Spintronics for image recognition : performance benchmarking via ultrafast data-driven simulations

Anatole MOUREAUX,[1, a)] Chloé CHOPIN,[1] Laurent JACQUES,[2] and Flavio ABREU ARAUJO[1, b)]
[1)]*Institute of Condensed Matter and Nanosciences, Université catholique de Louvain, Place Croix du Sud 1, 1348 Louvain-la-Neuve, Belgium*
[2)]*Institute for Information and Communication Technologies, Electronics and Applied Mathematics, Université catholique de Louvain, Avenue Georges Lemaître 4, 1348 Louvain-la-Neuve, Belgium*

We present a demonstration of image classification using a hardware-based echo-state network (ESN) that relies on spintronic nanostructures known as vortex-based spin-torque oscillators (STVOs). Our network is realized using a single STVO multiplexed in time. To circumvent the challenges associated with repeated experimental manipulation of such a nanostructured system, we employ an ultrafast data-driven simulation framework called the data-driven Thiele equation approach (DD-TEA) to simulate the STVO dynamics. We use this approach to efficiently develop, optimize and test an STVO-based ESN for image classification using the MNIST dataset. We showcase the versatility of our solution by successfully applying it to solve classification challenges with the EMNIST-letters and Fashion MNIST datasets. Through our simulations, we determine that within a large ESN the results obtained using the STVO dynamics as an activation function are comparable to the ones obtained with other conventional nonlinear activation functions like the reLU and the sigmoid. While achieving state-of-the-art accuracy levels on the MNIST dataset, our model's performance on EMNIST-letters and Fashion MNIST is lower due to the relative simplicity of the system architecture and the increased complexity of the tasks. We expect that the DD-TEA framework will enable the exploration of more specialized neural architectures, ultimately leading to improved classification accuracy. This approach also holds promise for investigating and developing dedicated learning rules to further enhance classification performance.

I. INTRODUCTION

The impact of artificial intelligence (AI) on various industries and our daily lives has been transformative. It has facilitated critical functions such as medical image recognition for early disease diagnosis and optimization of manufacturing processes[1,2]. However, the widespread use of AI has raised concerns about its energy consumption and environmental impact during resource-intensive training and large-scale inference phases[3]. To address these challenges, growing efforts to develop energy-efficient alternatives with cutting-edge performance are underway. Moreover, conventional computers, despite decades of miniaturization and optimization, are nearing their limits in computing power[4,5]. This observation prompted researchers around the world to explore unconventional computing architectures like quantum computers and neuromorphic systems inspired by the human brain[6].

In response to these challenges, novel approaches in artificial intelligence have emerged, aiming to surpass the limitations of conventional methods. One such paradigm is reservoir computing (RC), which involves injecting data into a neural reservoir: a random recurrent neural network comprising nonlinearly activated units. This process projects data into a higher-dimensional space, ensuring that the outputs of the reservoir are linearly separable. Notably, only the weights of the readout layer need optimization, leaving the reservoir weights unchanged. Furthermore, the readout training step can be performed linearly using a simple linear regression, hence significantly reducing the training time compared to traditional algorithms reliant on gradient descent and error backpropagation. As the learning step can be performed in a single shot over the whole training set, RC offers the advantage of finding a global minimum of the loss function linearly, ensuring determinism and reproducibility[7]. The fixed nature of the reservoir makes it well-suited for hardware implementations, enabling the creation of dedicated hardware processors to optimize performance and energy efficiency[8]. Intriguingly, research has demonstrated that such reservoirs can be implemented using a single nonlinear unit multiplexed in time without sacrificing performance[9–11]. This feature further enhances its potential for hardware-based solutions.

Efficient data classification through physical reservoir computing has been demonstrated in various works, utilizing optoelectronic systems[12,13] as well as nonlinear oscillators known as vortex-based spin-torque oscillators[14] (STVOs). STVOs are nanometer-sized magnetic tunnel junctions (MTJs) that rely on spintronics to operate[15]. Spintronics, also referred to as spin electronics, describe the physical phenomena related to the transport of spin in magnetic materials and nanostructures. When an electric current is injected into these STVOs, they exhibit stable oscillations of the electrical resistance. By utilizing an amplitude-modulated signal, these oscillations display a strong nonlinearity with respect to the signal amplitude[16], leading to oscillating voltages across the MTJ. STVOs offer numerous advantages, such as low power consumption, minimal noise and compatibility with complementary metal-oxide-semiconductor (CMOS) technology[17], making them attractive candidates for nonlinear data transformation in RC-based neuromorphic applications[18–21].

Despite the progress in STVO-based hardware reservoirs, a significant challenge lies in their development, optimization and testing due to the laborious and costly process of manipulating nanostructures repeatedly. To overcome these

[a)]anatole.moureaux@uclouvain.be
[b)]https://flavio.be

limitations, simulating STVOs and their complex dynamics proves to be a more feasible approach. Micromagnetic simulations[22] (MMS) offer highly accurate results but are time and energy-consuming, making them impractical for simulating STVO-based neuromorphic computing and optimization. A computationally lighter alternative is the Thiele equation approach (TEA), which employs a single ordinary differential equation for each space direction to model STVO dynamics[23]. However, TEA is only quantitatively accurate in the resonant regime *i.e.*, when the magnetization does not oscillate. TEA is merely qualitative in the steady and transient oscillatory states *i.e.*, when the magnetization undergoes stable oscillations or switches from a stable state to another one. The very latter states are however crucial for neuromorphic computing[24].

A recent solution to this challenge is the data-driven Thiele equation approach (DD-TEA) developed by Abreu Araujo *et al.*. This semi-analytical model accurately describes the STVO dynamics in both steady and transient states, achieved by using numerical parameters extracted from a small set of MMS[24]. Consequently, DD-TEA enables simulation with the same level of accuracy as MMS but with a significant acceleration factor of 9 orders of magnitude[25].

