Efficient Design of Hardware-Enabled Recurrent Neural Networks[*]

Bogdan Penkovsky, Laurent Larger, and Daniel Brunner

FEMTO-ST / Optics Dept., UMR CNRS 6174, Univ. Bourgogne Franche-Comté,
15B avenue des Montboucons, 25030 Besançon Cedex, France

(Dated: March 7, 2020)

In this work, we propose a new approach towards the efficient design of reservoir computing hardware. First, we adapt the reservoir input mask to the structure of the data via linear autoencoders. We therefore incorporate the advantages of dimensionality reduction and dimensionality expansion achieved by conventional and efficient linear algebra procedures of principal component analysis. Second, we employ evolutionary-inspired genetic algorithm techniques resulting in a highly efficient optimization of reservoir dynamics. We illustrate the method on the so-called single-node reservoir computing architecture, especially suitable for implementation in ultrahigh-speed hardware. The combination of both methods and the resulting reduction of time required for performance optimization of a hardware system establish a strategy towards machine learning hardware capable of self-adaption to optimally solve specific problems. We confirm the validity of those principles building RC hardware based on a field-programmable gate array.

I. INTRODUCTION

Machine learning (ML) development has drastically progressed during the last decade. To only name a few examples, now machines can accurately describe images [1], identify and recognize faces [2], recognize speech [3] and compose music [4, 5]. In 2017, AlphaGo Zero and AlphaZero with no prior domain knowledge have beaten the best human and machine players both in Go [6] and chess [7] games. These all are challenges which, until recently, were thought to remain reserved for the human intelligence only.

However, the efficiency of current ML methods is restricted by hardware, which in its turn is fundamentally limited by the minimal transistor size. Today's silicon industry is already close to that limit. Another potential issue is related to the fact the vast majority of ML hardware rely only on a single design: the Turing-von Neumann architecture. That results in the second, conceptual limitation: our machines are constrained by their implementation's design, and that design is mostly one. This computing architecture can be regarded as a single meta-algorithm guiding existing ML techniques in only one possible direction. A look at nature can be illuminating to illustrate what potential directions might be missing. Unlike biological organisms, our machines (1) have only centralized memory, while in the human brain, memory and algorithms are inseparable; (2) are much less energy efficient; (3) cannot efficiently self-adapt to unforeseen challenges.

While the majority of discourses quickly tend towards the philosophical realm, we will constrain ourselves to improving ML hardware computational effectiveness. A viable way to circumvent present limitations is by shifting the design paradigm away from Turing–von Neumann architectures. Furthermore, this shift may also give insight into questions related to self-adapting hardware.

In this paper, we propose a hardware architecture which can efficiently adapt its dynamical properties to unforeseen changes in data [8] and can be later extended to non-silicon physical systems. As a proof-of-concept, we prototype a self-adapting system in simulation, evaluate it on a speech recognition benchmark, and verify the validity of the approach using a field-programmable gate array (FPGA) hardware. The results achieved by FPGA-based RC with limited bit resolution closely match those obtained in simulations with noise.

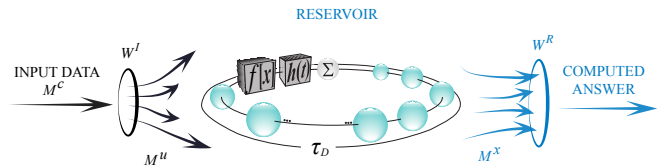### A. Reservoir computing. Single-node approach

Figure 1. **Single-node reservoir computing architecture**. The three core components are: input masking, reservoir transformation, and linear readout. First, the input data M^c are masked by multiplying the mask W^I, then the masked input M^u is transformed by reservoir's nonlinear delay dynamics. Finally, the answer is obtained by multiplying the readout W^R and the reservoir's state M^x.

In this work, we employ a so-called *single-node* reservoir computing (RC) architecture based on complex nonlinear delay dynamics [9]. Reservoir computing first appeared as a modification to recurrent neural networks training and was proposed independently in [10–12]. The core principle behind RC is a random mapping of the low-dimensional input information onto a higher-dimensional state space, where this information is expected to become linearly separable. Therefore, a linear readout should

[*] {bogdan.penkovskyi,laurent.larger,daniel.brunner}@femto-st.fr

arXiv:1805.03033v1 [cs.ET] 4 May 2018

be sufficient to interpret that information. As a consequence, instead of adapting the whole recurrent network, only the linear readout layer is trained.

