

PAPER • OPEN ACCESS

Stable attractors for neural networks classification via ordinary differential equations (SA-nODE)

To cite this article: Raffaele Marino *et al* 2024 *Mach. Learn.: Sci. Technol.* **5** 035087

View the [article online](#) for updates and enhancements.

You may also like

- [Assessing non-nested configurations of multifidelity machine learning for quantum-chemical properties](#)
Vivin Vinod and Peter Zaspel
- [Towards an End-to-End Method for High-Variance SERS Spectra Classification and Quantification](#)
Vincent Thibault and Jean-Francois Masson
- [Advancements in prostate zone segmentation: integrating attention mechanisms into the nnU-Net framework](#)
Aleksas Vaitulevius, Jolita Bernataviien, Jurgita Markeviut et al.



PAPER

OPEN ACCESS

RECEIVED
27 May 2024REVISED
7 September 2024ACCEPTED FOR PUBLICATION
24 September 2024PUBLISHED
1 October 2024

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Stable attractors for neural networks classification via ordinary differential equations (SA-nODE)

Raffaele Marino* , Lorenzo Buffoni , Lorenzo Chicchi, Lorenzo Giambagli and Duccio Fanelli

Dipartimento di Fisica e Astronomia, Università degli studi di Firenze, Via Giovanni Sansone, 1, 50019 Sesto Fiorentino (FI), Italy
* Author to whom any correspondence should be addressed.E-mail: raffaele.marino@unifi.it

Keywords: stable attractors, neural networks, ordinary differential equations

Abstract

A novel approach for supervised classification is presented which sits at the intersection of machine learning and dynamical systems theory. At variance with other methodologies that employ ordinary differential equations for classification purposes, the untrained model is *a priori* constructed to accommodate for a set of pre-assigned stationary stable attractors. Classifying amounts to steer the dynamics towards one of the planted attractors, depending on the specificity of the processed item supplied as an input. Asymptotically the system will hence converge on a specific point of the explored multi-dimensional space, flagging the category of the object to be eventually classified. Working in this context, the inherent ability to perform classification, as acquired ex post by the trained model, is ultimately reflected in the shaped basin of attractions associated to each of the target stable attractors. The performance of the proposed method is here challenged against simple toy models crafted for the purpose, as well as by resorting to well established reference standards. More precisely, we achieved an accuracy of 98.06% on the MNIST test set and 88.21% on the Fashion MNIST test set. Although this method does not reach the performance of state-of-the-art deep learning algorithms, it illustrates that continuous dynamical systems with closed analytical interaction terms can serve as high-performance classifiers.

1. Introduction

Machine learning (ML) [1, 2] and deep learning (DL) [3–5] stand at the forefront of global technology use and research [6]. In an age where data complexity is ever-increasing, these tools excel in uncovering novel patterns, thereby enabling the achievement of groundbreaking results. Their profound impacts are indeed felt across myriad of disciplines, from medicine to optimization, physics, partial differential equations and finance [7–14]. The escalating sophistication of ML and DL technologies poses pressing challenges, chiefly in the realm of interpretability [15], pointing to the need for novel technical solutions that could contribute to elucidate the process that underlies decision making, which is often obscured from our understanding.

In standard approaches to dense neural network training one seeks to optimise the weights that link pairs of neurons associated to adjacent layers. This is achieved by computing the gradient of the imposed loss with respect to the weights. This procedure amounts to operate in the so called direct space. Alternatively, the learning can be reformulated in reciprocal space, the spectral attributes (eigenvalues and eigenvectors) of the underlying transfer operators being the actual target of the optimisation. This procedure introduced in [16] allows for a substantial compression of the space of trainable parameters. Eigenvalues do provide a reliable ranking of the nodes, an observation which can be used, downstream of training, to prune unessential computing units and return a significantly compressed network, with almost identical classification abilities [17]. Neural networks, in their original conception, are static entities. As such, they do not incorporate time as a distinctive algorithmic element.

