Sparsity in Reservoir Computing Neural Networks

Claudio Gallicchio
Department of Computer Science
University of Pisa
Pisa, Italy
gallicch@di.unipi.it

Abstract—Reservoir Computing (RC) is a well-known strategy for designing Recurrent Neural Networks featured by striking efficiency of training. The crucial aspect of RC is to properly instantiate the hidden recurrent layer that serves as dynamical memory to the system. In this respect, the common recipe is to create a pool of randomly and *sparsely* connected recurrent neurons. While the aspect of sparsity in the design of RC systems has been debated in the literature, it is nowadays understood mainly as a way to enhance the efficiency of computation, exploiting sparse matrix operations.

In this paper, we empirically investigate the role of sparsity in RC network design under the perspective of the richness of the developed temporal representations. We analyze both sparsity in the recurrent connections, and in the connections from the input to the reservoir. Our results point out that sparsity, in particular in input-reservoir connections, has a major role in developing internal temporal representations that have a longer short-term memory of past inputs and a higher dimension.

Index Terms—Reservoir Computing, Echo State Networks, Short-term Memory, Sparse Recurrent Neural Networks

I. INTRODUCTION

Recurrent Neural Networks (RNNs) [1] are a fundamental tool for adaptive processing of dynamically evolving information, with excellent performance in fields such as time-series forecasting [2], machine translation [3], speech and text processing [4], [5], just to mention a few. An increasing number of works are analyzing the role of sparsity in the design of trained (dynamical) neural networks systems, for example through pruning [6] or re-wiring [7] connections. The characterization emerging from these studies is that having sparse connections between neurons is not only advantageous in computational terms - as it enables fast sparse matrix computations - but can also be beneficial to obtain a better performance in practice. Moreover, in the context of neurobiologically-inspired information processing systems, a sparse degree of connectivity between neurons has been shown to improve the quality of the developed internal representations [8]. Interestingly, the optimal amount sparsity in the numerical simulations matched observed properties of cerebellum-like circuits.

Reservoir Computing (RC) neural networks [9]–[11] represent an intriguing development in the field of RNNs. In RC, the recurrent hidden layer of a RNN is left untrained after initialization subject to asymptotic stability conditions

of the corresponding dynamical system. As a result, learning is applied only to a simple readout component with striking advantages in terms of required training times compared to fully trained RNNs. Pushing the involved algorithms towards extreme simplicity and efficiency makes the RC approach very well suited for real-world application scenarios featured by (possibly severe) resource constraints, such as neuromorphic hardware implementations [12] or cyber-physical systems where the learning modules are embedded at the edge [13].

A typical strategy in the design of RC networks is to setup the recurrent layer in a sparse way. The initial intuition was that sparsity in the recurrent untrained layer could enable a decoupling of state variables and hence richer representations [14]. Successively, several authors pointed out empirical evidences contrary to the initial intuition (see, e.g., [10], [15], [16]). Currently, the sparse design of reservoirs is commonly understood mainly as a way to speedup state computations, without a practical effect on the resulting performance. However, the impact of sparsity on the performance of RC neural networks has been typically studied limited to the recurrent connections only. In this paper, we intend to shed more light on the role of sparsity in RC by extending the analysis to both recurrent and input connections. Specifically, we empirically show the effect of recurrent and input sparsity in reservoirs, evaluated by means of short-term memory capacity and effective dimension of the resulting state trajectories.

The rest of this paper is structured as follows. We introduce the basics of RC methodology in Section II, discussing initialization and sparsity of reservoirs. Then, in Section III we present the concepts of short-term memory capacity and effective reservoir dimension. Our experimental analysis is described in Section IV. Finally, in Section V we draw our conclusions and sketch possible developments.

II. RESERVOIR COMPUTING NEURAL NETWORKS

Here we give a brief description of the RC design methodology for RNNs, focusing on the Echo State Network (ESN) [11], [14] model.

