



HAL
open science

Evolutionary Echo State Network: evolving reservoirs in the Fourier space

Sebastian Basterrech, Gerardo Rubino

► **To cite this version:**

Sebastian Basterrech, Gerardo Rubino. Evolutionary Echo State Network: evolving reservoirs in the Fourier space. IJCNN 2022 - 2022 International Joint Conference on Neural Networks, Jul 2022, Padua, Italy. pp.1-8, 10.1109/IJCNN55064.2022.9892892 . hal-03898605

HAL Id: hal-03898605

<https://inria.hal.science/hal-03898605>

Submitted on 14 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Evolutionary Echo State Network: evolving reservoirs in the Fourier space

Sebastian Basterrech *Faculty of Electrical Engineering and Computer Science
VŠB-Technical University of Ostrava, Ostrava, Czech Republic*

Sebastian.Basterrech@vsb.cz

ORCID: 0000-0002-9172-0155

Gerardo Rubino *INRIA Rennes – Bretagne Atlantique
Rennes, France*

Gerardo.Rubino@inria.fr

ORCID: 0000-0002-1712-0477

Abstract

The Echo State Network (ESN) is a class of Recurrent Neural Network with a large number of hidden-hidden weights (in the so-called reservoir). Canonical ESN and its variations have recently received significant attention due to their remarkable success in the modeling of non-linear dynamical systems. The reservoir is randomly connected with fixed weights that don't change in the learning process. Only the weights from reservoir to output are trained. Since the reservoir is fixed during the training procedure, we may wonder if the computational power of the recurrent structure is fully harnessed.

In this article, we propose a new computational model of the ESN type, that represents the reservoir weights in the Fourier space and performs a fine-tuning of these weights applying genetic algorithms in the frequency domain. The main interest is that this procedure will work in a much smaller space compared to the classical ESN, thus providing a dimensionality reduction transformation of the initial method. The proposed technique allows us to exploit the benefits of the large recurrent structure avoiding the training problems of gradient-based method. We provide a detailed experimental study that demonstrates the good performances of our approach with well-known chaotic systems and real-world data.

Index Terms

Recurrent Neural Networks, Echo State Networks, Genetic Algorithms, Reservoir Computing, Fourier Transform.

I. INTRODUCTION

A Neural Network (NN) with a recurrent topology (RNN) is a powerful tool for modeling time series and solving machine learning problems with sequential data. Recurrences (circuits) are relevant for the computational power of neural architectures; thanks to them, hidden states are able to store more information from input streams. Numerous works explain the distinct advantages presented by RNNs [1]–[4]: good theoretical properties, biologically plausibility, success in several real-world applications. However, it is still hard to properly define optimal network architectures in RNNs [2]. The difficulties are even higher when the network is large and the input signal contains long-term data dependencies [1]. In summary, the principal well-identified problems are: not guaranteed convergence, long training times, and difficulties in storing information in internal states for long time periods (long-range memory) [1], [3], [4].

A class of single-layer RNNs called Reservoir Computing (RC) has attracted considerable attention during the last years due to its remarkable success in the modeling of complex non-linear dynamical systems [2]. An RC model relies on the assumption that supervised adaptation of all hidden-hidden weights is not necessary [4]. A single-layer recurrent network (so-called *reservoir*) is driven by an input signal, and projects the input information into a high-dimensional space. Reservoir weights are randomly

initialized and scaled for satisfying technical conditions of dynamic stability [4]. A distinctive characteristic of an RC model is that the training algorithm only adjusts the readout structure, in general recurrence-free weights that map the reservoir projection to the output space. Nowadays, learning algorithms for adapting the recurrence-free parameters are fast and robust.

The research community has been actively studying methods for creating efficient reservoir models [2]. Specific cyclic topologies were studied in [5], [6], and pre-training techniques with self-organized strategies were analyzed in [7], [8]. In [9], [10], reservoir units with parametric activation functions were trained for maximizing the entropy using the Intrinsic Plasticity (IP) rule. Other approaches based on Evolutionary Algorithms (EAs) for optimizing hyper-parameters of the reservoir structure, such as number of reservoir units, scaling parameters, connectivity properties (density and spectral radius of the reservoir weights matrix), have also been introduced in the literature [11], [12]. EAs have also been used for optimizing interconnected reservoirs [13], [14], where GAs are applied for searching hyper-parameters of hierarchical reservoirs. Besides, Bayesian optimization for finding hyper-parameters of particular reservoirs was also analyzed in [15], [16]. Furthermore, a relatively new prominent research in the RC area consists of models with interconnected reservoirs. A tree structure where the tree-nodes satisfies the reservoir characteristics was developed in [17]. The weight connections among the reservoirs are set in order of creating a global contractive mapping. Furthermore, layered/hierarchical architectures composed of multi-layered reservoirs have also been studied in [13], [14], [18]–[21].

Motivations. Defining good hyper-parameters of ESNs is still often made by experience using a grid search strategy, partially brute force, and in general requires numerous trails [5]. Another drawback is that different reservoir matrices with the same hyper-parameters may produce substantially different results [9]. In addition, since the reservoir is fixed during the training procedure, the computational power of the model is not fully harnessed [1], [5].

Contributions. In this article, we introduce the EVOLUTIONARY Echo State Network (EvoESN), a novel computational model rooted in RC concepts. We combine ideas regarding dimensionality reduction using the Fourier transform for searching weights of RNNs (fully connected architectures) [22], optimization of network weights using an evolutionary search (such as in EVOLINO [23]), and the potential of ESNs for fast and robust temporal learning [4].

Given a reservoir with a fixed pattern of connectivity (as in standard ESNs), we apply an indirect encoding procedure for representing the reservoir weights into the Fourier space (specifically, by means of the Discrete Cosine Transform, DCT) [22], [24]. Then, we perform a fine-tuning of the reservoir weights applying genetic algorithms over the DCT coefficients. It is important to note that we are not looking for a good reservoir directly working with the weights, that becomes changing the RC paradigm and leaving the fast learning advantages of an RC model. Instead, we solve an optimization problem in the frequency domain. Using this Fourier representation drastically reduces the size of the searching space, and makes it possible to efficiently look for good reservoir weights. Our framework is a strong integration between learning in the RC paradigm and using EAs to improve the reservoir (to find *better* weights) following a global iterative process moving at each iteration between the weights space and the frequency space using the DCT transform and its inverse.