This study aims to showcase the capability of a STVO-based neural reservoir in achieving state-of-the-art accuracy for image classification. Firstly, we will introduce our implementation of a time-multiplexed echo state network (ESN) using a single STVO. The data preprocessing and training procedures will be detailed. Next, we will explain how we simulate the STVO dynamics using DD-TEA. Subsequently, we will evaluate the network's performance in classifying various types of images, such as handwritten digits from MNIST[26], characters from EMNIST-letters[27] and clothing items from Fashion MNIST[28] (FMNIST). The STVO dynamics will then be compared to more conventional forms of nonlinearity: the reLU and the sigmoid functions. Those studies will be carried for various number of virtual nodes in the ESN reservoir to demonstrate the potential of DD-TEA for hyperparameter optimization.

II. METHODS

A. Time-multiplexed reservoir computing

Reservoir computing employs a recurrent neural network (RNN) with random and fixed internal weights. Each reservoir node exhibits a nonlinear behavior based on its internal dynamics. The input layer directly feeds data into the reservoir, and the internal state of the N_θ nonlinear nodes are collected in an output state matrix. The output state matrix is then transformed into the final output using the weights of the readout layer. These readout layer weights are the only adjustable parameters and can be learned using linear regression over the entire training set during the training phase. This approach enables RC to discover a global optimum of the loss function while minimizing the number of parameters that need to be learned.

Time-multiplexing enables the replacement of the reservoir with a single nonlinear unit[9]. Instead of having multiple nodes spatially interconnected within the reservoir, this singular node effectively emulates a network of virtual nodes linked sequentially through time. This unique capability arises from the transient state of the node, functioning as a short-term memory that facilitates the connection of successive states over time. Consequently, this approach simplifies the experimental setup for easy manipulation and allows for the exclusion of spatial interactions among nodes during simulations by trading on the execution time[11,29].

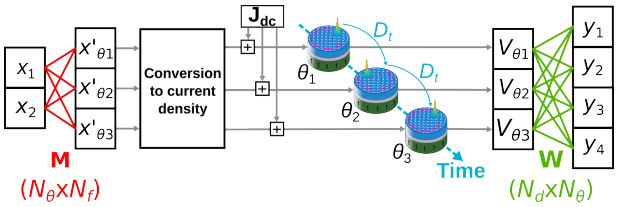

FIG. 1. Reservoir computing using a time-multiplexed, single chain connected echo state network. In this example, a 2-dimensional input $\mathbf{x}$ is sent into a reservoir of $N_\theta = 3$ virtual STVOs using the weights of the random matrix $\mathbf{M}$. Each virtual STVO is implemented by a single STVO delayed every D_t nanoseconds. The final output $\mathbf{Y}$ is reconstructed by linear combination from the output state matrix $\mathbf{V}$ using the learned matrix $\mathbf{W}$.

In our approach, images are injected into the virtual reservoir using a random connectivity matrix $\mathbf{M}$ with weights uniformly distributed between 1 and -1, defining the fixed weights of the input layer. These masked inputs are then injected sequentially into the single STVO every D_t nanoseconds. We choose D_t to be shorter than the characteristic time of the transient state of the STVO dynamics. The output state matrix $\mathbf{V}_{\text{out}}$ is sequentially recorded and mapped to the output matrix $\mathbf{Y}$ using the readout weights matrix $\mathbf{W}$ learned during the training phase (Fig. 1).

The virtual network lacks recurrent connections as it is directed towards future in a feed-forward fashion as represented in the example in Fig. 1. This is compensated by an artificial delay loop implemented by the matrix $\mathbf{M}$ in the input layer. This architecture is equivalent to the simple-chain ESN multiplexed in time presented in ref.[10]. Indeed, the STVO receives N_θ random linear combinations of the input features spaced in time by D_t nanoseconds, thus implementing a form of feedback[11]. Note that as N_θ is chosen much higher than the length of the input signal N_f, the time-multiplexing technique requires more execution time than for a spatially connected network.

B. Modeling the STVO dynamics

In a STVO, the energetically preferred ground state of the magnetization is identified by a non-uniform topological distribution referred to as vortex. The vortex results in a magnetization pattern that curls within the plane of the MTJ, while the core of the vortex displays magnetization oriented outward from the plane, as depicted in Fig. 2a. When the STVO is

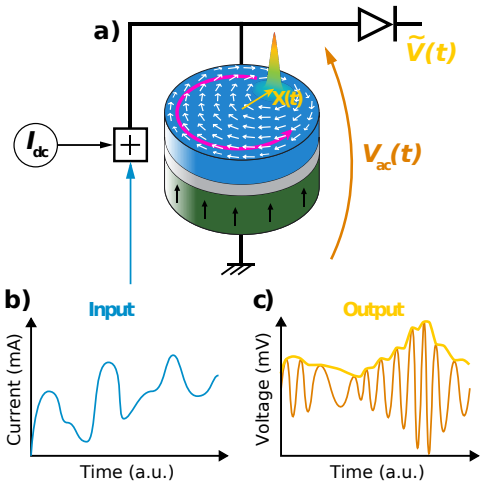

FIG. 2. **(a)** Radio-frequency oscillations of the vortex core (pink circular arrow) and of its reduced position $s(t)$ are triggered by **(b)** the injection of an amplitude-modulated signal in the STVO. **(c)** The non-linearity of $s(t)$ can be retrieved experimentally by recording the envelope $\tilde{V}(t)$ of the voltage $V_{\rm ac}(t)$ across the oscillator.

subjected to a time-varying amplitude-modulated input signal, such as the one shown in Fig. 2b, the vortex core undergoes stable circular radiofrequency oscillations within the plane of the STVO[14,30,31], represented by the pink arrow in Fig. 2a. The resistance of the STVO depends on the position $\mathbf{X}(t)$ of the vortex core, a phenomenon known as the tunnel magnetoresistance (TMR) effect[32]. Consequently, the electrical resistance of the STVO exhibits similar high-frequency dynamics, leading to an alternating voltage $V_{\rm ac}(t)$ across the oscillator due to Ohm's law, as seen in the orange curve in Fig. 2c. It's worth noting that the input signal must be combined with a dc current intensity $I_{\rm dc}$, referred to as the bias current intensity, to drive the dynamics of the vortex core.

