

Mod-DeepESN: Modular Deep Echo State Network

Zachariah Carmichael
(zjc2920@rit.edu)

Humza Syed
(hxs7174@rit.edu)

Stuart Burtner
(stb3444@rit.edu)

Dhiresha Kudithipudi
(dxkeec@rit.edu)

Neuromorphic AI Lab, Rochester Institute of Technology
1 Lomb Memorial Drive, Rochester, NY 14623 USA

Abstract

Neuro-inspired recurrent neural network algorithms, such as echo state networks, are computationally lightweight and thereby map well onto untethered devices. The baseline echo state network algorithms are shown to be efficient at solving small-scale spatio-temporal problems. However, they underperform for complex tasks that are characterized by multi-scale structures. In this research, an intrinsic plasticity-infused modular deep echo state network architecture is proposed to solve complex and multiple timescale temporal tasks. It outperforms state-of-the-art for time series prediction tasks.

Keywords: echo state networks (ESN); intrinsic plasticity; time series prediction; reservoir computing (RC)

Introduction

Echo state networks (ESNs) are recurrent rate-based networks that are proven to be efficient in solving spatio-temporal tasks (Jaeger et al., 2007). Studies of the associated spiking models have shown that the recurrent layer (*a.k.a.* reservoir layer) corresponds to the granular cells in the cerebellum and the readout layer corresponds to the Purkinjee cells (Yamazaki & Tanaka, 2007). These models are attractive because they offer lightweight, resilient, and conceptually simple networks with rapid training time, as training occurs only at the readout layer. The recurrent nature of these models introduces feedback signals similar to an innate form of fading memory.

In a recent resurgence to enhance the capabilities of ESNs, few research groups have studied deep ESNs. The premise of adding depth to an ESN is to support hierarchical representations of the temporal input and also capture the multi-scale dynamics of the input features. In prior literature, ESNs have been applied to several temporal tasks, such as speech processing, EEG classification, and anomaly detection (Soures et al., 2017; Jaeger et al., 2007). The Deep-ESN architecture proposed by Ma et al. (2017) consists of stacked reservoir layers and unsupervised encoders to efficiently exploit the temporal kernel property of each reservoir. In Gallicchio et al. (2017), multiple deep ESN architectures based on the shallow Leaky-Integrator ESN of Jaeger et al. (2007) are analyzed. Of the architectures, the DeepESN network achieved the best results with and without utilizing intrinsic plasticity on memory capacity experiments.

Previous works have shown that increasing the depth of reservoir networks is necessary to capture multi-scale dynamics of time series data as well as extract features with a higher order of complexity (Gallicchio et al., 2017). However, to extract richer features from the data, a wider architecture com-




Figure 1: High-level representation of the Mod-DeepESN architecture with the *Wide* topology.

prises greater reservoir real estate as each reservoir learns *distinct local features*. Therefore, we propose a modular deep ESN architecture, known as Mod-DeepESN, which utilizes multiple reservoirs with heterogeneous topology and connectivity to capture and integrate the multi-scale dynamical states of temporal data. The proposed architecture is studied for a mix of benchmarks and consistently performs well across different topologies when compared to the baseline.

Proposed Design

Architecture

The proposed Mod-DeepESN architecture consists of modular and deep reservoirs that can be realized in multiple topologies and connectivities. Specifically, four topologies are explored in this study: (i) *Wide* (shown in Figure 1) (ii) *Layered* (iii) *2x2 Criss-Cross* and (iv) *Wide+Layered*. The reservoirs that are connected to the observed input receive the vector $\mathbf{u}(t) \in \mathbb{R}^{N_U}$ at time t , which is projected via the input weight matrix $\mathbf{W}_{in} \in \mathbb{R}^{N_R \times C(\mathbf{C}_{u,1}) \times N_U}$ into each reservoir layer, where $C(x, j) = \sum_i [x_i = j]$ and $[\dots]$ are Iverson brackets. N_U is the total number of inputs, N_R is the fixed number of neurons within each reservoir, and lastly N_L is the total number of reservoirs. \mathbf{C} is the binary and triangular connectivity matrix that determines the feedforward connections between reservoirs and the input \mathbf{u} . For example, a '1' at position $\mathbf{C}_{u,2}$ indicates a feedforward connection from the input \mathbf{u} to the second reservoir. The output of the l th reservoir, $\mathbf{x}^{(l)} \in \mathbb{R}^{N_R}$, is computed using (1)