We demonstrate the utility of EvoESN by experiments on two well-known chaotic systems (the Mackey-Glass system and the Lorenz systems [23], [25]) and a real-world problem (Sunspot series [26]). Preliminary results of part of the content of this paper were briefly presented in the 2-page poster [27].

The rest of the article is organized as follows. In the next section we present the ESN framework. Our main contributions are the object of Section III, where we discuss the reservoir encoding and the evolutionary search approach. Experimental results are provided in Section IV. The article ends in Section V with a discussion about the presented work, its limitations and possible future research avenues.

II. PRELIMINARIES

A. The Echo State Network framework

The large recognition of ESNs comes from their record-breaking success in predicting observations through the Mackey-Glass attractor [25]. The most common application of ESNs is time series prediction, where the operation is divided into two phases. An initial training phase where the model is externally driven by an input signal and the parameters are trained using supervised learning, followed by an autonomous phase (so-called free-run prediction) where the input signal is provided by the model itself using previous predicted values [28].

Let us consider an ESN with K input neurons, $N > K$ hidden neurons and L output neurons. The dynamics of the model is described in terms of a discrete-time system with an input signal $\mathbf{u}(t) \in \mathbb{R}^K$ and a hidden state $\mathbf{x}(t) \in \mathbb{R}^N$. The recurrent state is computed according to the following recurrence [29]:

$$\mathbf{x}(t) = f(\mathbf{W}^{\text{in}}\mathbf{u}(t) + \mathbf{W}^{\text{r}}\mathbf{x}(t-1) + \mathbf{W}^{\text{fb}}\mathbf{y}(t-1) + \boldsymbol{\varepsilon}(t)), \quad (1)$$

where t is the discrete time index, \mathbf{W}^{in} is a $N \times K$ matrix with input-to-reservoir weights, the $N \times N$ matrix \mathbf{W}^{r} has the reservoir internal weights, \mathbf{W}^{fb} is a weight $N \times L$ matrix with the feedback connections going from the output neurons to the reservoir, and $\boldsymbol{\varepsilon}(t)$ is a noise vector with dimension N . Function $f(\cdot)$ is the activation function, usually an hyperbolic function, applied element-wise. Matrix \mathbf{W}^{fb} and vector $\boldsymbol{\varepsilon}$ are optional [25]. The model's output $\mathbf{y}(t) \in \mathbb{R}^L$ is computed by a readout layer defined as:

$$\mathbf{y}(t) = g(\mathbf{W}^{\text{out}}[\mathbf{u}(t); \mathbf{x}(t)]), \quad (2)$$

where \mathbf{W}^{out} is the $N \times (K+L)$ reservoir-to-output weight matrix, $g(\cdot)$ is a coordinate-wise function (usually an hyperbolic tangent or, in some cases, the identity) and $[\cdot; \cdot]$ denotes vector concatenation. The readout weights are trained using a classical tool such as Tikhonov regularization [4], [28]. In the case of an ESN with feedback connections (non-null \mathbf{W}^{fb}), during the training procedure the feedback values are *teacher-forced*: instead of feeding back the reservoir with the predicted value $\mathbf{y}(t)$, the real target is used as the feedback value [30]. In exploitation mode, the model's output $\mathbf{y}(t)$ is fed back to the reservoir.

Since the reservoir weights are frozen during learning, the initial reservoir design may have an impact on the model's performance. The most relevant hyper-parameters to be chosen are: reservoir size, input scaling factor, activation functions and reservoir spectral radius. Most research effort has gone into understanding the impact of those parameters. For more information, especially about the fundamental property of the reservoir dynamics (so-called *Echo State Property* (ESP)), we suggest [2], [4], [28], [31], [32].

B. Evolutionary Algorithms applied on RC models

There has been a rising interest in large-scale neural architecture engineering using evolutionary algorithms [33]. Several EAs have been developed to optimize NNs' architectures [23], [34], [35]. The RC field has also taken advantage of the recent advances in evolutionary computation. EAs have been employed for finding the hyper-parameters of an RC model, such as: sparsity, input scaling factor, injected noise, spectral radius, leaky rate and number of reservoir units. Particle Swarm Optimization (PSO) and Genetic Algorithms (GAs) have been used as metaheuristics for selecting optimal hyperparameters in [11], [12], [36]. Global parameters of an extended ESN (DRESN) were optimized using GAs for solving medical signal prognosis [37]. EAs have also been applied for evolving the topology of some RC models. An evolutionary pre-training procedure based on the extreme learning approach has been studied in [38], where the authors applied PSO for adjusting a subset of the reservoir weights. Recently, NEAT and HyperNEAT algorithms were also applied for evolving the pattern of connectivity of reservoirs in [39], [40]. In [40] the evolution of the reservoir topology was made considering a fitness function that combines the memory capacities of the model and its accuracy. HyperNEAT algorithm over a full recurrent topology has a significant computational cost [40]. GAs have also been used for optimizing the hyper-parameters of hierarchical reservoirs and for finding connection parameters between reservoir

structures [13], [14]. Furthermore, a microbial GA was used for optimizing multilayered architectures of interconnected reservoirs [13]. It was also employed for finding good architectures of single ESNs with leaky integrator units. Recently, a hybrid approach using swarm optimization and local search was applied for evolving hierarchical ESNs [41].

Differences between EvoESN and previous works. There are also other works that tune the reservoir weights using evolutionary techniques. We identify two relevant differences between our approach and the proposed in [40]. EvoESN performs evolution over a fixed reservoir topology and the fine-tuning is done only over the weights, while in [40] the evolution is applied in both structure and connection weights of the reservoirs using HyperNEAT. In addition, we perform the optimization of the weights in the Fourier space applying a considerable reduction of the searching space. As already mentioned above, there are also other works where was applied EA over the reservoir weights, e.g. microbial GA was used on a single ESN [13] and PSO in [38]. The difference between EvoESN and these previous works is the encoding technique, in the mentioned works the evolution was applied directly over the weight connections.