The RC approach achieves multiple objectives: (1) the training procedure is fast, and (2) is guaranteed to converge using conventional linear algebra techniques, and crucially for the development of novel computing systems, (3) the fixed nonlinear part of the network can practically be delegated to low-level hardware, i.e. physically existing dynamical systems, not limited to digital electronics. Reservoir computing is a computation paradigm that potentially addresses the issues of inherently fast and energy-efficient hardware. This is mainly due to its support of information processing directly on the very hardware level. Several experimental implementations of RC hardware are known: using digital-analog electronics [13, 14], electro-optical and all-optical systems [15–17], and spintronic nanoscillators [18]. Both numerical and experimental RC systems often beat state-of-the-art in speed (such as in speech recognition task [19]) and accuracy (e.g. time series prediction [20]).

The single-node approach to RC takes advantage of delay dynamics, which has a high-dimensional, mathematically speaking even infinite-dimensional phase space, and can be interpreted as a virtual network [13, 21]. Single-node RC is frequently implemented in hardware as it is a technologically efficient way to construct a nonlinear reservoir network. The benefit of the method is the ability to reduce the physical neural network's size to a single nonlinear unit, thereby resulting in a smaller number of dynamical parameters to control. Moreover, the single-node RC architecture is especially suitable for ultrahigh-speed photonic hardware implementations [19].

A general single-node RC architecture is schematically represented in Fig. 1. Input information M^c is masked by the input mask W^I and then, mapped as M^u on a high-dimensional state space of a delay reservoir. Then, the reservoir's nonlinear response creates a state matrix M^x. The final answer is obtained by a linear readout, i.e. by multiplying matrix W^R. An introduction to the single-node RC can be found in [13, 19, 22].

B.   Input streamlining

In [23] it was suggested that information processing in the brain (e.g. in the primary visual cortex) is performed in three stages: first, input projection into principal feature dimensions, second, redundant information filtering, and finally third, higher-level information processing. Motivated by that strategy, we propose automatic feature weighting via redundant information filtering, to enhance the conventional random masking of RC. That is achieved by principal component analysis (PCA), a technique that constructs linear combinations of input features.

First, we apply PCA to remove the dimensions with the lowest variance, i.e. to compress the input data.

That allows us to focus on the input data's most relevant structure. Then, an inverse to compression operation, dimensionality expansion, is employed to restore the shape of the inputs. These two linear operators, compression and expansion, partially remove irrelevant feature information, such as noise. Therefore, PCA acts as a particular case of *autoencoders* [24]. Finally, random masking conventional to RC is performed in order to map the information onto a higher dimensional state space of the reservoir.

C.   Self-adapting reservoir dynamics

Due to the simplified training step in RC, the main action in system performance optimization is dynamical system parameters exploration. While the single-node RC method's complexity is much reduced, this can still be a substantial bottleneck for real-world applications since each problem may require a different set of dynamical parameters. For instance, a set of optimized speech recognition system parameters is potentially different from that of handwriting recognition. Therefore, quick parameter search is crucial when adapting RC to a new task.

Another case when dynamical parameters optimization is essential is testing new RC substrates when there is no prior RC parameter estimate. This becomes even more relevant when choosing between several alternative dynamical RC systems that differ in materials. For example, performing speech recognition in a bucket as in [25], which liquid is more suitable? A mixture, if then in which proportions, at which temperature, how deep would be an optimal reservoir?

As illustrated, one typically deals with multidimensional parameter optimization. There are two quite contrary approaches towards RC parameters optimization which are currently prevailing. One is a so-called parameter fine-tuning (essentially, trial and error) method. Here, one relies on often partially heuristic arguments why a certain set of starting parameters might be well suited. From that point one searches for the nearest, potentially only local, performance optimum. Although this ad-hoc practice can be frequently observed in the ML community, it does not guarantee an optimal parameter combination.

An opposite case is a more systematic optimization approach, the grid search (GS) technique. The method consists in an exhaustive search of all parameter combinations under certain constraints. GS has its advantage in guaranteeing the identification of a global parameter optimum, provided that the parameter search grid is sufficiently dense. In addition, the technique provides a multidimensional error landscape, giving insight into the structure of the parameter space and through that potentially into the relationship between task and computing system. However, GS may take substantial optimization time due to the exponentially increasing amount of data

points with each additional optimized parameter. Therefore, in practice GS is often limited to three-four scanned parameters.

In the present work, we provide a strategy how to create a data-driven self-adapting reservoir dynamics by employing an evolutionary selection-inspired technique known as genetic algorithm (GA) [26]. The GA optimization method can be regarded as a sweet spot between the two optimization extrema: it provides a systematic search while dramatically reducing the number of trials, making GA especially suitable for RC hardware design.

