

Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless telecommunication

Herbert Jaeger*, Harald Haas, International University Bremen, Bremen, Germany.

*to whom correspondence should be addressed, email: h.jaeger@iu-bremen.de

Abstract. *We present a method for learning nonlinear systems, echo state networks (ESNs). ESNs employ artificial recurrent neural networks in a way that has recently been proposed independently as a learning mechanism in biological brains. The learning method is computationally efficient and easy to use. On a benchmark task of predicting a chaotic time series, accuracy is improved by a factor of 2400 over previous techniques. The potential for engineering applications is illustrated by equalizing a communication channel, where the signal error rate is improved by two orders of magnitude.*

Nonlinear dynamical systems abound in the sciences and in engineering. If one wishes to simulate, predict, filter, classify or control such a system, one needs an executable system model. However, it is often infeasible to obtain analytical models. In such cases one has to resort to blackbox models, which ignore the internal physical mechanisms and instead only reproduce the outwardly observable input-output behavior of the target system.

If the target system is linear, efficient methods for blackbox modeling are available. Most technical systems however become nonlinear if operated at higher operational points (that is, closer to saturation). Although this might lead to cheaper and more energy-efficient designs, it is not done because the resulting nonlinearities cannot be harnessed. Many biomechanical systems use their full dynamic range (up to saturation) and thereby become lightweight and energy-efficient – and thoroughly nonlinear.

Here we present an approach to learning blackbox models of nonlinear systems, echo state networks (ESNs). An ESN is an artificial recurrent neural network (RNN). RNNs are characterized by feedback ("recurrent") loops in their synaptic connection pathways. They can maintain an ongoing activation even in the absence of input and thus exhibit dynamic memory. Biological neural networks are typically recurrent. Like biological neural networks, an artificial RNN can learn to mimick a target system – in principle with arbitrary accuracy (1). Several learning algorithms are known (2–4) that incrementally adapt the synaptic weights of an RNN in order to tune it toward the target system. These algorithms have not been widely employed in technical applications because of slow convergence and suboptimal solutions (5,6). The ESN approach differs from these methods in that a large RNN is used (order of 50 to 1000 neurons, previous techniques typically use 5 to 30 neurons) and in that only the synaptic connections from the RNN to the output readout neurons are modified by learning (previous techniques tune all synaptic connections, Fig. 1). Because there are no cyclic dependencies between the trained readout connections, training an ESN becomes a simple linear regression task.




Fig. 1. (A) Schema of previous approaches to RNN learning. **(B)** Schema of ESN approach. Solid bold arrows: fixed synaptic connections, dotted arrows: adjustable connections. Both approaches aim at minimizing the error $d(n) - y(n)$, where $y(n)$ is the network output and $d(n)$ is the "teacher" time series observed from the target system.

We illustrate the ESN approach on a task of chaotic time series prediction (7). The Mackey-Glass system (8) is a standard benchmark system for time series prediction studies. It generates a subtly irregular time series (Fig. 1A). The prediction task has two steps: (i) use an initial teacher sequence generated by the original MGS to learn a blackbox model M of the generating system, (ii) use M to predict the value of the sequence some steps ahead.




Fig. 2. (A) Prediction output of the trained ESN (dotted) overlaid with the correct continuation (solid). (B) Learning the MG attractor. Three sample activation traces of internal neurons are shown. They "echo" the teacher signal $d(n)$. After training, the desired output is re-created from the echo signals through output connections (dotted arrows) whose weights w_i are the result of the training procedure.

First we created a random RNN with 1000 neurons (called the "reservoir") and one output neuron. Importantly, the output neuron was equipped with random connections that project back into the reservoir (Fig. 2B). A 3000 step teacher sequence $d(1), \dots, d(3000)$ was generated from the MGS equation and fed into the output neuron. This excited the internal neurons through the output feedback connections. After an initial transient, they started to exhibit systematic individual variations of the teacher sequence (Fig. 2B).

The fact that the internal neurons display systematic variants of the exciting external signal is constitutional for ESNs: the internal neurons must work as "echo functions" for the driving signal. Not every randomly generated RNN has this property, but it can effectively be built into a reservoir (supporting online text).

It is important that the "echo" signals be richly varied. This was ensured by a sparse interconnectivity of 1% within the reservoir: this condition lets the reservoir decompose into many loosely coupled subsystems, establishing a richly structured reservoir of excitable dynamics.

After time $n = 3000$, output connection weights w_i ($i = 1, \dots, 1000$) were computed (dashed arrows in Fig. 1B) from the last 2000 steps $n = 1001, \dots, 3000$ of the training run such that the training error $\text{MSE}_{\text{train}} = 1/2000 \sum_{n=1001}^{3000} \left(d(n) - \sum_{i=1}^{1000} w_i x_i(n) \right)^2$ was minimized [$x_i(n)$: activation of the i -th internal neuron at time n]. This is a simple linear regression.

With the new w_i in place, the ESN was disconnected from the teacher after step 3000 and left running freely. A bidirectional dynamical interplay of the network-generated output signal with the internal signals $x_i(n)$ unfolded. The output signal $y(n)$ was created from the internal neuron activation signals $x_i(n)$ through the trained connections w_i , by $y(n) = \sum_{i=1}^{1000} w_i x_i(n)$. Conversely, the internal signals were "echoed" from that output signal through the fixed output feedback connections (supporting online text).

For testing, a 84-step continuation $d(3001), \dots, d(3084)$ of the original signal was computed for reference. The network output $y(3084)$ was compared with the correct continuation $d(3084)$. Averaged over 100 independent trials, an normalized root mean square error

$$\text{NRMSE} = \left(\sum_{j=1}^{100} (d_j(3084) - y_j(3084))^2 / 100 \sigma^2 \right)^{1/2} \approx 10^{-4.2}$$

was obtained [d_j and y_j teacher

and network output in trial j , σ^2 variance of MGS signal], improving the best previous techniques (9–15), which used training sequences of length 500 to 10000, by a factor of 700. If the prediction run was continued, deviations typically became visible after about 1300 steps (Fig. 2A). With a refined variant of the learning method (7) the improvement factor rises to 2400. Models of similar accuracy were also obtained for other chaotic systems (supporting online text).

The main reason for the jump in modelling accuracy is that ESNs capitalize on a massive short-term memory. We showed analytically (16) that under certain conditions an ESN of size N may be able to "remember" in the order of the last N inputs. This information is more massive than the information used in other techniques (supporting online text).

