

Special Issue on Echo State Networks and Liquid State Machines

Seeking plausible models for brain computation has been a continuing effort in neuroscience, computer science, biophysics and machine learning. Quite generally speaking, there are two routes toward understanding brains. In a bottom-up way, one can attempt to re-create artificial brain structures from empirical observations, whose emerging dynamics are then studied by simulation, searching for dynamical patterns that can be understood in terms of information processing. Conversely, one can proceed top-down, starting from computational metaphors taken from computer science or signal processing and control engineering, and try to synthesize artificial brain modules from these principles whilst relating them to what is known about biological brains. Whatever route one takes, at some point one has to introduce a computational mechanism, some mathematical information processing principle. A large number of such principles have been considered – we just mention logical calculi, Turing computation, "cybernetic" regulation mechanisms, energy-minimizing and particle dynamics motivated by information theory and statistical physics, field theories and pattern-forming nonlinear PDEs, or chaotic attractor dynamics. Originating from outside neuroscience, these information processing mechanisms often are difficult to connect to known neural processing mechanisms. There is however a subset of such principles of very elementary nature whose main learning and activation phenomena can readily be mapped to biology – and indeed have been partially motivated by natural neural systems. We would count in this category the perceptron (trained with the Widrow-Hoff rule), Hopfield networks, and self-organizing maps. The elementary nature of these models makes them amenable to mathematical analysis and invites mapping them to biological brains in many ways. We view echo state networks and liquid state machines, the heroes of this special issue, as a further member in this family of versatile basic computational metaphors with a clear biological footing.

Circuits of neurons in the brain are recurrently connected, and recurrently connected circuits are obviously needed for all computational and cognitive tasks that require temporal integration of information. However it has turned out to be difficult to understand how synapses within such recurrent circuit should be modified in order to improve the computational performance of the system. One radical solution is to view the recurrent circuit as a generic device (in analogy to the kernel of a support vector machine), and to concentrate learning efforts on the training of linear readouts from the recurrent circuit. This strategy has the advantage that training cannot get stuck in local minima of the mean-squared-error function, and that generalization of learned responses to new instances of the task is optimally supported (due to the low VC-dimension of linear readouts in comparison with nonlinear learning devices). This strategy had in fact already been suggested in Rosenblatt's book on perceptrons [Rosenblatt, 1962], as a side remark in connectionist AI [Hinton 1985] and also in neuroscience [Buonomano and Merzenich, 1995]. The resulting new computational model has been investigated more rigorously under the name of echo states networks (ESNs) and liquid state machines (LSMs). The hallmark of both models is that they compute using a large, distributed, nonlinear dynamical recurrent network with fixed weights, called a reservoir, with adaptation restricted to the readout. This greatly simplifies the training in practical

applications, and avoids the biological implausibility of multilayer gradient descent optimization of previous recurrent neural networks. These two models had been designed independently, with different application types and different parameter regimes in mind. Theoretical results on LSMs are quite general, and have been formulated within the mathematical frameworks of dynamical system theory and filtering theory. Hence they apply in particular also to ESNs. But since the primary goal of the development of LSMs was to provide a biologically plausible paradigm for computations in generic cortical microcircuits, applications of this model have only been explored for circuits of spiking neurons with a biologically characteristic large amount of internal noise. In contrast, ESNs have been designed to provide high performance for a number of engineering tasks, and have been primarily applied to recurrent artificial neural networks without internal noise, which are better suited for such tasks. Theoretical results for ESNs (e.g., on the connection between network structure and memory time span) have contributed more complex results that are valid for simpler classes of recurrent circuits (in particular for circuits consisting of linear neurons). As to terminology, "reservoir computing" appears to be increasingly used as the generic name for LSMs, ESNs, and variants thereof which are being created.

During the last few years, a number of research groups have started to explore these new computational paradigms from different angles. This special issue can be seen both as a snapshot and a resume of the current state of the art. One research question that is common to most of this research is the question which recurrent circuits (or more generally: which dynamical systems) are optimal for a given range of tasks. This research can be seen as a parallel to ongoing research on support vector machines in machine learning, where one asks which kernels are optimal for a given range of application tasks (although the situation there is similar insofar as there one also gets already very good performance with a fixed generic kernel, for example an RBF-kernel).