II.   METHODS

A.   Enhanced masking via PCA

In our approach, we adapt the input mask to the data structure particular to each task. A new input mask $\mathbf{W}_2^I$ is constructed as a superposition of three linear operators, compression $\mathbf{W}_c$, decompression $\mathbf{W}_c^\intercal$, and conventional (random) masking $\mathbf{W}^I$:

$$\mathbf{W}_2^I = \mathbf{W}^I \cdot \mathbf{W}_c^\intercal \cdot \mathbf{W}_c, \qquad (1)$$

where the transposed matrix pair $\mathbf{W}_c$ and $\mathbf{W}_c^\intercal$ is calculated using the standard unsupervised dimensionality reduction technique of principal component analysis [27, 28]; $\mathbf{W}^I$ is randomly generated as usual for RC. The resulting mask $\mathbf{W}_2^I$ remains fixed during all RC experiments for the given task. Furthermore, as it optimizes input information content, it is not optimized for a particular set of dynamical reservoir dynamics.

Principal component analysis (PCA) can be described as follows. First, an isolated subset of data is selected such that it reflects the data distribution of the whole dataset. Then, a covariance matrix $\mathbf{\Sigma} \in \mathbb{R}^{M \times P_0}$ is constructed using said subset. Here, M is the input dimensionality and P_0 is the total number of feature vectors in the subset. During the next step, a new matrix $\mathbf{W} \in \mathbb{R}^{M \times M}$ is obtained such that columns in $\mathbf{W}$ are eigenvectors of $\mathbf{\Sigma}$ sorted by decreasing magnitude of corresponding eigenvalues. This is achieved via *singular value decomposition*. Finally, a compression matrix $\mathbf{W}_c \in \mathbb{R}^{M' \times M}$ is constructed by selecting the first $M' < M$ vector-columns of $\mathbf{W}$, the principal components, and transposing the resulting matrix.

The superposition of compression $\mathbf{W}_c$ and decompression $\mathbf{W}_c^\intercal$ operators is an autoencoder. The autoencoder $\mathbf{W}_c^\intercal \cdot \mathbf{W}_c$ facilitates general data structure extraction by learning efficient data representation, while $\mathbf{W}^I$ helps to map the input data onto a higher dimensional state space. Therefore, the new operator $\mathbf{W}_2^I$ can be interpreted as a mask made more sensitive to the most relevant features in the input data, rather than the commonly employed simple random feature mapping via $\mathbf{W}^I$.

Another implication for ML hardware implementation of our architecture is tackling the input bottleneck. By decomposing $\mathbf{W}_2^I$ into two independent steps, first compression $\mathbf{W}_c$ and second masking and decompression $\mathbf{W}^I \cdot \mathbf{W}_c^\intercal$, both steps can be performed by different units. One can therefore preprocess the information by a simplistic special unit according to the first step. The information ultimately to be injected into the physical reservoir is then compressed at ratio M/M'.

B.   Reservoir computing

Prior to the reservoir transformation (Fig. 1), input sample matrix $\mathbf{M}^c$, consisting of L input feature vectors $\mathbf{c}(n), n = 1 \ldots L$, is first masked by multiplying an input mask $\mathbf{W}_2^I$ and then, temporally encoded:

$$
\begin{aligned}
\mathbf{W}_2^I \mathbf{c}(n) &= \mathbf{W}_2^I \left(c_1(n), c_2(n), \ldots, c_M(n) \right) \\
&= (u_1(n), u_2(n), \ldots, u_N(n)) \\
&= (u(t+\theta), u(t+2\theta), \ldots, u(t+N\theta)),
\end{aligned}
\qquad (2)
$$

where M is the input data dimensionality, and N is the reservoir network size. $\mathbf{W}_2^I$ is calculated from Eq. (1) with $\mathbf{W}^I \in \mathbb{R}^{N \times M}$, $(N > M)$ having weights randomly drawn from $\{-0.4; 0; 0.4\}$ (30% connectivity) and remaining fixed for all experiments. Finally, the temporally-encoded input signal $u(t)$ is kept constant in-between times $(t + i\theta, t + (i+1)\theta)$, $i = 1 \ldots N$, corresponding to the temporal separation between the virtual nodes [13]. This temporal encoding technique is sometimes called a *sample-and-hold* operation.

The temporal information input signal $u(t)$ is subsequently processed by the delayed-feedback nonlinear system reservoir (Eq. (3)). The choice of this particular reservoir dynamics model was motivated by its recent implementation as a substrate for numerous photonic RC devices [13, 16, 19], and can often be described by the low-pass delay-differential equation

$$\tau \dot{x}(t) = -x(t) + f\left(x(t - \tau_D) + \rho u(t)\right). \qquad (3)$$

In our case, we employ $f(x) = \beta \sin^2(x + \Phi_0)$ as nonlinearity. The nonlinear dynamics parameters τ, β, Φ_0, and ρ are subject to optimization while delay time $\tau_D = 6$ is kept constant in our experiments. Moreover, similar bandpass-filtered systems can be easily implemented in in electro-optical substrates [29].