We now illustrate the approach in a task of practical relevance, namely, the equalization of a wireless communication channel (7). The essentials of equalization are as follows. A sender wants to communicate a symbol sequence $s(n)$. This sequence is first transformed into an analog envelope signal $d(n)$, then modulated on a high-frequency carrier signal and transmitted, then received and demodulated into an analog signal $u(n)$ which is a corrupted version of $d(n)$. Major sources of corruption are (i) noise (thermal or due to interfering signals), (ii) multipath propagation, which leads to a superposition of adjacent symbols (intersymbol interference), and nonlinear distortion induced by operating the sender's power amplifier in the high gain region. In order to avoid the latter, the actual power amplification is run well below the maximum amplification possible, thereby incurring a significant loss in energy efficiency which is clearly undesirable in cellphone and satellite communications. The corrupted signal $u(n)$ is then passed through an equalizing filter whose output $y(n)$ should restore $u(n)$ as closely as possible to $d(n)$. Finally, the equalized signal $y(n)$ is converted back into a symbol sequence. The quality measure for the entire process is the fraction of incorrect symbols finally obtained (symbol error rate).

In order to compare the performance of an ESN equalizer with standard techniques, we took a channel model for a nonlinear wireless transmission system from a study (17) that compared three customary nonlinear equalization methods, namely a linear decision feedback equalizer (DFE), which is actually a nonlinear method; a Volterra DFE; and a bilinear DFE. The model equation featured intersymbol interference across 10 consecutive symbols, a second-order and a third-order nonlinear distortion, and additive white Gaussian noise. All methods investigated in that study had 47 adjustable parameters and used sequences of 5000 symbols for training. In order to make the ESN equalizer comparable with the equalizers studied in (17), we took ESNs with a reservoir of 46 neurons (which is small for the ESN approach), which yielded 47

adjustable parameters (the 47th comes from a direct connection from the input to the output neuron).

We carried out numerous learning trials (7) to obtain ESN equalizers, using an online learning method (a version of the recursive least square algorithm known from linear adaptive filters) to train the output weights on 5000 step training sequences. We chose an online adaptation scheme (17) here because the methods in (17) were online adaptive, too, and because wireless communication channels mostly are time-varying, such that an equalizer must adapt to changing system characteristics. The entire learning-testing procedure was repeated for signal-to-noise ratios ranging from 12 to 32 db. Fig. 3 compares the average symbol error rates obtained to the results reported in (17), showing an improvement of two magnitudes for high signal-to-noise ratios.




Fig. 3. Using an ESN for nonlinear channel equalization: results. Plot shows signal error rate SER vs. signal-to-noise ratio SNR. **a.** linear DFE, **b.** Volterra DFE, **c.** bilinear DFE (**a – c** taken from (18)). **d.** Blue line represents average ESN performance with randomly generated "reservoirs" (error bars: variation across networks). **e.** Green line indicates performance of best network chosen from the networks averaged in **d** (error bars: variation across learning trials).

For tasks with multi-channel input and/or output, the ESN approach can be accommodated simply by adding more input or output neurons (19,20).

ESNs can be taken to all basic tasks of signal processing and control, including time series prediction, inverse modeling, pattern generation, event detection and classification, modeling distributions of stochastic processes, filtering and nonlinear control (16,18,19,20). Because a single learning run takes only a few seconds (or minutes for very large datasets and networks), engineers can test out variants at a high turnover rate, a crucial factor for practical usability.

ESNs have been developed from a mathematical and engineering perspective, but exhibit typical features of biological RNNs: a large number of neurons, recurrent pathways, sparse random connectivity and local modification of synaptic weights. The idea of using randomly connected RNNs to represent and memorize dynamic input in network states has frequently been explored in specific contexts, for instance artificial intelligence models of associative memory (21),

models of prefrontal cortex function in sensory-motor sequencing tasks (22), models of birdsong (23), models of the cerebellum (24) or general computational models of neural oscillators (25). Many different learning mechanisms were considered, mostly within the RNN itself. The contribution of the ESN is to have elucidated the mathematical properties of large RNNs such that they can be used with a linear, trainable readout mechanism for general blackbox modeling. An approach essentially equivalent to ESNs, liquid state networks (26,27), has been developed independently to model computations in cortical microcircuits. Recent findings in neurophysiology suggest that the basic ESN/liquid state network principle seems not uncommon in biological networks (28-30) and could eventually be exploited to control prosthetic devices by signals collected from a collective of neurons (31).

References and notes

1. K.-I. Funahashi, Y. Nakamura, *Neural Networks* **6**, 801 (1993).
2. D. Zipser, R. J. Williams, *Neural Comput.* **1**, 270 (1989).
3. P. J. Werbos, *Proc. IEEE* **78**(10), 1550 (1990).
4. L.A. Feldkamp, D. V. Prokhorov, C. F. Eagen, F. Yuan, in *Nonlinear Modeling: Advanced Black-Box Techniques*, J. A. K. Suykens, J. Vandewalle, J., Eds. (Kluwer 1998) pp. 29-54.
5. K. Doya, in *The Handbook of Brain Theory and Neural Networks*, M.A. Arbib, Ed. (MIT Press 1995) pp. 796-800.
6. H. Jaeger, "Tutorial on training recurrent neural networks" (GMD-Report 159, German National Research Institute for Computer Science, 2002).
ftp://borneo.gmd.de/pub/indy/publications_herbert/CompleteTutorialTechrep.pdf
7. Methods are available as supporting material on *Science Online*.
8. M.C. Mackey, L. Glass, *Science* **197**, 287(1977).
9. J. Vesanto, in *Proc. WSOM'97* (1997).
<http://www.cis.hut.fi/projects/monitor/publications/papers/wsom97.ps.zip>
10. L. Chudy, I. Farkas, *Neural Network World* **8**, 481 (1998).
11. H. Bersini, M. Birattari, G. Bontempi, in *Proc. IEEE World Congr. on Computational Intelligence (IJCNN '98)* pp. 2102-2106 (1997).
ftp://iridia.ulb.ac.be/pub/lazy/papers/IridiaTr1997-13_2.ps.gz
12. T. M. Martinetz, S. G. Berkovich, K. J. Schulten, *IEEE Trans. Neural Networks* **4**, 558 (1993).
13. Yao, X. & Liu. Y. *IEEE Trans. Neural Networks* **8**, 694 (1997).
14. F. Gers, D. Eck, J. F. Schmidhuber, "Applying LSTM to time series predictable through time-window approaches" (IDSIA-IDSIA-22-00, 2000).
<http://www.idsia.ch/~felix/Publications.html>
15. J. McNames, J. A. K. Suykens, J. Vandewalle, *Int. J. Bifurcation Chaos* **9**, 1485 (1999).