$$\mathbf{x}^{(l)}(t) = (1 - a^{(l)})\mathbf{x}^{(l)}(t-1) + a^{(l)} \tanh \left(\mathbf{W}_{res}^{(l)} \mathbf{i}^{(l)}(t) + \hat{\mathbf{W}}_{res}^{(l)} \mathbf{x}^{(l)}(t-1) \right) \quad (1)$$

where $\mathbf{i}^{(l)}(t)$ is computed using (2).

$$\mathbf{i}^{(l)}(t) = \begin{cases} \mathbf{u}(t) & l = 1 \\ \mathbf{x}^{(l-1)}(t) & l > 1 \end{cases} \quad (2)$$

Equation (1) is also referred to as the state transition function in the literature. $\mathbf{W}_{res}^{(l)} \in \mathbb{R}^{N_R \times N_R}$ is a feedforward weight matrix that connects two reservoirs in the *Layered*, 2×2 *Criss-Cross*, and *Wide+Layered* topologies, while $\hat{\mathbf{W}}_{res}^{(l)} \in \mathbb{R}^{N_R \times N_R}$ is a recurrent weight matrix that connects intra-reservoir neurons. The per-layer leaky parameter $a^{(l)}$ controls the leakage rate in a moving exponential average manner. The non-linear activation function for the recurrent neurons in this work is the hyperbolic tangent. Note that the bias vectors are left out of the formulation for simplicity. The state of the Mod-DeepESN network is defined as the concatenation of the output of each reservoir, *i.e.* $\mathbf{x}(t) = (\mathbf{u}(t), \mathbf{x}^{(1)}(t), \dots, \mathbf{x}^{(N_L)}(t)) \in \mathbb{R}^{N_U + N_L N_R}$. The matrix of all states is denoted by $\chi(\mathbf{s}) = (\mathbf{x}(0), \dots, \mathbf{x}(N_t - 1)) \in \mathbb{R}^{(N_U + N_L N_R) \times N_t}$ where N_t is the number of time steps in the time series. Finally, the output of the network for the duration N_t is computed as a linear combination of the Mod-DeepESN state matrix using (3).

$$\mathbf{y}(\mathbf{s}) = \mathbf{W}_{out} \chi(\mathbf{s}) \quad (3)$$

The matrix $\mathbf{W}_{out} \in \mathbb{R}^{N_Y \times (N_U + N_L N_R)}$ contains the feedforward weights between reservoir neurons and the N_Y output neurons. Ridge regression using the Moore-Penrose pseudo-inverse is used to solve for optimal \mathbf{W}_{out} , shown in (4).

$$\mathbf{W}_{out} = \mathbf{y}(\mathbf{s}) \cdot (\chi(\mathbf{s}) \cdot \chi(\mathbf{s})^T + \beta \mathbb{I})^{-1} \quad (4)$$

The formulation includes a regularization term β , and \mathbb{I} is the identity matrix. This computation is performed only during the training phase of the model.

Mod-DeepESN Topologies

For all the proposed topologies, (i) *Wide* (ii) *Layered* (iii) 2×2 *Criss-Cross* and (iv) *Wide+Layered*, the input layer is directly connected to the output layer; this connection captures sensitive spatial features and improves the network performance experimentally. Additionally, each reservoir is connected to the output layer.

Figure 1 illustrates the *Wide* topology. Distinct connections are created between the input layer and each reservoir, *i.e.* the input weights are not shared between reservoirs. As a result, an ensemble of shallow ESNs emerges, each of which captures varying dynamics of the input data locally.

The *Layered* topology, shown in Figure 2(a), presents the input to the first reservoir with each successive reservoir receiving feedforward input from its predecessor. By stacking reservoirs in this fashion, stateful dynamics of the system are integrated with features that increase in complexity with depth.