The non-linearity of the STVO dynamics can also be observed in the time-dependent reduced position of the vortex core, denoted as $s(t) = ||\mathbf{X}(t)||/R$, with R being the radius of the oscillator[24]. This value represents the normalized distance between the vortex core and the center of the STVO. Experimentally, this is analogous to considering the envelope of the voltage $\tilde{V}(t)$ across the STVO, as illustrated by the yellow curve in Fig. 2c.

The reduced position $s(t)$ can be analytically modeled within the DD-TEA[24,25] framework using Eq. (1). In this equation, D_t (in nanoseconds) represents the sampling rate constant of the input signal and the rate at which the masked signal is sent into the STVO. The parameters $\alpha(t)$ and $\beta(t)$ are functions dependent on the amplitude of the input signal at time t. Additionally, the parameter $n(t)$ represents the level of complexity accounted for in the model and is also a function of the input signal amplitude at time t. To ensure the physical consistency of the model, α, β, and n are fitted to values obtained through MMS[24,25]. The reduced position $s(t)$ incorporates $s(t-D_t)$ to capture the transient state of the STVO when it is positioned between two different pieces of input. In the context of reservoir computing, the model described in Eq. (1) serves as the internal transfer function implemented by the virtual STVOs in the reservoir[8,12,13,21,29,33].

$$s(t) = \frac{s(t-D_t)}{\sqrt[n(t)]{\left(1+\frac{s(t-D_t)^{n(t)}}{\alpha(t)/\beta(t)}\right)\exp\left(-n(t)\alpha(t)D_t\right) - \frac{s(t-D_t)^{n(t)}}{\alpha(t)/\beta(t)}}} \tag{1}$$

C. Pre-treating the data

The following sections describe how the classification of handwritten digits from the MNIST dataset can be performed by simulating STVO-based reservoir computing. The MNIST dataset consists in 60,000 training images $\mathbf{x}_i$ of digits between 0 and 9, and 10,000 testing images. The images are arrays of 28×28 grayscale pixels characterized by a value between 0 and 255.

The images of the training set are first normalized between 0 and 1 (see Fig. 3a), and then flattened into 1-D vectors. Then, their dimension is reduced from 784 to N_f values using principal components analysis (PCA)[34,35]. Selecting the $N_f = 44$ most significant components as in (Fig. 3b) allows explaining 80.33% of the variance in the training set. The weights of these components are later used to extract the principal components from the images of the testing set.

A reservoir containing N_θ nodes is then designed with only one STVO using the time-multiplexing technique[21,29,36]. The reservoir is a recurrent neural network characterized by a mask $\mathbf{M}$ of dimension $(N_\theta \times N_f)$ that plays the role of connectivity matrix. In our implementation, $\mathbf{M}$ is generated randomly. Its values are uniformly distributed between -1 and 1. Each sample of input data $\mathbf{x}_i$ is masked using $\mathbf{M}$ in order to spread the input in the network (see Eq. (2) and Fig. 4).

$$\underbrace{\mathbf{x}_i'}_{(N_\theta\times 1)} = \underbrace{\mathbf{M}}_{(N_\theta\times N_f)}\underbrace{\mathbf{x}_i^{\rm T}}_{(N_f\times 1)} \tag{2}$$

The response of the STVO is influenced by three parameters. The first one is the bias current intensity $I_{\rm dc}$ to which is added the input signal[25]. The second parameter is the amplitude of the input signal, spanning over a voltage range ΔV. This parameter can also be used to tune the dynamics of the

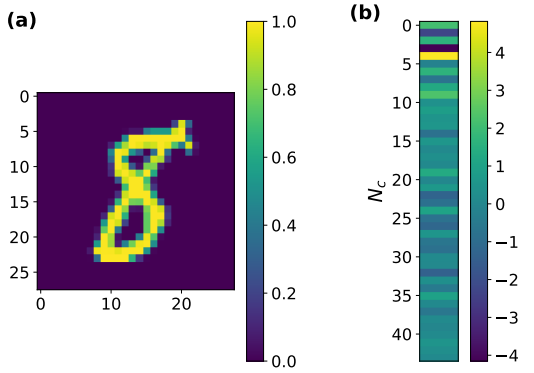

FIG. 3. **(a)** A sample from the MNIST training dataset, noted $\mathbf{x}_i$. **(b)** Projection of the sample on the first 44 components extracted using PCA. Note that $\mathbf{x}_i$ is a line vector. The scales represent the intensity of the pixels and of the signal after projection on the PCA components.

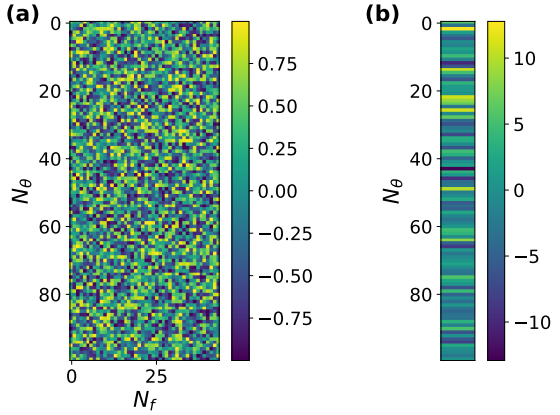

FIG. 4. **(a)** Example of a mask $\mathbf{M}$ representing a reservoir of $N_\theta = 100$ nodes. **(b)** The masked input $\mathbf{x}'_i$ from Fig. 3. The scales represent the value of the mask elements and the intensity of the masked signal.

oscillator[25]. The third parameter is the sampling rate constant D_t introduced with Eq. (1). The signal carried by $\mathbf{x}'_i$ is scaled in the $[-\Delta V/2, \Delta V/2]$ range, then added to I_{dc} and finally sampled accordingly to D_t as represented in Fig. 5.