III. EVOESN: EVOLUTIONARY ECHO STATE NETWORK

A. Dimensionality reduction of reservoir weights applying a Fourier-type transform

Genetic encoding is a fundamental aspect of EAs, with special interest when algorithms are applied for optimizing large complex structures. A direct encoding (that is, a one-to-one mapping where every weight is represented in the gene) is clearly unsuitable for large-scale networks [24]. Furthermore, it has been shown that naive crossover operations present deficiencies when used for combining recurrent networks [42]. In [22], an indirect encoding procedure for NNs was introduced, where the network weights are represented in the frequency domain applying a Fourier transform, more precisely, the type-II DCT of the weight matrix. The evolutionary search is conducted in a reduced space composed by the coefficients of the low frequencies only [24]. This leads to a significant dimensionality reduction. The approach has been successfully applied over rather small fully connected RNNs, on three benchmark problems of the area of control theory: pole balancing, ball throwing and octopus arm [22]. The authors have also applied this procedure on computer vision problems and in reinforcement learning [24]. In [43], the authors studied the trade-off between non-linear mappings and memory capacity of ESNs. The authors applied FFTs to estimate the deviation from linearity in the reservoir activations. Furthermore, the state transition structure of the reservoir was analyzed in the frequency domain in [44].

In our case, we also exploit the well-known benefits of Fourier transform for compressing data information. This efficient indirect encoding allows us to represent large reservoir matrices in a reduced space using only the low frequency coefficients of the transform. As a consequence, instead of having to tune hundreds of thousands or millions of reservoir weights, we will look for hundreds or thousands of Fourier coefficients, which implies a huge gain in complexity.

Discrete Cosine Transforms (DCTs). Let $s = (s_0, s_1, \dots, s_{J-1})$ be a sequence of real numbers having size J . In the sequel, when we say DCT we mean the normalized DCT transform of type-II. The DCT S of s (we write $S = \text{DCT}(s)$) is also a sequence of J reals $S = (S_0, S_1, \dots, S_{J-1})$, defined as follows:

$$S_0 = \frac{1}{\sqrt{J}} \sum_{j=0}^{J-1} s_j,$$

and for $\ell = 1, 2, \dots, J - 1$,

$$S_\ell = \sqrt{\frac{2}{J}} \sum_{j=0}^{J-1} s_j \cos\left(\frac{\pi\ell(2j+1)}{2J}\right).$$

The inverse transform, going from S to s , denoted IDCT, is given by the following expression: for $\ell = 0, 1, \dots, J - 1$,

$$s_\ell = \frac{S_0}{\sqrt{J}} + \sqrt{\frac{2}{J}} \sum_{j=1}^{J-1} S_j \cos\left(\frac{\pi j(2\ell + 1)}{2J}\right).$$

This IDCT transform is also called normalized DCT of type-III. So, $S = \text{DCT}(s)$ and $s = \text{IDCT}(S)$.

DCT satisfies some properties that make the method suitable for encoding large recurrent structures and performing GAs. Observe that both mappings are linear, so they are continuous. Also, observe that the fact that one is inverse of the other means that they define one-to-one mappings from the space of vectors having size J to itself.

The DCT transform has a nice “energy compaction” property, informally meaning that if we keep in $S = \text{DCT}(s)$ just a first few terms (called “low frequency” terms) and replace the rest by 0s, obtaining a new vector S' , then, reconstructing the initial vector from S' using the IDCT has a small error. That is, $\text{IDCT}(S') \approx s$. This characteristic of the DCT is strongly used in signal compression, and is a key element of our approach in this paper.

B. Evolving reservoirs in the Fourier space

Recall that in the canonical ESN, we select a subset of weights in the reservoir to be randomly initialized and that the rest is set to 0. We denoted by N the number of neurons in the reservoir; let M be the number of weights randomly chosen, $M < N^2$ (for instance, $M \approx 0.2N^2$ [4], [31]).

Let us denote by \mathbf{p} the positions of the sampled weights in the reservoir matrix \mathbf{W}^r , a vector having size M . In our approach, we will keep fixed the vector \mathbf{p} but we will modify the corresponding weights during the procedure. For that reason, we call these weights the “unfrozen” ones, and we call the remaining $N^2 - M$ weights the frozen-to-zero ones. After choosing an arbitrary order on those unfrozen weights, the k th coordinate in \mathbf{p} corresponds to some element (i, j) of \mathbf{W}^r , that is, $p_k = (i, j)$. Define now vector $\boldsymbol{\nu}$ having dimension M , containing the value of those weights, that is, $\nu_k = \mathbf{W}_{p_k}^r = \mathbf{W}_{i,j}^r$. The proposed method moves the vector $\boldsymbol{\nu}$ to the frequency domain using the DCT, and there, a GA will produce new vectors, always in the frequency space. Call \mathcal{P} this population of candidates. Each vector in \mathcal{P} (denoted by $\boldsymbol{\alpha}$) will be transported to the network weights using the IDCT, occupying the fixed positions given in \mathbf{p} . The ESN training procedure (e.g. ridge regression over the readout structure) will then be run, and the resulting network evaluated using a validation data set. The training can be done using the techniques already mentioned in Section II-A. This will allow us to select the best individuals of \mathcal{P} . If its quality is not good enough, the best elements in \mathcal{P} are used to start a new generation of candidates through GAs. The selection procedure used to find the best element in \mathcal{P} must move each C -dimensional vector $\boldsymbol{\alpha}$ to the weights space, in order to evaluate the corresponding ESN. For this purpose, we first extend $\boldsymbol{\alpha}$ with zeros, to obtain a vector with size M . This procedure is called *padding*, which is necessary if $C < M$ because the input and output of an inverse DCT have equal dimensions. The dimension C of the frequency space is a control parameter. The “energy compaction” property of DCTs makes that using some C value not too small, those coefficients will have enough information about the corresponding weights vectors, in spite of having much less components.

Figure 1 visualizes the global EvoESN procedure. Summarizing, the padding operation is performed if necessary (if $C < M$), then sent to the weight space as a new feasible solution through the inverse DCT. The ESN is evaluated, after training, and its error is used in the evolutionary operations. Aiming to increase the readability, we describe the proposed algorithm into two pseudo-codes. The evolutionary search is presented in pseudo-code 1, and the computation of the fitness value of each feasible solution by the evolutionary algorithm is given in pseudo-code 2.