Figure 2(b) depicts the 2×2 *Criss-Cross* topology, which comprises a total of 4, or $N_L = N^2$ ($N = 2$), reservoirs. This model exhibits dense feedforward connectivity between




Figure 2: (a) *Layered* topology, (b) 2×2 *Criss-Cross* topology, and (c) *Wide+Layered* topology of Mod-DeepESN.

all reservoirs, but maintains reservoir-local recurrent connections. The system states are hierarchically integrated at multiple depths from disparate inputs.

The *Wide+Layered* topology, presented in Figure 2(c), deviates from the 2×2 *Criss-Cross* topology through greater sparsity in its feedforward connectivity; dynamical states of the input are integrated in discrete pathways that resemble the *Layered* topology.

Weight Initialization

As weights in ESN reservoirs are fixed, recurrent weights must be initialized to maintain stable reservoir states, especially in deep ESNs where inputs to a network propagate through multiple reservoirs. As mathematically supported in the ESN literature, deep ESNs need to satisfy the *Echo State Property* for such stability. Accordingly, each $\hat{\mathbf{W}}_{res}^{(l)}$ is initialized using a uniform distribution and scaled such that (5) is satisfied.

$$\max_{1 \leq l \leq N_L} \rho \left((1 - a^{(l)}) \mathbb{I} + a^{(l)} \hat{\mathbf{W}}_{res}^{(l)} \right) < 1 \quad (5)$$

The spectral radius of the matrix, *i.e.* the magnitude of the largest eigenvalue, is computed by the function $\rho(\cdot)$. We substitute a hyperparameter $\hat{\rho}$ for the value 1 to fine-tune model reservoirs for each experiment.

The remaining weight matrices are initialized from a uniform distribution such that the Euclidean distance between weights is equivalent to either σ_{in} or σ_l , *i.e.* $\|\mathbf{W}_{in}\|_2 = \sigma_{in}$ and $\|\mathbf{W}_{res}^{(l)}\|_2 = \sigma_{res}$. σ_{in} and σ_l are tunable hyperparameters.

Rather than normalizing weights using the L_2 norm or spectral radius, Xavier initialization (Glorot & Bengio, 2010) can be used to determine the distribution of weights. In this method, any set of weights \mathbf{W} can be drawn from a normal distribution as shown by (6) and (7).

$$\sigma_{xav} = \sqrt{2/(n_{in} + n_{out})} \quad (6)$$

$$\mathbf{W} \sim \mathcal{N}(\mu = 0, \sigma_{xav}^2) \quad (7)$$

The distribution is parameterized with a mean of $\mu = 0$ and a standard deviation σ_{xav} , where n_{in} is the number of inputs and n_{out} is the number of outputs that the weight matrix is connecting. Additionally, this initialization can be made modular to isolate which sets of weights utilize a Xavier initialization and which use either L_2 or spectral radius normalization.

All modes of initialization additionally use a sparsity hyperparameter that determines the probability that each weight is nullified. Specifically, s_{in} determines sparsity for \mathbf{W}_{in} , \hat{s}_l for each $\hat{\mathbf{W}}_{res}^{(l)}$, and s_l for each $\mathbf{W}_{res}^{(l)}$.

Intrinsic Plasticity (IP)

In this section, we study how a neurally-inspired IP mechanism can enhance the performance of the Mod-DeepESN. Originally proposed by Schrauwen et al. (2008), the IP rule introduces a gain and bias term to the $\tanh(\cdot)$ function: $\tanh(x) \rightarrow \tanh(gx + b)$ where g is the gain and b is the bias. The update rules are given by (8) and (9) where

$$\Delta b = -\eta \left(-\frac{\mu}{\sigma^2} + \left(\frac{\tilde{x}}{\sigma^2} \right) (2\sigma^2 + 1 - \tilde{x}^2 + \mu\tilde{x}) \right) \quad (8)$$

$$\Delta g = \eta/g + \Delta b x \quad (9)$$

x is the net weighted sum of inputs to a neuron, \tilde{x} is the tanh activation of x , i.e. $\tilde{x} = \tanh(gx + b)$, and η is the learning rate. The hyperparameters σ and μ are the standard deviation and mean, respectively, of the target Gaussian distribution. In a pre-training phase, the learned parameters are initialized as $b = 0$ and $g = 1$, and are updated iteratively according to (8) and (9). The application of such results in the minimization of the Kullback-Leibler (KL) divergence between the empirical output distribution and the target Gaussian distribution (Schrauwen et al., 2008).