An RC network is a neural information processing system that treats data in the form of (temporal) sequences. Architecturally, the neural network is composed by a hidden recurrent layer called *reservoir*, and an output layer called *readout*. Fig. 1 illustrates the building blocks of a typical RC network.

In what follows, we denote the number of reservoir neurons, i.e., the reservoir dimension, by N, and the state of the
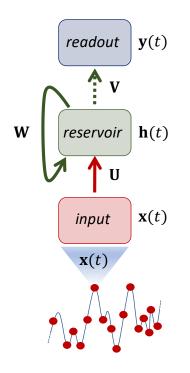
Fig. 1: Architecture of an RC neural network. The dotted arrow indicates trained connections.

reservoir system at time t by $\mathbf{h}(t) \in \mathbb{R}^N$. This state is evolved by following a state update equation:

$$\mathbf{h}(t) = \tanh(\mathbf{U}\mathbf{x}(t) + \mathbf{W}\mathbf{h}(t-1)), \qquad (1)$$

where $\mathbf{x}(t) \in \mathbb{R}^M$ is the M-dimensional input at time t, $\mathbf{U} \in \mathbb{R}^{N \times M}$ is the input weight matrix, modulating the influence of the external input on the current state, and $\mathbf{W} \in \mathbb{R}^{N \times N}$ is the recurrent weight matrix, which controls the impact of previous state on the current state. The state is typically set to a zero vector as initial condition, i.e. $\mathbf{h}(0) = \mathbf{0} \in \mathbb{R}^N$. Note that here we dropped from (1) the reference to bias terms to focus the analysis on the external stimulating input signal alone. Both weight matrices $\mathbf{U}$ and $\mathbf{W}$ remain untrained after initialization (see Section II-A).

The reservoir system is coupled with a linear readout layer that computes an L-dimensional output at each time step, i.e. $\mathbf{y}(t) \in \mathbb{R}^L$, as an affine transformation of the reservoir state:

$$\mathbf{y}(t) = \mathbf{V}\mathbf{h}(t) + \mathbf{b}, \qquad (2)$$

where $\mathbf{V} \in \mathbb{R}^{L \times N}$ is a readout weight matrix and $\mathbf{b} \in \mathbb{R}^L$ is a bias vector (that assumes a constant unitary input bias for the readout). The readout parameters are the only ones that undergo a training process, typically in closed-form fashion by using pseudo-inversion [9].

A. Initialization of Reservoirs

The fundamental characterization of RC neural networks is that all the reservoir parameters remain untrained after initialization. Such initialization is performed in agreement to asymptotic stability conditions expressed by the Echo State Property (ESP) [14], [17], [18], which essentially require to control the magnitude of the weights in $\mathbf{U}$ and $\mathbf{W}$. Usually, both the input weights in $\mathbf{U}$ and the recurrent weights in $\mathbf{W}$ are randomly drawn from a uniform distribution in $[-1, 1]$. After that, the elements in $\mathbf{U}$ are re-scaled by a factor ω_{in}, which takes the role of input scaling. The weights in $\mathbf{W}$ are re-scaled to control the largest absolute eigenvalue, i.e., the spectral radius ρ, typically to a value smaller than 1 [14].

The design strategy of the reservoir topology (i.e., the way in which the reservoir neurons are connected among each other) has been subject of several studies in literature (see, e.g., [19], [20]). While some of the proposed reservoir organizations can be beneficial in specific application circumstances, a random and sparse topological organization of the reservoir is the architecture of choice in general cases. This is the focus of our analysis in this paper.

