ESANN 2011 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 27-29 April 2011, i6doc.com publ., ISBN 978-2-87419-044-5. Available from http://www.i6doc.com/en/livre/?GCOI=28001100817300.

Reservoir regularization stabilizes learning of Echo State Networks with output feedback

René Felix Reinhart and Jochen Jakob Steil

Research Institute for Cognition and Robotics, Bielefeld University
Universitätsstr. 25, 33615 Bielefed - Germany
{freinhar, jsteil}@cor-lab.uni-bielefeld.de

Abstract. Output feedback is crucial for autonomous and parameterized pattern generation with reservoir networks. Read-out learning can lead to error amplification in these settings and therefore regularization is important for both generalization and reduction of error amplification. We show that regularization of the inner reservoir network mitigates parameter dependencies and boosts the task-specific performance.

1 Introduction

Output feedback is a crucial ingredient for reservoir computing applications like autonomous pattern generation [1, 2], parameterized pattern generation [3, 4], where piecewise constant inputs additionally modulate otherwise autonomously generated patterns, and also bidirectional pattern association [5, 6]. Output feedback introduces a trainable feedback loop which can potentially amplify small deviations from the target pattern. In this context, "sticking to a learned pattern" is often called stability without precise definition. We refer to this notion of stability throughout the text and formalize output feedback stability later on. Although stability in this sense can be achieved using heuristics [1], regularization of the read-out learning [2] or online learning techniques like BPDC [5] and FORCE learning [7], the success of learning and thus stability depends strongly on the choice of the read-out learning and regularization parameters. We show that a generic and efficient regularization approach for the inner reservoir makes read-out learning more robust to parameter variations.

The method we pursue here is based on the recently introduced reservoir regularization (RR) for input-driven recurrent networks [8]. Very recently, Jaeger took a similar approach in order to compile the functionality of an external controller into the reservoir [9]. The idea is to reimplement a previously harvested state sequence with the smallest weights possible, i.e. introducing a Gaussian prior distribution with zero mean for the reservoir weights while constraints demand to implement the desired state sequence. Since the dynamics are reimplemented for a wide range of input patterns with minimized weight norm [8], reservoir regularization provides a generic, task-independent criterion for a good reservoir and a means to prefer one reservoir to another. The minimized weight norm is particularly useful for reservoirs with output feedback: We show that networks with smaller weight norm are more robust against erroneous feedback but also allow to relax regularization constraints of the read-out learning which in turn increases the task-specific performance.

ESANN 2011 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 27-29 April 2011, i6doc.com publ., ISBN 978-2-87419-044-5. Available from http://www.i6doc.com/en/livre/?GCOI=28001100817300.

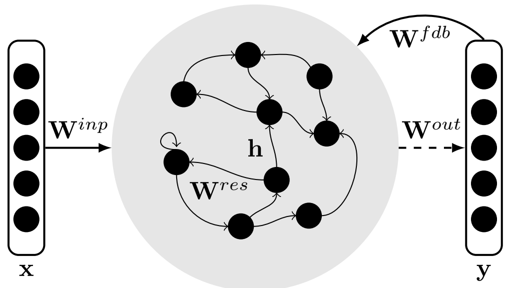

Fig. 1: Echo State Network with output feedback connections $\mathbf{W}^{fdb}$.

2 Echo State Network (ESN) with output feedback

We focus on ESN architectures with output feedback as depicted in Fig. 1. The network architecture comprises a recurrent reservoir network of nonlinear neurons which are connected randomly with small connection strengths. We consider the discrete and synchronous recurrent network dynamics

$$\mathbf{h}(k+1) = \sigma\left(\mathbf{W}^{inp}\mathbf{x}(k) + \mathbf{W}^{res}\mathbf{h}(k) + \mathbf{W}^{fdb}\mathbf{y}(k)\right) \quad (1)$$

