Perspectives and challenges for recurrent neural network training

Marco Gori, *University of Siena, Italy.*

Barbara Hammer, *Clausthal University of Technology, Germany.*

Pascal Hitzler, *Wright State University, Dayton, Ohio.*

Guenther Palm, *University of Ulm, Germany.*

Recurrent neural networks (RNNs) offer flexible machine learning tools which share the learning abilities of feedforward networks and which extend their expression abilities based on dynamical equations. Hence, they can directly process complex spatiotemporal data and model complex dynamic systems. Since temporal and spatial data are present in many domains such as processing environmental time series, modelling the financial market, speech and language processing, robotics, bioinformatics, medical informatics, etc., RNNs constitute promising candidates for a variety of applications. Further, their rich dynamic repertoire as time dependent systems makes them suitable candidates for modelling brain phenomena or mimicking large-scale distributed computations and argumentations. Thus, RNNs carry the promise of efficient biologically plausible signal processing models optimally suited for a wide area of industrial applications on the one hand and an explanation of cognitive phenomena of the human brain on the other hand.

Despite these facts, however, the design of efficient training methods for RNNs as well as their mathematical investigation with respect to reliable information representation and generalization ability when dealing with complex data structures is still a challenge. It has led to diverse approaches and architectures including specific training modes such as echo and liquid-state-machines, backpropagation decorrelation, or long short term memory, specific architectures such as recursive and graph networks, and hybrid systems at the borderline of symbolic and subsymbolic processing such as the core method. Interestingly, very heterogeneous domains are included and contributions to the area of RNNs stem from very different fields of research including, for example, logic, iterated function systems, and biological networks.

The aim of this special issue is to collect recent work developed in the field of recurrent information processing, which bridges the gap between different approaches and which sheds some light on canonical solutions or principled problems which occur in the context of recursive information processing when considered across the disciplines. This idea was born during a Dagstuhl seminar entitled ‘Recurrent Neural Networks- Models, Capacities, and Applications’ which took place in 2008 and which centered around the connection of RNNs to biological systems on the one side and logical models on the other side, gathering together experts in all three domains. This volume contains five papers which were accepted out of

Vol. 18 No. 5,
doi:10.1093/jigpal/jzp042 Advance Access published 26 August 2009

more than ten submissions to the special issue after a review process. These contributions cover a wide area of recent developments in the context of recurrent neural network training.

The first contribution is connected to the topic of deep learning. Daan Wiestra et al. consider the problem of reinforcement learning for partially observable Markov decision processes, a topic of high relevance to training systems in complex environments such as robotics or strategy optimization. Unlike classical models which are often based on policy iteration and the Bellman principle, a direct policy gradient is considered and the learning technique is adapted to recurrent neural networks which are capable of representing complex nonlinear state-based policies. Thereby, integration of the specially tailored long short term memory architecture allows to learn complex policies which include long term dependencies this way.

The following two contributions center around reservoir computing, one of the most promising paradigms of RNN training in the last years which borrows ideas from neuroscience to arrive at very efficient and powerful recurrent systems. Basically, a rich recurrent reservoir is combined with a simple trainable linear read-out, yielding surprisingly powerful recurrent processors. Both approaches deal with the important question of how the reservoir, which is simply random in the original proposal of echo state machines, can be adapted to the task at hand to yield a better accuracy and better suited network topologies. The work by Benjamin Roeschies and Christian Igel proposes specially tailored evolutionary algorithms to adapt the reservoir topology and connectivity to the given learning task. Thereby, they develop a technique for vectorial optimization of the two contradicting goals of high accuracy and small network topology, such that the user can pick the appropriate solution from the Pareto front. Ali Ajderi Rad et al. take a more direct way and propose a method to directly optimize the reservoir based on gradient techniques. Thereby, they avoid the problem of a usually huge number of free parameters by representing the connection matrices in terms of Kronecker products of small matrices which allow to control the sparseness and the weights in terms of only few directly adaptable parameters.

The last two contributions in this special issue take an even broader perspective of RNNs in the context of cognitive science by linking RNNs to biological modelling in neuroscience, on the one hand, and to explicit integrated neuro-symbolic systems, on the other hand. The approach of Claudius Gros and Gregor Kaczor investigates the relation of external sensory stimuli and self-sustained activity in autonomous recurrent networks as local models of brain activity. The latter are considered at an intermediate time level as quasi-stationary dynamical systems with transient states, and the role of external stimuli in the transient states is demonstrated in concrete experiments. These studies have the potential of concrete mathematical models for brain activities e.g. in the context of object perception. Patrick Simon and Thad Polk propose a complex model at the opposite end of cognitive science which relates recurrent networks to symbolic models. Based on leaky integrators as building blocks, they propose a concrete architecture which can encode both, symbolic phenomena (such as finite state machines and logical inference) as well as subsymbolic phenomena (such as noisy behavior or pattern recognition) by appropriate parameter settings. Both models have in common that they provide very interesting recurrent architectures and insights into their capability which can help to solve challenging problems, but only first steps towards their trainability are made based on these findings.

Altogether, the contributions in this special issue lay the foundation for powerful training tools of next generation RNN networks. These models carry the promise of relevant building blocks for the solution of future challenges such as deep learning in complex environments, modeling the human brain, or neuro-symbolic integration.