To fill this gap, neural ordinary differential equations (Neural ODEs) [18, 19] have been proposed as a prime example of the intersection between dense neural networks and dynamical systems. Neural ODEs illustrate that the information flow within a neural network can be seen as the evolution of a system over

time. Unlike traditional networks that have a fixed number of hidden layers, these models parametrize the derivative of the hidden state, using a neural network to determine the layer's depth dynamically. The model's output is calculated using a differential equation solver, allowing it to adapt to different inputs and balance numerical precision with computational speed. The training process is enhanced by a scalable method for backpropagation through any ordinary differential equation (ODE) solver that implements a strategy known as Adjoint Method [20]. A correspondence can be drawn between Neural ODEs and Skip—Connection Neural Networks [21, 22] where the skip connection acts from one layer to the following one. Notably, the weights should be the same for every step for the equivalence to hold. Relevant is also the setting of Liquid Time-Constant Networks [23] which implement the biologically inspired idea of an intrinsic relaxation time of activation and continuous time recurrent neural network [24–26].

Working in this setting, we here propose a variant of the Neural ODEs which accommodates for an analytical expression of the computing non linear kernel. Further, and by leveraging on the spectral formulation of the coupling operator, the untrained model is *a priori* furnished with a set of stationary stable attractors. Classifying implies shaping the basin of the planted attractors, in such a way that items belonging to different categories, supplied as an input, will be eventually directed towards distinct targets. These latter will uniquely flag the class of pertinence of the object to be eventually classified. The ability of the system to cope with a given classification task will be therefore inherited by the characteristics of the basin of attractions, as sculpted upon training. In our view, this paves the way for a better conceptual grasping of the overall decision making process, in terms of a genuine dynamical system that flows in time, by funnelling different input towards their deputed destination site.

More concretely, we will focus on image classification tasks. The network which defines the backbone of the examined dynamical model is made of N nodes, N denoting the number of pixels of the images to be classified. Each node is assigned with a double well potential and decorated with a scalar and continuous state variable which can explore the landscape of the assumed potential. Nodes are then sensing its nearest neighbours, as stipulated by a linear coupling term. The web of inter-nodes connections is stored in a $N \times N$ adjacency weighted matrix, whose elements represent the target of the training process. The above matrix is formulated in reciprocal domain: a subset of suitably tailored eigenvectors is assigned to its kernel and define the attractive poles for the globally coupled dynamics. The eigenvalues populate a specific interval of the real axis that yields stability of the crafted attractors, as dictated by a linear stability analysis. In essence the model to be trained is an extended collection of interacting units subject to contrasting tendencies: local reactions force the system towards the minima of the potential, while global, spatially (across nodes) extended, interactions make the system to evolve towards complex heterogeneous stable states. The algorithmic scheme proposed will be denoted in short SA-nODE, Stable Attractors for Neural Networks Classification via Ordinary Differential Equations. The computational complexity of SA-nODE, after the training process is complete, is quantified as $O(N^2)$, resulting from simple matrix-vector multiplication during the integration of the dynamics, where N represents the size of the input. Additionally, SA-nODE offers an advantage in terms of memory storage. Like nODE models, SA-nODE has a fixed number of parameters that do not increase with the depth of the neural network.

After having established the theoretical background for the proposed methodology, we will turn to presenting an extensive collection of tests. These include exploring the impact of random perturbations on the algorithm's performance. In conclusion, this manuscript present: (i) a dynamical system where stable attractors are embedded into the dynamics; (ii) how to use spectral training techniques for transforming such a dynamical system in a image classification algorithm. The paper is organized as follows: In section 2, we will formulate the model, including the strategy to embed stable attractors. Section 3 details our experimental setup and explains the training approach employed. In section 4, we present results from a synthetic dataset composed of letters of size 7×7 [27]. Section 5 demonstrates the performance of the algorithm on two well-known benchmark datasets, namely MNIST and Fashion MNIST. In section 6, we interpret the implications of our results and conclude our manuscript.

2. The model

In this section, we will introduce the basic principles that underlie the functioning of SA-nODE. As anticipated, the proposed classification model is a learnable autonomous dynamical system, composed of N interacting neurons, which can be cast in the general form:

$$\dot{\vec{x}}(t) = \vec{F}(\vec{x}), \quad (1)$$

where $\vec{x} \in \mathbb{R}^N$, and $\vec{F}(\vec{x}) = -\vec{\nabla}V(\vec{x}(t), a, \gamma) + \beta\mathbf{A}\vec{x}(t)$. The vector \vec{x} has entries of $O(1)$, while $\dot{\vec{x}}(t)$ represents the derivative with respect to time of $\vec{x}(t)$. $V(\vec{x}(t), a, \gamma) : \mathbb{R}^N \rightarrow \mathbb{R}$ is a scalar field, i.e. a potential; with $\vec{\nabla}$ we

define the gradient operator. The adjacency matrix describing the graph is $\mathbf{A} \in \mathbb{R}^{N \times N}$. Such a matrix encapsulates information on existing (weighted and sign sensitive) interactions between the nodes of the network. β, γ and a are parameters of the system.