The output $\mathbf{y}_i$ of the reservoir is then computed accordingly to Eq. (3) by feeding the scaled signal from Fig. 5 into Eq. (1) (Fig. 6).

$$\mathbf{y}_i := \Psi(\mathbf{x}'_i) := \sigma_{\text{STVO}}(\mathbf{M}\mathbf{x}_i^{\text{T}}) \tag{3}$$

D. Training the readout layer

This section describes the training procedure of the readout layer. The outputs $\mathbf{y}_i$ of the reservoir for the $60,000$ samples of the MNIST training set are computed in parallel accordingly to the procedure described in the previous section. These outputs $\mathbf{y}_{i,\text{train}}$ are then stacked horizontally to form the output matrix $\mathbf{V}_{\text{out, train}}$ of dimension $(N_\theta \times 60,000)$. The weight matrix $\mathbf{W}$ is then computed following Eq. (4). $\mathbf{T}_{\text{train}}$ (dimensions $(10 \times 60,000)$) is the target matrix containing the $60,000$ categorical targets l_i (*i.e.*, a column vector of 0s, except at the index of the corresponding category where it is equal to 1) of the training set stacked horizontally. This step consists in a linear regression between the outputs of the reservoir and the desired targets over the whole training set. Standard algorithms such as ridge regression can be used, so that the regularization parameter can be tuned accordingly to the problem to avoid overfitting. In our case, the Moore-Penrose pseudo-inverse matrix $\mathbf{V}^{\dagger}_{\text{out, train}}$ [37,38] is considered. This linear regression acts as the only training process in the whole framework.

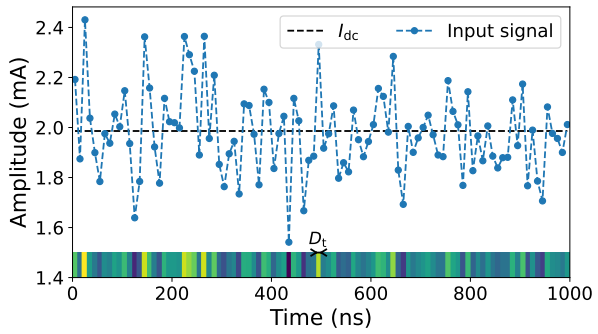

FIG. 5. The masked input $\mathbf{x}'_i$ scaled accordingly to I_{dc}, ΔV and D_t.

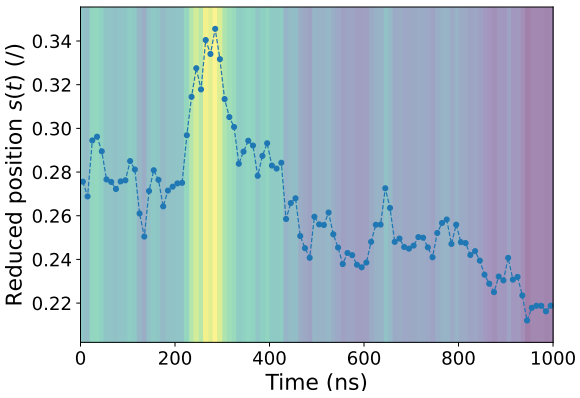

FIG. 6. Output $\mathbf{y}_i$ of the STVO for the input signal in Fig. 5.

$$\underbrace{\mathbf{W}}_{(10\times N_\theta)} = \underbrace{\mathbf{T}_{\text{train}}}_{(10\times 60,000)} \times \underbrace{\mathbf{V}^{\dagger}_{\text{out, train}}}_{(60,000\times N_\theta)} \tag{4}$$

E. Testing and classifying unseen data

To infer the classification of unseen data, the samples from the testing set are processed in the same way as during the training process (PCA, masking, scaling and STVO treatment). The outputs $\mathbf{y}_{i,\text{test}}$ of the STVO for all the samples of

the testing set are stacked horizontally to form the inference output matrix $\mathbf{V}_{\text{out, test}}$ of shape $(N_\theta \times 10,000)$.

$$\underbrace{\mathbf{Y}_{\text{test}}}_{(10\times 10{,}000)} = \underbrace{\mathbf{W}}_{(10\times N_\theta)} \times \underbrace{\mathbf{V}_{\text{out, test}}}_{(N_\theta \times 10{,}000)} \tag{5}$$

The estimated categories $\mathbf{Y}_{\text{test}}$ are obtained with Eq. (5). Finally, the classification of each data sample $\hat{\mathbf{T}}_{\text{test}}$ is obtained with Eq. (6), where the argmax function is performed column-wise on $\mathbf{Y}_{\text{test}}$.

$$\underbrace{\hat{\mathbf{T}}_{\text{test}}}_{(1\times 10{,}000)} = \text{argmax} \underbrace{(\mathbf{Y}_{\text{test}})}_{(10\times 10{,}000)} \tag{6}$$

The overall process is depicted in Fig. 10 with a focus on digit 8 and can be summarized by Eq. (7).

$$\hat{\mathbf{T}}_{\text{test}} = \text{argmax}\left(\mathbf{T}_{\text{train}}\left(\left[\sigma_{\text{STVO}}\left(\mathbf{M}\mathbf{x}^{\text{T}}_{i,\text{train}}\right)\right]_1^{60{,}000}\right)^{\dagger}\left[\sigma_{\text{STVO}}\left(\mathbf{M}\mathbf{x}^{\text{T}}_{i,\text{test}}\right)\right]_1^{10{,}000}\right) \tag{7}$$

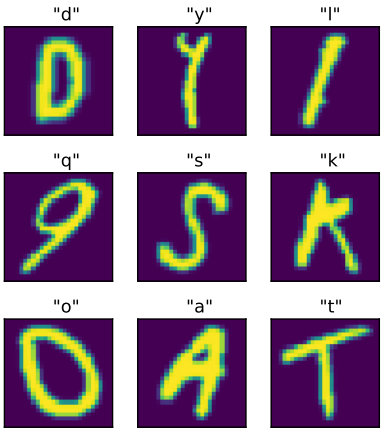

FIG. 7. Some images randomly selected in the EMNIST-letters training set.