16. H. Jaeger, "The echo state approach to analysing and training recurrent neural networks" (GMD-Report 148, German National Research Institute for Computer Science 2001).
ftp://borneo.gmd.de/pub/indy/publications_herbert/EchoStatesTechRep.pdf
17. V. J. Mathews, J. Lee, in *Advanced Signal Processing: Algorithms, Architectures, and Implementations V (Proc. SPIE Vol. 2296)*, (SPIE, San Diego, CA, 1994), pp. 317-327.
18. H. Jaeger, in *Advances in Neural Information Processing Systems 15*, S. Becker, S. Thrun, K. Obermayer, Eds. (MIT Press, Cambridge, MA, 2003) pp. 593-600
19. H. Jaeger, "Short term memory in echo state networks" (GMD-Report 152, German National Research Institute for Computer Science 2002).
ftp://borneo.gmd.de/pub/indy/publications_herbert/STMEchoStatesTechRep.pdf
20. J. Hertzberg, H. Jaeger, F. Schönherr, in *Proc. 15th Europ. Conf. on Art. Int. (ECAI 02)*, F. van Harmelen, Ed., (IOS Press, Amsterdam, 2002), pp. 708-712.
<http://www.ais.fhg.de/~schoenhe/papers/ECAI02.pdf>
21. G. E. Hinton, in *Parallel Models of Associative Memory*, G. E. Hinton, J. A. Anderson, Eds. (Erlbaum, Hillsdale, NJ, 1981) pp.161-187.
22. D. G. Beiser, J. C. Houk, *J Neurophysiol* **79**: 3168 (1998).
23. S. Dehaene, J.-P. Changeux, J.-P. Nadal, *Proc. Natl. Acad. Sci. USA* **84**, 2727 (1987).
24. M. Kawato, in *The Handbook of Brain Theory and Neural Networks*, M. Arbib, Ed. (MIT Press, 1995), pp. 172-178.
25. K. Doya, S. Yoshizawa, *Neural Networks* **2**, 375 (1989).
26. W. Maass, T. Natschläger, H. Markram, *Neural Computat.* **14**, 2531 (2002).
27. W. Maass, T. Natschläger, H. Markram, in *Computational Neuroscience: A Comprehensive Approach*, J. Feng, Ed. (Chapman & Hall/CRC, 2003), 575-605. [Article](#)
28. G. B. Stanley, F. F. Li, Y. Dan, *J. Neurosci.* **19**, 8036 (1999). [Article](#)
29. G. B. Stanley, *Neurocomputing* **38-40**, 1703 (2001).
30. W. M. Kistler, Ch. I. de Zeeuw, *Neural Comput.* **14**, 2597 (2002). [Article](#)
31. S. Mussa-Ivaldi, *Nature* **408** (6810) 361 (2000).
32. The first author thanks Thomas Christaller for unfaltering support and Wolfgang Maass for friendly co-operation. International patents are claimed by Fraunhofer AIS (PCT/EP01/11490).

Supporting Online Material

Methods

Network equations

The ESNs described in the article are standard discrete-time sigmoid networks. The state update equation is $\mathbf{x}(n+1) = \tanh(\mathbf{W}\mathbf{x}(n) + \mathbf{w}_{in}u(n+1) + \mathbf{w}_{fb}y(n) + v(n))$ [\mathbf{W} : $N \times N$ matrix of internal connection weights; \mathbf{w}_{in} : N -size vector of input connection weights; \mathbf{w}_{fb} (optional): weight vector for feedback connections from the output neuron to the reservoir; $v(n)$ (optional): noise vector; the \tanh is applied elementwise]. The output equation for a single-output network is $y(n) = \tanh(\mathbf{w}_{out}(\mathbf{x}(n), u(n)))$ [\mathbf{w}_{out} : $(N+1)$ -size vector of weights of connections to the output neuron; \mathbf{w}_{out} contains the w_i used in the article text]. The output nonlinearity \tanh is optional. In the Mackey-Glass experiment it was present, whereas in the channel equalization study we used linear output units. In order to make a random reservoir qualify as an echo state network, it must exhibit certain damping properties. They can be ascertained by ensuring that the weight matrix \mathbf{W} has a spectral radius smaller than unity (SI).

The Mackey-Glass system

Preparation of training and testing data. The MG delay differential equation $dx/dt = 0.2x(t-\tau)/(1+x(t-\tau)^{10} - 0.1x(t))$ [τ : delay; here we used $\tau = 17$, the value standardly employed in most of the MGS prediction literature] was simulated using the `dde23` solver for delay differential equations from the commercial toolbox Matlab. This solver allows one to specify the absolute accuracy; it was set to $1e-16$. A stepsize of 1.0 was used (which is finer than the stepsize automatically selected by `dde23`). The resulting time series were shifted by -1 and passed through a \tanh function, whereby they fell into a range of $(-0.5, 0.3)$. From all data series thus generated, the first 1000 steps were discarded to get rid of initial washouts.

Network setup. To construct the ESN, a 1000×1000 reservoir weight matrix \mathbf{W} was generated with 1% connectivity and random weights drawn from a uniform distribution over $(-1, 1)$, then rescaled to spectral radius 0.8. Output feedback connection weights were randomly selected from a uniform distribution over $(-1, 1)$. An auxiliary input unit was attached with similar random connections to feed in a constant bias input of size 0.2.

Training and testing. Length of training sequence was 3000, first 1000 steps were discarded to wash out initial transient, echo signals $\mathbf{x}(n)$ were sampled from remaining 2000 steps. Noise $v(n)$ of size $1.0e-10$ was added on network units during sampling, a trick to increase stability of the trained network, see comments in (1). The error to be minimized was

$$\text{MSE}_{train} = 1/2000 \sum_{n=1001}^{3000} (\tanh^{-1} d(n) - \mathbf{w}_{out}\mathbf{x}(n))^2$$
. (The \tanh^{-1} was left out in the formula given in the article because it would have required an explanation of the network equations). The weight vector \mathbf{w}_{out} that minimizes this MSE is made from the regression weights of the linear

regression of the $\tanh^{-1} d(n)$ values on the $\mathbf{x}(n)$ values. Any method for computing linear regressions can be used to obtain \mathbf{w}_{out} . We employed the pseudoinverse routine `pinv` from Matlab. The obtained training error was $MSE_{train} \approx 1.2E-15$. For testing, 100 independent instantiations of the MGS time series were teacher-forced on the trained network for 2000 steps, then the network was left running freely and the network's continuation was compared to the true continuation after 84 steps to obtain an average normalized root mean square error for the 84 step prediction of $NRMSE_{84} = \left(\sum_{i=1}^{100} (d(+84) - y(+84))^2 / 100 \sigma^2 \right)^{1/2} \approx 0.000025$ [$d(+84)$ correct continuation, $y(+84)$ network prediction, σ^2 variance of MG signal]. Fig. 2A in the article shows an overlay of a 2500 step free running continuation of the trained network with the continuation of the original MGS series.