$$\mathbf{y}(k+1) = \mathbf{W}^{out}\mathbf{h}(k). \quad (2)$$

$\mathbf{x}, \mathbf{h}$ and $\mathbf{y}$ are the input, reservoir and output neurons, where $\mathbf{h}$ is obtained by applying sigmoid activation functions component-wise. We subsume all connections to the reservoir in $\mathbf{W}^{net} = \left(\mathbf{W}^{inp}\, \mathbf{W}^{res}\, \mathbf{W}^{fdb}\right)$. Note that in case of autonomous pattern generation there are no input neurons, i.e. $\mathbf{W}^{inp}\mathbf{x}(k) \equiv 0$.
Read-out Learning: Originally, learning of ESNs is restricted to the supervised optimization of the read-out weights $\mathbf{W}^{out}$ in order to infer a desired input-to-output mapping from a set of training examples. The network is driven by external inputs and the reservoir states $\mathbf{h}(k)$ as well as the desired output series $\mathbf{t}(k)$ are recorded for $k = 1, \ldots, K$ in a reservoir state matrix $\mathbf{H} = (\mathbf{h}(1), \ldots, \mathbf{h}(K))^T \in \mathbb{R}^{K \times dim(\mathbf{h})}$ and target matrix $\mathbf{T} = (\mathbf{t}(1), \ldots, \mathbf{t}(K))^T$, respectively. The procedure of collecting the reservoir states $\mathbf{H}$ is called *harvesting states*. The optimal read-out weights are then determined by the least squares solution $(\mathbf{W}^{out})^T = \left(\mathbf{H}^T\mathbf{H} + \alpha\mathbb{1}\right)^{-1}\mathbf{H}^T\mathbf{T}$, where the Tikhonov factor $\alpha \geq 0$ regularizes the read-out weights $\mathbf{W}^{out}$.
Output feedback and teacher-forcing: Without output feedback ($\mathbf{W}^{fdb} \equiv \mathbf{0}$), only inputs $\mathbf{x}(k)$ drive the reservoir dynamics and training the read-out weights $\mathbf{W}^{out}$ does not affect the reservoir state. If $\mathbf{W}^{fdb} \neq \mathbf{0}$, the network state is also driven by the outputs $\mathbf{y}(k)$ (compare (1) and Fig. 1). During learning, the outputs $\mathbf{y}(k)$ are typically *teacher-forced* [1, 2, 3, 4, 6], i.e. clamped to the desired output sequence $\mathbf{t}(k)$. In exploitation mode, the estimated outputs $\mathbf{y}(k)$ are fed back into the reservoir which we refer to as *output feedback driven* (OFD) mode. The trained feedback loop parameterized by $\mathbf{W}^{out}$ and $\mathbf{W}^{fdb}$ can lead to an error amplification when the network is output feedback driven [1, 2].

ESANN 2011 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 27-29 April 2011, i6doc.com publ., ISBN 978-2-87419-044-5. Available from http://www.i6doc.com/en/livre/?GCOI=28001100817300.

Output feedback stability: We formalize the concept of stability for ESNs with output feedback. The reservoir shall implement a desired (output) behavior which we can formulate in terms of the network state trajectory $\mathbf{h}(k)$: Consider a teacher-forced network state sequence $\mathbf{h}(k)$ with external inputs $\mathbf{x}(k)$ and $\mathbf{t}(k)$ for $k = -\infty, \ldots, 0$. Then, duplicate the network and run one copy further with teacher-forced inputs and outputs. The output neurons of the second copy are released from the teacher signal yielding the OFD sequence $\mathbf{h}_{fdb}(k)$. The output feedback driven system is said to be *output feedback stable* (OFS) if it remains ϵ-close to the teacher-forced sequence, i.e. $||\mathbf{h}(k) - \mathbf{h}_{fdb}(k)|| < \epsilon$ for $k = 1, \ldots, K$ and any small $\epsilon > 0$. Note that this stability criterion is closely related to the task-specific performance and ties both concepts together: Teacher-forced networks that can not fit desired outputs are not OFS, and non-OFS networks fail to reproduce desired outputs. The OFS criterion is stricter than orbital stability and monitors the precise temporal coupling of input and network dynamics.
Read-out regularization: Whereas Tikhonov regularization is typically meant to improve generalization, the preference for small weights also reduces the potential for error amplification [1, 2]. The Tikhonov factor α can therefore be tuned in order to reduce the potential error amplification on the one hand, but to still implement the desired pattern on the other hand. This can be accomplished by conducting a (grid) search over α for each network, where in the test condition the network is output feedback driven (see Algorithm 1 and [2]).
Reservoir Regularization (RR): In addition to the read-out regularization, we propose to also regularize the reservoir in order to mitigate the dependency on the read-out regularization parameter α. The main idea is to express a preference for small weights conditioned on the constraint to reimplement a previously harvested network state sequence [8]. Adopting the constraint optimization problem in [8] to reservoirs with output feedback, we obtain the regularized weights $(\mathbf{W}^{net}_{reg})^T = \left(\mathbf{S}^T\mathbf{S} + \beta\mathbb{1}\right)^{-1}\mathbf{S}^T\mathbf{A}$, where $\mathbf{S} = (\mathbf{s}(1), \ldots, \mathbf{s}(K))^T$ with $\mathbf{s}(k)^T = (\mathbf{x}(k)^T, \mathbf{h}(k)^T, \mathbf{t}(k)^T)$ are the harvested states together with the inputs $\mathbf{x}(k)$ and target outputs $\mathbf{t}(k)$. The reservoir targets $\mathbf{A} = (\mathbf{a}(2), \ldots, \mathbf{a}(K+1))^T$ with $\mathbf{a}(k+1) = \mathbf{W}^{inp}\mathbf{x}(k) + \mathbf{W}^{res}\mathbf{h}(k) + \mathbf{W}^{fdb}\mathbf{y}(k)$ are the neural activities before application of non-linear activation functions σ one time step ahead. The solution is similar to read-out learning with Tikhonov regularization, where β cor-

