

Memory capacity of nonlinear recurrent networks: Is it informative?

Giovanni Ballarin¹, Lyudmila Grigoryeva^{1,2}, Juan-Pablo Ortega³

Abstract

The total memory capacity (MC) of linear recurrent neural networks (RNNs) has been proven to be equal to the rank of the corresponding Kalman controllability matrix, and it is almost surely maximal for connectivity and input weight matrices drawn from regular distributions. This fact questions the usefulness of this metric in distinguishing the performance of linear RNNs in the processing of stochastic signals. This note shows that the MC of random nonlinear RNNs yields arbitrary values within established upper and lower bounds depending just on the input process scale. This confirms that the existing definition of MC in linear and nonlinear cases has no practical value.

Keywords: memory capacity, nonlinear recurrent neural networks, reservoir computing, echo state networks.

1 The problem

Main goal. The main focus of this note is the notion of total memory capacity ((T)MC), which has been used for decades for measuring the ability of recurrent neural networks (RNNs) to reproduce past inputs. The results in [Ball 24] for linear (random) RNNs show the practical drawbacks of this metric at the time of comparing different linear RNN architectures. This note shows that these negative conclusions also hold for the MC of nonlinear RNNs. The MC is demonstrated to be not invariant with respect to the scale of the inputs, which can be tuned so that the corresponding MC equals arbitrary values within the bounds known in the literature. This fact has not been reported earlier, begs the question about how valuable some of the existing contributions based on the MC are, and signals the need for other memory capacity notions inherently immune to the reported ambiguity.

Context and motivation. Recurrent neural networks (RNNs) are a popular architecture class for modeling data with time dependence or sequential structure and have been widely studied [Elma 90, Hoch 97, Pasc 13]. Different measures of *memory capacity* (MC) for RNNs have been proposed and discussed in the literature [Havi 19, Li 21], which aim at assessing the memory properties of a certain model via measuring its capability to reconstruct past inputs from the model's current states. [Jaeg 02] introduced a definition of memory capacity that was originally adapted to a specific class of RNN models with random internal weights, namely, the so-called echo state networks (ESNs) [Matt 94, Jaeg 04], which are a popular family of RNNs in the reservoir computing (RC) literature. These RNN architectures can be seen as, in general, nonlinear state-space models with parameters of the state dynamics *randomly*

¹Mathematics and Statistics Division, University of St. Gallen, Rosenbergstrasse 22, CH-9000 St. Gallen, Switzerland. Giovanni.Ballarini@unisg.ch, Lyudmila.Grigoryeva@unisg.ch

²Honorary Associate Professor, Department of Statistics, University of Warwick, Coventry CV4 7AL, UK. Lyudmila.Grigoryeva@warwick.ac.uk

³Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, 21 Nanyang Link, Singapore 637371. Juan-Pablo.Ortega@ntu.edu.sg

sampled from fixed regular probability distributions and the parameters of the observation equation subject to the training.

In this note, we consider the recurrent neural networks that admit the nonlinear state-space representation given by

$$\mathbf{x}_t = F(\mathbf{x}_{t-1}, z_t) := \varphi(A\mathbf{x}_{t-1} + \mathbf{C}z_t + \boldsymbol{\xi}), \quad (1.1)$$

$$\mathbf{y}_t := h(\mathbf{x}_t) = W^\top \mathbf{x}_t, \quad (1.2)$$

for $t \in \mathbb{Z}_-$, where $F : \mathbb{R}^N \times \mathbb{R} \rightarrow \mathbb{R}^N$ is called the state map, $h : \mathbb{R}^N \rightarrow \mathbb{R}^m$ is the readout map, $\mathbf{x}_t \in \mathbb{R}^N$ are the states, $z_t \in \mathbb{R}$ are the inputs, $\mathbf{y}_t \in \mathbb{R}^m$ are the outputs, and $m, N \in \mathbb{N}$. The map $\varphi : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is obtained by componentwise application of an activation function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$. Whenever this map is the identity, (1.1)-(1.2) correspond to the linear ESN model studied in [Ball 24]. In this note, we assume φ to be \mathcal{C}^1 -continuous, bounded, and non-constant. The fixed reservoir (connectivity) matrix $A \in \mathbb{M}_N$, the input mask $\mathbf{C} \in \mathbb{R}^N$, and the input shift $\boldsymbol{\xi} \in \mathbb{R}^N$ are randomly sampled independently of inputs, while $W \in \mathbb{M}_{N,m}$ is estimated via minimizing the empirical risk associated to a chosen loss function. Additionally, we impose that the reservoir matrix is normalized such that $\|A\|_2 < 1$, with $\|\cdot\|_2$ the spectral norm, which is a sufficient condition for the existence and uniqueness of the solutions of (1.1)-(1.2) (the so-called Echo State Property (ESP)).