In order to test whether the trained network indeed generates a chaotic time series (as opposed to a long limit cycle behavior that might superficially look like a chaotic sequence), the largest Lyapunov exponent of the network-generated output signal was empirically estimated. To this end, the trained network was first run for 1500 time steps in teacher-forced mode, after which it was in state $\mathbf{x}(1500)$. From this state, it was left running freely for 317 steps, and the resulting output sequence $y(1501), \dots, y(1817)$ was recorded. In addition, the state $\mathbf{x}(1500)$ was perturbed by a uniform noise vector of size $1e-7$, and the network was left running freely starting now from the perturbed state for 317 steps, and the resulting output sequence $y'(1501), \dots, y'(1817)$ was recorded. The exponential divergence rate λ between these two sequences was estimated by computing $d_1 = \| [y(1601) \dots y(1617)] - [y'(1601) \dots y'(1617)] \|$, $d_2 = \| [y(1801) \dots y(1817)] - [y'(1801) \dots y'(1817)] \|$, $\lambda = \log(d_2/d_1) / 200$. The subsequence length 17 that goes into the computation of d_1 and d_2 was chosen because it corresponds to roughly one full "loop" of the attractor. Averaging over 100 such trials gave an estimated Lyapunov exponent of 0.0058. An analog empirical estimate of the Lyapunov exponent of the time series obtained from the `dde23` solver yielded a value of 0.0059. Both values agree with estimates for the MG Lyapunov exponent found in the literature, which vary (considerably) around 0.006. Therefore, the network indeed generated a chaotic time series with a degree of chaoticity similar to the original system.

The ESN equalizer

Data preparation. We took the channel model from (S2). The channel is modeled by a linear system with memory length 10 followed by a memoryless noisy nonlinearity. The input to the channel is an i.i.d. random sequence $d(n)$ with values from $\{-3, -1, 1, 3\}$. The input-output equation of the linear channel is given by $q(n) = 0.08 d(n+2) - 0.12 d(n+1) + d(n) + 0.18 d(n-1) - 0.1 d(n-2) + 0.09 d(n-3) - 0.05 d(n-4) + 0.04 d(n-5) + 0.03 d(n-6) + 0.01 d(n-7)$. The noisy nonlinearity is given by $u(n) = q(n) + 0.036 q(n)^2 - 0.011 q(n)^3 + v(n)$, where $v(n)$ is an i.i.d. Gaussian noise with zero mean adjusted in power to yield signal-to-noise ratios ranging from 16 to 32 db.

Network setup. 46-neuron ESN reservoir network weight matrices \mathbf{W} were randomly generated with a connectivity of 20%, and nonzero connection weights drawn from a uniform distribution over $(-1, 1)$. The resulting weight matrix was rescaled to a spectral radius of 0.5. An input neuron was attached with connection weights drawn from a uniform distribution over $(-0.025, 0.025)$.

The output neuron was a linear neuron here, so the output equation of the network is $y(n) = \mathbf{w}_{out}(\mathbf{x}(n), u(n))$.

Input and teacher signals given to the network. The input given to the ESN equalizer was not the raw corrupted signal $u(n)$ but a shifted version $u(n) + 30$. The teacher output signal was $d(n - 2)$, so the learning task consisted in minimizing the square error $(y(n) - d(n - 2))^2$. Note that this implies a delay in the equalizer response of 4 time steps, because the signal $u(n)$ depends on d values up to time $n + 2$. The equalizers investigated in (S2) had response delays ranging from 8 to 21.

The online learning algorithm. We used the following plain textbook version of the RLS algorithm (S3):

Initialization: Create a diagonal auxiliary matrix $\Psi^{-1}(0)$ of size 47 by 47, with large diagonal entries of size 10^{10} . Create an initial all-zero output weight vector $\mathbf{w}_{out}(0)$ of size 47 and an initial all-zero network state $\mathbf{x}(0)$. Define a forgetting rate λ (we used $\lambda = 0.998$).

Input into update cycle n : Previous output weight estimates $\mathbf{w}_{out}(n - 1)$ of size 47; network state + input vector $(\mathbf{x}(n), u(n) + 30) =: \mathbf{v}(n)$ of size 47; one-dimensional teacher output $d(n - 2) =: d^+(n)$; the auxiliary matrix $\Psi^{-1}(n - 1)$.

Output after update cycle n : Network output $y(n)$, update $\mathbf{w}_{out}(n)$ of output weight vector, update of auxiliary matrix $\Psi^{-1}(n)$.

Computation steps within one update cycle:

1. $\mathbf{u}(n) = \Psi^{-1}(n - 1) \mathbf{v}(n)$ [comment: this \mathbf{u} is not related to the input u].
2. $\mathbf{k}(n) = \frac{1}{\lambda + \mathbf{v}(n)^T \mathbf{u}(n)} \mathbf{u}(n)$ [comment: T indicates transpose]
3. $y(n) = \mathbf{w}_{out}(n - 1)^T \mathbf{v}(n)$
4. $e(n) = d^+(n) - y(n)$
5. $\mathbf{w}_{out}(n) = \mathbf{w}_{out}(n - 1) + \mathbf{k}(n) e(n)$
6. $\Psi^{-1}(n) = \lambda^{-1} (\Psi^{-1}(n - 1) - \mathbf{k}(n) [\mathbf{v}(n)^T \Psi^{-1}(n - 1)])$

During the first 100 steps of the 5000-step training run, the RLS update was disabled to allow the reservoir network to wash out initial transients resulting from the arbitrary starting state.

Learning and testing, suite 1 (curve **d** in Fig. 3). We carried out 20 independent learning trials for each signal-to-noise ratio (SNR) 12, 16, 20, 24, 28, 32. A single learning trial consisted of creating a random 46-neuron reservoir and random input connection weights as indicated above, setting the noise $v(n)$ to yield a prescribed signal-to-noise ratio, running the RLS algorithm on the random network for 5000 update cycles, freezing the weights and letting the trained equalizer run on a test sequence. The equalized signal $y(n)$ from the test run was converted into a 4-symbol sequence by equidistant thresholding (symbol "1" was chosen if $y(n) < 2$, "2" if $2 \leq y(n) < 0$, etc.).