The articles of this special issue fall into three groups. In the first group one finds contributions which investigate LSMs as models of natural neural systems, on various levels of abstraction. *Yamazaki and Tanaka* study the cerebellum as an LSM and state evidence that the granular layer acts as the reservoir, while the Purkinje cells act as the readout neurons, unlike the established analogy with the perceptron. *Joshi* maps LSM modules (with empirically plausible dynamical parameters) to subtasks of a more complex behavioral task that involves working memory and decision making, and shows that quite different subtasks can be learned using the same reservoir and different readouts. *Lazar, Pipa and Triesch* combine spike timing dependent plasticity and intrinsic plasticity to maintain homeostasis of neuronal activity that stabilizes the LSM and benefits the recognition of temporal patterns in time series. Finally, *Legenstein and Maass* ask which aspects of recurrent circuits of spiking neurons are relevant for their computational performance, and derive two measures (one of which evaluates to what extent a given neural circuit shares properties with an RBF-kernel, the other is a VC dimension measure that estimates generalization capability) which predict the computational capability of a given circuit of spiking neurons.

The largest group of articles deals with reservoir analysis and optimization on a more abstract level, without a direct claim of biological modeling or with a particular application in the focus. *Jaeger et al* study ESNs made of leaky integrator neurons, present basic stability conditions, investigate parameter optimization by stochastic gradient descent, and demonstrate the usefulness of leaky integrator ESNs in test cases that require long time constants and insensitivity to time warped patterns. *Steil* uses the intrinsic plasticity of neurons to propose a new local, unsupervised adaptation rule for in-reservoir connections that improves the richness of its dynamics from an information point of view. Another approach to optimize reservoir dynamics is taken by *Xue, Yang and Haykin*, who implement lateral inhibition structures in a modular ESN to improve the richness of the reservoir dynamics. *Ozturk and Principe* propose a new readout for ESNs which use a special type of linear associative memory called MACE to achieve high specificity in dynamical pattern recognition applications. *Verstraeten et al* present a Lyapunov exponent based method to evaluate the richness of the reservoir, and compare reservoir computing methods (LSM, ESN and RNN with a backpropagation decorrelation rule) in a set of experiments using their reservoir computing toolbox.

The final group of articles gives a small sample of a variety of engineering and data analysis applications that are currently emerging. *Venayagamoorthy* applies ESNs to monitoring a multi-machine power system, demonstrating improved performance with much simpler training when compared with time delay neural networks. *Skowronski and Harris* bring ESNs to speech recognition and show improved performance with respect to Hidden Markov Models in low signal to noise ratio regimes. *Tong et al* use ESNs in language modeling to learn grammatical structure and show that their performance is similar to Elman networks, even though the ESN does not train the recurrent connections.

Reservoir computing is a very young field. This special issue gives witness of its fertility. We would like to conclude by pointing out two themes which we perceive as main directions for research in the next few years. First, not much is known at this early stage about the relationship between the dynamical characteristics of a given task (biological or engineering) and reservoir properties. Here we expect both a growth of analytical insight, and the development of practical methods to optimize a reservoir toward the task at hand, by design or by unsupervised adaptation. We furthermore predict a confluence of research on general properties of naturally occurring graphs (scale free networks etc), which have so far only been tested in their role as communication networks, research on dynamical properties (such as criticality) of circuits that employ various types of connectivity graphs, and research on computational properties of the resulting reservoirs. Secondly, once we know which connectivity structure and dynamical properties are optimal for the execution of particular types of tasks, research on the design of “optimal” computation systems for complex tasks, such as processing multiscale high-dimensional data (e.g. speech, video input streams, or robot sensor data), will move to the design and analysis of networks of interconnected reservoirs (e.g. hierarchies of reservoirs with different time constants). This second line of research is likely to also provide new insight into the computational role of individual brain areas within a recurrent network of interconnected brain areas that collaborate on the solution of complex cognitive and motor control tasks in the brain.

We would like to thank a large number of referees, whose comments have contributed to the improvement of the articles in this volume.

The Guest Editors

Herbert Jaeger
Wolfgang Maass
Jose Principe

REFERENCES:

[Hinton, 1981]

Hinton, G.E. Implementing semantic networks in parallel hardware. In: Hinton, G.E. and Anderson, J.A. (eds.), *Parallel Models of Associative Memory*, Erlbaum, Hillsdale, NJ, 1981, 161-187

[Rosenblatt, 1962]

Rosenblatt, F. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Spartan Books, 1962

[Buonomano and Merzenich, 1995]

Buonomano, D.V. and Merzenich, M.M. Temporal information transformed into a spatial code by a neural network with realistic properties. Science 267: 1028-1030, 1995