---

Algorithm 1 Grid search for read-out regularization

Require: training data $(\mathbf{x}(k), \mathbf{t}(k))$ with $k = 1, \ldots, K$, grid of α's, network
Require: harvest teacher-forced states $\mathbf{H}$
1: **for** $\alpha \in \{10^{-6}, 5 \times 10^{-6}, 10^{-5}, 5 \times 10^{-5}, \ldots, 1.5, 2\}$ **do**
2: read-out learning by $(\mathbf{W}^{out})^T = \left(\mathbf{H}^T\mathbf{H} + \alpha\mathbb{1}\right)^{-1}\mathbf{H}^T\mathbf{T}$ using α
3: output feedback driven testing yields task-specific error $E^{task}(\alpha) = \frac{1}{K}\sum_{k=1}^{K} ||\mathbf{t}(k) - \mathbf{y}(k)||^2$
4: **end for**
5: **return** trained reservoir network using α^{opt} that yields smallest error E^{task}

ESANN 2011 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 27-29 April 2011, i6doc.com publ., ISBN 978-2-87419-044-5. Available from http://www.i6doc.com/en/livre/?GCOI=28001100817300.

responds to the Tikhonov factor which trades off reimplementation constraints against weight norm. RR is conducted before read-out learning, i.e. reservoir states are harvested, the reservoir is regularized and then Algorithm 1 is applied.

3 Reservoir regularization of ESNs with output feedback

In this section, we show that the proposed reservoir regularization facilitates learning in Echo State Networks with output feedback in two complementary scenarios. All results are averaged over 200 independent trials. The reservoirs comprise 50 neurons with $tanh(\cdot)$ activation functions. The weight matrix $\mathbf{W}^{net}$ is randomly initialized according to a uniform distribution in range $[-0.5, 0.5]$. The spectral radius of the reservoir submatrix $\mathbf{W}^{res}$ is scaled to 0.9.

3.1 Learning inverse kinematics

We first consider a combined prediction and sequence transduction task from the robotics domain. The ESN learns the inverse kinematics of an arm with two degrees of freedom from a training sequence (see Fig. 2), i.e. the network has to map Cartesian end effector coordinates to joint angles. We use a circular and a figure eight-like motion as training and test data (compare Fig. 2). Each sequence comprises $K = 500$ samples. We represent joint angles in a coordinate representation for reasons of normalization, i.e. $\mathbf{y}(q_1, q_2) = (sin(q_1), cos(q_1), sin(q_2), cos(q_2))^T$.

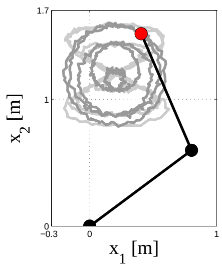

Fig. 2: Robot arm with two degrees of freedom. Training and test patterns are shown in gray and light gray, respectively.