The result of RC $\mathbf{y}(n)$ is computed as:

$$\mathbf{y}(n) = \mathbf{W}^R \mathbf{x}(n), \qquad (4)$$

where vector $\mathbf{x}(n) = (x(t+\theta), x(t+2\theta), \ldots, x(t+N\theta))$ is the decoded nonlinear reservoir response (Eq. 3). The linear readout weights $\mathbf{W}^R$ are obtained on a computer from previously processed data samples (Eqs (2)-(3)) using the ridge regression:

$$\mathbf{W}^R = (\mathbf{M}_x \cdot \mathbf{M}_x^\intercal + \lambda \cdot \mathbf{I})^{-1}(\mathbf{M}_x \cdot \mathbf{T}^\intercal), \qquad (5)$$

GA meta-parameter	Value
Crossover rate	88%
Mutation rate	12%
Crossover prob.	50%
Bit mutation prob.	20%
Population size	20
Total no. of generations	40

Table I. Summary of GA parameters. Total no. of RC training: $40 \times 20 = 800$ evaluations.

where $\lambda \ll 1$ is a small regularization constant. $\mathbf{M}_x \in \mathbb{R}^{N \times Q}$ is a feature matrix of concatenated horizontally state vectors $\mathbf{x}(n)$. $\mathbf{T} \in \mathbb{R}^{K \times Q}$ is a teacher matrix, Q denotes output dimensionality, and K the number of training feature vectors. For classification tasks, the teacher is a *one-hot encoded* matrix, i.e. consists of target answer vectors $\mathbf{y}^{\text{tgt}} \in \mathbb{R}^{K \times 1}$ where the only nonzero elements correspond to the correct class label.

C. Parameter self-optimization

Genetic algorithms (GAs) is a family of evolutionary-inspired techniques. Our particular GA implementation can be summarized as follows. Optimized parameters are encoded in binary chromosomes [30]. That encoding scheme allows for a simple definition of the crossover and mutation operators as bit exchange and bit flipping, respectively.

Parameters of GA (meta-parameters) are summarized in Table I. The crossover (mutation) rate is the offsprings percentage in a new generation by crossover (mutation). The probability of bit exchange in crossover is 50%, while the probability of bit flip is 20% during mutation. The number of chromosomes is 20 in a single generation.

D. Hardware implementation

We employ an Artix-7 (XC7A100T) FPGA chip as a digital hardware substrate for RC. The implemented architecture (Fig. 2) is a pipeline of three components working in parallel: masking, delayed-feedback dynamics, and the readout. The asynchronous communication between the modules implements a data strobe-like communication protocol [31].

First, the information input is compressed on a computer using matrix $\mathbf{W}_c$ precomputed via PCA. The resulting data are transferred to the FPGA via a USB cable using the serial UART protocol. The masking block is implemented on FPGA and realizes both masking and decompression $\mathbf{W}^I \cdot \mathbf{W}_c^{\mathsf{T}}$ as a single matrix-vector multiplication operation. The data are then transferred to the reservoir block which simulates the delay dynamics of Eq. (3) using the second-order Heun's method.

The FPGA implements the 16-bit fixed-point arithmetic, thus introducing quantization (digitization) noise.
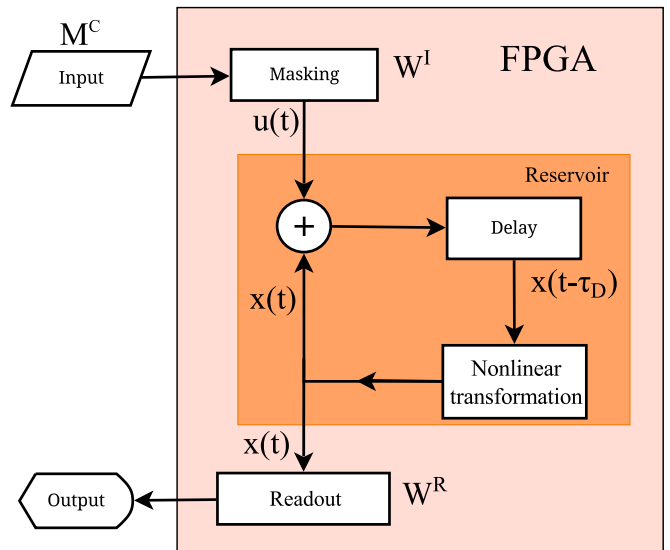

Figure 2. **FPGA-based standalone RC architecture** implements all essential RC blocks: masking, reservoir, and readout (cf. Fig. 1). Arrows represent 16 bit wide asynchronous communication buses.