Over the last two decades, memory capacity has received much attention, especially in the RC literature. The idea behind existing classical definitions of MC is that the ability of the network states to contain information about previous independent and identically distributed (i.i.d.) inputs can be quantified by the correlation between the outputs of the network and its past inputs and underpins existing contributions for linear [Herm 10, Damb 12, Bara 14, Coui 16], echo state shallow [Whit 04, Fark 16, Verz 19], and deep architectures [Gall 18]. Extensions for temporally dependent inputs have also been recently proposed [Damb 12, Grig 14, Grig 16a, Grig 16b, Marz 17, Gono 20].

Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space on which random variables are defined and consider a variance-stationary stochastic input $\mathbf{z} : \Omega \rightarrow \mathbb{R}^{\mathbb{Z}_-}$ and the associated covariance-stationary output state process $\mathbf{x} : \Omega \rightarrow (\mathbb{R}^N)^{\mathbb{Z}_-}$. Following [Jaeg 02, Gono 20, Ball 24], the τ -lag memory capacity of the state-space system with respect to \mathbf{z} , with $\tau \in \mathbb{N}$, is defined as

$$\text{MC}_\tau := 1 - \frac{1}{\text{Var}(z_t)} \min_{W \in \mathbb{M}_{N,m}} \mathbb{E} [(z_{t-\tau} - W^\top \mathbf{x}_t)^2]. \quad (1.3)$$

Under the assumption that the joint process $(\mathbf{x}_t, z_{t-\tau})_{t \in \mathbb{Z}_-}$ is also covariance-stationary and $\Gamma_{\mathbf{x}} := \text{Cov}(\mathbf{x}_t, \mathbf{x}_t)$ is nonsingular, MC_τ has the following closed-form expression (see [Gono 20])

$$\text{MC}_\tau = \frac{\text{Cov}(z_{t-\tau}, \mathbf{x}_t) \Gamma_{\mathbf{x}}^{-1} \text{Cov}(\mathbf{x}_t, z_{t-\tau})}{\text{Var}(z_t)}. \quad (1.4)$$

The *total memory capacity* (MC) of an ESN is then given by

$$\text{MC} := \sum_{\tau=0}^{\infty} \text{MC}_\tau \quad (1.5)$$

and a fundamental result [Damb 12, Gono 20] is that for i.i.d. inputs it necessarily holds that

$$1 \leq \text{MC} \leq N \quad (1.6)$$

for any linear or nonlinear ESN model.

Contribution. We complement the results of [Ball 24], where the memory properties of *linear* ESN models were characterized, and its invariance under specific choices of the ESN architecture was proved. In this note, we demonstrate that a direct implication of defining memory as in (1.3) and (1.5) is that

the memory of a given *nonlinear* ESN model can be tuned to *attain both sides of the memory bounds* (1.6) just by changing the scale of the input. More specifically, we demonstrate that the change in the variance of the inputs may lead to different capacities. This implies that the MC is actually not only a function of the RNN architecture and is strictly *task-specific*, meaning that its properties also rely on the nature of the input process, and hence, MC serves as a poor measure to discriminate between the signal processing properties of different RNN architectures.