2.1. Planting the attractors

As anticipated, we postulate a double well potential, expressed as $V(\vec{x}(t), a, \gamma) = \gamma(\vec{x}^2(t) - a^2 \vec{\mathbf{1}})^T (\vec{x}^2(t) - a^2 \vec{\mathbf{1}})$. In this formulation, each element of \vec{x}^2 is denoted as x_i^2, i ranging from 1 to N . In the uncoupled limit ($\beta = 0$), each node asymptotically approaches one of the two available equilibria, respectively locates in $\pm a$. The uncoupled model displays hence 2^N distinct attractors, i.e. the whole set of independent combinations obtained by permuting the two above equilibria across the collection of N available nodes. These latter are no longer solutions of the examined model when β is made different from zero. In the following, we choose $\beta = \frac{1}{\sqrt{N}}$. This scaling is in line with the model proposed in [28, 29], where the authors assume that the entries of \mathbf{A} are independently and identically distributed as $O(1)$ variables from a standard Gaussian distribution. The value of γ , which sets the strength of the potential, will be self-consistently adjusted through training. It should be emphasized that a prototypical double well potential is found to emerge in models relevant to computational neuroscience, as follow the intertwined interaction between distinct families of excitatory and inhibitory neurons [30]. In this respect, the simplified model that we have here formulated unlocks a perspective view full of captivating biomimetic implications.

As previously indicated, the matrix \mathbf{A} is a critical component as it incorporates the structural details of the network’s topology. Assume \mathbf{A} to be represented as $\mathbf{A} = \Phi \Lambda \Phi^{-1}$, where Φ belongs to the set of real matrices $\mathbb{R}^{N \times N}$ and Φ^{-1} signifies its inverse. Within this framework, the matrix Φ comprises the eigenvectors of \mathbf{A} , arranged as its columns. Meanwhile, the matrix Λ , also in $\mathbb{R}^{N \times N}$, is a diagonal matrix containing \mathbf{A} ’s eigenvalues. For ease of reference, we denote the columns of Φ as $\vec{\phi}^{(l)}$, where l ranges from 1 to N . This approach of decomposing the interaction matrix mirrors the spectral approach to machine learning, as outlined in [16, 17, 31–33]. The whole idea of SA-nODE is to *a priori* plant a congruous number of non linear stable attractors within the full untrained model. Indeed the $k = 1, \dots, K$ columns of matrix Φ , where K stands for the total number of classes to be eventually categorised, correspond to our identified attractors.

To reach our goal we assign the entries of the selected vectors $\vec{\phi}^{(k)}$ to take values $\pm a$, namely the positions of the fixed points of the a -spatial dynamics ($\beta = 0$). In formulae, we require $\left(\vec{\phi}^{(k)}\right)_i = \pm a, \forall i = 1, \dots, N$.

Further, we place the target vectors $\vec{\phi}^{(k)}$ in the kernel of \mathbf{A} . In other words, the matrix Λ has eigenvalues $\lambda^{(k)} = 0, \forall k = 1, \dots, K$, for satisfying the property that the corresponding eigenvectors sits in the kernel of \mathbf{A} .

Thus, in light of the above prescriptions, $\vec{x}_{st}^k = \vec{\phi}^{(k)}, k = 1, \dots, K$ are stationary solutions ($\dot{\vec{x}} = 0$) of the examined model: when the solution aligns to \vec{x}_{st}^k the linear inter-nodes coupling term disappear, by construction. Further the individual components of \vec{x}_{st}^k are chosen in such a way $-\vec{\nabla} V(\vec{x}_{st}^k, a, \gamma) = 0$, thus yielding the sought condition $\dot{\vec{x}} = \vec{F}(\vec{x}_{st}^k) = 0$.

However, merely planting attractors is insufficient. Our objective is to enforce asymptotic stability, following the procedure detailed above.