IV. EXPERIMENTS

To measure the computational power of the presented approach, we have selected three well-known benchmark problems: Mackey-Glass system, Lorenz system and monthly sunspot index data. Due to their

interesting properties both chaotic systems have been largely studied in the RC community [6], [23], [25], [43], [45]–[47]. In addition, monthly Sunspot data has information of the magnetic field in the Sun’s outer regions. It also has been analyzed using RC models in at least [6], [14], [48]. An open source monthly sunspot series is available at [26].

A. Benchmark problems

The Mackey-Glass system (MGS). As far as we know, ESN is still the reference RNN for modeling the MGS chaotic time series [23], [25], defined by the equation:

$$\frac{\partial y}{\partial t} = \frac{\alpha y(t - \tau)}{1 + y(t - \tau)^\beta} - \gamma y(t),$$

where the parameters are usually $\alpha = 0.2$, $\beta = 10$, $\gamma = 0.1$. For these parameters’ values, the system has a chaotic attractor if the delay τ satisfies $\tau > 16.8$ [25]. We study the system with a delay $\tau = 17$, the most commonly used value. For this problem we replicate the experimental setting used in [25]. The first 1000 time steps were discarded to washout the initial transients. The training length was 3000 time steps. The remaining 2084 time steps were used for testing. From those, the first 2000 steps were used in a teacher-forced mode, then the network was in exploitation mode.

Lorenz attractor. The sequence data set is based on the following equations:

$$\frac{\partial x}{\partial t} = \sigma(y - x), \quad \frac{\partial y}{\partial t} = rx - y - xz, \quad \frac{\partial z}{\partial t} = xy - bz.$$

The parameters used are $r = 28$, $b = 8/3$, $\sigma = 10$, and the step size is 0.01. We replicated the pre-processing procedure for sequence x and the experimental setting described in [25]. The sequence was rescaled by a factor of 0.01, and the results are presented with data scaled back. There was a washout of 1000 time steps, and the training data size had 6000 time steps. For testing, we used 1000 points in a teacher-forced mode and 600 time steps in exploitation mode.

Monthly sunspot series. We studied the sunspot index from Jan. 1749 to Dec. 2021, a total of 3276 data points, available at [26]. In this problem, we replicated the experimental setting investigated in [6]. With a minor difference that our sequence is slightly larger than the series studied in [6]. Initial washout of 100 samples, 1600 points for training, 500 points of validation and rest of the sequence for testing.

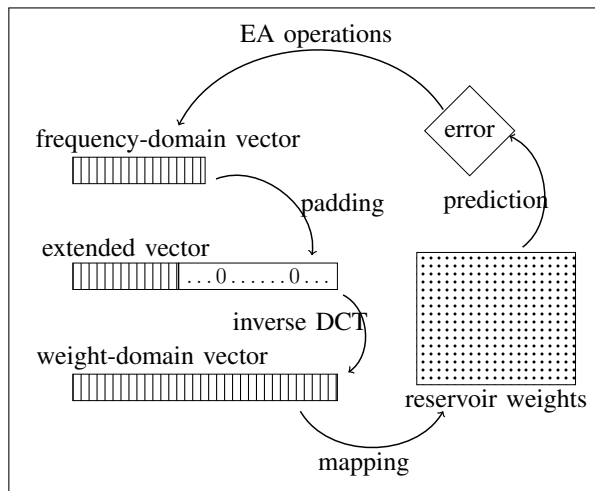


Figure 1: High-level visualization of the EvoESN model.

Algorithm 1: Evolutionary search for a fine-tuning the reservoir weights in the Fourier space.

```

// Initialization of  $I$  individuals
// in GA population
1 Random initialization of  $\mathcal{P} = \{\alpha^{(1)}, \dots, \alpha^{(I)}\}$ ;
2 Compute fitness for each  $\alpha^{(i)}$  using Algorithm 2;
  // Apply evolutionary search
3 while convergence criterion is not satisfied do
4   | Selection of best individuals;
5   | Mutation operation;
6   | Crossover operation;
7   | Compute fitness for each  $\alpha^{(i)}$  using Algorithm 2;
8 end while

```

Algorithm 2: Evaluation of a frequency-domain C-vector α .

```

1 if  $C < M$  then
2   | Padding operation for  $\alpha$ ;
3 end if
  // Mapping from frequency domain
  // to weight domain
4  $\nu = \text{IDCT}(\alpha)$ ;
5 Assign  $\nu$  values to  $\mathbf{W}^r$  according to positions in  $p$ ;
6 Train readout weights;
7 Return ESN evaluation error of the given individual;

```

B. Hyper-parameter selection and experimental setup

The selection of the most important hyper-parameters of the proposed EvoESN model was made following a grid search strategy. In the case of the ESNs we also used suggestions from the literature that are described below according to each specific benchmark problem. For all the benchmark problems we applied a grid search in the interval $[0.1, 1.4]$ for finding a good spectral radius value. Note that we consider spectral radius also larger than 1. As a consequence, the ESP is not guaranteed and the reservoir can exhibit unstable dynamics. This strategy was also applied in [9], [12], [43]. In the MGS and Lorenz problems some parameters, such as injected noise, input bias, activation functions, and scaling feedback connections, were directly taken from the seminal article [25].

MGS problem. We adopted the network configuration employed in [25]. We designed a reservoir with 1000 units and the hyperbolic tangent activation function. The three weight matrices \mathbf{W}^{in} , \mathbf{W}^r and \mathbf{W}^{fb} were sampled from the uniform distribution on $(-1, 1)$. According to our results, a good candidate for the spectral radius of the reservoir matrix is 0.8, a result coherent with [25]. Left figure in Fig. 2 shows the relationship between spectral radius and reservoir sparsity. The curves were produced with the average of testing errors computed over 30 independent simulations. Last, the following parameters were taken from [25]: injected uniform noise was scaled by 10^{-10} , and input bias had value 0.2.