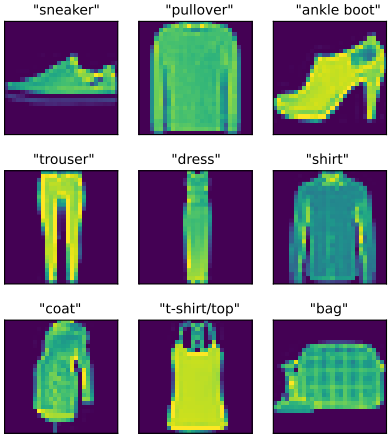

FIG. 8. Some images randomly selected in the FMNIST training set.

III. GENERALIZABILITY

To assess the generalizability of our implementation of simulated hardware reservoir computing, the training and inference processes described in the previous part are applied on two additional image datasets. The first one is the EMNIST-letters dataset: an extension of MNIST to handwritten letters[27]. It consists of 124,800 training images and 20,800 testing images, classified into 26 classes corresponding to the 26 letters in the alphabet (see Fig. 7). Each class contains lower and upper case letters.

The second dataset is Fashion MNIST (FMNIST[28]) and is composed of 60,000 training images and 10,000 testing images, all extracted from the Zalando website. The images are separated into 10 classes representing clothing items: {*t-shirt/top*, *trouser*, *pullover*, *dress*, *coat*, *sandal*, *shirt*, *sneaker*, *bag*, *ankle boot*}. Some examples are shown in Fig. 8.

IV. PERFORMANCE EVALUATION

The performance of our time-multiplexed STVO-based ESN at classifying images from the MNIST, EMNIST-letters and FMNIST datasets will be assessed as follows. A STVO-based ESN will be trained with each dataset accordingly to the procedure described in the previous section. The number of PCA components selected will be chosen to ensure an explained variance of 80%. This allows to reduce the dimension of the data from 784 to $N_f = 38$ (resp. $N_f = 24$) for the EMNIST-letters (resp. FMNIST) dataset. The accuracy (the ratio of correctly classified samples in the testing dataset) and the normalized root-mean-square error (NRMSE) between each predicted output $\mathbf{y}_i$ and the corresponding target $\mathbf{t}_i$ defined as $\text{NRMSE} := \sqrt{\overline{(\mathbf{y}_i - \mathbf{t}_i)^2}}/\overline{\mathbf{t}_i}$ will be averaged over the entire testing set of each dataset. To decrease the influence of the randomness of $\mathbf{M}$, the metrics will be averaged over 10 iterations using different randomly generated versions of $\mathbf{M}$.

The process will be repeated for a increasing number of

TABLE I. Physical parameters for the DD-TEA model

Parameter	Value	
D_t	50	ns
I_{dc}	2	mA
ΔV	125	mV

virtual neurons in the reservoir. This will show the potential of DD-TEA for performing extensive parametric studies and optimizing such hyperparameters. The physical parameters plugged in DD-TEA will be fixed to the values presented in Table I. Finally, the performance will be compared to the same architecture with more conventional nonlinear transfer functions such as the reLU and the sigmoid functions. This will be achieved by artificially replacing Eq. (1) by the corresponding expressions in the simulations. The identity function will also be tested as an activation function to assess the role of the nonlinearity in the learning ability of the ESN.

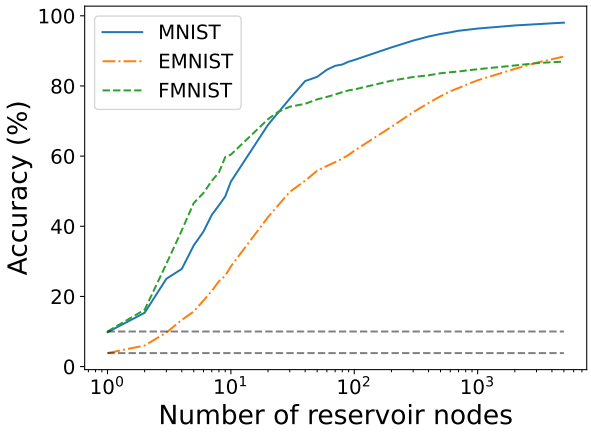

FIG. 9. Accuracy of the STVO-based ESN for the MNIST, EMNIST-letters and FMNIST datasets. The dashed lines represent the random choice accuracy levels linked to the number of categories in each dataset ($100\%/26 = 3.85\%$ and $100\%/10 = 10\%$).

V. RESULTS

The accuracy and the NRMSE of our STVO-based ESNs is depicted for the three datasets in Fig. 11. Without surprise, the accuracy (resp. NRMSE) increases (resp. decreases) with an increasing number of nodes in the reservoirs. As a matter of fact, the number of nodes is directly proportional to the number of tunable parameters in the model (*i.e.*, the number of elements in the learned matrix $\mathbf{W}$).

The shaded area in each plot represents the range where the input data $\mathbf{x}$ is projected into a smaller dimension space due to the low amount of virtual neurons in the reservoir. Indeed, if $N_\theta < N_f$, then the size of $\mathbf{x}'$ is smaller than the original input $\mathbf{x}$, hence compressing information instead of projecting it into a higher dimension space. Above this limit, the absence of nonlinearity in the identity function leads to a saturation of the performance. Considering a linear transfer function represented by $\sigma(\mathbf{x}') = k\mathbf{x}'$, and by assuming that $[k\mathbf{M}]^\dagger [k\mathbf{M}] \approx \mathbf{I}$, then Eq. (7) reduces to Eq. (8), which is independent of $\mathbf{M}$ and hence of the reservoir internal structure itself.

$$\hat{\mathbf{T}}_{\text{test}} = \text{argmax}\left(\mathbf{T}_{\text{train}} \left(\mathbf{X}^{\text{T}}_{\text{train}}\right)^\dagger \mathbf{X}^{\text{T}}_{\text{test}}\right) \tag{8}$$

In all the cases, the accuracy and the NRMSE seems to tend to a common value for all the nonlinear activation functions. This implies that the STVO dynamics nonlinearity is equivalent to that of these more conventional functions. This also indicates that the nature of the nonlinearity is not determinant in the case of an ESN when the number of learnable parameters becomes high. However, any form of nonlinearity is still required to improve the quality of the data classification when $N_\theta > N_f$ as showed by the high performance gap between the nonlinear functions and the identity function.