Making the connections among reservoir neurons sparse has the fundamental practical advantage to reduce the cost of state update operations in (1). Actually, for densely connected reservoirs (and assuming $N >> M$) the cost of state updating scales as $\mathcal{O}(N^2)$, i.e. quadratically with the reservoir size. A first approach to make the reservoir sparsely connected would be to impose a (small) fixed percentage, say C, of non-zero weights in the involved weight matrices matrices. Although reducing the running times in practice for smaller reservoirs, this approach would asymptotically scale as $\mathcal{O}(N^2 C/100)$, hence still quadratically with the reservoir size. A more effective approach, which is adopted in this paper, is to fix the number, say χ_R, of incoming recurrent connections for each reservoir unit. This indeed makes the state update cost as small as $\mathcal{O}(N \chi_R)$, i.e. scaling only linearly with the number of neurons in the reservoir. A similar strategy can be adapted for the setup of the input connections. In this case, to ensure that each input dimension is actually forwarded to the reservoir, we fix the number of outgoing connections from each input units, denoted as χ_I. The sparse architectural reservoir setup used in this paper is exemplified in Figure 2. Notice that in this case, every row of $\mathbf{W}$ has exactly χ_R non-zero values, and every column of $\mathbf{U}$ has exactly χ_I non-zero elements, with both χ_R and χ_I being not greater than N.

III. SHORT-TERM MEMORY AND EFFECTIVE RESERVOIR SPACE DIMENSION

The role of the recurrent reservoir system is to embed the input time-series into an internal "state" representation, given by the activation of the reservoir neurons over time. Here we analyze the quality of such internal reservoir representation by quantifying its short-term memory and effective dimension.

Short-term Memory Capacity (MC) [21] tests the ability of a recurrent neural system to reconstruct its driving input time-series from the transient state dynamics. More in detail, the reservoir is driven by a uni-dimensional time-series, $x(t), t = 1, 2, \ldots$, and different readout units are trained to recall progressively delayed versions of the input. I.e., the i-th
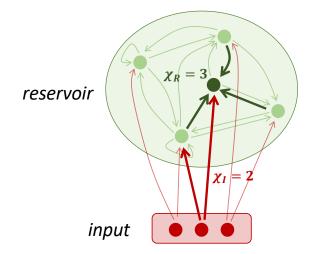
Fig. 2: Illustration of sparsity in input to reservoir and recurrent reservoir connections. χ_R indicates the number of incoming recurrent connections for each reservoir unit. χ_I indicates the number of outgoing connections from each input unit.

readout unit $y_i(t)$ should approximate $x(t-i)$. The MC of an RC network is then quantified as follows:

$$MC = \sum_{i=1}^{\infty} \frac{cov^2(x(t-i), y_i(t))}{\sigma^2(x(t-i))\sigma^2(y_i(t))},\qquad(3)$$

i.e., as the sum of squared correlation coefficients of the delayed input and reconstructed signals.

Effective Dimension (N_{eff}) [8], [22] is a measure of the number of orthogonal directions in the neuronal system's state trajectory over time. While the evolution of the reservoir system in (1) is described by an N-dimensional state vector $\mathbf{h}(t)$, the actual reservoir trajectory lies into a lower-dimensional manifold whose dimension can be quantified as follows:

$$N_{eff} = \frac{(\sum_{i=1}^{N} \lambda_i)^2}{\sum_{i=1}^{N} \lambda_i^2},\qquad(4)$$

where λ_i, $i = 1, 2, \ldots, N$, denote the eigenvalues of the covariance matrix of the reservoir state activation over time. When measured for a reservoir under the driving influence of an external time-series, (4) gives an estimate of the number of directions of reservoir state variability that are (linearly) uncorrelated along the observed trajectory.

IV. EXPERIMENTAL ANALYSIS

We measured the short-term memory (MC) and the effective reservoir dimension (N_{eff}) introduced in Section III for RC networks varying the amount of recurrent and input connections. Our experimental settings are described in Section IV-A, while the results are reported in Section IV-B.