Parameter	Min value	Max value	Resolution step
τ	$7.8 \cdot 10^{-3}$	0.99	2^{-7}
β	-3.98	3.98	2^{-6}
Φ_0	0	π	2^{-6}
ρ	-3.88	3.88	2^{-3}

Table II. Parameter ranges used for GA search

The impact of quantization noise could be strongly reduced by an implementation based on a floating-point module. However, the downsides of the floating-point FPGA implementation are more consumed programmable logic area and potentially slower processing rates. To demonstrate the practical applicability to other RC realizations, we stick to the less accurate fixed-point representation natively supported by our hardware.

During the training step, the system is run without the readout component. The resulting dynamics is sent to a computer where the readout matrix $\mathbf{W}^R$ is obtained with Eq. (5). During the testing step, to avoid model over-fitting, we utilize a separate testing dataset, i.e. data neither used in FPGA training, nor in model optimization.

III. RESULTS

To benefit from the underlying recurrent network, we apply RC to a time-dependent signal, human speech. The benchmark employed in this paper is a speech recognition task based on the clean isolated digits subset of Aurora-II database [32] (2412 samples). Following the established speech recognition paradigm, we model the dynamics of the inner ear and utilize Lyon model *cochleagrams* [33] as
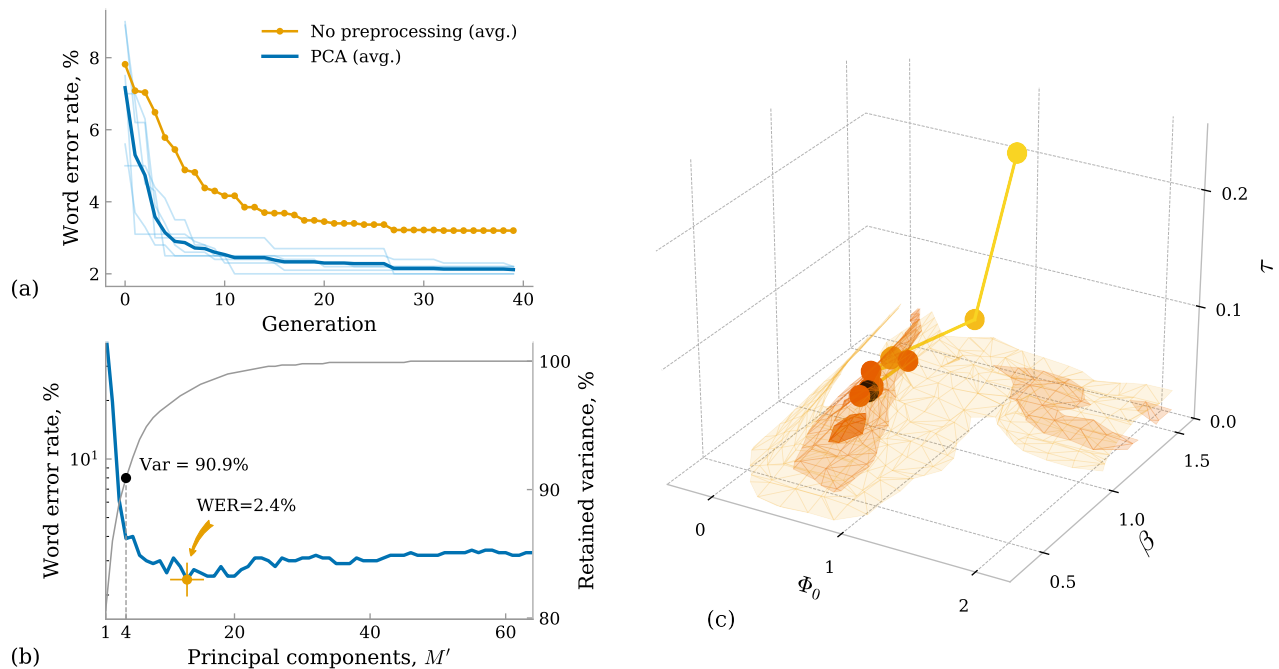
Figure 3. (a) **Genetic algorithm convergence averaged over 6 individual runs**. Individual genetic algorithm (GA) runs with PCA preprocessing are illustrated by narrow light-blue lines. The PCA preprocessing not only provides better accuracy (bold blue line) but also faster GA convergence and better overall performance than without PCA (orange dotted). Both experiments were conducted for a network of $N = 600$ virtual neurons. (b) **Selecting the number of principal components**. Principal component analysis predicts four as the minimal principal components number for current dataset. That corresponds to the retained variance of 90.9%. The most accurate value WER = 2.4% is achieved at 13 principal components (orange crosshair) with 97.6% retained variance. The small number of principal components ($M' = 13$) compared to the original number of channels ($M = 64$) indicates that the space containing human speech is sparse. Fixed (suboptimal) dynamical parameters are $\tau = 5 \cdot 10^{-3}$, $\tau_D = 6$, $\beta = 0.8$, $\Phi_0 = 0.3$, $\rho = 1.5$ (Eq. (3)). (c) **Projection of multidimensional error surfaces in 3D parameter space**. The fixed parameter is $\rho = 1.5$. The error landscape in the three parameter dimensions (Φ_0, β, τ) is characterized by extensive grid search (11,340 data points). An example evolution of the best chromosomes in each GA generation is visualized with circles converging to a local error minimum after ~ 20 generations. The final parameter obtained by this GA run is WER = 2.5% marked in black color. The nested error isosurfaces correspond to WER = $10, 5, 3\%$, respectively. The darkest orange volume contains the absolute error minimum of WER = 1.9%.