The comparison of the accuracy of the classification using the STVO-based ESN for the three datasets is represented in Fig. 9. It can be seen that for a very low number of virtual nodes in the reservoir, the accuracy reaches the random choice (*i.e.*, 10% for MNIST and FMNIST, and 3.85% for EMNIST-letters). The accuracy eventually reaches an accuracy of 98% on MNIST and an accuracy of 88.4% (resp. 86.9%) on the EMNIST-letters (resp. FMNIST) dataset with 5000 nodes in the ESN. We believe that the relatively low score achieved in FMNIST is due to the lack of convolutional treatment of the input features, something that is known to be of prior importance for complex images classification. Concerning EMNIST-letters, we suspect that the lower quality of the recognition is due to the presence of lower case and upper case characters in each category, which can be regarded as the presence of two distinct datasets with the same labels. This assumption is consolidated by computing the average intraclass variance in each dataset. The variance of each class $\mathbf{X}$ is computed accordingly to $\text{Var}_{\text{intra}}(\mathbf{X}) = \left(\sum_i (x_i - \overline{\mathbf{X}})^2\right)/|\mathbf{X}|$ and averaged over all the classes for each dataset. It is equal to 3452.9 in MNIST, 4408.7 in EMNIST-letters and 3401 in FMNIST. The average intraclass variance is about 28% bigger in EMNIST-letters than in MNIST and FMNIST.

For comparison, the top accuracy reached on MNIST known to this day is 99.87%[39]. While our solution performs better than other reservoir computing-based approaches (in software and hardware)[40–42], other works indicate that higher performance could be achieved by increasing further the number of nodes in the reservoir[43]. The top accuracy on EMNIST-letters and FMNIST are respectively 95.96% and 96.91%[44,45]. However, these works are based on neural networks far more complex than our, involving many specialized layers such as convolutional layers and attention modules. Therefore, these solutions are not suited for hardware integration within a neuromorphic system.

VI. CONCLUSION

We designed an echo-state network based on hardware nonlinear spintronics oscillators called STVOs. This neuromorphic system aims at solving cognitive tasks more efficiently than conventional software solutions. The system has a simple architecture and only requires the learning of a small part of the involved parameters. Moreover, the learning can be performed linearly using the Moore-Penrose pseudoinverse, as all the nonlinear treatment of the data takes place in a fixed and random reservoir.

To overcome the limitations inherent to experimentally developing, optimizing and testing such a nanostructured system, we used ultrafast simulations of the STVOs. This model yields an explicit expression for the STVO dynamics that we used as neuronal activation function. This approach allows to simulate an entire STVO-based ESN accurately at ultra-high speed compared to previously available frameworks.

We showed that a STVO-based ESN is able to classify images of handwritten digits from the MNIST dataset with state-of-the-art accuracy levels. Indeed, the nonlinearity of STVO dynamics is equivalent to more conventional nonlinearities such as the reLU and the sigmoid functions. Moreover, the system is generalizable to other problems such as the EMNIST-letters and the fashion MNIST datasets. We suspect that better performance can be obtained on the two latter tasks by designing a more specialized architecture, for example involving convolutional data treatment before the reservoir[46].

The use of the DD-TEA framework allowed to carry on extensive studies of the performance with respect to the number of reservoir nodes. This could be reproduced with other hyperparameters such as the input and output shapes of the data or the properties of the input matrix $\mathbf{M}$. Physical operating parameters having an influence on the STVO dynamics could also be optimized by doing so, allowing an efficient preparation of an experimental setup.

We expect that STVO-based ESNs can be successfully used for other tasks, such as time series prediction. Moreover, we think that DD-TEA could allow to develop and test other neural architectures involving STVOs, and investigate other types of learning rules to perform complex machine learning tasks with increased efficiency.

VII. ACKNOWLEDGEMENTS

Computational resources have been provided by the Consortium des Équipements de Calcul Intensif (CÉCI), funded by the Fonds de la Recherche Scientifique de Belgique (F.R.S.-FNRS) under Grant No. 2.5020.11 and by the Walloon Region.