Lorenz problem. For this task we also followed the hyper-parameters selection studied in [25]. Input scaling and data normalization was done following [25]. We used a reservoir with 600 units, with the hyperbolic tangent activation function. Right figure in Fig.2 shows the testing error for different values of reservoir spectral radius and reservoir sparsity. Curves were created averaging the results of 30 independent simulations. We observed that the best spectral radius candidates were close to 1. The reservoir matrix

had 20% of non-zero weights. Then, we set up a spectral radius to 0.97 (the same value was selected in [25]). The input and reservoir matrices were initialized with the uniform distribution over $(-1, 1)$. Uniform injected noise had values $\pm 10^{-7}$ [25]. The feedback connections were sampled following a uniform distribution in $(-4, 4)$ (this particular range was suggested in [25]).

Sunspot series problem. We adopted the setting used in [6], and we also analyzed some parameters using grid search. An important difference in this problem with the previous task, is that the goal is to predict only one single time-step ahead (as it was studied in [6]). In previous tasks we applied free run prediction over large testing time windows. For this task we studied only networks without feedback connections. Left figure in Fig. 3 shows the relationship between spectral radius and sparsity of the reservoir. Besides, for this problem we also analyzed leaky units in the range $[0.01 : 0.05 : 0.5]$. For the regularization parameter we evaluated the values 10^{-i} with $i \in \{3, 4, 5, 6, 7, 8\}$, and selected 10^{-5} . The selected input scaling was 0.1 after evaluation over the grid $\{0.01, 0.02, 0.03, 0.05, 0.07, 0.09, 0.1, 0.15, 0.2, 0.3, 0.5\}$.

Computational aspects. The IDCT was chosen with the orthogonal norm using the `fftpack` Python package. We conducted GA as EA, and it was implemented using the `DEAP` package [49]. We analyzed the impact of C (number of coefficients in the frequency-domain vector) in a size range between 25 and 750. Besides, different numbers of generations were evaluated: 70, 150, 300, 600. After carrying out multiple runs of the GA, we selected the hyper-parameters in order to have relatively fast convergence and to diminish the probability of being trapped in a local minima. Here we present the results for the following GA parameters. We used a Gaussian mutation with parameter `mutpb`=0.15, two-point crossover, probability of mating two individuals of 0.5, number of individuals participating in each tournament = 3, and population size = 20. The readout was computed using a ridge regression with a regularization factor of 10^{-9} (MGS task) and 10^{-6} (Lorenz task) [25], and 10^{-5} (Sunspot task).

C. Results analysis

Accuracy evaluation. It is common to see the NRMSE metric on experiments conducted with MGS and Lorenz data [23], [25], [43]. Therefore, for comparison purposes we also used NRMSE for these two tasks, and for the sunspot series problem we use NMSE as it was applied in [6]. We denote by $\text{NMRSE}(H)$ the error computed in the testing phase using the residuals over *the whole horizon* H (by residual we mean the difference between target and predicted values). We also present another metric denoted by NRMSE_H following the notation in [23], [25]. NRMSE_H is computed in the testing phase after H autonomous steps. The error is computed only using the residual at the H time step. NRMSE_{84} has become the standard metric for evaluating the MGS prediction [23].

Accuracy comparison. Table I presents a non-exhaustive overview of the reported state-of-the-art results. Previous studies have analyzed the problems using different error metrics. Then, the last column in the table shows the information about the error measure. Furthermore, we also implemented the canonical ESN, the ESN with leaky integrator units [50], and one variation of ESN with feedback connections introduced in [25]. In [25], the authors describe a learning approach for correcting the bias introduced by the teacher-forcing schema in the case of feedback connections. Jaeger *et al.* obtained the best performance for MGS prediction using this particular learning schema introduced as *refined learning method*. Let *ESN with FB (I)* denote the ESN with feedback connections using the teacher-forcing schema, and *ESN with FB (II)* denote the ESN with feedback connections using the refined learning method.

Results. We partially investigated the impact of the parameter C . Some results are shown in the right figure in Fig. 3 where the prediction error, in a logarithmic scale, is plotted against time (generation number), for different values of C . The fitness function was computed using 300 points of generalization on 12 independent prediction tasks. The curves show the evolution of the average of NRMSE across the 12 independent experiments. They also show that the accuracy is increasing with C and with time. See that if $C \geq 200$, for instance, then after generation 50, the performance of the method becomes roughly independent of C . The results are consistent, the curves behave as expected, and globally, the procedure

shows its efficiency once the frequency-domain vectors are large enough. Figure 4 shows results over Lorenz data with $C = 150$. There is a group of 6 figures. The top-left figure has an example of Lorenz testing data. Next graphics show EvoESN predictions (of a randomly selected individual) for different generations of the GA. In a few iterations, the model is able to reproduce with significant accuracy a large segment of the Lorenz attractor. Figure 5 shows the absolute error for different C values on 600 time units in the Lorenz system task. For the three values of C the results outperform the ones reported in [25] using different variations of ESNs. Note that the red curve ($C = 300$) has a better accuracy than the other two curves ($C = 50$ and $C = 150$) until the time horizon $t \approx 300$. For longer time horizons than $t = 300$ both the red and blue curves seem to have similar behavior. The explanation probably lies in that the fitness function used in the GA algorithm was computed using 200 points of generalization, i.e. the reservoir weights are optimized for predicting the next 200 points. After that threshold limit the models behaves similarly (independently of C). Further possible improvements will be looked for by studying different fitness functions in the evolutionary search.

Table II summarizes the results of experiments over MGS. In the case of EvoESN, we present the results for 500 coefficients, and 70 and 300 generations in the GA algorithm. A total of 50 randomly created ESNs were evaluated. We show the average of the error produced by independent runs. EvoESN reaches much lower accuracy than the obtained with ESNs using our implementation. EvoESN also reaches the results reported in [25]. However, even though the variance is very low, we cannot say that the difference is significant with the error reported in [25]. We expect to do further studies using statistical analysis and a full evaluation of the GA parameters, in order to check if the difference is statistically significant.