64-dimensional inputs to the reservoir described by Eqs (2)-(3).

We start our experimentation with unoptimized reservoir dynamics parameters and perform GA search. Table II summarizes parameter ranges selected with respect to *physically meaningful* RC dynamics. For instance, the delay system Eq. (3) is π-periodic because of the nonlinear function $f(x) = \sin^2(x)$, therefore we restrict $\Phi_0 \in [0; \pi]$; parameter τ cannot be large with respect to delay time τ_D, otherwise the system's complexity is substantially reduced; finally, β and ρ cannot be large otherwise the system will bifurcate away from useful dynamics. Otherwise, we do not provide any knowledge common to RC implementations (such as edge-of chaos), i.e. dynamics parameters are self-adapted to the speech recognition task.

To evaluate the classification accuracy during GA and GS optimizations, a so-called two-fold cross-validation is employed. First, training is performed on a group of 500 digits and another group of 500 digits is used for validation. Then, the roles are reversed, i.e. training is performed on the second group and validation, on the first one. Finally, a separate dataset of 1000 digits is used for testing the FPGA implementation. The remaining 412 samples are employed for PCA. Each individual step involved in our procedure is therefore carried out on a unique data-set, ensuring that findings can be transferred to applications with a typical continuous stream of input data. As an error measure (loss function) we utilize word error rate (WER), i.e. the ratio between errors and total number of evaluated samples.

To run the experiments quickly, we keep the reservoir size at a moderate value of $N = 600$ nodes; though our hardware could support substantially larger systems. The genetic algorithm efficiently converges to optimal dynamics settings, see Fig. 3(a), orange dotted curve. The best obtained parameters are $\tau = 7 \cdot 10^{-2}, \beta = -1.69, \Phi_0 = -1.33, \rho = 1.5$ with WER = 3.8%.

In the next step, we study the impact of dimensionality reduction on the classification accuracy. With the

help of PCA, we decrease the number of input dimensions by removing the principal components corresponding to the smallest eigenvalue magnitudes, i.e. containing the redundant information. General practice in PCA is to reduce the number of principal components so that at least 90% of variance is preserved. Therefore, we anticipate that the minimal number of principal components that can be used with these data is four (90.9% variance). This hypothesis is confirmed in Fig. 3(b), where the error sharply increases when the number of principal components goes below four. Principal component analysis shows that the best result in Fig. 3(b) is obtained for 13 principal components (97.6% of variance). Thus, we may conclude that the remaining $64 - 13 = 51$ principal components carry 2.4% of non-essential information such as noise. By removing those principal components we are able to effectively filter the residual information, which improves the recognition accuracy (Fig. 3(b)). In the rest of our experiments we reduce the number of input dimensions to seven principal components, preserving thus 94.9% of variance. That corresponds to $64/7 \simeq 9$ times compression rate. In our case, where data transfer is serial, this compression rate substantially reduces the transmission time to FPGA processing unit. Furthermore, ROM memory capacity, containing coefficients for the input masking, is reduced 9 times comparing to conventional masking without compression. Crucially, according to Fig. 3(b), recognition performance is hardly affected by this stronger compression.

The selected dimensionality reduction consistently improves the overall accuracy (Fig. 3(a), thick blue curve). The parameters obtained by the GA $\tau = 7.8125 \cdot 10^{-3}$, $\tau_D = 6$, $\beta = -1.09375$, $\Phi_0 = -3.3125$, $\rho = 1.5$ result in an optimal performance of WER = 2.1%. Moreover, PCA preprocessing also helps the GA to converge faster. This can be explained by the fact that the space of sounds (and therefore, cochleagrams) is sparse with respect to the words pronounced by humans, hence the majority of the sound space is populated by information only weakly correlated to the information content. Therefore, a significant part of information contained in cochleagrams is redundant.