A. Settings

We used a uni-dimensional signal as driving input for the reservoir (i.e., $M = 1$). To maximally test the intrinsic quality of reservoir representations, we used iid randomly sampled inputs $x(t)$ from a uniform distribution (in $[-0.8, 0.8]$). The length of the generated input time-series was 6000, and the number of reservoir neurons was fixed to $N = 100$. To compute MC, we used the first 5000 time-steps as training set[1], using the remaining 1000 time-steps to assess the MC score. The total number of delays used for the computation of (3) was 200, which is in practice sufficient to account for all the non-negligible contributions for 100-dimensional reservoirs. The last 1000 time-steps of the dataset were also used to compute the effective reservoir dimension N_{eff} (see (4)).

In our experiments, we used RC networks with spectral radius $\rho = 0.9$ and input scaling $\omega_{in} = 1$. While this setup is of common use in RC practice, we also ran preliminary experiments with other choices of these hyper-parameters, finding that the outcomes are not qualitatively different. We varied both the number of recurrent connections (χ_R) and of input connections (χ_I) from 1 to 100 (with step of 1). For each configuration we averaged the results over 50 reservoir realizations.

B. Results

The achieved values of MC and N_{eff} in correspondence of the possible sparsity settings (values of χ_R and χ_I) are shown in Fig. 3. We can draw two major observations from the results. First, the number of input connections has a decisive impact on both the short-term memory and the effective reservoir dimension of the networks. Indeed, maximally sparse input connections, with $\chi_I = 1$, achieved the highest performances. Interestingly, simply propagating the input to all the reservoir neurons degrades the performance sensibly. Second, the role of sparsity in recurrent connections seems to be much less important. In fact, the trend in Fig. 3 indicates that for a given input connectivity, the achieved results are not much sensible to the exact number of recurrent connections (after a minimum number has been exceeded).

The results are further detailed in Fig. 4, which shows the best result for each choice of input (resp. recurrent) connectivity in Fig. 4(a) (resp. Fig. 4(b)), as well as the results achieved for maximally sparse input connectivity, i.e. for $\chi_I = 1$, in Fig. 4(c). Figs. 4(a)-(b) confirm the already observed trends. On the one hand the performance of the RC networks tends to deteriorate for less sparse input weight matrices. On the other hand, a modest number of recurrent connections is already sufficient to achieve a performance not far from the highest possible one. For RC networks with $\chi_I = 1$ (Fig. 4(c)), both MC and N_{eff} saturate for fairly small values of χ_R, without appreciable differences for settings with more than 20 recurrent connections per reservoir neuron.

[1]We used pseudo-inversion to train the readout, discarding the first 1000 time-steps as initial transient.
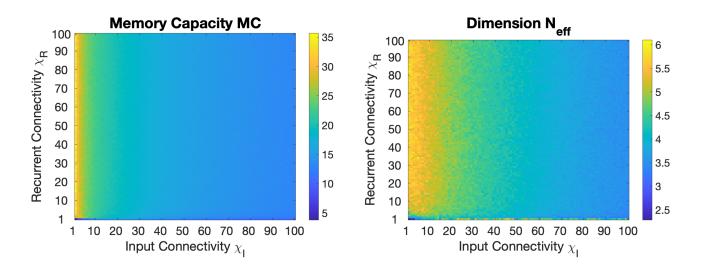
Fig. 3: Short-term Memory Capacity (MC) and effective dimension (N_{eff}) of RC networks. Results corresponds to $N = 100$ reservoir neurons, spectral radius $\rho = 0.9$, and input scaling $\omega_{in} = 1$. Recurrent (χ_R) and input (χ_I) connectivity varied from 1 to 100 with step 1. For each of the 10000 configurations the results are averaged over a number of 50 reservoir realizations.
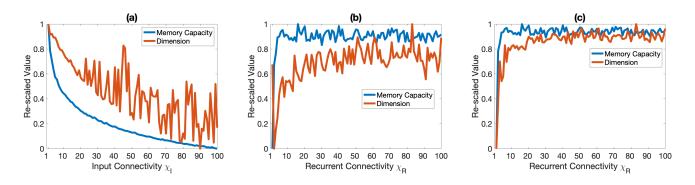

Fig. 4: Short-term Memory Capacity (MC) and effective dimension (N_{eff}) of 100 units RC networks, detailed for: **(a)** best results for increasing input connectivity; **(b)** best results for increasing recurrent connectivity; **(c)** results for maximally sparse input connections ($\chi_I = 1$) and increasing recurrent connectivity. Results are re-scaled to $[0, 1]$.