2.2. Enforcing linear stability

To investigate the linear stability, we introduce a perturbation $\delta \vec{x}$ around $\vec{\phi}^{(k)}$. Each component of the perturbation is defined as follow: $\delta x_i = x_i - \bar{\phi}_i^{(k)}$, with δx_i representing a small deviation from the state $\bar{\phi}_i^{(k)}$. The model, i.e. equation (1), is then linearized around $\vec{\phi}^{(k)}$, leading to the following expression for each component:

$$\begin{aligned} \delta \dot{x}_i &= f\left(\bar{\phi}_i^{(k)}\right) + f'\left(\bar{\phi}_i^{(k)}\right) \delta x_i + \\ &\beta \sum_{j=1}^N A_{ij} \bar{\phi}_j^{(k)} + \beta \sum_{j=1}^N A_{ij} \delta x_j + O\left(\delta x_i^2\right), \end{aligned} \tag{2}$$

where $f\left(\bar{\phi}_i^{(k)}\right) = -\left.\frac{\partial V(\vec{x}, a, \gamma)}{\partial x_i}\right|_{x_i = \bar{\phi}_i^{(k)}}$ and $f'\left(\bar{\phi}_i^{(k)}\right)$ identifies the derivative with respect to the variable x_i .

Higher-order terms in δx_i are neglected in this treatment. In (2), the term $f\left(\bar{\phi}_i^{(k)}\right) = f(\pm a) = 0$ by

construction, as well as $\beta \mathbf{A} \vec{\phi} = \vec{0}$. Equation (2) can, therefore, be further simplified to:

$$\delta \dot{x}_i = f'(\pm a) \delta x_i + \beta \sum_{j=1}^N A_{ij} \delta x_j. \quad (3)$$

To proceed with the linear stability analysis, we aim to express δx_i in the basis of eigenvectors of \mathbf{A} . In this context, δx_i with $\sum_{\alpha} c_{\alpha} \phi_i^{(\alpha)}$, with $\mathbf{A} \vec{\phi} = \Lambda \vec{\phi}$, where, we recall, Λ is diagonal.

By substituting δx_i with the expression above, we obtain an equation for the evolution of coefficients c_{α} . This equation is given by:

$$\sum_{\alpha=1}^N \phi_i^{(\alpha)} \left\{ \dot{c}_{\alpha} - \left(f'(\pm a) + \beta \lambda^{(\alpha)} \right) c_{\alpha} \right\} = 0, \quad \forall \alpha. \quad (4)$$

The condition for a stable solution of the differential equation $\dot{c}_{\alpha} = (f'(\pm a) + \beta \lambda^{(\alpha)}) c_{\alpha}$ is thus provided by $\lambda^{(\alpha)} < \frac{8a^2\gamma}{\beta}, \forall \alpha$. From the solution of the differential equation, we can immediately observe that if

$\lambda^{(\alpha)} < \frac{8a^2\gamma}{\beta}, \forall \alpha$, then $\vec{\phi}^{(k)}$ is a stable attractor of the dynamics, meaning that perturbations decay

exponentially with time around it. Conversely, if there exists at least one $\lambda^{(\alpha)} > \frac{8a^2\gamma}{\beta}$ then the attractor $\vec{\phi}^{(k)}$ becomes unstable, as a small perturbation grows exponentially with time. Moreover, in the case of stability, we can define $t_s = \frac{1}{\min |f'(\pm a) + \beta \lambda|}$, which indicates the characteristic time scale for an imposed perturbation to vanish.

In conclusion, we have examined the linear stability of the imposed attractors, the eigenvectors of the coupling matrix assigned to its kernel. The analysis yields an upper bound for the set of eigenvalues $\lambda^{(\alpha)}$, a condition that should be enforced for the attractors to prove stable to external perturbations.