Table III shows the results for the Lorenz problem. The NRMSE_{84} error has the same order of magnitude as obtained by Jaeger *et al.* using a feedback connection and a refined learning method [25]. It also shows the absolute error over a 600 time horizon, that was studied in [25]. Table IV presents the test errors for the Sunspot series problem. EvoESN outperforms the results reported in [6], and also the results obtained by our implementations of ESN with leaky-integrators. We ran 50 independent simulations with the same ESN architecture, and the standard deviation of error across simulation has been really low ($\approx 1 \times 10^{-6}$) independently of the leaky rate. Standard deviation values for EvoESN experiments are shown in the table. Regarding complexity, EvoESN has a higher cost than standard ESNs. Note that the GA performs, at each generation, training of the readouts using linear regression, computation of reservoir spectral radius and of inverse DCT. Other GA operations such as crossover, selection and mutation are less costly than the mentioned above. However, it seems that a few iterations are enough for achieving much lower errors with respect to ESNs.

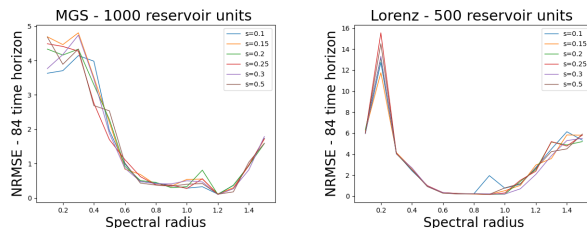


Figure 2: Impact of spectral radius versus reservoir sparsity.

V. CONCLUSIONS

In this article, we have presented the EvoESN, an RC-based model where the reservoir weights are not set randomly, but are obtained by applying GA, integrated with the classical training of the ESN readouts. The proposed algorithm projects the reservoir weights into the frequency domain using a DCT. The optimization is performed in a frequency space whose dimension is much smaller than that of the non-null reservoir weights. This leads to a significant dimensionality reduction, in turn producing a significant computer time reduction, while keeping the quality of the optimization procedure. These are consequences

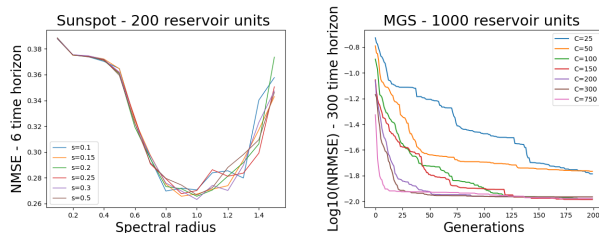


Figure 3: Impact of hyper-parameters. Left figure shows the relationship between spectral radius and sparsity in the Sunspot series problem. Right figure shows the impact of the number of coefficients C in the MGS problem.

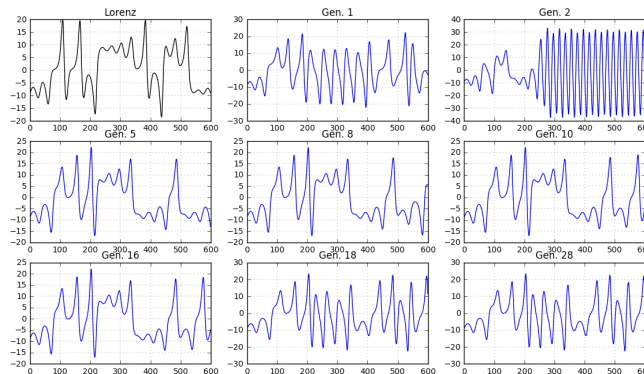


Figure 4: Prediction of the Lorenz attractor: convergence of the proposed method according to the GA generations.

of the properties of Fourier-like transforms. Since we use an indirect encoding the dimension of the searching space is basically independent of the reservoir size, an additional advantage of the proposed technique. Another advantage is that the fine-tuning of the reservoir weights is made using the fitness function of the evolutionary search, which can be different than the one used for training the readouts. In addition, by applying GA we avoid the difficulties associated with the training of RNNs. The quality of these first experiments clearly show the potential of the approach, and the first work to be done in the close future is to push the procedure towards better performances by a careful analysis of both weight space encoding and evolutionary search procedures.

Table I: Accuracy comparison with other RC methods.

Problem	Reference	Model	Error	Metric
MGS	H. Jaeger <i>et al.</i> [25]	ESN-FB	$\approx 10^{-4.2}$	NRMSE ₈₄
	Schmidhuber <i>et al.</i> [23]	Evolino	$\approx 1.9 \times 10^{-3}$	NRMSE ₈₄
	J. Steil [51]	BPDC	0.318	NMSE
	A. Ferreira <i>et al.</i> [12]	RCDESIGN	0.00051	NRMSE
	C. Gallicchio <i>et al.</i> [47]	RDES	1.112×10^{-8}	MSE
Lorenz	Jaeger <i>et al.</i> [25]	ESN-FB	≈ -4.27	$\log_{10}(\text{NRMSE}_{84})$
Sunspot	Rodan <i>et al.</i> [6]	ESN	≈ 0.1042	NMSE
	Rodan <i>et al.</i> [6]	DLR, DLRB, SCR	0.1039	NMSE

REFERENCES

- [1] J. Martens and I. Sutskever, "Learning Recurrent Neural Networks with Hessian-Free Optimization," in *Proceeding of the 28th International Conference on Machine Learning*. Bellevue, WA, USA: JMLR.org, 2011, pp. 1033–1040.
- [2] G. Tanaka, T. Yamane, J. B. Héroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, and A. Hirose, "Recent advances in physical reservoir computing: A review," *Neural Networks*, vol. 115, pp. 100–123, 2019.

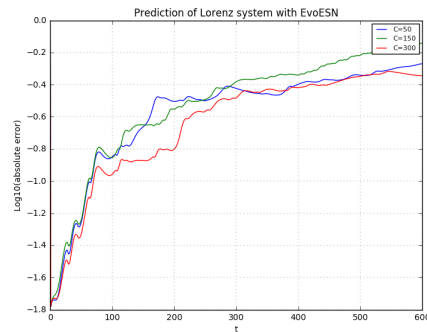


Figure 5: Prediction of the Lorenz attractor for different values of C . Type of visualization taken from [25]

Table II: Comparison between different types of ESNs and EvoESN (with $C = 500$) for the MSG prediction task.