V. CONCLUSIONS

We have empirically analyzed the performance of RC neural networks in relation to sparsity of input and recurrent connections. Our results indicate that under commonly used reservoir configurations, the number of non-zero connections can play a decisive role in determining the richness of the developed representations. In particular, while a modest number of recurrent connections is already sufficient to achieve good performance, we found that maximally sparse input to reservoir connections lead to the best results both in terms of short-term memory and in terms of effective dimension of the state manifold. Overall, our analysis points out a simple rule of thumb for shaping reservoir weight matrices in case of uni-dimensional driving time-series: (i) connect the input to just one reservoir neuron, and (ii) set a small number of incoming recurrent connections ($\approx 20\%$) for each reservoir neuron.

The study presented in this paper can be seen as preparatory to opening further and deeper lines of research. First of all, the role of sparsity can be investigated in synergy with *structured* (rather than random) recurrent reservoir topologies, such as those based on cyclic [20] or small-world [23] connections. Similarly, the study can be extended towards *deep* RC neural networks [24], [25], where multiple reservoir layers are connected in a pipeline. In this case, the sparsity of input connections for higher layers has the even more intriguing role of modulating the extent of signal propagation between consecutive internal representations. Neuromorphic hardware implementations [12], [26]–[28] of deep recurrent neural systems are an important example of a domain where such insights can be capitalized in practice. Under a broader perspective, and outside the RC world, the analysis presented here pointed out that a sparse setting of RNN connections brings advantages even before learning of the non-zero connections. How these architectural advantages can be further exploited by (supervised or unsupervised) *training* is another exciting open research question.

REFERENCES

[1] J. F. Kolen and S. C. Kremer, *A field guide to dynamical recurrent networks*. John Wiley & Sons, 2001.

[2] N. Laptev, J. Yosinski, L. E. Li, and S. Smyl, "Time-series extreme event forecasting with neural networks at uber," in *International Conference on Machine Learning*, vol. 34, 2017, pp. 1–5.

[3] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Advances in neural information processing systems*, 2014, pp. 3104–3112.

[4] A. Graves, A.-r. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *2013 IEEE international conference on acoustics, speech and signal processing*. IEEE, 2013, pp. 6645–6649.

[5] R. Nallapati, B. Zhou, C. Gulcehre, B. Xiang *et al.*, "Abstractive text summarization using sequence-to-sequence rnns and beyond," *arXiv preprint arXiv:1602.06023*, 2016.

[6] S. Narang, E. Elsen, G. Diamos, and S. Sengupta, "Exploring sparsity in recurrent neural networks," *ICLR 2017. arXiv preprint arXiv:1704.05119*, 2017.

[7] G. Bellec, D. Kappel, W. Maass, and R. Legenstein, "Deep rewiring: Training very sparse deep networks," *ICLR 2018. arXiv preprint arXiv:1711.05136*, 2018.

[8] A. Litwin-Kumar, K. D. Harris, R. Axel, H. Sompolinsky, and L. Abbott, "Optimal degrees of synaptic connectivity," *Neuron*, vol. 93, no. 5, pp. 1153–1164, 2017.

[9] M. Lukoševičius and H. Jaeger, "Reservoir computing approaches to recurrent neural network training," *Computer Science Review*, vol. 3, no. 3, pp. 127–149, 2009.