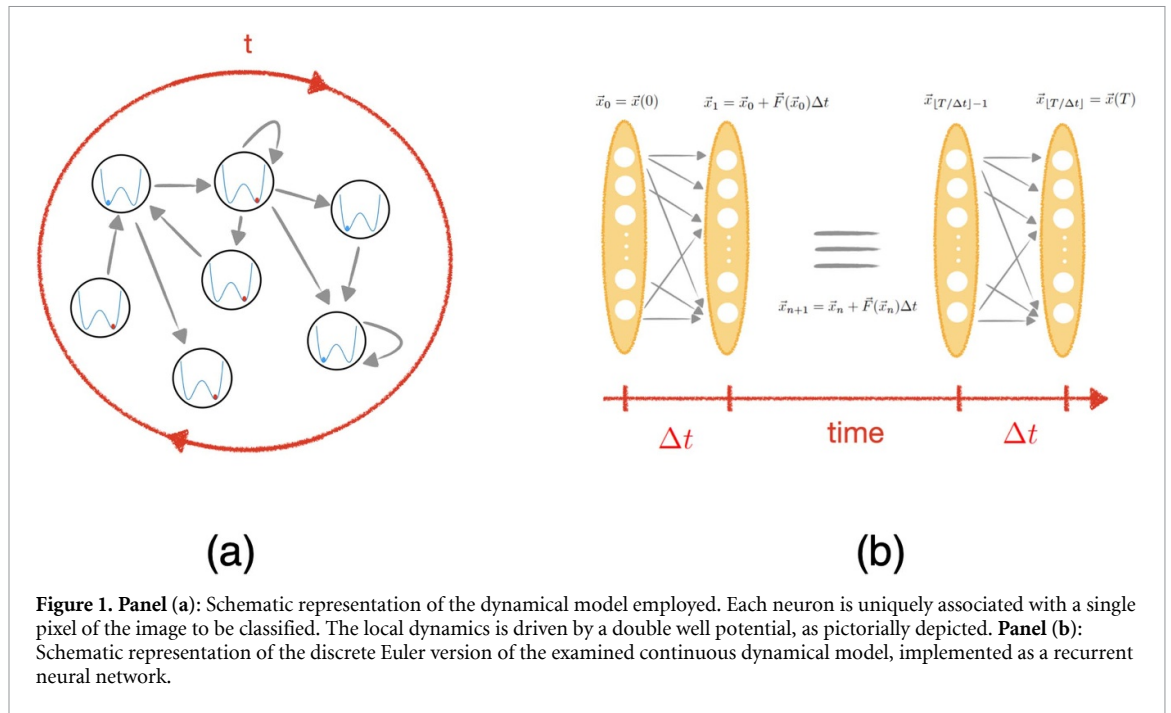
3. Training the model and the metric of convergence

Following the analysis carried out above, we are now in a position to introduce the training process that the algorithm must undergo to learn to classify images. Without loss of generality, we fix the parameters a at 0.5 and let γ to be trainable. From the observations made in the previous section, we constrain the eigenvalues of the matrix during the learning process to take any value in the range $(-\infty, \frac{8a^2\gamma}{\beta})$. As an initial condition for γ we opt for a value equal to 0.4, while trainable eigenvalues are initialized as $\mathcal{N}(-5, 1)$, i.e. a normal distribution with mean equal to -5 and variance equal to 1.

With this stated, we construct a statistical learning model where only the components of the matrix Φ , which are not the embedded eigenvectors, as well as the non-zero eigenvalues can be self-consistently learned. In other words, the first K columns of matrix Φ , those identifying the embedded attractors, are not to be learned, just as the first K eigenvalues of the matrix Λ , which are fixed to be zero by construction. As an initial condition for the non-embedded eigenvectors, we opt for orthogonal random eigenvectors.

As our system evolves over time, the values of \vec{x} are updated. The updating is performed by an Euler algorithm implemented as a recurrent neural network, following the scheme depicted in figure 1. Identical layers made of N nodes are mutually linked by a linear coupling matrix \mathbf{A} , computed via its spectral decomposition, which contains the parameters to be trained. Non linear filters acting on each node follow the interaction term here postulated, i.e. the symmetric double well potential. Time flows along the horizontal axis, two successive layers of the imposed feedforward deep architecture being separated by a finite amount, $\Delta t \ll 1$. By operating in this framework, enables us to resort to standard optimization tools, as routinely employed within the machine learning community, to train the proposed dynamical model. By choosing Δt sufficiently small we make sure that the trained discrete model behaves like its continuous counterpart. In concrete terms, we therefore train an ensemble made of continuous and linearly coupled ordinary differential equations, which display a set of assigned asymptotic attractors, to operate as genuine classifier. In the annexed appendix A, we consider a variant of the integration method that makes use of a Runge–Kutta scheme. The underlying idea is identical to what previously illustrated, except for the fact that for the deployment on an equivalent recurrent neural network one needs to accommodate for multiple evaluations of the updated state (namely account for additional layers), as following the prescription of the employed numerical algorithm. Tests on the convergence of the algorithm are enclosed in the aforementioned appendix A.