Notation. \mathbb{R}^+ denotes the set of all nonnegative real numbers. For an m by n \mathbb{R} -valued matrix A we write $\|A\|_2$ to denote its spectral norm. For any m by n \mathbb{R} -valued matrix A over \mathbb{R} , we use the notation $\lfloor A \rfloor$ for its minimal absolute entry that is $\lfloor A \rfloor := \min_{1 \leq i \leq m, 1 \leq j \leq n} |A_{ij}|$. The infinity norm of a vector $\mathbf{x} \in \mathbb{R}^N$ is given by $\|\mathbf{x}\|_\infty = \max_{1 \leq i \leq N} |x_i|$. Let $\mathbf{X} : \Omega \rightarrow B$ be a random variable with $(B, \|\cdot\|_B)$ a normed space endowed with a σ -algebra (for example, its Borel σ -algebra). We write $\text{ess sup}_{\omega \in \Omega} \{\|\mathbf{X}(\omega)\|_B\} = \inf \{b \in \mathbb{R}^+ \mid \|\mathbf{X}\|_B \leq b \text{ almost surely}\}$ and $\text{ess inf}_{\omega \in \Omega} \{\|\mathbf{X}(\omega)\|_B\} = \sup \{b \in \mathbb{R}^+ \mid \|\mathbf{X}\|_B \geq b \text{ almost surely}\}$. A discrete-time stochastic process is a map of the type:

$$\begin{aligned} \mathbf{z} : \mathbb{Z} \times \Omega &\longrightarrow \mathbb{R}^n \\ (t, \omega) &\longmapsto \mathbf{z}_t(\omega), \end{aligned}$$

such that, for each $t \in \mathbb{Z}$, the assignment $\mathbf{z}_t : \Omega \rightarrow \mathbb{R}^n$ is a random variable. For each $\omega \in \Omega$, we will denote by $\mathbf{z}(\omega) := \{\mathbf{z}_t(\omega) \in \mathbb{R}^n \mid t \in \mathbb{Z}\}$ the realization or the sample path of the process \mathbf{z} . We denote $\|\mathbf{z}\|_{L^\infty} := \text{ess sup}_{\omega \in \Omega} \{\sup_{t \in \mathbb{Z}} \{\|\mathbf{z}_t(\omega)\|_\infty\}\}$ and $\lfloor \mathbf{z} \rfloor_{L^\infty} := \text{ess inf}_{\omega \in \Omega} \{\inf_{t \in \mathbb{Z}} \{\min_{i \in \{1, \dots, n\}} |(\mathbf{z}_t(\omega))_i|\}\}$.

2 The result

The core idea of our main result is to show that the memory capacity of nonlinear ESNs is driven by the relationship between the input mask \mathbf{C} and the input process $\{z_t\}_{t \in \mathbb{Z}_-}$. Therefore, the properties of $\{z_t\}_{t \in \mathbb{Z}_-}$, and, in particular, its variance, play a key role in our demonstration. We consider the system (1.1)-(1.2), where, for simplicity, we set $\boldsymbol{\xi} = \mathbf{0}$ and assume a particular type of activation function φ , which conforms with most of the common choices of RNN architectures. In the remainder of the note, without loss of generality, we assume that A and \mathbf{C} are sampled from some regular distributions and kept fixed. Our goal is to show that for any sampled ESN, one can find appropriate input scalings that make its memory take any value between its theoretical maximum and minimum in (1.6).

Definition 2.1 (Sigmoid activation function). We say that an activation function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is a *sigmoid* function whenever there exists some $a \in \mathbb{R}$ such that $|\varphi(x)| \leq a$ (a -bounded), for all $x \in \mathbb{R}$, $\varphi \in \mathcal{C}^1(\mathbb{R})$ with $\varphi'(x) \geq 0$ for all $x \in \mathbb{R}$.

Definition 2.2 (Piecewise smooth sigmoid). A sigmoid activation function φ is said to be a δ -linear D -saturating sigmoid, with $\delta, D \in \mathbb{R}^+$, $\delta < D$, if the following conditions hold:

- (i) $\varphi(x) = x$ for all $x \in (-\delta, \delta)$;
- (ii) $\varphi(x) = -1$ for all $x \in (-\infty, -D)$ and $\varphi(x) = 1$ for all $x \in (D, +\infty)$.

This choice of activation function allows us to treat inputs of significantly small or large magnitudes. The conditions (i)-(ii) in Definition 2.2, which specify the linear behavior of φ in the δ -neighborhood of 0 and for inputs of absolute value above some D , also simplify the analysis of the situations which are typical for the numerical evaluation of activation functions. More precisely, we notice that due to the finite numerical precision of computers, the computation of φ for close to 0 and very large values in the domain renders the behavior compatible with Definition 2.2.