To better illustrate the GA search (case of PCA preprocessing), we perform grid search (GS) along the three most significant parameter dimensions, i.e. Φ_0, β, and τ, crucially forced to use much coarser resolution: less than 25 points per dimension already result in a total of 11,340 points in parameter space. The exhaustive GS in all four parameter dimensions with resolution comparable to the one we used in GA would take 5,106 times longer than GA. If, as here for our case, a single GA run in our implementation takes around an hour, GS would take more than seven months. Adding an additional (fifth) parameter dimension scanned along e.g. 100 points, would immediately increase the GS time to 59 years. Both GA and GS can be parallelized, but in case of GS, parallelization cannot overcome the exponential growth off necessary resources. That clearly highlights the advantage of GA over
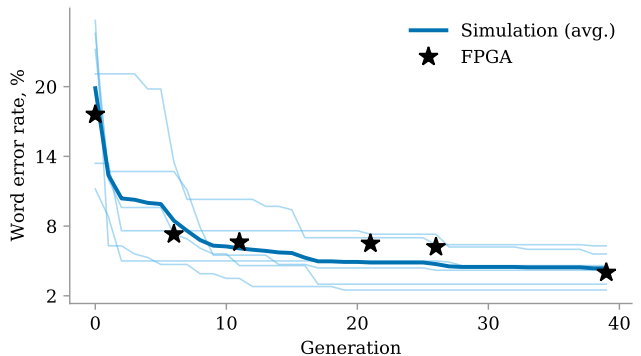

Figure 4. **Genetic algorithm convergence in hardware**. Simulation with noise (lines), average over 6 runs in bold; and FPGA realizations (stars).

exhaustive GS.

Figure 3(c) reveals error isosurfaces in the three-dimensional parameter space obtained as a result of GS. The isosurfaces present nested objects corresponding to WERs = 10%, 5%, and 3%. Error rates obtained from the GA search are visualized with circles. It can be seen that the topmost circle is a result of the random search (zero generation), corresponding WER $\simeq$ 16%. Then, as GA is efficiently converging, the circles are quickly approaching an acceptable local minimum. Although Fig. 3(c) illustrates the GA search in only three dimensions, GA is simultaneously optimizing parameters in all four dimensions.

Finally, we apply the GA technique to an FPGA-based RC. Before we actually implement RC on FPGA, we accurately estimate an optimal parameter set offline on a computer. In order to take into account the quantization noise in FPGA with limited bit resolution, we simulate the dynamics Eq. (3) adding the white noise of level $1.2 \cdot 10^{-4}$ to the dynamical variable $x(t)$, the delay term $x(t - \tau_D)$, and the result of nonlinear transformation f. The noise level corresponds to the quantization noise on an actual FPGA device. Additionally, to better model the behavior of our hardware, in the beginning of numerical experiment we add the noise of the same magnitude to masking coefficients and also we repeat the procedure with readout coefficients right after training.

Figure 4 shows the results simulated on a computer GA (lines). Due to the noise in the experiment, the accuracy of classification has degraded overall. We then select parameters from a GA optimization creating results close to the average performance convergence. These are used for the RC implementation in the FPGA. Computational results obtained fully autonomously by the FPGA correspond to the black stars. They excellently match the average convergence obtained from our offline model, thereby validating our approach.

IV. CONCLUSION

We have proposed a technique towards practical application of reservoir computing (RC). The technique consists of two components: data-driven input mask optimization and efficient dynamical parameter optimization. We have illustrated those methods and their strong positive impact on the speech recognition. The advanced input masking reduced the input data to be transfered to the device by 9 times and lowered the average classification error by 1.7%, a 1.8 fold improvement. The improved parameter optimization reduced the number of iterative optimization steps by 5,106 times when compared to exhaustive grid search.

We took advantage of the fact that the exact RC model was known in advance and were able to run genetic algorithm (GA) optimizations offline on a PC. We took into account the hardware's quantization noise and therefore added an additional white noise term to the model. Finally, we have built an FPGA RC confirming our evolutionary technique. This illustrates how our method can be applied to various physically existing RC systems where noise is inevitably present.

Another significant benefit of GA is that the method could be applied even when the exact model of the optimized system was unknown. This would enable RC optimization online, i.e. directly on the actual hardware, as it was done e.g. in [34]. In this work we have shown that evolutionary-inspired optimization can significantly reduce the time to adapt RC dynamics to an unforeseen task. We leave the implementation of online GA to the future investigations as the next logical step towards self-adapting hardware [35].

ACKNOWLEDGMENTS

[1] Y. LeCun, Y. Bengio, and G. E. Hinton, Nature **521**, 436 (2015).

[2] B. Amos, B. Ludwiczuk, and M. Satyanarayanan, *Open-Face: A general-purpose face recognition library with mobile applications*, Tech. Rep. (CMU-CS-16-118, CMU School of Computer Science, 2016).

[3] A. Graves, A. Mohamed, and G. Hinton, Icassp , 6645 (2013).

