

# Multiobjective Optimization of Echo State Networks for multiple motor pattern learning

André Frank Krause\*, Volker Dürr†, Bettina Bläsing\* and Thomas Schack\*

\*Faculty of Sport Science, Dept. Neurocognition & Action email: andre\_frank.krause@uni-bielefeld.de

†Faculty of Biology, Dept. for Biological Cybernetics email: volker.duerr@uni-bielefeld.de

all authors: Cognitive Interaction Technology, Center of Excellence

University of Bielefeld, D-33615 Bielefeld, Germany

**Abstract**—Echo State Networks are a special class of recurrent neural networks, that are well-suited for attractor-based learning of motor patterns. Using structural multi-objective optimization, the trade-off between network size and accuracy can be identified. This allows to choose a feasible model capacity for a follow-up full-weight optimization. It is shown to produce small and efficient networks, that are capable of storing multiple motor patterns in a single net. Especially the smaller networks can interpolate between learned patterns using bifurcation inputs.

## I. INTRODUCTION

Neural networks are biological plausible models for pattern generation and learning. A straight-forward way to learn motor patterns is to store them in the dynamics of recurrent neuronal networks. For example, Tani [1] argued that this distributed storage of multiple patterns in a single network gives good generalisation compared to local, modular neural network schemes [2]. In [3] it was shown that it is not only possible to combine previously stored motor patterns to generate new ones, but also to establish an implicit functional hierarchy by using leaky integrator neurons with different time constants in a single network. This network can then generate and learn sequences by use of stored motor patterns and combine them to form new, complex behaviours. Tani [3] uses back-propagation through time (BPTT, [4]), that is computationally complex and rather biologically implausible. Echo State Networks (ESNs, [5]) are a special kind of recurrent neuronal networks. They are very easy and fast to train compared to classic, gradient-based training methods.

The general idea behind ESNs is to have a large, fixed, random reservoir of recurrently and sparsely connected neurons. Only a linear readout layer that taps this reservoir needs to be trained.

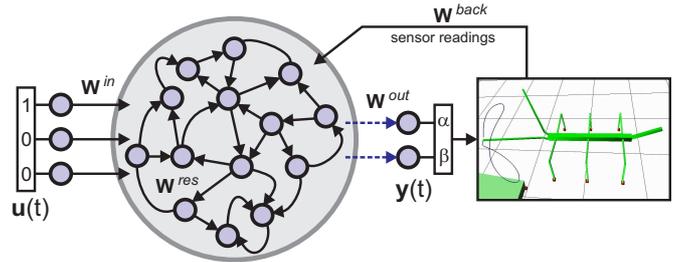


Fig. 1. General structure of an echo state network. Solid arrows indicate fixed, random connections, while dotted arrows are trainable readout connections. The output  $[\alpha, \beta]$  sets the joint angles of a bi-articular manipulator, e.g., a bio-inspired active tactile sensor. Joint angles are fed back via the back-projection weight matrix  $\mathbf{W}^{back}$ .

Typically, the structural parameters of ESNs, for example the reservoir size and connectivity, are chosen manually by experience and task demands. This may lead to suboptimal and unnecessary large reservoir structures for a given problem. Smaller ESNs may be more robust, show better generalisation, be faster to train and computationally more efficient. Here, multi-objective optimization is used to automatically find good network structures and explore the trade-off between network size and network error.

## II. ECHO STATE NETWORK

A basic, discrete-time ESN with a sigmoid activation function was implemented in Matlab©2009b. The purpose of this ESN was to control the joints of a bi-articular manipulator that could serve as a bio-inspired, active tactile sensor. The overall goal was to use the input to the ESN to set the tactile sampling pattern as desired. The state update equations used are:

$$\begin{aligned} \mathbf{y}(n) &= \mathbf{W}^{out} \mathbf{x}(n) \\ \mathbf{x}(n+1) &= \tanh(\mathbf{W}^{res} \mathbf{x}(n) + \mathbf{W}^{in} \mathbf{u}(n+1) + \\ &\quad \mathbf{W}^{back} \mathbf{y}(n) + \nu(n)) \end{aligned} \quad (1)$$