The test run was stopped after the 10th symbol error was encountered or after 10^7 steps, whichever occurred first. The symbol error rate (SER) was estimated by the quotient *number of incorrectly re-converted symbols / length of test run*.

Learning and testing, suite 2 (curve **e** in Fig. 3). The reservoir that gave the best performance in the SNR = 32 condition in suite 1 was used in all trials of a second suite of learning trials. The variance found for each SNR in curve **e** from Fig. 3 results from differences in the training and test sequences.

Supporting Material Text

Refined version of the learning method

When the ESN method is used with feedback from the output neuron (as in the Mackey-Glass example), the reservoir network is "forced" by the correct teacher output during the presentation of the training data. When the trained network is later exploited for generating output autonomously, it is receiving its own output signal through the feedback connections instead of the correct signal. Because the self-generated output slightly deviates from the correct one, the dynamics of the reservoir state will be slightly different in training vs. autonomous signal generation. If during training the output neurons could be trained using a reservoir with a dynamics closer to the one encountered at exploitation time, modelling accuracy could be improved. This idea leads to an improved, three-stage version of the training method, which will now be described for the Mackey-Glass prediction task:

1. First stage: an ESN is trained on the Mackey-Glass prediction task exactly as in the basic method, with a teacher time series $d(n)$ obtained from a numerical simulation of the MG equation. This results in a preliminary set of output weights \mathbf{w}_{out}^0 .
2. Second stage: A new teacher time series $d'(n)$ is computed from $d(n)$ and the weights \mathbf{w}_{out}^0 , as follows. Put $d'(1) = d(1)$. In order to compute $d'(n)$ for $n > 1$, start network in random state $\mathbf{x}(1)$. To obtain $d'(2)$, write $d(1)$ into the output neuron and update the network once according to $\mathbf{x}(2) = \tanh(\mathbf{W}\mathbf{x}(1) + \mathbf{w}_{fb}d(1))$ and $y(2) = \tanh(\mathbf{w}_{out}^0(\mathbf{x}(2)))$. Put $d'(2) = y(2)$. Iterate: to obtain $d'(n)$, write $d(n-1)$ into the output neuron, update according to $\mathbf{x}(n) = \tanh(\mathbf{W}\mathbf{x}(n-1) + \mathbf{w}_{fb}d(n-1))$ and $y(n) = \tanh(\mathbf{w}_{out}^0(\mathbf{x}(n)))$, put $d'(n) = y(n)$. Thus, $d'(n)$ is the sequence of one-step predictions of the network trained in the first stage (except for $d'(1)$).
3. Third stage: Use $d'(n)$ as a new teacher signal and retrain the ESN with $d'(n)$ instead of $d(n)$, obtaining the final set of output weights \mathbf{w}_{out}^1 .

Because $d'(n)$ incorporates some effects of the autonomous network update dynamics, using it as a teacher in stage 3 brings the network dynamics closer to the autonomous network dynamics in exploitation runs, and thereby yields an altogether more accurate model.

The generation of a synthetic training sequence $d'(n)$ from $d(n)$ can be iterated by generating $d''(n)$ from $d'(n)$ and \mathbf{w}_{out}^1 as $d'(n)$ was generated from $d(n)$ and \mathbf{w}_{out}^0 , etc. However, further such

"relaxation stages" beyond $d(n)$ did not result in further improvements of model accuracy in the Mackey-Glass task (but sometimes did in other chaotic time series prediction tasks not reported here).

When the basic learning method is used to train chaotic attractor predictors, adding noise during sampling time is often helpful to obtain stable network dynamics. This was found to be unnecessary in the refined version of the method.

The improvements afforded by this re-training method can further be enhanced by a standard trick of machine learning: averaging over different models. We demonstrate this on the MGS prediction task (delay 17). Twenty ESN reservoirs of 1000 units each were randomly created, with the same scaling parameters as reported for the basic version. A 3000 step teacher sequence $d(n)$ was computed from the MG equations as described in the methods section. For each of the 20 ESNs, the three-stage refined version of the learning method was carried out independently using $d(n)$, yielding 20 sets of output weights. As in the basic method, network state data from an initial washout period of 1000 steps were discarded from computing the linear regression of the output weights in stages 1 and 3.

For exploitation, the 20 ESN models were combined through averaging their output. Concretely, in each update step n , the 20 networks were updated independently. Their averaged output was fed back into each of the networks. The sequence of averaged outputs was considered as the total model output. The NRMSE_{84} was computed as described above. This entire procedure was repeated 10 times, yielding 10 values of NRMSE_{84} , each corresponding to a model of 20 combined 1000-neuron ESNs. The average of the \log_{10} 's of the 10 NRMSE_{84} 's was -5.09 (stddev 0.25). Compared to the previously achievable accuracy for this prediction task of $\log_{10}(\text{NRMSE}_{84}) = -1.7$, this is an improvement of about 3.4 orders of magnitude, or a factor of 2450. Fig. S1 illustrates the prediction error development on long prediction runs.




Figure S1: (A) Prediction of the MG (delay 17) attractor using the refined ESN method. Blue: correct continuation, green: network prediction. (B) \log_{10} of absolute prediction error. Blue: basic method (averaged over 10 networks, 10 predictions each). Red: refined method (one 20-network model, averaged over 20 predictions). Green lines indicate the divergence rate of the original system, that is, the function $\exp(\lambda_1)$, where λ_1 is the largest Lyapunov exponent of the MG system. The x -axis shows network updates (= MG time units) in both plots.

The echo state property

Intuitively, a RNN which is driven by an external signal $u(n)$ has the echo state property if the activations $x_i(n)$ of the RNN neurons are systematic variations of the driver signal $u(n)$. More formally, this means that for each internal unit x_i there exists an "echo function" e_i , such that, if the network has been run for an indefinitely long time in the past, the current state can be written as $x_i(n) = e_i(u(n), u(n-1), u(n-2), \dots)$. We gave several nontrivial alternative definitions of this condition and algebraic characterizations of which network weight matrices \mathbf{W} lead to networks having the echo state property (SI). For practical purposes it suffices to fix the spectral radius of \mathbf{W} below unity to ensure the echo state property.

The echo state property is crucial for making the ESN learning method work. Consider some target system of the NMA type (= Nonlinear Moving Average). The current output $d(n)$ of an NMA system is a function of the input history, that is, $d(n) = f(u(n), u(n-1), u(n-2), \dots)$. Due to the echo state property, the ESN learning approach boils down to approximate the nonlinear system function f by a linear combination of the echo functions, where the linear combination weights are the trained output connection weights: $f \approx \sum_i w_i e_i$. Written out, this becomes $d(n) = f(u(n), u(n-1), \dots) \approx \sum_i w_i e_i(u(n), u(n-1), \dots)$. The equalizer system considered in the article is of type NMA.