[10] B. Schrauwen, D. Verstraeten, and J. Van Campenhout, "An overview of reservoir computing: theory, applications and implementations," in *Proceedings of the 15th european symposium on artificial neural networks. p. 471-482 2007*, 2007, pp. 471–482.

[11] H. Jaeger and H. Haas, "Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication," *science*, vol. 304, no. 5667, pp. 78–80, 2004.

[12] L. Larger, M. C. Soriano, D. Brunner, L. Appeltant, J. M. Gutiérrez, L. Pesquera, C. R. Mirasso, and I. Fischer, "Photonic information processing beyond turing: an optoelectronic implementation of reservoir computing," *Optics express*, vol. 20, no. 3, pp. 3241–3249, 2012.

[13] D. Bacciu, P. Barsocchi, S. Chessa, C. Gallicchio, and A. Micheli, "An experimental characterization of reservoir computing in ambient assisted living applications," *Neural Computing and Applications*, vol. 24, no. 6, pp. 1451–1464, 2014.

[14] H. Jaeger, "The echo state approach to analysing and training recurrent neural networks-with an erratum note," *Bonn, Germany: German National Research Center for Information Technology GMD Technical Report*, vol. 148, no. 34, p. 13, 2001.

[15] Y. Xue, L. Yang, and S. Haykin, "Decoupled echo state networks with lateral inhibition," *Neural Networks*, vol. 20, no. 3, pp. 365–376, 2007.

[16] C. Gallicchio and A. Micheli, "Architectural and markovian factors of echo state networks," *Neural Networks*, vol. 24, no. 5, pp. 440–456, 2011.

[17] I. B. Yildiz, H. Jaeger, and S. J. Kiebel, "Re-visiting the echo state property," *Neural networks*, vol. 35, pp. 1–9, 2012.

[18] C. Gallicchio, "Chasing the echo state property," in *Proceedings of ESANN*, 2019, pp. 667–672.

[19] T. Strauss, W. Wustlich, and R. Labahn, "Design strategies for weight matrices of echo state networks," *Neural computation*, vol. 24, no. 12, pp. 3246–3276, 2012.

[20] A. Rodan and P. Tino, "Minimum complexity echo state network," *IEEE transactions on neural networks*, vol. 22, no. 1, pp. 131–144, 2010.

[21] H. Jaeger, *Short term memory in echo state networks*. GMD-Forschungszentrum Informationstechnik, 2001, vol. 5.

[22] L. F. Abbott, K. Rajan, and H. Sompolinsky, "Interactions between intrinsic and stimulus-evoked activity in recurrent neural networks," *The dynamic brain: an exploration of neuronal variability and its functional significance*, pp. 1–16, 2011.

[23] Y. Kawai, J. Park, and M. Asada, "A small-world topology enhances the echo state property and signal propagation in reservoir computing," *Neural Networks*, vol. 112, pp. 15–23, 2019.

[24] C. Gallicchio, A. Micheli, and L. Pedrelli, "Deep reservoir computing: A critical experimental analysis," *Neurocomputing*, vol. 268, pp. 87–99, 2017.

[25] ——, "Design of deep echo state networks," *Neural Networks*, vol. 108, pp. 33–47, 2018.

[26] J. Moughames, X. Porte, M. Thiel, G. Ulliac, M. Jacquot, L. Larger, M. Kadic, and D. Brunner, "Three dimensional waveguide-interconnects for scalable integration of photonic neural networks," *arXiv preprint arXiv:1912.08203*, 2019.

[27] M. Freiberger, P. Bienstman, and J. Dambre, "Towards deep physical reservoir computing through automatic task decomposition and mapping," *arXiv preprint arXiv:1910.13332*, 2019.

[28] J. Partzsch and R. Schuffny, "Analyzing the scaling of connectivity in neuromorphic hardware and in models of neural networks," *IEEE transactions on neural networks*, vol. 22, no. 6, pp. 919–935, 2011.