Attaining the lower MC bound. We now prove that such networks, when applied to i.i.d. inputs with sufficiently high variance, attain minimal memory, namely the lower bound in (1.6). Consider a process $\{z_t\}_{t \in \mathbb{Z}_-}$ with $z_t = \sigma \zeta_t$ where $\{\zeta_t\}_{t \in \mathbb{Z}_-}$ is a sequence of i.i.d. Rademacher random variables, which take values -1 and 1 with equal probability, and $\sigma > 0$. We work with dense input mask vectors \mathbf{C} , which we formalize in the following assumption.

Assumption 1. Define $\underline{c} := \lfloor \mathbf{C} \rfloor$ as the minimum absolute entry of input mask \mathbf{C} and assume $\underline{c} > 0$.

Remark 2.3. In practice ESN models are often constructed by drawing the entries of the reservoir matrix A from *sparse* distributions. The distribution of the entries of input mask \mathbf{C} is more rarely chosen to be sparse. Assumption 1 immediately holds whenever the entries of \mathbf{C} are sampled from any continuous distribution. More generally, \mathbf{C} needs to be sampled from the law that has no atom at 0 or, equivalently, $\min_{1 \leq i \leq N} \mathbb{P}(C_i = 0) = 0$, which is not a restrictive assumption and is satisfied in most practical scenarios.

We start by observing that for any \mathbf{C} under Assumption 1 one can require that $\sigma > D/\underline{c}$. It is easy to see that in this case

$$\lfloor \mathbf{Cz} \rfloor_{L^\infty} = \lfloor \mathbf{C} \sigma \text{sign}(\zeta) \rfloor_{L^\infty} = \lfloor \mathbf{C} \rfloor \sigma \geq \sigma \underline{c} > D.$$

Additionally, one can use Definition 2.2 of the piecewise smooth sigmoid φ to show that for any $t \in \mathbb{Z}_-$

$$\varphi(\mathbf{C}z_t) = \varphi(\mathbf{C} \sigma \text{sign}(\zeta_t)) = \begin{cases} \varphi(-CD/\underline{c}) & \text{if } z_t = -\sigma, \\ \varphi(+CD/\underline{c}) & \text{if } z_t = +\sigma, \end{cases}$$

which leads to the following definition.

Definition 2.4 (Extreme states of the ESN). Given an ESN constructed using the piecewise smooth sigmoid activation φ , the reservoir matrix A , and the input mask vector \mathbf{C} that satisfies Assumption 1, let $\{\zeta_t\}_{t \in \mathbb{Z}_-}$ be a sequence of i.i.d. random variables with unit variance. Given a scaling $\sigma > D/\underline{c}$ and the associated inputs $\{z_t\}_{t \in \mathbb{Z}_-}$, $z_t = \sigma \zeta_t$, $t \in \mathbb{Z}_-$, for which

$$\varphi(\mathbf{C}z_t) = \begin{cases} \varphi(-CD/\underline{c}) =: \mathbf{x}_- & \text{if } z_t = -\sigma, \\ \varphi(+CD/\underline{c}) =: \mathbf{x}_+ & \text{if } z_t = +\sigma, \end{cases}$$

we call $\mathbf{x}_+, \mathbf{x}_- \in \{-1, +1\}^N$ *extreme states* of the ESN associated to the input $\{\zeta_t\}_{t \in \mathbb{Z}_-}$.

The following result shows that, whenever the ESN state dynamics are defined by (1.1)-(1.2), one can always choose a scaling σ of the input such that the ESN states are all extreme.

Lemma 2.5. Let $\mathbf{z} = \{z_t\}_{t \in \mathbb{Z}_-}$, $z_t \in \{-1, +1\}$, be a sequence of i.i.d. rescaled Rademacher random variable of standard deviation σ . Let $\mathbf{x} = \{\mathbf{x}_t\}_{t \in \mathbb{Z}_-}$, $\mathbf{x}_t \in \mathbb{R}^N$, be the states associated to inputs $\{z_t\}_{t \in \mathbb{Z}_-}$ of the ESN with D -saturating sigmoid and randomly sampled A and \mathbf{C} . If $\sigma > \frac{\sqrt{N}\|A\|_2 + D}{\underline{c}}$ with $\underline{c} = \lfloor \mathbf{C} \rfloor$, then for any $t \in \mathbb{Z}_-$

$$\mathbf{x}_t = \begin{cases} \mathbf{x}_-, & \text{if } z_t = -\sigma, \\ \mathbf{x}_+, & \text{if } z_t = +\sigma. \end{cases}$$