[4] D. D. Johnson, in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 10198 LNCS (2017) pp. 128–143.

[5] D. D. Johnson, R. M. Keller, and N. Weintraut, Eighth International Conference on Computational Creativity (ICCC'17) , 151 (2017).

[6] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, and L. Sifre, Nature Publishing Group **550**, 354 (2017).

[7] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. Lillicrap, K. Simonyan, and D. Hassabis, , 1 (2017).

[8] J. Torresen, Field Programmable Logic and Application , 1 (2004).

[9] M. C. Soriano, J. García-Ojalvo, C. R. Mirasso, and I. Fischer, Reviews of Modern Physics **85**, 421 (2013).

[10] H. Jaeger, German National Research Center for Information Technology GMD Technical Report **148:34** (2001).

[11] J. Steil, in *2004 IEEE International Joint Conference on Neural Networks (IEEE Cat. No.04CH37541)*, Vol. 2 (IEEE, 2004) pp. 843–848.

[12] W. Maass, T. Natschläger, and H. Markram, Neural Comput **14**, 2531 (2002).

[13] L. Appeltant, M. C. Soriano, G. V. D. Sande, J. Danckaert, S. Massar, J. Dambre, B. Schrauwen, C. R. Mirasso, and I. Fischer, Nature Communications **2**, 466 (2011).

[14] M. C. Soriano, S. Ortin, L. Keuninckx, L. Appeltant, J. Danckaert, L. Pesquera, and G. der Sande, IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS **26**, 388 (2015).

[15] L. Larger, M. C. Soriano, D. Brunner, L. Appeltant, J. M. Gutierrez, L. Pesquera, C. R. Mirasso, and I. Fischer, Optics express **20**, 3241 (2012).

[16] D. Brunner, M. C. Soriano, C. R. Mirasso, and I. Fischer, Nature Communications **4**, 1364 (2013).

[17] K. Vandoorne, P. Mechet, T. Van Vaerenbergh, M. Fiers, G. Morthier, D. Verstraeten, B. Schrauwen, J. Dambre, and P. Bienstman, Nature communications **5**, 3541 (2014).

[18] J. Torrejon, M. Riou, F. A. Araujo, S. Tsunegi, G. Khalsa, D. Querlioz, P. Bortolotti, V. Cros, K. Yakushiji, A. Fukushima, H. Kubota, S. Yuasa, M. D. Stiles, and J. Grollier, Nature **547**, 428 (2017).

[19] L. Larger, A. Baylón-Fuentes, R. Martinenghi, V. S. Udaltsov, Y. K. Chembo, and M. Jacquot, Physical Review X **7**, 011015 (2017).

[20] H. Jaeger, Science **304**, 78 (2004).

[21] L. Larger, B. Penkovsky, and Y. Maistrenko, Physical Review Letters **111**, 054103 (2013).

[22] J. Schumacher, H. Toutounji, and G. Pipa, in *Artificial Neural Networks*, Vol. 4, edited by K. N. Koprinkova-Hristova P., Mladenov V. (Springer, Cham, 2015) artificial ed., pp. 63–90.

[23] R. Miikkulainen, J. A. Bednar, Y. Choe, and J. Sirosh, Psychology of Learning and Motivation - Advances in Research and Theory **36**, 257 (1997).

[24] T. Manning-Dahan, *PCA and Autoencoders*, Available at: http://tylermd.com/pdf/pca_ae.pdf. Accessed 26 Apr. 2018.

[25] C. Fernando and S. Sojakka, in *Proceedings of the 7th European Conference on Artificial Life* (2003) pp. 588–597.

[26] K. Deb, Sadhana **24**, 293 (1999).

[27] K. Pearson, The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science **2**, 559 (1901).

[28] H. Hotelling, Journal of Educational Psychology **24**, 417 (1933).

[29] L. Larger, B. Penkovsky, and Y. Maistrenko, Nature Communications **6**, 7752 (2015).

[30] The employed GA library was https://hackage.haskell.org/package/GA.

[31] B. Penkovsky, *Theory and Modeling of Complex Nonlinear Delay Dynamics Applied to Neuromorphic Computing*, Ph.D. thesis, Universite de Bourgogne Franche-Comte (2017).

[32] H.-G. Hirsch and D. Pearce, Isca Itrw Asr2000 , 181 (2000).

[33] M. Slaney, Apple Technical Report #13 , 1 (1988).

[34] A. Thompson, Evolvable Systems: From Biology to Hardware **1259**, 390 (1996).

[35] C. Aporntewan and P. Chongstitvatana, in *Proceedings of the 2001 Congress on Evolutionary Computation (IEEE Cat. No.01TH8546)*, Vol. 1 (IEEE) pp. 624–629.