Model	Characteristics	Error	Metric
ESN	without FB, non-linear readout	0.1770	NRMSE (84H)
ESN	without FB, linear readout	0.29857	NRMSE (84H)
ESN	with FB (I), linear readout	0.19620	NRMSE (84H)
ESN	with FB (I), non-linear readout	0.19376	NRMSE (84H)
ESN	with FB (II), non-linear readout	0.07684	NRMSE (84H)
EvoESN	300 generations	$\approx 1.075 \times 10^{-4}$	NRMSE ₈₄
EvoESN	70 generations	4.485×10^{-4}	NRMSE (84H)

Table III: Results for the Lorenz attractor prediction.

Model	C	Error	Metric
EvoESN50		-3.4666	\log_{10} NRMSE ₈₄
EvoESN100		-3.6014	\log_{10} NRMSE ₈₄
EvoESN150		-4.1379	\log_{10} NRMSE ₈₄
EvoESN50		-0.2704	\log_{10} absolute error (600H)
EvoESN150		-0.1407	\log_{10} absolute error (600H)

Table IV: Results for the Sunspot prediction.

Model	Characteristics	Error	Metric
ESN	without FB	0.085358	NMSE
LI-ESN	leaky rate 0.1	0.033120	NMSE
LI-ESN	leaky rate 0.7	0.031816	NMSE
EvoESN	C 50, NGEN 150	0.0290556 (8.1×10^{-4})	NMSE
EvoESN	C 75, NGEN 150	0.0291978 (1.9×10^{-4})	NMSE
EvoESN	C 100, NGEN 150	0.0292394 (7.4×10^{-5})	NMSE
EvoESN	C 200, NGEN 150	0.0293263 (7.7×10^{-5})	NMSE
EvoESN	C 300, NGEN 150	0.0309722 (4.4×10^{-5})	NMSE

- [3] J. Schmidhuber, “Deep learning in neural networks: An overview.” *Neural Networks*, vol. 61, pp. 85–117, 2015.
- [4] M. Lukoševičius and H. Jaeger, “Reservoir Computing Approaches to Recurrent Neural Network Training,” *Computer Science Review*, vol. 3, pp. 127–149, 2009.
- [5] A. Rodan and P. Tiño, “Simple deterministically constructed cycle reservoirs with regular jumps,” *Neural Computation*, vol. 24, pp. 1822–1852, 2012.
- [6] A. Rodan and P. Tiño, “Minimum Complexity Echo State Network,” *IEEE Transactions on Neural Networks*, vol. 22, pp. 131–144, 2011.

- [7] M. Lukoševičius, “Reservoir Computing and Self-Organized Neural Hierarchies,” Ph.D. dissertation, School of Engineering and Science, Jacobs University, december 2011.
- [8] S. Basterrech, C. Fyfe, and G. Rubino, “Self-organizing Maps and Scale-invariant Maps in Echo State Networks,” in *11th International Conference on Intelligent Systems Design and Applications, ISDA 2011, Córdoba, Spain, November 22-24, 2011*, November 2011, pp. 94–99.
- [9] B. Schrauwen, M. Wardermann, D. Verstraeten, J. J. Steil, and D. Stroobandt, “Improving Reservoirs using Intrinsic Plasticity,” *Neurocomputing*, vol. 71, pp. 1159–1171, March 2008.
- [10] J. J. Steil, “Online reservoir adaptation by intrinsic plasticity for backpropagation-decorrelation and echo state learning,” *Neural Networks*, 2007.
- [11] A. Ferreira and T. Ludermir, “Comparing evolutionary methods for reservoir computing pre-training,” in *Neural Networks (IJCNN), The 2011 International Joint Conference on*, July 2011, pp. 283–290.
- [12] A. A. Ferreira, T. B. Ludermir, and R. R. B. De Aquino, “An Approach to Reservoir Computing Design and Training,” *Expert Syst. Appl.*, vol. 40, no. 10, pp. 4172–4182, Aug. 2013.
- [13] M. Dale, “Neuroevolution of hierarchical reservoir computers,” in *Proceedings of the Genetic and Evolutionary Computation Conference*, ser. GECCO ’18. New York, NY, USA: ACM, 2018, pp. 410–417.
- [14] Q. Ma, L. Shen, and G. W. Cottrell, “DeePr-ESN: A deep projection-encoding Echo State Network,” *Information Sciences*, vol. 511, pp. 152–171, 2020.
- [15] J. R. Maat, N. Gianniotis, and P. Protopapas, “Efficient Optimization of Echo State Networks for Time Series Datasets,” in *2018 International Joint Conference on Neural Networks (IJCNN)*, 2018, pp. 1–7.
- [16] G. Trierweiler Ribeiro, J. Guilherme Sauer, N. Fraccanabbia, V. Cocco Mariani, and L. dos Santos Coelho, “Bayesian optimized echo state network applied to short-term load forecasting,” *Energies*, vol. 13, no. 9, 2020.
- [17] C. Gallicchio and A. Micheli, “Tree Echo State Networks,” *Neurocomputing*, vol. 101, p. 319–337, 02 2013.
- [18] —, “Deep echo state network (deepesn): A brief survey,” *CoRR*, vol. abs/1712.04323, 2017.
- [19] —, “Reservoir topology in deep echo state networks,” in *ICANN’2019*, I. V. Tetko, V. Kůrková, P. Karpov, and F. Theis, Eds. Cham: Springer International Publishing, 2019, pp. 62–75.
- [20] —, “Echo state property of deep reservoir computing networks,” *Cognitive Computation*, vol. 9, no. 3, pp. 337–350, Jun 2017.
- [21] C. Gallicchio, A. Micheli, and L. Pedrelli, “Deep reservoir computing: A critical experimental analysis,” *Neurocomputing*, vol. 268, no. C, pp. 87 – 99, 2017, advances in artificial neural networks, machine learning and computational intelligence.
- [22] J. Koutník, F. Gomez, and J. Schmidhuber, “Evolving neural networks in compressed weight space,” in *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation*, ser. GECCO ’10. New York, NY, USA: Association for Computing Machinery, 2010, pp. 619–626.
- [23] J. Schmidhuber, D. Wierstra, M. Gagliolo, and F. Gomez, “Training Recurrent Networks by Evolino,” *Neural Networks*, vol. 19, pp. 757–779, 2007.
- [24] J. Koutník, G. Cuccu, J. Schmidhuber, and F. Gomez, “Evolving large-scale neural networks for vision-based reinforcement learning,” in *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation*. New York, NY, USA: Association for Computing Machinery, 2013, pp. 6–10.
- [25] H. Jaeger and H. Haas, “Harnessing Nonlinearity: Predicting Chaotic Systems and Saving Energy in Wireless Communication,” *Science*, vol. 304, pp. 78–80, April 2004.
- [26] SILSO World Data Center, “The international sunspot number,” *International Sunspot Number Monthly Bulletin and online catalogue*, 1749-2021.
- [27] S. Basterrech and G. Rubino, “Evolving Reservoir Weights in the Frequency Domain,” in *2021 Genetic and Evolutionary Computation Conference Companion (GECCO ’21 Companion)*. ACM, 2021.
- [28] A. Hart, J. Hook, and J. Dawes, “Embedding and approximation theorems for echo state networks,” *Neural Networks*, vol. 128, pp. 234 – 247, 2020.
- [29] H. Jaeger, “The “echo state” approach to analysing and training recurrent neural networks,” German National Research Center for Information Technology, Tech. Rep. 148, 2001.
- [30] R. F. Reinhart and J. J. Steil, “Regularization and stability in reservoir networks with output feedback,” *Neurocomputing*, vol. 90, 2012.
- [31] M. Lukoševičius, “A Practical Guide to Applying Echo State Networks,” in *Neural Networks: Tricks of the Trade*, ser. Lecture Notes in Computer Science, G. Montavon, G. Orr, and K.-R. Müller, Eds. Springer Berlin Heidelberg, 2012, vol. 7700, pp. 659–686.
- [32] S. Basterrech, “Empirical analysis of the necessary and sufficient conditions of the echo state property,” in *International Joint Conference on Neural Networks, IJCNN’17*. Anchorage, AK, USA: IEEE Press, 2017, pp. 888–896.
- [33] T. Elsken, J. H. Metzen, and F. Hutter, “Neural Architecture Search: A Survey,” *Journal of Machine Learning Research*, vol. 20, pp. 1–21, 2019.
- [34] A. Rawal and R. Miikkulainen, “Evolving Deep LSTM-based Memory networks using an Information Maximization Objective,” in *Genetic and Evolutionary Computation Conference (GECCO 2016), Colorado, USA*, 2016.
- [35] K. O. Stanley, J. Clune, J. Lehman, and R. Miikkulainen, “Designing neural networks through neuroevolution,” *Nature Machine Learning*, vol. 1, pp. 24–35, 2019.
- [36] A. T. Sergio and T. B. Ludermir, “PSO for Reservoir Computing Optimization,” in *Artificial Neural Networks and Machine Learning-ICANN 2012*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2012, vol. 7552, pp. 685–692.
- [37] S. Zhong, X. Xie, L. Lin, and F. Wang, “Genetic algorithm optimized double-reservoir echo state network for multi-regime time series prediction,” *Neurocomputing*, vol. 238, no. C, pp. 191–204, may 2017.
- [38] S. Basterrech, E. Alba, and V. Snášel, “An Experimental Analysis of the Echo State Network Initialization Using the Particle Swarm Optimization,” in *Nature and Biologically Inspired Computing (NaBIC), 2014 Sixth World Congress on*, July 2014, pp. 214–219.
- [39] K. C. Chatzidimitriou and P. A. Mitkas, “A neat way for evolving echo state networks,” in *Proceedings of the 2010 Conference on ECAI 2010: 19th European Conference on Artificial Intelligence*. NLD: IOS Press, 2010, pp. 909–914.