Proof. By construction, it holds that

$$\lfloor \mathbf{Cz} + \mathbf{Ax} \rfloor_{L^\infty} \geq \lfloor \mathbf{Cz} \rfloor_{L^\infty} - \lceil \mathbf{Ax} \rceil_{L^\infty} \geq \underline{c} \sigma - \sqrt{N} \|A\|_2.$$

Hence, whenever $\sigma > (\sqrt{N}\|A\|_2 + D)/\underline{c}$, then $\lfloor \mathbf{Cz} + \mathbf{Ax} \rfloor_{L^\infty} > D$. Thus, by Definition 2.2 when $z_t > 0$, $t \in \mathbb{Z}_-$, it follows that $\varphi(\mathbf{C}z_t + \mathbf{Ax}_t) = \mathbf{x}_+$ for any possible state \mathbf{x}_t . The same applies to the case $z_t < 0$, $t \in \mathbb{Z}_-$, with respect to \mathbf{x}_- . Therefore, whenever $\sigma > (\sqrt{N}\|A\|_2 + D)/\underline{c}$, the ESN state space consists of the set of extremal states almost surely. \square

Remark 2.6. Notice that although using a different bound that satisfies $\lceil \mathbf{Ax} \rceil_{L^\infty} \leq \bar{a}N$ with $\bar{a} = \|A\|_\infty$ one obtains a less sharp larger lower bound for σ , it is observed to produce more stable empirical illustrations.

Attaining the upper MC bound. We now show that the states of ESNs that are constructed using the i.i.d. inputs with sufficiently small variance are equivalent to those of a linear recurrent network (or linear ESN), which yields provably maximal memory. We use again the i.i.d. rescaled Rademacher sequence $\{z_t\}_{t \in \mathbb{Z}}$ of standard deviation σ and notice that, for $\bar{c} := \|\mathbf{C}\|_\infty$, it holds

$$[\mathbf{C}\mathbf{z}]_{L^\infty} \leq \|\mathbf{C}\|_\infty [\mathbf{z}]_{L^\infty} = \bar{c}\sigma.$$

Lemma 2.7. *Let $\{z_t\}_{t \in \mathbb{Z}}$ be a sequence of i.i.d. rescaled Rademacher random variables of standard deviation σ . Let $\{\mathbf{x}_t\}_{t \in \mathbb{Z}}$ be the states associated with inputs $\{z_t\}_{t \in \mathbb{Z}}$ of the nonlinear ESN with δ -linear sigmoid, A and \mathbf{C} drawn from random distributions such that $\|A\|_2 < 1$. If $\sigma < \frac{\delta(1-\|A\|_2)}{\bar{c}}$, with $\bar{c} = \|\mathbf{C}\|_\infty$, then $\{\mathbf{x}_t\}_{t \in \mathbb{Z}}$ are equivalent to the states of the linear ESN model with the state dynamics defined as*

$$\mathbf{x}_t := A\mathbf{x}_{t-1} + \mathbf{C}z_t.$$

Proof. We begin by bounding the maximal possible entry of the states of the ESN. We write that

$$\begin{aligned} [\mathbf{x}]_{L^\infty} &:= \operatorname{ess\,sup}_{\omega \in \Omega} \left\{ \sup_{t \in \mathbb{Z}_-} \{\|\mathbf{x}_t(\omega)\|_\infty\} \right\} = \operatorname{ess\,sup}_{\omega \in \Omega} \left\{ \sup_{t \in \mathbb{Z}_-} \left\{ \left\| \sum_{j=0}^{\infty} A^j \mathbf{C} z_{t-j}(\omega) \right\|_\infty \right\} \right\} \\ &\leq \operatorname{ess\,sup}_{\omega \in \Omega} \left\{ \sup_{t \in \mathbb{Z}_-} \left\{ \sum_{j=0}^{\infty} \|A^j \mathbf{C} z_{t-j}(\omega)\|_\infty \right\} \right\} \\ &\leq \operatorname{ess\,sup}_{\omega \in \Omega} \left\{ \sup_{t \in \mathbb{Z}_-} \left\{ \sum_{j=0}^{\infty} \|A\|_2^j \cdot \|\mathbf{C} z_{t-j}(\omega)\|_2 \right\} \right\} \\ &\leq \operatorname{ess\,sup}_{\omega \in \Omega} \left\{ \sup_{t \in \mathbb{Z}_-} \left\{ \sqrt{N} \sum_{j=0}^{\infty} \|A\|_2^j \cdot \|\mathbf{C} z_{t-j}(\omega)\|_\infty \right\} \right\} \\ &\leq \operatorname{ess\,sup}_{\omega \in \Omega} \left\{ \sup_{t \in \mathbb{Z}_-} \left\{ \sqrt{N} \sum_{j=0}^{\infty} \|A\|_2^j \cdot \bar{c}\sigma \right\} \right\} \leq \frac{\sqrt{N}}{1 - \|A\|_2} \bar{c}\sigma, \end{aligned}$$