Reservoir Regularization stabilizes dynamics: We first show that RR stabilizes the output feedback driven dynamics for a wide range of read-out regularization parameters α. Therefor, we calculate the difference $E^{state} = \frac{1}{K \cdot dim(\mathbf{h})} ||\mathbf{H} - \mathbf{H}_{fdb}||^2$ between the state trajectories $\mathbf{H}$ in the teacher-forced case (the network outputs are set to the target values $\mathbf{t}(k)$) and the output feedback driven case $\mathbf{H}_{fdb}$ (the predicted outputs $\mathbf{y}(k)$ are fed back into the reservoir). Fig. 3 (left) reveals that without RR the network state trajectory $\mathbf{H}$ in the teacher-forced case is only achieved in the OFD case if the read-out regularization α is strong enough (E^{state}_{ini}). Note that E^{state} directly corresponds to the OFS criterion and that the large errors in Fig. 3 (left) and Fig. 4 (left) for $\alpha < 10^{-4}$ correspond to non OFS networks. RR mitigates this dependency on α and makes reservoirs with output feedback more robust against the choice of this parameter (E^{state}_{reg} shown for $\beta \in \{10^{-6}, 10^{-4}, 10^{-2}\}$ in Fig. 3 (left)). β forces the reservoir weights to have a smaller norm (compare Fig. 3 (right)) which damps the amplification of deviations at the output. With RR, the network is more robust against deviations of the output trajectory.

Reservoir Regularization improves performance: Fig. 4 (left) clearly points out the connection between learning, regularization and stability: output feedback yields to increased training errors for small α's without RR. The

ESANN 2011 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 27-29 April 2011, i6doc.com publ., ISBN 978-2-87419-044-5. Available from http://www.i6doc.com/en/livre/?GCOI=28001100817300.

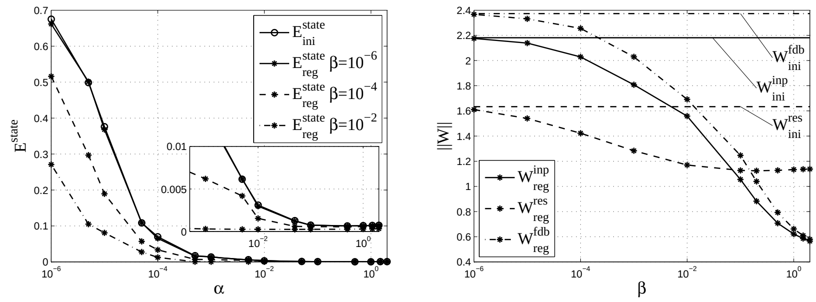

Fig. 3: Deviation E^{state} of the output feedback driven state sequence from the teacher-forced state sequence for the training scenario as function of the read-out regularization parameter α (left). Matrix norm of weights $\mathbf{W}^{inp}$, $\mathbf{W}^{res}$ and $\mathbf{W}^{fdb}$ as function of the RR parameter β (right). Subscripts indicate whether the networks where regularized (*reg*) before read-out training, or not (*ini*).

regularized reservoir weights reduce the effective gain of the system and cause the grid search Algorithm 1 to yield smaller values for α (see Fig. 4 (right)). This in turn improves the task-specific performance while keeping the network stable (compare Fig. 3 (left) and Fig. 4 (left)). We observe a highly significant correlation (Spearman rank correlation test with significance level 0.01) between the change of the read-out regularization parameter $\Delta\alpha = \alpha_{reg}^{opt} - \alpha_{ini}^{opt}$ and the change of the task-specific error $\Delta E^{task} = E_{reg}^{task} - E_{ini}^{task}$ (plotted using ranks for $\Delta\alpha$ in Fig. 4 (right)). This means that with RR, the read-out regularization can be reduced which consequently results in smaller task-specific errors. The combination of both small task-specific errors and regularized reservoir weights minimizes the risk of error amplification due to the output feedback loop.