where  $\mathbf{u}$ ,  $\mathbf{x}$  and  $\mathbf{y}$  are the activation of the input, reservoir and output neurons, respectively.  $\nu(n)$  adds a small amount of uniformly distributed noise to the activation values of the reservoir neurons. This tends to stabilize solutions, especially in models that use output feedback for cyclic attractor learning [6].  $\mathbf{W}^{in}$ ,  $\mathbf{W}^{res}$ ,  $\mathbf{W}^{out}$  and  $\mathbf{W}^{back}$  are the input, reservoir, output and back-projection weight matrices. All matrices are sparse, randomly initialised, and stay fixed, except for  $\mathbf{W}^{out}$ . The weights of this linear output layer are learned using offline batch training. During training, the teacher data is forced into the network via the back-projection weights (teacher forcing), and internal reservoir activations are collected (state harvesting). After collecting internal states for all training data, the output weights are directly calculated using ridge regression. Ridge regression uses the Wiener-Hopf solution  $\mathbf{W}^{out} = \mathbf{R}^{-1}\mathbf{P}$  and adds a regularization term (Tikhonov regularization):

$$\mathbf{W}^{out} = (\mathbf{R} + \alpha^2\mathbf{I})^{-1}\mathbf{P} \quad (2)$$

where  $\alpha$  is a small number,  $\mathbf{I}$  is the identity matrix,  $\mathbf{R} = \mathbf{S}'\mathbf{S}$  is the correlation matrix of the reservoir states and  $\mathbf{P} = \mathbf{S}'\mathbf{D}$  is the cross-correlation matrix of the states and the desired outputs. Ridge regression leads to more stable solutions and smaller output weights, compared to ESN training using the Moore-Penrose pseudoinverse. A value of  $\alpha = 0.08$  was used for all simulations in this paper.

### III. MULTI-OBJECTIVE NETWORK STRUCTURE OPTIMIZATION

Multi-objective optimization (MO) is a tool to explore trade-offs between conflicting objectives. In the case of ESN optimization, the size of the reservoir versus the network performance is the main trade-off. In MO, the concept of dominance replaces the concept of a single optimal solution in traditional optimization. A solution dominates another, if strictly one objective value is superior and all other objectives are at least equal to the corresponding objective values of another solution. Following this definition, multiple (possibly infinite) non-dominated solutions can exist, instead of a single optimal solution. The set of non-dominated or pareto-optimal solutions is called the pareto front of the multi-objective problem. The goal of MO is to find a good approximation of the true (best) pareto front, but usually MO algorithms converge to a local (sub-optimal) pareto front due to complexity of the problem and computational constraints.

Usually, the structural parameters of an ESN are chosen manually by experience and task demands. Here, the full set of free network parameters was optimized using MO. The MO was performed with the function 'gamultiobj' from the *Matlab Genetic Algorithm and Direct Search (GADS) Toolbox*, that implements a variant of the 'Elitist Non-dominated Sorting Genetic Algorithm version II' (NSGA-II algorithm, [7]). The network structure was encoded into the genotype as a seven-dimensional vector of floating-point numbers. The first six structural parameters were the sparsity and weight range of the input-, reservoir- and back-projection weights. The seventh parameter was the number of reservoir neurons. The search range of the algorithm was constrained to  $[0, 1]$  for the sparsity values, to  $[-5, 5]$  for the weight values and to  $[1, 100]$  for the reservoir size ( $[1, 500]$  for the 4-pattern problem). The optimization was started with a population size of 1000 and converged after around 120 generations. In each iteration of the MO, all genomes were decoded into network structures, the networks were trained and then simulated with random initial activations for 1000 frames per pattern. In order to neglect the initial transient behaviour, the first 50 iterations of network output were rejected. The network output and the training patterns are usually not in-phase. The best match between training pattern and network output was searched by phase-shifting both output time courses by  $\pm 50$  frames relative to the training pattern and calculating the mean Manhattan distance across all pairs of data points. The training error was then defined as the smallest distance found in that range. The acceptable error threshold (fig.2) is expressed as the percentage of the amplitude of the training patterns, that is 1.0 units for all patterns. The pareto front for a circular pattern (fig.2a) reveals that even very small networks are capable of learning and generating two sine waves with identical frequency and  $90^\circ$  phase shift. The smallest network found had only 3 reservoir neurons. Including the two output neurons, the overall network size was 5. In comparison, 7 neurons are required for this task when using gradient-based learning methods [8]. Network size increases with the complexity of the motor pattern, and especially when having to store multiple patterns in a single network. Storing 4 patterns in a single network required 166 reservoir neurons to reach an error below 5% (fig.2d).