where we used that $\|A\|_2 < 1$. Finally, notice that whenever $\sigma < \delta(1 - \|A\|_2)/(\sqrt{N}\bar{c})$, then necessarily $[\mathbf{x}]_{L^\infty} < \delta$ which yields that for the δ -linear sigmoid $\varphi(A\mathbf{x}_t + \mathbf{C}z_t) = A\mathbf{x}_t + \mathbf{C}z_t$ for all $t \in \mathbb{Z}_-$ as required. \square

Remark 2.8. Note that the proof of Lemma 2.7 can be generalized to work for any input sequence that is almost surely bounded. For the sake of simplicity and to be consistent with our proof of Lemma 2.5, here we have limited ourselves to the case of rescaled Rademacher processes. We also note that, in simulations, a looser bound, i.e. one where the \sqrt{N} factor is omitted and thus $\sigma < \delta(1 - \|A\|_2)/\bar{c}$, still yields robust empirical results.

Total Memory Saturation. We are now ready to state our main theoretical result regarding the total memory capacity of nonlinear ESN models.

Theorem 2.9. *Let $\{z_t\}_{t \in \mathbb{Z}}$ be a sequence of i.i.d. rescaled Rademacher random variables of standard deviation σ . Let a ESN model be defined as in (1.1)-(1.2). Let φ be a δ -linear D -saturating sigmoid activation, and let A and \mathbf{C} be drawn from random distributions. Assume that $\|A\|_2 < 1$ and that Assumption 1 holds. Let $\{\mathbf{x}_t\}_{t \in \mathbb{Z}}$ be the states associated with inputs $\{z_t\}_{t \in \mathbb{Z}}$ of the ESN. Then, there exist positive constants $\bar{\sigma}_{A,C}$ and $\underline{\sigma}_{A,C}$ such that*

(i) *If $\sigma > \bar{\sigma}_{A,C}$, then $\text{MC} = 1$, almost surely.*

(ii) If $\sigma < \underline{\sigma}_{A,C}$, then $\text{MC} = N$, almost surely.

Proof. (i) We use Lemma 2.5 to show that $\text{MC}_0 = 1$, while $\text{MC}_\tau = 0$ for all $\tau \geq 1$ whenever $\sigma > \bar{\sigma}_{A,C}$ with $\bar{\sigma}_{A,C} := \frac{\sqrt{N}\|A\|_2 + D}{\underline{c}}$, where $\underline{c} = \lfloor C \rfloor$. Without loss of generality, for $\tau = 0$ set $W_0 = \sigma \mathbf{e}_1$ if the first entry of \mathbf{x}_+ is 1, and $W_0 = -\sigma \mathbf{e}_1$ if $(\mathbf{x}_+)_1 = -1$, with $\mathbf{e}_1 \in \mathbb{R}^N$ the first canonical basis vector. By construction and Lemma 2.5, $W_0^\top \mathbf{x}_+ = \sigma$ and $W_0^\top \mathbf{x}_- = -\sigma$, and hence

$$\min_{W_0 \in \mathbb{R}^N} \mathbb{E} \left[(z_t - W_0^\top \mathbf{x}_t)^2 \right] = 0.$$

When $\tau > 1$, Lemma 2.5 implies that $z_{t-\tau}$ and $\mathbf{x}_t = \varphi(Cz_t)$ are independent random variables for any $t \in \mathbb{Z}_-$. Therefore, since $\mathbb{E}[z_{t-\tau}] = 0$,

$$\min_{W_\tau \in \mathbb{R}^N} \mathbb{E} \left[(z_{t-\tau} - W_\tau^\top \mathbf{x}_t)^2 \right] = \mathbb{E}[z_{t-\tau}^2] + \min_{W_\tau \in \mathbb{R}^N} \mathbb{E}[(W_\tau^\top \mathbf{x}_t)^2] = \sigma^2,$$

and hence $\text{MC}_\tau = 0$ for all $\tau \geq 1$. This shows that $\text{MC} = 1$ almost surely.