3.2 Autonomous generation of a unit circle

To complement the results for the input-driven case, we also show that RR improves the stability in case of autonomous pattern generation. In this scenario, the networks have no input neurons and shall produce a unit circle pattern at their outputs autonomously. 5 periods of a unit circle $\mathbf{y}(k) = (sin(\omega k), cos(\omega k))^T$ with $\omega = \frac{2\pi}{200}$ are presented for training, where the first 4 periods are used as wash-out phase. For testing, the networks have to reproduce the circle running freely for 10^6 steps within an error margin of 5%. We neglect phase lags typically occurring during long-term recursive prediction, i.e. monitor the orbital stability of the output pattern $\mathbf{y}(k)$ with respect to the target pattern $\mathbf{t}(k)$. Without reservoir regularization, only 75% of the networks comply with this error bound for 10^6 steps, whereas 98, 5% of the networks satisfy this tight criterion if the reservoir is regularized with $\beta = 10^{-2}$. This confirms the previous result that RR improves the task-specific performance and makes ESN learning more robust in scenarios where output feedback is crucial.

ESANN 2011 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 27-29 April 2011, i6doc.com publ., ISBN 978-2-87419-044-5. Available from http://www.i6doc.com/en/livre/?GCOI=28001100817300.

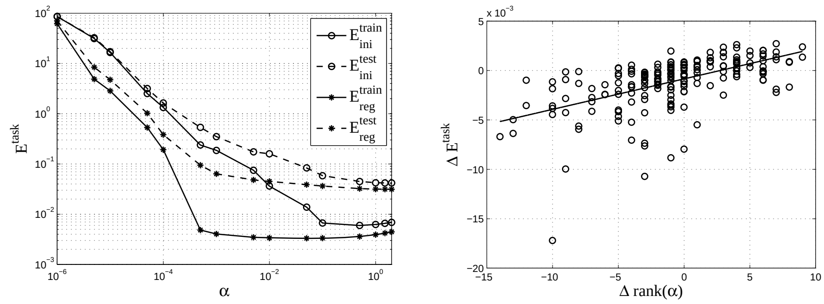

Fig. 4: Task-specific error E^{task} on the training and test set as function of the read-out regularization α (left). Change of the task-specific error due to reservoir regularization as function of the changed read-out regularization parameter α for the training scenario (right). $\beta = 10^{-2}$ is used in both plots.

4 Conclusion

Reservoir regularization enables robust offline training of ESNs with output feedback. RR mitigates the sensible balance of task-specific performance and read-out regularization by minimizing the norm of the weights that excite the reservoir. This in turn allows to increase the weight norm of the read-out weights and consequently improves the task-specific performance and stability. Although reservoir regularization introduces extra computational costs and an extra parameter, the gained robustness of learning is a necessary prerequisite for the convenient application of reservoir networks with output feedback.

References

[1] H. Jaeger, M. Lukosevicius, D. Popovici, and U. Siewert. Optimization and applications of echo state networks with leaky-integrator neurons. *Neural Networks*, 20(3):335–352, 2007.

[2] F. wyffels, B. Schrauwen, and D. Stroobandt. Stable output feedback in reservoir computing using ridge regression. In *ICANN*, pages 808–817, 2008.

[3] F. wyffels and B. Schrauwen. Design of a central pattern generator using reservoir computing for learning human motion. In *ECSIS Symp. on LAB-RS*, pages 118–122, 2009.

[4] 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*, volume 6, pages 11–21. Springer, 2009.

[5] R. F. Reinhart and J. J. Steil. Recurrent neural associative learning of forward and inverse kinematics for movement generation of the redundant PA-10 robot. In *ECSIS Symp. on LAB-RS*, pages 35–40, 2008.

[6] E. Antonelo, B. Schrauwen, and J. V. Campenhout. Generative modeling of autonomous robots and their environments using reservoir computing. *Neural Processing Letters*, 26(3):233–249, 2007.

[7] D. Sussillo and L. F. Abbott. Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63:544–557, 2009.

[8] R. F. Reinhart and J. J. Steil. A constrained regularization approach for input-driven recurrent neural networks. *Differ. Equ. Dyn. Syst.* DOI: 10.1007/s12591-010-0067-x.

[9] Herbert Jaeger. Reservoir self-control for achieving invariance against slow input distortions. Technical Report 23, Jacobs University, Oct. 2010.