### IV. FULL OPTIMIZATION OF THE NETWORK WEIGHTS

From the pareto front of the two-pattern task, four candidate network structures were selected and opti-

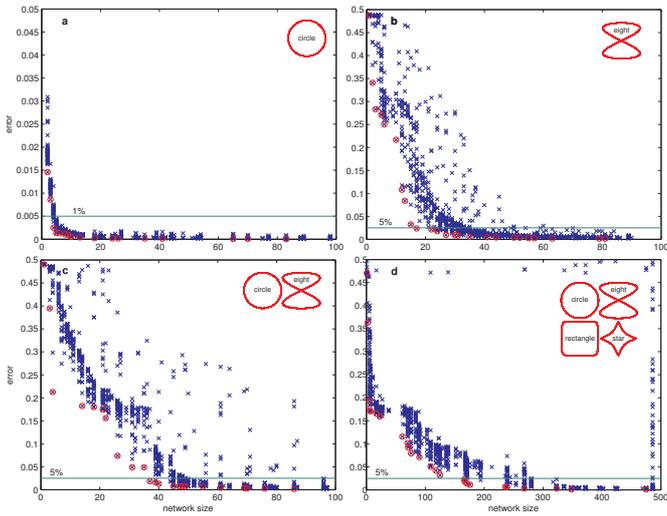


Fig. 2. Minimum reservoir size depends on task complexity. All panels show a set of pareto-optimal solutions (red circles) and the final population (blue crosses). (a) Learning a simple, circular pattern. All networks with 3 or more neurons show an error below 1%. (b) Pareto-front for the figure eight pattern. Learning this pattern requires a notably larger reservoir. Please note the different scaling of the error compared to the easier circle task. Networks with 17 or more neurons have an error below 5%. (c) Storing two motor patterns (circle and figure-eight) as cyclical attractors in a single network requires 37 or more reservoir neurons for errors below 5%. (d) Simultaneous learning of four patterns required 166 neurons.

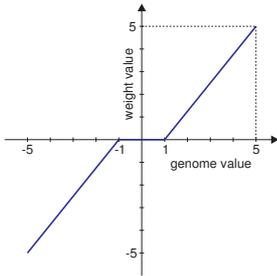


Fig. 3. Threshold function that decodes genome values into weight values, preserving sparse weight coding.

mized further, using a single-objective genetic algorithm. This time, all network weights except the output layer were fully optimized. The output layer was still trained by ridge regression. An initial random population of 200 parents was created from the network structure information of the selected candidate solutions with 4, 14, 26 and 37 reservoir neurons. Network weights were constrained to  $[-5, 5]$  and decoded from the genome with a threshold function that preserves sparsity. The threshold function sets a weight to zero, if the genome value is between -1 and 1, see fig.3.