Genetic Algorithm

To fine-tune network hyperparameters, a genetic algorithm is employed (Bäck et al., 2000). The evolution is executed with a population size of 50 individuals for 50 generations, and evaluated in a tournament comprising 3 randomly selected individuals. Individuals mate with a crossover probability of 50% and mutate with a probability of 10% to form each successive generation. This evolution is used to tune all model hyperparameters as well as the Mod-DeepESN topology type.

Experiments & Results

Two datasets, the chaotic Mackey-Glass time series generated using the fourth-order Runge-Kutta method (RK4) and a daily minimum temperature series (Time Series Data Library, 1981–1990), are used to evaluate Mod-DeepESN as they present variable nonlinearity at multiple time scales. To analyze the effectiveness of the Mod-DeepESN model, root mean squared error (RMSE), normalized RMSE (NRMSE), and mean absolute percentage error (MAPE) are computed as shown in (10), (11), and (12), respectively.

$$\text{RMSE} = \sqrt{\frac{1}{N_t} \sum_{t=1}^{N_t} [\mathbf{u}(t) - \hat{\mathbf{u}}(t)]^2} \quad (10)$$

$$\text{NRMSE} = \sqrt{\frac{\sum_{t=1}^{N_t} [\mathbf{u}(t) - \hat{\mathbf{u}}(t)]^2}{\left(\sum_{t=1}^{N_t} [\mathbf{u}(t) - \bar{\mathbf{u}}]^2 \right)}} \quad (11)$$

$$\text{MAPE} = \frac{1}{N_t} \sum_{t=1}^{N_t} \frac{|\mathbf{u}(t) - \hat{\mathbf{u}}(t)|}{\mathbf{u}(t)} \times 100\% \quad (12)$$




Figure 3: (a) The Mackey-Glass time series and (b) the Melbourne daily minimum temperature time series.

$\hat{\mathbf{u}}(t)$ is the predicted time series value at time step t , and $\bar{\mathbf{u}}$ is the average value of the time series over the N_t time steps. Each reported result is averaged over 10 runs using the best hyperparameters found from the genetic algorithm search.

The Mackey-Glass dataset, shown in Figure 3(a), is split into 8,000 training samples and 2,000 testing samples to forecast 84 time steps in advance. To reduce the influence of initial reservoir states, the first 100 predictions from the network are discarded in a washout period.

The daily minimum temperature series dataset, shown in Figure 3(b), comprises data from Melbourne, Australia (1981–1990), and is split into 2,920 training samples and 730 testing samples. A washout period of 30 steps is realized during model evaluation for the one-step ahead prediction.

In each experiment, $\alpha^{(l)}$ is set to the same value to reduce the hyperparameter search space. We denote this value as α in Tables 1 and 2. Additionally, a value in the N_L column with parenthesis is used to indicate the reservoir width of the 2×2 *Criss-Cross* and *Wide+Layered* topologies. If Xavier initialization is used in place of initialization by spectral radius or L_2 normalization, X is reported. Hyperparameters are chosen independently of the topology, as the experimental performance consistently surpasses that of the best hyperparameters found per topology. Each topology is evaluated with and without IP to determine its efficacy, but only the best result using IP is reported. Baseline ESN results for both datasets are retrieved from Ma et al. (2017). The Mod-DeepESN implementation makes use of the scikit-learn (Pedregosa et al., 2011), SciPy (Jones et al., 2001), NumPy (Oliphant, 2006), and DEAP (Fortin et al., 2012) software libraries.

The Mackey-Glass experimental results demonstrate that model fitness depends on the Mod-DeepESN topology. An IP pre-training phase greatly reduces the network error and performs best while evaluated with the *Wide* topology. The proposed architecture outperforms all baseline models with the exception of the Deep-ESN (Ma et al., 2017).

Table 1: Mackey-Glass Time Series 84-Step Prediction Results (Scores and Standard Deviations).