VIII. REFERENCES

[1] W. Sun, B. Zheng, and W. Qian, in *Medical imaging 2016: computer-aided diagnosis*, Vol. 9785 (SPIE, 2016) pp. 241–248.
[2] D. S. Leme, S. A. da Silva, B. H. G. Barbosa, F. M. Borém, and R. G. F. A. Pereira, Computers and electronics in agriculture **156**, 312 (2019).
[3] P. Dhar, Nat. Mach. Intell. **2**, 423 (2020).
[4] M. V. Wilkes, ACM SIGARCH Computer Architecture News **23**, 4 (1995).
[5] D. J. Frank, R. H. Dennard, E. Nowak, P. M. Solomon, Y. Taur, and H.-S. P. Wong, Proceedings of the IEEE **89**, 259 (2001).
[6] C. Mead, Proceedings of the IEEE **78**, 1629 (1990).
[7] B. Schrauwen, D. Verstraeten, and J. Van Campenhout, in *Proceedings of the 15th european symposium on artificial neural networks. p. 471-482 2007* (2007) pp. 471–482.
[8] G. Tanaka, T. Yamane, J. B. Héroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, and A. Hirose, Neural Networks **115**, 100 (2019).
[9] L. Appeltant, M. C. Soriano, G. Van der Sande, J. Danckaert, S. Massar, J. Dambre, B. Schrauwen, C. R. Mirasso, and I. Fischer, Nature communications **2**, 468 (2011).
[10] S. Ortín, M. C. Soriano, L. Pesquera, D. Brunner, D. San-Martín, I. Fischer, C. Mirasso, and J. Gutiérrez, Scientific reports **5**, 14945 (2015).
[11] M. Borghi, S. Biasi, and L. Pavesi, Scientific Reports **11**, 15642 (2021).
[12] L. Larger, M. C. Soriano, D. Brunner, L. Appeltant, J. M. Gutiérrez, L. Pesquera, C. R. Mirasso, and I. Fischer, Optics express **20**, 3241 (2012).
[13] Y. Paquot, F. Duport, A. Smerieri, J. Dambre, B. Schrauwen, M. Haelterman, and S. Massar, Scientific reports **2**, 287 (2012).
[14] V. Pribiag, I. Krivorotov, G. Fuchs, P. Braganca, O. Ozatay, J. Sankey, D. Ralph, and R. Buhrman, Nature physics **3**, 498 (2007).
[15] Y. Gaididei, V. P. Kravchuk, and D. D. Sheka, International Journal of Quantum Chemistry **110**, 83 (2010).
[16] F. Abreu Araujo, C. Chopin, and S. de Wergifosse, Scientific Reports **12**, 10605 (2022).
[17] K. Yogendra, D. Fan, and K. Roy, IEEE Transactions on Magnetics **51**, 1 (2015).
[18] J. Torrejon, M. Riou, F. Abreu Araujo, S. Tsunegi, G. Khalsa, D. Querlioz, P. Bortolotti, V. Cros, K. Yakushiji, A. Fukushima, *et al.*, Nature **547**, 428 (2017).
[19] M. Romera, P. Talatchian, S. Tsunegi, F. Abreu Araujo, V. Cros, P. Bortolotti, J. Trastoy, K. Yakushiji, A. Fukushima, H. Kubota, *et al.*, Nature **563**, 230 (2018).
[20] D. Marković, N. Leroux, M. Riou, F. Abreu Araujo, J. Torrejon, D. Querlioz, A. Fukushima, S. Yuasa, J. Trastoy, P. Bortolotti, *et al.*, Applied Physics Letters **114** (2019).
[21] F. Abreu Araujo, M. Riou, J. Torrejon, S. Tsunegi, D. Querlioz, K. Yakushiji, A. Fukushima, H. Kubota, S. Yuasa, M. D. Stiles, *et al.*, Scientific reports **10**, 328 (2020).
[22] J. Leliaert and J. Mulkers, Journal of Applied Physics **125** (2019).
[23] A. Thiele, Physical Review Letters **30**, 230 (1973).
[24] F. Abreu Araujo, C. Chopin, and S. de Wergifosse, arXiv preprint arXiv:2206.13596 (2022).
[25] A. Moureaux, S. de Wergifosse, C. Chopin, J. Weber, and F. Abreu Araujo, arXiv preprint arXiv:2301.11025 (2023).
[26] L. Deng, IEEE Signal Processing Magazine **29**, 141 (2012).
[27] G. Cohen, S. Afshar, J. Tapson, and A. Van Schaik, in *2017 international joint conference on neural networks (IJCNN)* (IEEE, 2017) pp. 2921–2926.
[28] H. Xiao, K. Rasul, and R. Vollgraf, arXiv preprint arXiv:1708.07747 (2017).
[29] L. Larger, A. Baylón-Fuentes, R. Martinenghi, V. S. Udaltsov, Y. K. Chembo, and M. Jacquot, Physical Review X **7**, 011015 (2017).
[30] D. C. Ralph and M. D. Stiles, Journal of Magnetism and Magnetic Materials **320**, 1190 (2008).
[31] K. Y. Guslienko, O. V. Sukhostavets, and D. V. Berkov, Nanoscale research letters **9**, 1 (2014).
[32] S. Yuasa and D. Djayaprawira, Journal of Physics D: Applied Physics **40**, R337 (2007).
[33] M. Lukoševičius and H. Jaeger, Computer science review **3**, 127 (2009).
[34] M. Ringnér, Nature biotechnology **26**, 303 (2008).
[35] S. Karamizadeh, S. M. Abdullah, A. A. Manaf, M. Zamani, and A. Hooman, Journal of Signal and Information Processing **4**, 173 (2013).
[36] M. Riou, F. Abreu Araujo, J. Torrejon, S. Tsunegi, G. Khalsa, D. Querlioz, P. Bortolotti, V. Cros, K. Yakushiji, A. Fukushima, *et al.*, in *2017 IEEE International Electron Devices Meeting (IEDM)* (IEEE, 2017) pp. 36–3.
[37] P. Courrieu, arXiv preprint arXiv:0804.4809 (2008).