There are two other major classes of nonlinear dynamical systems, NAR ("Nonlinear Auto-Regressive") and NARMA ("Nonlinear Auto-Regressive Moving Average") systems. NAR systems have system functions of the kind $d(n) = f(d(n-1), d(n-2), \dots)$ and NARMA systems of the kind $d(n) = f(u(n), u(n-1), u(n-2), \dots, d(n-1), d(n-2), \dots)$. The Mackey-Glass attractor is a NAR system. In NAR systems, the next system output is determined by previous output values. This was effected in the MGS study reported in the article through the output feedback connections. In such systems, the teacher output $d(n)$ is approximated through the echo functions by $d(n) = f(d(n-1), d(n-2), \dots) \approx \sum_i w_i e_i(d(n-1), d(n-2), \dots)$. The article contained no NARMA example, which however can be learnt in a similar way through ENSs featuring both input neurons and output feedback connections, through

$$d(n) = f(u(n), u(n-1), \dots, d(n-1), d(n-2), \dots) \approx \sum_i w_i e_i(u(n), u(n-1), \dots, d(n-1), d(n-2), \dots).$$

Comparison with previous approaches to MGS prediction

The MGS system with delay parameter $\tau = 17$ is a standard benchmark system in research on time series prediction. Almost every available technique for nonlinear system modelling and prediction has been tested on the MGS system. Virtually all previous techniques (except a few which used other versions of RNNs) rely on Takens' theorem (S4). This theorem implies that a chaotic system observed in one variable $y(n)$ can in principle be perfectly predicted from a few previous support points, that is, the next value $y(n+1)$ is determined by a few previous output values $y(n), y(n-l), y(n-2l), \dots, y(n-kl)$, where l is the distance between support points and k is related to the embedding dimension. Typical values used for k for the MGS are 7 – 10, (7 is the theoretical minimum required). The values taken for l depend on the sampling frequency of the MGS signal. Using this basic approach, the prediction task boils down to estimate from training

data a prediction function f which produces the next value from the support values by $y(n+1) = f(y(n), y(n-l), \dots, y(n-kl))$. Methods differ in the mathematical representation of f , which may be expressed for instance by methods of fuzzy splines, feedforward neural networks, radial basis function networks, Volterra expansions, bilinear expansions, support vector machines, or weighted combinations of previously observed subsequences in the training data ("direct methods"). If carried out with care, these approaches invariably yield NRMSE prediction errors for the 84-step prediction in the range of $10^{-1.2}$ to $10^{-1.7}$ (S5-S11). In contrast, ESNs do not try to estimate such a prediction function and thus do not ultimately rely on Takens' theorem. The much better precision achieved with ESNs suggest that it is the reliance on Takens' theorem which limited the precision of previous approaches. Specifically, the inevitable numerical inaccuracy entailed by rounding errors implies that the information about the current MGS system state contained in the few support points of traditional techniques is poorer than the massive information that is retained in an ESN "echo" network state about a long previous history (S12).

Applying the ESN method to predicting other chaotic attractors

The Lorenz attractor

The Lorenz attractor is governed by the three-dimensional ordinary differential equation $dx/dt = p(y - x)$, $dy/dt = -xz + rx - y$, $dz/dt = xy - bz$. We used Lorenz' original parameters $p = 10$, $r = 28$, and $b = 8/3$, for which one obtains a chaotic attractor with a largest Lyapunov exponent of $\lambda_{1 \text{ Lorenz}} \approx 0.906$, which thus exhibits a much stronger chaoticity than the Mackey-Glass system. Training and testing data were prepared by solving the Lorenz equation with the Runge-Kutta (4,5) method, stepsize 0.02 (employing Matlab's `ode45` solver). The stepsize 0.02 was also taken for the ESN update interval, thus one network update covers 0.02 units of original Lorenz time. Following most other learning studies of this attractor, we used only the x -coordinate trajectory for training and testing. The x -coordinate time series thus obtained from the Lorenz equation were rescaled by a factor of 0.01 to obtain series $x'(n)$ to be used for training and testing.

We compared the basic training method, as described in the main text, with the refined method outlined in the SOM.

For one training-test run of the original method, a single network of 500 units was created (spectral radius 0.97, connectivity 1/50, input connection weights sampled from uniform distribution over $(-1, 1)$, output feedback connection weights sampled from uniform distribution over $(-4, 4)$). It was trained on 6000 update steps on $x'(n)$, with uniform noise sampled from $(-1e-8, 1e-8)$ added to states during updates. Data from the first 1000 updates were discarded. The trained network was tested on 20 independent prediction tasks, each consisting of first running the network through 1000 teacher-forced updates and then letting it run freely to predict the next 600 steps. All error computations were done with data scaled back to the original Lorenz scale. 84-step prediction errors NRMSE_{84} were obtained from the 20 prediction tasks and taken to \log_{10} . This train-test procedure was repeated for 10 independently created networks on independently created data. The \log_{10} NRMSE_{84} had an average of -2.33 (stddev 0.54 across the 10 networks). Fig. S2B shows the average absolute error development.

For the refined method, 20 randomly created networks of 500 units each were trained on 6000 sample points with one relaxation stage. Network scaling parameters were the same as in the basic method. No noise was added to states during sampling. Data from the first 1000 cycles were discarded. The resulting 20-network model was tested by running it through 20 prediction tasks on independent test data. Each prediction task was carried out by first running the model in teacher-forced mode through an initial washout period of 1000 steps and then letting it run freely to predict the continuation for another 600 steps. For the 84-step prediction, an error of $\log_{10} \text{NRMSE}_{84} = -4.27$ was obtained, about two orders of magnitude better than with the basic method. Fig. S2A shows a typical prediction, and Fig. S2B shows the average absolute error development on the 600 steps. The prediction reaches its maximal deviation from the correct continuation after about 600 network updates, corresponding to 12 Lorenz time units.

The accuracy relative to the degree of chaos is quite similar to the accuracy seen in the (refined) Mackey-Glass prediction: noticing that the Mackey-Glass predictions reach their maximum average departure from the correct continuation after about 1700 time units (see Fig. S1) we find $12 \text{ Lorenz time units} * \lambda_1 \text{ Lorenz} = 10.9 \approx 10.2 = 1700 \text{ MG time units} * \lambda_1 \text{ MackeyGlass}$.