- [40] F. Matzner, "Neuroevolution on the edge of chaos," in *Proceedings of the Genetic and Evolutionary Computation Conference*, ser. GECCO'17. New York, NY, USA: ACM, 2017, pp. 465–472.
- [41] J. Long, S. Zhang, and C. Li, "Evolving Deep Echo State Networks for Intelligent Fault Diagnosis," *IEEE Transactions on Industrial Informatics*, vol. 16, no. 7, pp. 4928–4937, 2020.
- [42] J. Lehman, J. Chen, J. Clune, and K. O. Stanley, "Safe mutations for deep and recurrent neural networks through output gradients," in *GECCO'18*. New York, NY, USA: Association for Computing Machinery, 2018, pp. 117–124.
- [43] J. B. Butcher, D. Verstraeten, B. Schrauwen, C. R. Day, and P. W. Haycock, "Reservoir Computing and Extreme Learning Machines for Non-linear Time-series Data Analysis," *Neural Networks*, vol. 38, pp. 76–89, feb 2013.
- [44] P. Aceituno, G. Yan, and Y.-Y. Liu, "Tailoring Echo State Networks for Optimal Learning," *iScience*, vol. 23, p. 101440, 08 2020.
- [45] H. Jaeger, "Controlling recurrent neural networks by conceptors," Jacobs University, Tech. Rep. 31, 2014, available at: <https://arxiv.org/pdf/1403.3369.pdf>.
- [46] —, "Using conceptors to manage neural long-term memories for temporal patterns," *Journal of Machine Learning Research*, vol. 18, pp. 1–43, 2017.
- [47] C. Gallicchio and A. Micheli, "Architectural and Markovian factors of Echo State Networks," *Neural Networks*, vol. 24, no. 5, pp. 440–456, 2011.
- [48] C. Sun, S. Hong, M. Song, Y. Zhou, Y. Sun, D. Cai, and H. Li, "TE-ESN: Time Encoding Echo State Network for Prediction Based on Irregularly Sampled Time Series Data," in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI-21)*, 05 2021.
- [49] F.-A. Fortin, F.-M. De Rainville, M.-A. Gardner, M. Parizeau, and C. Gagné, "DEAP: Evolutionary algorithms made easy," *Journal of Machine Learning Research*, vol. 13, pp. 2171–2175, jul 2012.
- [50] H. Jaeger, M. Lukoševičius, D. Popovici, and U. Siewert, "Optimization and applications of Echo State Networks with leaky-integrator neurons," *Neural Networks*, vol. 20, no. 3, pp. 335–352, 2007.
- [51] J. J. Steil, "Backpropagation-Decorrelation: online recurrent learning with $O(N)$ complexity," *In Proceedings of IJCNN'04*, vol. 1, 2004.