(ii) Let now

$$\underline{\sigma}_{A,C} = \frac{\delta(1 - \|A\|_2)}{\sqrt{N}\bar{c}},$$

where $\bar{c} = \|C\|_\infty$. Then, by Lemma 2.7, for $\sigma < \underline{\sigma}_{A,C}$ the nonlinear ESN with φ a δ -linear D -saturating sigmoid activation is equivalent to a linear ESN model. As proven in [Gono 20], linear ESNs have maximal memory almost surely, therefore $\text{MC} = N$ almost surely. \square

Simulation. To give an empirical illustration of Theorem 2.9, we provide in Figure 1 the results of a simulation aimed at computing the total memory capacity of different nonlinear ESN models with i.i.d. rescaled Rademacher inputs and dense input matrix.¹ These clearly show that, while for some types of reservoir matrices it is indeed possible to achieve practically maximal memory, as the variance of the input process increases, all models eventually drop to the lower bound of $\text{MC} = 1$.²

3 Conclusion

In this work, we have shown that without making any precise assumptions on the properties of the inputs of a state space system other than their i.i.d. character, it is *a priori* impossible to establish a nonlinear memory characterization that is more precise than just its boundedness within the range $1 \leq \text{MC} \leq N$. The definition of MC due to [Jaeg 02], which is often studied in the literature on reservoir computing methods, depends on the specific nature of the data. To be more precise, the variance of inputs was shown to directly control whether any fixed nonlinear ESN has maximal or minimal total memory capacity, based on whether it behaves approximately like a linear model or, rather, its state space consists only of two extremal states.

These results, taken together with those developed in [Ball 24] for linear recurrent networks, effectively spell out the limits of the MC concept. Our conclusion is that memory capacity should not be considered a useful metric to evaluate the inherent “qualities” or “capabilities” of an ESN or RNN model as it is merely a statistical quantity that may offer some insight into the application of a given model to a given input sequence.

Acknowledgments: GB thanks the Center for Doctoral Studies in Economics and the University of Mannheim, thanks to which a significant part of this work was developed over the course of doctoral

¹The simulation code is publicly available at <https://github.com/Learning-of-Dynamic-Processes/memorycapacity>.

²See [Ball 24] for an in-depth discussion of why, even in the linear ESN limit, it is generally not possible to obtain precise memory estimates by simulation due to intrinsic numerical issues.




Figure 1: Plot of estimated total memory capacity of nonlinear ESN models, $\widehat{MC}(\sigma)$, as a function of scaled Rademacher standard deviation, $\sigma \in [\underline{\sigma}_{A,C}, \bar{\sigma}_{A,C}]$. The horizontal axis shows the value of σ in logarithmic scale. Connectivity matrix $A = (A_{ij}) \in \mathbb{M}_N$ – with $N = 30$ and $\|A\|_2 = 0.95$ – sampled from (a) orthogonal ensemble (via Gram-Schmidt decomposition of entry-wise standard Gaussian matrix); (b) conditioning sparse Gaussian ensemble with sparsity degree 0.1 and conditioning number 0.7; (c) standard Gaussian ensemble. In all panels $c_i \sim \text{i.i.d. } \mathcal{N}(0, 1)$. For each value of σ , we use 10^5 Monte Carlo replications to estimate $\widehat{MC}(\sigma)$.

studies. GB also thanks the University of St. Gallen for the hospitality over the years. LG thanks the hospitality of Nanyang Technological University, where part of this work was completed. JPO thanks the hospitality of the University of St. Gallen, where part of this work was completed and acknowledges partial financial support from the School of Physical and Mathematical Sciences of the Nanyang Technological University, Singapore.