Figure S2: (A) Prediction of the Lorenz attractor using the refined ESN method. The first coordinate of the attractor is shown. Blue: correct continuation, green: network prediction. (B) \log_{10} of absolute prediction error. Blue: basic method (averaged over 10 networks, 20 predictions each), red: refined method. The x -axis shows network updates in both plots. Divide by 50 to obtain Lorenz timescale.

A conspicuous feature in Fig. S2B is the initial sharp error rise in the first 40-50 network updates. This results from the fact that the network was teacher-forced by the correct signal before the free-running prediction starts. When the network is decoupled from the forcing teacher at update step 1, it relaxes from a forced mode into its autonomous dynamics, where it settles after about 40-50 updates. A similar relaxation with an initial error rise can be seen in Fig. S1B, although it is less pronounced there.

The Lorenz attractor has been used often to demonstrate time series prediction methods. Unfortunately, prediction studies using the Lorenz attractor vary in important characteristics (number of training points, sampling rates, prediction horizons, error criteria, dimensions used for training), which makes a comparison difficult. The best Lorenz prediction study (S13) known to the authors manage to keep track of the Lorenz time series through 4 "switches" of the attractor from one lobe to the next, while our ESN model (refined version) stays with the true continuation for 5-11 such switches (mean 7.0 on 20 visually inspected predictions, stddev 1.6). Considering the fact that roughly such switching events occur after 2 Lorenz time units on average, in order to

stay with the correct continuation for one more such switch one needs an additional accuracy factor of $\exp(2 * \lambda_{1 \text{ Lorenz}}) \approx 6$. A prediction method A that stays with the correct continuation for 3 more lobe switches than another prediction method B would therefore have to be about $6^3 \approx 1e2.2$ times more accurate (in terms of absolute distance between predicted and correct trajectories).

The Mackey-Glass attractor with delay 30

Besides the MG attractor with a delay of 17, the other widely investigated version of the MG attractor is defined by putting the delay to 30. This attractor still is only mildly chaotic ($\lambda_{\max} \approx 0.007$), but the resulting time series is phenomenologically much richer than for $\tau = 17$ and therefore more difficult to learn. We applied the basic and the refined version of ESN training to this task.

Training and testing data were prepared by simulating the attractor with a stepsize of 1.0, using the `dde23` solver of Matlab with absolute accuracy set to $1e-16$. The resulting time series were subsampled by 6, such that one increment in the resulting time series $d(n)$ corresponds to 6 time units of the original MG equation, a common setting in MG modeling studies. These data $d(n)$ were brought into a range suitable for training by putting $d'(n) = (0.3 * \tanh(d(n) - 1)) + 0.2$.

For testing the basic learning method, we created 10 random ESNs (500 neurons each, connectivity 1/50, spectral radius 0.9) with a single input and a single output neuron. Input connection weights were sampled from a uniform distribution over $(-0.2, 0.2)$, and output feedback weights from a uniform distribution over $(-1, 1)$. The networks were individually trained with the basic method on a 3000-step sequence $d'(n)$ with a constant "bias" input of size 1.0. While the network was driven by the teacher data, uniform noise $v(n)$ sampled from $(-0.0008, 0.0008)$ was added on the network states. Data from the first 1000 steps were discarded before computing the output weights. Figure S3B shows the average absolute prediction error for predictions lasting 150 network updates. For all test error computations, the network outputs were first retransformed back to the original MG scaling and compared with the originally scaled MG signals $d(n)$.

For testing the refined method, ten independently created ESNs (700 neurons each, one input neuron, one output neuron, spectral radius 0.85, connectivity 1/70) were employed in the refined learning method. Input and output feedback connection weights were randomly assigned like in the basic method. The model was trained by teacher-forcing it through 3000 data points, of which the first 1000 were discarded before computing the regression of the output weights on the network states. Figure S3B shows the average absolute prediction error on independent test data. A typical prediction is shown in Fig. S3A. Both graphics demonstrate that the model loses track of the original after about 120 iterations, corresponding to 720 MG time units.




Figure S3: (A) A sample prediction with the refined ESN model. Blue: true continuation, green: model prediction. (B) Average absolute prediction error of ESN models for iterated prediction. For the basic method, the average is taken over 10 networks and 20 predictions each, for the refined method over 50 predictions of a single model made from 10 combined ESNs. x -axis corresponds to network updates in both plots. Multiply by 6 to obtain MG time units.

From a perspective of nonlinear dynamics, the delay = 30 case of the MG attractor is more difficult to model than the delay = 17 case not because of different degrees of chaos (they are almost the same) but because of a higher embedding dimension of the delay-30 attractor (S14), namely, 4 vs 3. Intuitively this means that the trajectory of the delay-30 attractor "lives" in a 4-dimensional space, as opposed to a 3-dimensional space for the delay-17 case. This has dramatic effects on learnability: if training information from 2000 points is dispersed in a 4-dimensional space as opposed to a 3-dimensional space, it covers state space volume much more sparsely. In order to obtain the same covering density as in the delay-17 case, $2000^{4/3} \approx 25,000$ training points would have been needed to ensure a similar model precision. In addition, the geometrical shape of the delay-30 attractor is much more involved than the delay-17 attractor (not discussed here; intuitively, it is much more "curly"), which again would require a relatively higher density of training points to obtain a similar model precision as in the delay-17 case. Both effects taken together render the delay-30 MG attractor a formidable modeling task.

The best approach to modeling this attractor we are aware of (S15) obtains a root mean square error for a prediction of 120 MG time units of 0.04, which corresponds to a \log_{10} NRMSE₁₂₀ of -0.15 . Our corresponding figures for the \log_{10} NRMSE₁₂₀ are -0.84 on average for the single nets trained with the basic method (stddev over the individual 10 networks' \log_{10} NRMSE₁₂₀: 0.083) and -1.42 for the refined method.

The Laser time series

One of the time series used in the 1994 time series prediction competition (S16) organized by the Santa Fe Institute was an empirical time series obtained from measuring a pulsing laser. This data set has been tackled with virtually every time series prediction method available. Figure S4A shows the data as they were published for the competition; the task was to predict the next 100 steps.

Using the refined method, we trained a collective of 20 ESNs with 300 neurons each on the original data set. The 200-step prediction obtained is shown in Fig. S4B. This prediction matches

the first high spikes very well, accurately catches the breakdown event, but is out of phase in the ensuing buildup of the oscillation. Very similar outcomes have been observed by a number of other researchers; no-one could catch the phase in the final buildup section correctly. The reasons for this have been elucidated (S17), as follows. The training dataset contains exactly one instance of a breakdown event like the one in the prediction interval (red box in Fig. S4A). The antecedent oscillations of this event in the training data match the oscillations right at the end of the training data and the first few steps into the prediction interval extremely closely. Figure S4C shows an overlay of the event in the training data with the event of the true continuation to illustrate this finding. Assuming that the physical laser is a deterministic system, the best prediction is therefore just to replicate the continuation from the similar section in the training data. This is what the ESN model does, as illustrated in Fig. S4D. Thus, the ESN model has "found out" that it faces essentially a pattern matching problem.

