocha, L.M., Costa, E., Harvey, I., Coutinho, A. (eds.) *ECAL 2007. LNCS (LNAI)*, vol. 4648, pp. 895–904. Springer, Heidelberg (2007)
- Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Local measures of information storage in complex distributed computation. *Information Sciences* 208, 39–54 (2012)
- Lukovševičius, M.: A practical guide to applying echo state networks. In: Montavon, G., Orr, G.B., Müller, K.-R. (eds.) *Neural Networks: Tricks of the Trade*, 2nd edn. LNCS, vol. 7700, pp. 659–686. Springer, Heidelberg (2012a)

- Lukovševičius, M.: Self-organized reservoirs and their hierarchies. In: Villa, A.E.P., Duch, W., Érdi, P., Masulli, F., Palm, G. (eds.) ICANN 2012, Part I. LNCS, vol. 7552, pp. 587–595. Springer, Heidelberg (2012b)
- Lukovševičius, M., Jaeger, H.: Reservoir computing approaches to recurrent neural network training. *Computer Science Review* 3(3), 127–149 (2009)
- Maass, W., Joshi, P., Sontag, E.D.: Computational aspects of feedback in neural circuits. *PLOS Computational Biology* 3(1), e165 (2007)
- Manjunath, G., Tino, P., Jaeger, H.: Theory of Input Driven Dynamical Systems. In: dice.ucl.ac.be, pp. 25–27 (April 2012)
- Martens, J., Sutskever, I.: Learning recurrent neural networks with hessian-free optimization. In: Proceedings of the 28th International Conference on Machine Learning, vol. 46, p. 68. Omnipress Madison, WI (2011)
- Martinetz, T., Schulten, K.: A “neural-gas” network learns topologies. *Artificial Neural Networks* 1, 397–402 (1991)
- Mitchell, M., Hraber, P.T., Crutchfield, J.P.: Revisiting the edge of chaos: Evolving cellular automata to perform computations. *Complex Systems* 7, 89–130 (1993)
- Obst, O., Boedecker, J., Asada, M.: Improving Recurrent Neural Network Performance Using Transfer Entropy. *Neural Information Processing Models and Applications* 6444, 193–200 (2010)
- Obst, O., Boedecker, J., Schmidt, B., Asada, M.: On active information storage in input-driven systems. preprint 1303.5526v1, arXiv.org (2013)
- Ozturk, M.C., Xu, D., Principe, J.C.: Analysis and design of echo state networks. *Neural Computation* 19(1), 111–138 (2007)
- Prokopenko, M., Lizier, J.T., Obst, O., Wang, X.R.: Relating Fisher information to order parameters. *Physical Review E* 84(4), 041116 (2011)
- Riedmiller, M., Braun, H.: A direct adaptive method for faster backpropagation learning: the rprop algorithm. In: IEEE International Conference on Neural Networks, vol. 1, pp. 586–591 (1993)
- Rissanen, J.: Modeling by shortest data description. *Automatica* 14(5), 465–471 (1978)
- Rolls, E.T., Deco, G.: *The Noisy Brain - Stochastic Dynamics as a Principle of Brain Function*. Oxford University Press (2010)
- Rumelhart, D., Hinton, G., Williams, R.: Learning representations by back-propagating errors. *Nature* 323(6088), 533–536 (1986)
- Schmidhuber, J., Wierstra, D., Gagliolo, M., Gomez, F.: Training recurrent networks by evoluno. *Neural Computation* 19(3), 757–779 (2007)
- Schrauwen, B., Wardermann, M., Verstraeten, D., Steil, J.J., Stroobandt, D.: Improving reservoirs using intrinsic plasticity. *Neurocomputing* 71(7-9), 1159–1171 (2008)
- Schreiber, T.: Measuring information transfer. *Physical Review Letters* 85(2), 461–464 (2000)
- Seung, H.S.: *Connectome: How the Brain’s Wiring Makes Us Who We Are*. Houghton Mifflin Harcourt, New York (2012)
- Sporns, O.: *Networks Of the Brain*. The MIT Press (2011)
- Sussillo, D., Barak, O.: Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. *Neural Computation* 25(3), 626–649 (2013)
- Tino, P., Rodan, A.: Short term memory in input-driven linear dynamical systems. *Neurocomputing* (2013)
- Triesch, J.: A gradient rule for the plasticity of a neuron’s intrinsic excitability. In: Duch, W., Kacprzyk, J., Oja, E., Zadrożny, S. (eds.) ICANN 2005. LNCS, vol. 3696, pp. 65–70. Springer, Heidelberg (2005)

- Voegtlin, T.: Recursive self-organizing maps. *Neural Networks* 15(8-9), 979–991 (2002)
- Werbos, P.J.: Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE* 78(10), 1550–1560 (1990)
- Williams, P.L., Beer, R.D.: Information dynamics of evolved agents. *From Animals to Animats* 11, 38–49 (2010)
- Williams, R.J., Zipser, D.: A learning algorithm for continually running fully recurrent neural networks. *Neural Computation* 1(2), 270–280 (1989)