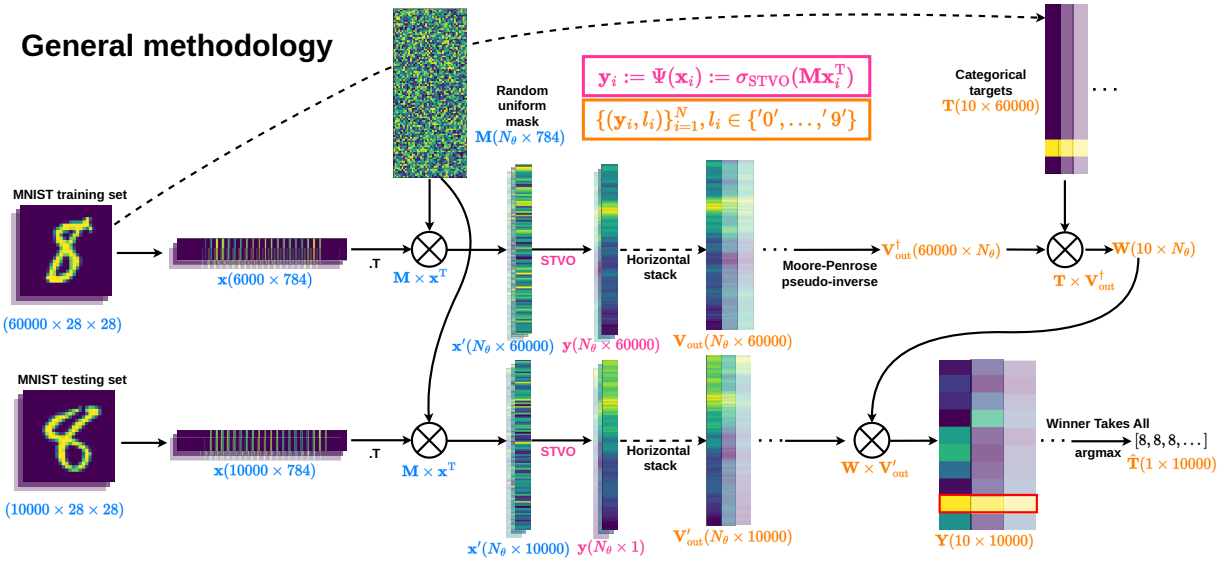

FIG. 10. Summary of the training and testing and inference processes focused on images of handwritten 8 digits from the MNIST dataset.

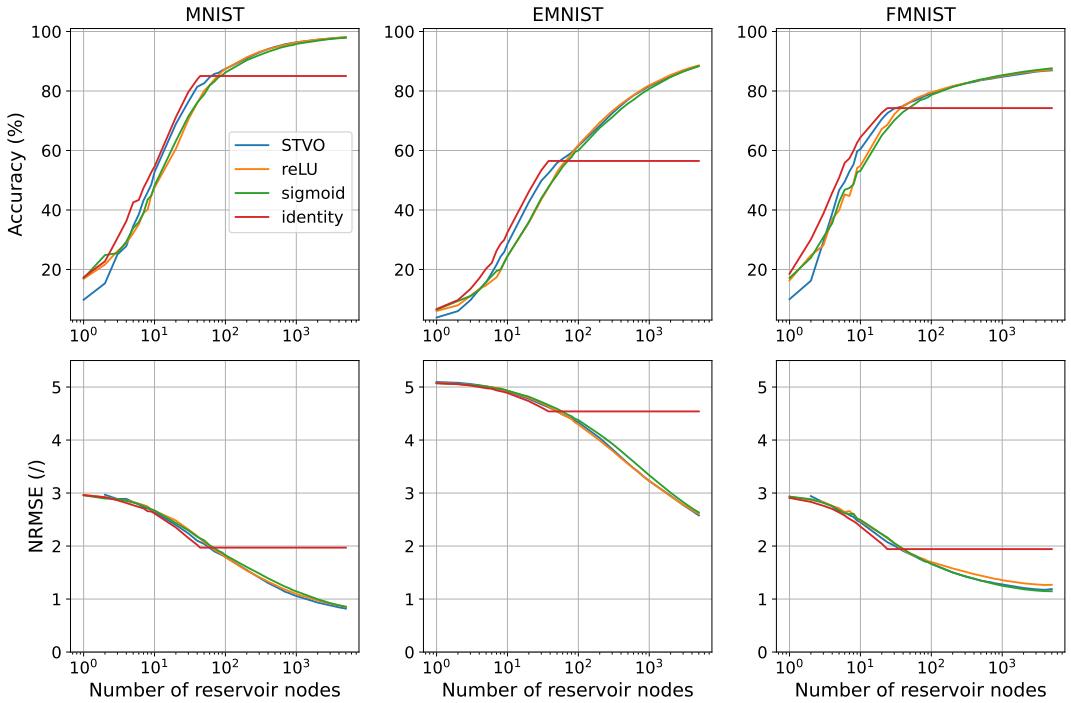

FIG. 11. Accuracy and NRMSE of an time-multiplexed ESN for the MNIST, EMNIST-letters and FMNIST datasets for an increasing number of nodes. The nonlinear transfer function of the reservoir nodes is set with the STVO dynamics and the reLU and sigmoid functions. Linear nodes are also tested using the identity function. The shaded area represents the data compression regime.

[38] J. C. A. Barata and M. S. Hussein, Brazilian Journal of Physics **42**, 146 (2012).
[39] A. Byerly, T. Kalganova, and I. Dear, Neurocomputing **463**, 545 (2021).
[40] S. D. Gardner, M. R. Haider, L. Moradi, and V. Vantsevich, in *2021 IEEE International Midwest Symposium on Circuits and Systems (MWSCAS)* (IEEE, 2021) pp. 255–258.
[41] R. Zhu, A. Loeffler, J. Hochstetter, A. Diaz-Alvarez, T. Nakayama, A. Stieg, J. Gimzewski, J. Lizier, and Z. Kuncic, in *International Conference on Neuromorphic Systems 2021* (2021) pp. 1–4.
[42] S. D. Gardner and M. R. Haider, IEEE Journal on Flexible Electronics **1**, 185 (2022).
[43] N. Schaetti, M. Salomon, and R. Couturier, in *2016 IEEE Intl conference on computational science and engineering (CSE) and IEEE Intl conference on embedded and ubiquitous computing (EUC) and 15th Intl symposium on distributed computing and applications for business engineering (DCABES)* (IEEE, 2016) pp. 484–491.
[44] P. Jeevan, K. Viswanathan, and A. Sethi, arXiv preprint arXiv:2205.14375 (2022).
[45] M. S. Tanveer, M. U. K. Khan, and C.-M. Kyung, in *2020 25th International Conference on Pattern Recognition (ICPR)* (IEEE, 2021) pp. 4789–4796.
[46] Z. Tong and G. Tanaka, in *2018 24th International Conference on Pattern Recognition (ICPR)* (IEEE, 2018) pp. 1289–1294.