Method		N_L	N_R	β	α	$\hat{\rho}$	σ_{in}	σ_l	s_{in}	\hat{s}_l	s_l	IP	RMSE $\times e-3$	NRMSE $\times e-3$	MAPE $\times e-3$
Baseline	ESN ¹	1	-	-	-	-	-	-	-	-	-	-	43.7 \pm 6.31	201 \pm 29.1	7.03 \pm 1.27
	ϕ -ESN ²	2	-	-	-	-	-	-	-	-	-	-	8.60 \pm 1.63	39.6 \pm 7.49	1.00 \pm .213
	R ² SP ³	2	-	-	-	-	-	-	-	-	-	-	27.2 \pm 4.27	125 \pm 19.6	1.00 \pm .213
	MESM ⁴	7	-	-	-	-	-	-	-	-	-	-	12.7 \pm 2.50	58.6 \pm 11.5	1.91 \pm .422
	Deep-ESN ⁵	3	-	-	-	-	-	-	-	-	-	-	1.12 \pm .187	5.17 \pm .861	.151 \pm .031
This Work	<i>Wide</i>	3	256	2e-8	.6	X	.1	X	.1	.1	.7	N	20.2 \pm 11.5	48.7 \pm 21.6	8.11 \pm 1.07
	<i>Layered</i>	3	256	2e-8	.6	X	.1	X	.1	.1	.7	N	57.8 \pm 20.2	96.2 \pm 14.1	19.9 \pm 3.34
	<i>Criss-Cross</i>	4 (2)	256	2e-8	.6	X	.1	X	.1	.1	.7	N	56.1 \pm 21.6	54.4 \pm 17.3	8.94 \pm .302
	<i>Wide+Layered</i>	6 (3)	256	2e-8	.6	X	.1	X	.1	.1	.7	N	41.1 \pm 5.15	55.4 \pm 3.74	11.2 \pm .320
	<i>Wide</i>	3	256	2e-8	.6	X	.1	X	.1	.1	.7	Y	7.22 \pm .623	27.5 \pm 4.00	5.55 \pm .088

Table 2: Daily Minimum Temperature Series Single-Step Prediction Results (Scores and Standard Deviations).

Method		N_L	N_R	β	α	$\hat{\rho}$	σ_{in}	σ_l	s_{in}	\hat{s}_l	s_l	IP	RMSE $\times e-3$	NRMSE $\times e-3$	MAPE $\times e-3$
Baseline	ESN ¹	1	-	-	-	-	-	-	-	-	-	-	501 \pm 3.70	139 \pm 1.02	39.5 \pm .237
	ϕ -ESN ²	2	-	-	-	-	-	-	-	-	-	-	493 \pm 3.86	141 \pm 1.10	39.6 \pm .374
	R ² SP ³	2	-	-	-	-	-	-	-	-	-	-	495 \pm 3.55	137 \pm .982	39.3 \pm .434
	MESM ⁴	7	-	-	-	-	-	-	-	-	-	-	478 \pm 3.39	136 \pm .967	37.7 \pm .336
	Deep-ESN ⁵	2	-	-	-	-	-	-	-	-	-	-	473 \pm 2.77	135 \pm .791	37.0 \pm .214
This Work	<i>Wide</i>	2	1024	7e-4	1	X	.4	X	.6	.3	.6	N	473 \pm 1.21	135 \pm .370	38.6 \pm .155
	<i>Layered</i>	2	1024	7e-4	1	X	.4	X	.6	.3	.6	N	470 \pm 1.11	134 \pm .307	38.2 \pm .113
	<i>Criss-Cross</i>	4 (2)	1024	7e-4	1	X	.4	X	.6	.3	.6	N	472 \pm .523	135 \pm .167	38.7 \pm .064
	<i>Wide+Layered</i>	4 (2)	1024	7e-4	1	X	.4	X	.6	.3	.6	N	471 \pm .814	135 \pm .224	38.6 \pm .154
	<i>Wide+Layered</i>	4 (2)	1024	7e-4	1	X	.4	X	.6	.3	.6	Y	459 \pm .236	132 \pm .111	37.1 \pm .074

In the daily minimum temperature time series experiment, the *Wide+Layered* Mod-DeepESN outperforms every baseline model, and further reduces the error with the inclusion of an IP pre-training phase. Varying the topology within this experiment is less impactful than within the Mackey-Glass experiment, but deeper models consistently yield better performance.