References

- [Ball 24] G. Ballarin, L. Grigoryeva, and J.-P. Ortega. “Memory of recurrent networks: Do we compute it right?”. *Journal of Machine Learning Research*, Vol. 25, No. 243, pp. 1–38, 2024.
- [Bara 14] P. Barancok and I. Farkas. “Memory capacity of input-driven echo state networks at the edge of chaos”. In: *Proceedings of the International Conference on Artificial Neural Networks (ICANN)*, pp. 41–48, 2014.
- [Coui 16] R. Couillet, G. Wainrib, H. Sevi, and H. T. Ali. “The asymptotic performance of linear echo state neural networks”. *Journal of Machine Learning Research*, Vol. 17, No. 178, pp. 1–35, 2016.
- [Damb 12] J. Dambre, D. Verstraeten, B. Schrauwen, and S. Massar. “Information processing capacity of dynamical systems”. *Scientific reports*, Vol. 2, No. 514, 2012.
- [Elma 90] J. E. Elman. “Finding structure in time”. *Cognitive Science*, Vol. 14, No. 2, pp. 179–211, 1990.
- [Fark 16] I. Farkas, R. Bosak, and P. Gergel. “Computational analysis of memory capacity in echo state networks”. *Neural Networks*, Vol. 83, pp. 109–120, 2016.
- [Gall 18] C. Gallicchio. “Short-term memory of Deep RNN”. 2018.
- [Gono 20] L. Gonon, L. Grigoryeva, and J.-P. Ortega. “Memory and forecasting capacities of nonlinear recurrent networks”. *Physica D*, Vol. 414, No. 132721, pp. 1–13., 2020.

- [Grig 14] L. Grigoryeva, J. Henriques, L. Larger, and J.-P. Ortega. “Stochastic time series forecasting using time-delay reservoir computers: performance and universality”. *Neural Networks*, Vol. 55, pp. 59–71, 2014.
- [Grig 16a] L. Grigoryeva, J. Henriques, L. Larger, and J.-P. Ortega. “Nonlinear memory capacity of parallel time-delay reservoir computers in the processing of multidimensional signals”. *Neural Computation*, Vol. 28, pp. 1411–1451, 2016.
- [Grig 16b] L. Grigoryeva, J. Henriques, and J.-P. Ortega. “Reservoir computing: information processing of stationary signals”. In: *Proceedings of the 19th IEEE International Conference on Computational Science and Engineering*, pp. 496–503, 2016.
- [Havi 19] D. Haviv, A. Rivkind, and O. Barak. “Understanding and controlling memory in recurrent neural networks”. In: *Proceedings of the 36th International Conference on Machine Learning*, pp. 2663–2671, 2019.
- [Herm 10] M. Hermans and B. Schrauwen. “Memory in linear recurrent neural networks in continuous time.”. *Neural Networks*, Vol. 23, No. 3, pp. 341–55, apr 2010.
- [Hoch 97] S. Hochreiter and J. Schmidhuber. “Long short-term memory”. *Neural Computation*, Vol. 9, No. 8, pp. 1735–1780, 1997.
- [Jaeg 02] H. Jaeger. “Short term memory in echo state networks”. *Fraunhofer Institute for Autonomous Intelligent Systems. Technical Report.*, Vol. 152, 2002.
- [Jaeg 04] 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.
- [Li 21] Z. Li, J. Han, W. E, and Q. Li. “On the curse of memory in recurrent neural networks: Approximation and optimization analysis”. In: *ICLR*, pp. 1–43, 2021.
- [Marz 17] S. Marzen. “Difference between memory and prediction in linear recurrent networks”. *Physical Review E*, Vol. 96, No. 3, pp. 1–7, 2017.
- [Matt 94] M. Matthews and G. Moschytz. “The identification of nonlinear discrete-time fading-memory systems using neural network models”. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, Vol. 41, No. 11, pp. 740–751, 1994.
- [Pasc 13] R. Pascanu, C. Gulcehre, K. Cho, and Y. Bengio. “How to construct deep recurrent neural networks”. *arXiv*, dec 2013.
- [Verz 19] P. Verzelli, C. Alippi, and L. Livi. “Echo State Networks with self-normalizing activations on the hyper-sphere”. *Scientific Reports*, Vol. 9, No. 1, p. 13887, dec 2019.
- [Whit 04] O. White, D. Lee, and H. Sompolinsky. “Short-term memory in orthogonal neural networks”. *Physical Review Letters*, Vol. 92, No. 14, p. 148102, apr 2004.