The Genetic Algorithm (GA) options were set to ranked roulette wheel selection, 20 elitist solutions, 80% crossover probability with scattered crossover and self adaptive mutation. Other options were left at their default

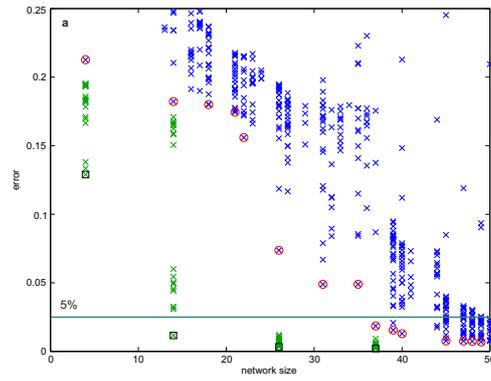


Fig. 4. Subsequent full-weight matrix optimization improves performance. Additional optimization of the four best networks of the two-pattern task with a reservoir size of 4, 14, 26 and 37 neurons. Starting from the best multi-objective solution, 20 GA runs were performed. a) Green crosses indicate the best fitness values of each run. Black squares indicate the overall best solutions that were found.

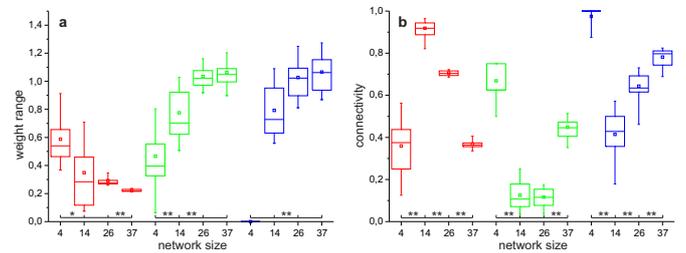


Fig. 5. Optimal weight range and connectivity depends on reservoir size. Network structure after full-weight optimization of the selected networks from fig.4. a) Weight range of all non-zero weights of the reservoir (red), the back-projection weights (green) and the input weights (blue). b) Connectivity (percentage of non-zero weights). Boxplots show 5%, 25%, 50%, 75% and 95% quantiles of  $N=20$  datapoints. \*  $p < 0.05$ ; \*\*  $p < 0.01$

values (see GADS toolbox, Matlab2009b). The GA-optimization was repeated 20 times for each network size. Fig. 4a shows the improvement in performance compared to the MO structure optimization run. A small network with only 14 reservoir neurons could reproduce the learned patterns with an error of 2.3%. Weight range and connectivity after optimization was analysed with an unpaired Wilcoxon rank sum test. Significant differences in connectivity and weight range were found (fig. 4b) with a clear trend for smaller reservoir weights and less reservoir connectivity with increasing network sizes. Both input- and back-projection weights tend to increase with reservoir size (fig. 4a). Although standard ESNs usually have full connectivity for input- and back-projection weights, evolutionary optimization seems to favor sparse connectivity for smaller networks, when given the choice (fig. 4b).

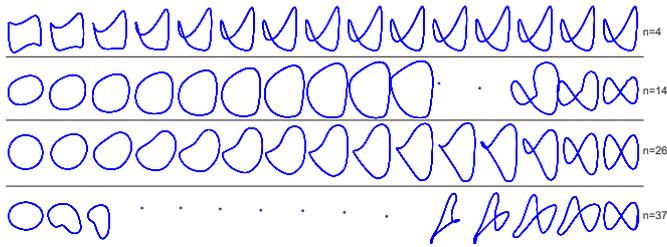


Fig. 6. Dynamic behaviour of selected networks with different reservoir sizes (blue trajectories). Shifting the dynamics of the networks by gradually changing the two input values from 1.0 to 0.0 and from 0.0 to 1.0, respectively, in 15 steps. Changing the input to the network causes a slow morphing between the two learned patterns, allowing the generation of new patterns that were not explicitly trained. Especially the small networks keep stable with no chaotic regions. Larger networks tend to converge to fixed points for input values other than zero or one (blue dots).

## V. DYNAMIC NETWORK BEHAVIOUR

Most of the smaller networks show an unexpected behaviour. They are able to interpolate between the learned patterns, generating novel, not explicitly trained outputs. Fig. 6 shows the dynamical responses from the fittest networks of section IV. The first input value was changed gradually in 15 steps from 1.0 to 0.0, while the second input was changed from 0.0 to 1.0. A gradual morphing from the circular to the figure-eight pattern can be observed. It is surprising, that already a small ESN with six reservoir neurons can store two different patterns. Larger networks tend to converge to fixed points for input values other than the trained ones. This interpolation effect might be applied to complex and smooth behaviour generation for neural network controlled robots.