It is important to note that the best Mod-DeepESN models use Xavier initialization for multiple weight matrices. While this does not necessarily suggest that spectral radius or L_2 normalization are inferior initialization methods, it does indicate that a desirable initialization can be achieved by using a non-parametrized Gaussian distribution that reduces the hyperparameter search space. Furthermore, the best-performing Mod-DeepESN models, *Wide* and *Wide+Layered*, incorporate distinct, non-coinciding channels from input layer to output layer, and consistently outperform the other topologies.

Conclusions

In this work, we have shown the efficacy of a deep and modular echo state network architecture with varying topology. Prediction error on multiple datasets is reduced by isolating the integration of dynamical input states to disparate computational pathways, and by incorporating IP in a pre-training phase. Xavier initialization of weights decreases the complexity of hyperparameter tuning and performs well consistently across all proposed topologies. By combining these mechanisms with

a genetic algorithm, the Mod-DeepESN outperforms several baseline models on non-trivial time series forecasting tasks.

Acknowledgments

The authors would like to thank the members of the RIT Neuromorphic AI Lab for their valuable feedback on this work.

References

- Bäck, T., Fogel, D. B., & Michalewicz, Z. (2000). *Evolutionary computation 1: Basic algorithms and operators* (Vol. 1). CRC press.
- Butcher, J. B., Verstraeten, D., Schrauwen, B., Day, C., & Haycock, P. (2013). Reservoir computing and extreme learning machines for non-linear time-series data analysis. *Neural networks*, 38, 76–89.
- Fortin, F.-A., De Rainville, F.-M., Gardner, M.-A., Parizeau, M., & Gagné, C. (2012, jul). DEAP: Evolutionary algorithms made easy. *Journal of Machine Learning Research*, 13, 2171–2175.
- Gallicchio, C., & Micheli, A. (2011). Architectural and markovian factors of echo state networks. *Neural Networks*, 24(5), 440–456.
- Gallicchio, C., Micheli, A., & Pedrelli, L. (2017). Deep reservoir computing: a critical experimental analysis. *Neurocomputing*, 268, 87–99.
- Glorot, X., & Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics* (pp. 249–256).
- Jaeger, H., Lukoševičius, M., Popovici, D., & Siewert, U. (2007). Optimization and applications of echo state networks with leaky-integrator neurons. *Neural networks*, 20(3), 335–352.
- Jones, E., Oliphant, T., Peterson, P., et al. (2001). *SciPy: Open source scientific tools for Python*. Retrieved from <http://www.scipy.org/>
- Ma, Q., Shen, L., & Cottrell, G. W. (2017). Deep-esn: A multiple projection-encoding hierarchical reservoir computing framework. *arXiv preprint arXiv:1711.05255*.
- Malik, Z., Hussain, A., & Wu, Q. M. J. (2016, 06). Multilayered echo state machine: A novel architecture and algorithm. , 47, 1-14.
- Oliphant, T. E. (2006). *A guide to numpy* (Vol. 1). Trelgol Publishing USA.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Schrauwen, B., Wardermann, M., Verstraeten, D., Steil, J. J., & Stroobandt, D. (2008). Improving reservoirs using intrinsic plasticity. *Neurocomputing*, 71(7-9), 1159–1171.
- Soures, N., Hays, L., & Kudithipudi, D. (2017, May). Robustness of a memristor based liquid state machine. In *2017 international joint conference on neural networks (ijcnn)* (p. 2414-2420).
- Time Series Data Library. (1981–1990). *Daily minimum temperatures in melbourne, australia*. Retrieved from <https://datamarket.com/data/set/2324/daily-minimum-temperatures-in-melbourne-australia-1981-1990>
- Yamazaki, T., & Tanaka, S. (2007). The cerebellum as a liquid state machine. *Neural Networks*, 20(3), 290 - 297.

¹Leaky-integrator ESN (Jaeger et al., 2007).

² ϕ -ESN (Gallicchio & Micheli, 2011).

³R²SP (Butcher et al., 2013).

⁴Multi-layered echo state machine (MESM) (Malik et al., 2016).

⁵Deep-ESN (Ma et al., 2017).