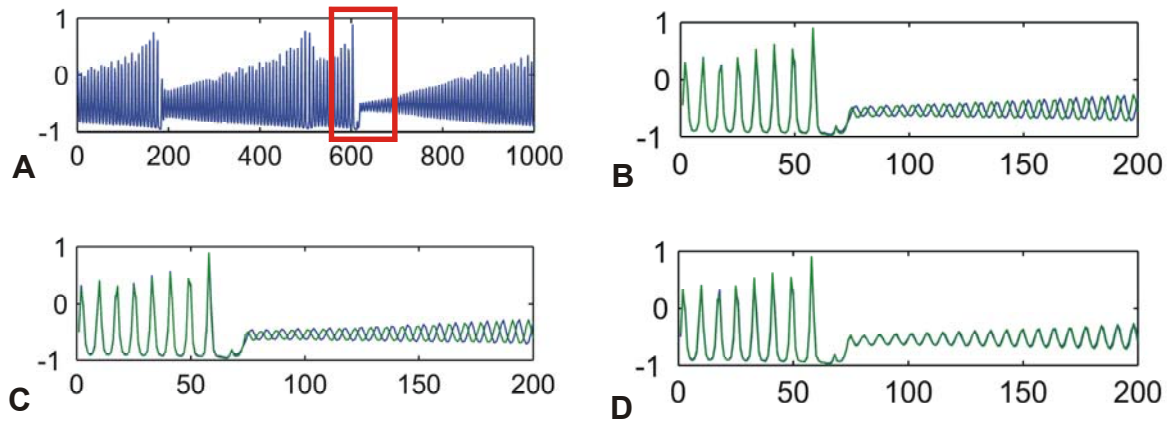


Figure S4: Predicting the Laser time series. (A) Training data. (B) ESN prediction (green) vs. true continuation (blue) after time 1000. (C) Situation in training data (green) [= data from red box in A] vs. true continuation (blue) after time 1000. (D) Situation in training data [= data from red box in A] (blue) vs. ESN prediction (green) after time 1000.

This problem exhibits specific difficulties not present in the other examples. First, the original data carry substantial numerical roundoff noise (they are measured with 8 bit precision), which has prompted most researchers to preprocess the dataset by interpolation/smoothing. We used the raw data. Second, the time series has two very different timescales: a fast oscillation and a slow buildup-breakdown cycle. Third, the breakdown event that has to be predicted occurs only once in the training data. All of these difficulties are not serious if one treats the problem as a pattern matching task right away and solves it by detecting the fact that the immediate antecedent of the prediction interval occurs already earlier in the training data, and re-using the original continuation. This is what many successful solutions do. However, if this valuable insight is not made or not used, the three difficulties become formidable. From a machine learning perspective, especially the third difficulty is noteworthy: we face here an instance of the "single shot learning" problem, that is, a model has to capture some phenomenon from a single presentation during training.

Supporting Online Material References

- S1. H. Jaeger, "The echo state approach to analysing and training recurrent neural networks" (GMD-Report 148, German National Research Institute for Computer Science 2001).
ftp://borneo.gmd.de/pub/indy/publications_herbert/EchoStatesTechRep.pdf
- S2. V. J. Mathews, J. Lee, in *Advanced Signal Processing: Algorithms, Architectures, and Implementations V (Proc. SPIE Vol. 2296)*, (SPIE, San Diego, CA, 1994), pp. 317-327.
- S3. B. Farhang-Boroujeny, *Adaptive Filters: Theory and Applications* (Wiley, New York 1998).
- S4. F. Takens, in *Dynamical Systems and Turbulence*, D.A. Rand, L.-S. Young, Eds. (Springer Lecture Notes in Mathematics **898**, 1991), pp. 366-381.
- S5. J. McNames, J. A. K. Suykens, J. Vandewalle, *Int. J. Bifurcation Chaos* **9**, 1485 (1999).
- S6. J. Vesanto, in *Proc. WSOM'97* (1997).
<http://www.cis.hut.fi/projects/monitor/publications/papers/wsom97.ps.zip>
- S7. L. Chudy, I. Farkas, *Neural Network World* **8**, 481 (1998).
- S8. H. Bersini, M. Birattari, G. Bontempi, in *Proc. IEEE World Congr. on Computational Intelligence (IJCNN '98)* pp. 2102-2106 (1997).
ftp://iridia.ulb.ac.be/pub/lazy/papers/IridiaTr1997-13_2.ps.gz.
- S9. T. M. Martinez, S. G. Berkovich, K. J. Schulten, *IEEE Trans. Neural Networks* **4**, 558 (1993).
- S10. Yao, X. & Liu. Y. *IEEE Trans. Neural Networks* **8**, 694 (1997).
- S11. F. Gers, D. Eck, J. F. Schmidhuber, "Applying LSTM to time series predictable through time-window approaches" (IDSIA-IDSIA-22-00, 2000).
<http://www.idsia.ch/~felix/Publications.html>.
- S12. H. Jaeger, "Short term memory in echo state networks" (GMD-Report 152, German National Research Institute for Computer Science 2002).
ftp://borneo.gmd.de/pub/indy/publications_herbert/STMEchoStatesTechRep.pdf
- S13. J. McNames, PhD thesis, Stanford University 1999.
<http://www.ece.pdx.edu/~mcnames/Publications/Dissertation.pdf>
- S14. J. D. Farmer, *Physica 4D* (1982).
- S15. L.A. Feldkamp, D. V. Prokhorov, C. F. Eagen, F. Yuan, in *Nonlinear Modeling: Advanced Black-Box Techniques*, J.A.K. Suykens, J. Vandewalle, Eds., (Kluwer 1998), pp. 29-54.
<http://mywebpages.comcast.net/dvp/bpaper.pdf>
- S16. A. S. Weigend, N. A. Gershenfeld (Eds.), *Time Series Prediction: Forecasting the Future and Understanding the Past* (Addison-Wesley, Reading, Mass., 1992).
- S17. J. Kohlmorgen, K.-R. Müller, *Neural Processing Letters* **7**(1), 43 (1998).
<http://www.first.gmd.de/persons/Kohlmorgen.Jens/publications/datasetA.pdf>