## VI. CONCLUSIONS

Using MO, good candidate network structures can be selected as starting points for a followup whole-network optimization and fine-tuning using genetic algorithms. The resulting pareto front helps to identify small and sufficiently efficient networks that are able to store multiple motor patterns in a single network. This distributed storage of motor behaviours as attractor states in a single net is in contrast to earlier, local module based approaches. *"If sequences contain similarities and overlap, however, a conflict arises in such earlier models between generalization and segmentation, induced by this separated modular structure."* [3]. By choosing a feasible model capacity, over-fitting and the risk of unwanted - possibly chaotic - attractor states is reduced. Also, with the right choice of the network

size, an interesting pattern interpolation effect can be evoked. Instead of using a classic genetic algorithm for fine-tuning of the network weights, new, very fast and powerful black box optimisation algorithms [9] [10] could further increase network performance and allow to find even smaller networks for better generalisation. ESNs can be used for direct control tasks ( see [11]) and scale well with a high number of training patterns and motor outputs [12]. A more complex simulation, for example of a humanoid robot, will show if direct, attractor-based storage of parameterized motor patterns is flexible enough for complex behaviour generation.

## REFERENCES

- [1] J. Tani, M. Itob, and Y. Sugitaa, "Self-organization of distributedly represented multiple behavior schemata in a mirror system: reviews of robot experiments using rnnpb," *Neural Networks*, vol. 17, pp. 1273 – 1289, 2004.
- [2] M. Haruno, D. M. Wolpert, and M. Kawato, "Mosaic model for sensorimotor learning and control," *Neural Computation*, vol. 13(10), pp. 2201–2220, 2001.
- [3] Y. Yamashita and J. Tani, "Emergence of functional hierarchy in a multiple timescale neural network model: A humanoid robot experiment," *PLoS Computational Biology*, vol. 4 (11), pp. –, 2008.
- [4] P. Werbos, "Backpropagation through time: what it does and how to do it," in *Proceedings of the IEEE*, vol. 78(10), 1990, pp. 1550–1560.
- [5] H. Jäger and H. Haas, "Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication," *Science*, vol. 304, pp. 78 – 80, 2004.
- [6] H. Jäger, M. Lukosevicius, D. Popovici, and U. Siewert, "Optimization and applications of echo state networks with leaky integrator neurons," *Neural Networks*, vol. 20(3), pp. 335–352, 2007.
- [7] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *IEEE Transactions on Evolutionary Computation*, vol. 6, No. 2, pp. 182–197, 2002.
- [8] B. A. Pearlmutter, "Learning state space trajectories in recurrent neural networks," *Neural Computation*, vol. 1, pp. 263–269, 1989.
- [9] O. Kramer, "Fast blackbox optimization: Iterated local search and the strategy of powell." in *The 2009 International Conference on Genetic and Evolutionary Methods (GEM'09)*, 2009, in press.
- [10] J. A. Vrugt, B. A. Robinson, and J. M. Hyman, "Self-adaptive multimethod search for global optimization in real-parameter spaces." *Evolutionary Computation, IEEE Transactions on*, vol. 13(2), pp. 243–259, 2008.
- [11] A. F. Krause, B. Bläsing, V. Dürr, and T. Schack, "Direct control of an active tactile sensor using echo state networks," in *Human Centered Robot Systems. Cognition, Interaction, Technology*, ser. Cognitive Systems Monographs, H. Ritter, G. Sagerer, R. Dillmann, and M. Buss, Eds., vol. 6. Berlin Heidelberg: Springer-Verlag, 2009, pp. 11–21.
- [12] H. Jäger, "Generating exponentially many periodic attractors with linearly growing echo state networks," IUB, technical report 3, 2006.