The parameter space for optimization resides in $\mathbb{R}^{(\mathcal{N}+1)(\mathcal{N}-K)+1}$, where K , we recall, denotes the number of classes in our classification problem, and the singular dimension relates to the parameter γ . In practice, we



take a dataset $\mathcal{D} = (\vec{x}, y)^{(j) \in [1, \dots, |\mathcal{D}|]}$ of size $|\mathcal{D}|$, where $\vec{x}^{(j)}$ is an input datum, and $y^{(j)}$ represents the target, i.e. a number that ranges from 1 to K . This target y is mapped into our attractor $\vec{\phi}^{(y^{(j)})}$ of the dynamics. In other words, the target $y = 1$ is the first eigenvector $\vec{\phi}^{(1)}$, the target $y = 2$ is the second eigenvector $\vec{\phi}^{(2)}$ and so on. Each analysed dataset is split into a training and a test sets.

The objective of the training is to minimize the loss function $\mathcal{L} = \frac{1}{|\mathcal{D}|} \sum_{j=1}^{|\mathcal{D}|} (\vec{x}^{*(j)} - \vec{\phi}^{(y^{(j)})})^T (\vec{x}^{*(j)} - \vec{\phi}^{(y^{(j)})})$. \vec{x}^* is the value of $\vec{x}^* = \vec{x}(T)$, when T is large enough, i.e. the dynamical system in (1) reaches its stationary state. In other words, we give as an initial condition to our dynamical system the input datum, i.e. $\vec{x}(0) = \vec{x}^{(j)}$, and let it evolve for a time T large enough. As anticipated, and in the minimal scheme here discussed, the evolution of the system is performed using Euler' algorithm, i.e. $\vec{x}_{n+1} = \vec{x}_n + \vec{F}(\vec{x}_n)\Delta t$, where Δt is the integration step and n is the index for the n th iteration. The maximum number of iterations is fixed *a priori*. Given a time T , and chosen the value for Δt , the total number of iterations is defined as $n_{\max} = \lfloor \frac{T}{\Delta t} \rfloor$. Reached the time T , the value of $\vec{x}^{*(j)} = \vec{x}^{(j)}(T)$ is used for optimizing the loss function. The optimization process can be facilitated using Adam algorithm [34] (with default learning rate).

Once optimal weight configuration is attained via Adam, we can proceed to evaluate the dynamical system's performance for classifications on test set elements. Importantly, these test set elements are a subset of the dataset \mathcal{D} that the system has never encountered during the training process. The number of epochs for each analysis in this manuscript is fixed at 700 and mini-batch size equal to 250. The hyper-parameters were systematically optimized to achieve the best performance. The code for reproducing the result can be downloaded from [35].

To evaluate the effectiveness of our algorithm, we must compute a metric defining the system's performance. We opt to compare the output images with their corresponding targets. To this end, we compute the overlap between the final state and the corresponding target. Mathematically, we define such overlap as:

$$\begin{aligned}
 m_f &= \frac{1}{|\mathcal{D}_{\text{test}}|} \sum_{j=1}^{|\mathcal{D}_{\text{test}}|} \frac{(\vec{x}^{*(j)})^T \vec{\phi}^{(y^{(j)})}}{Na^2} \\
 &= \frac{1}{|\mathcal{D}_{\text{test}}|} \sum_{j=1}^{|\mathcal{D}_{\text{test}}|} \frac{1}{Na^2} \sum_{i=1}^N x_i^{*(j)} \vec{\phi}_i^{(y^{(j)})}, \tag{5}
 \end{aligned}$$

where $|\mathcal{D}_{\text{test}}|$ is the cardinality of the test set. m_f can be seen as a sort of averaged magnetization of the system, as well described in disordered systems [36].

As alternative measure of the convergence to the deputed attractors, we introduce the Mean Squared Error (MSE), which can be monitored over time t . The MSE is defined as:

$$\text{MSE}(t) = \frac{1}{B} \sum_{m=1}^B \left(\vec{x}(t)^{(m)} - \vec{\phi}(y^{(m)}) \right)^T \left(\vec{x}(t)^{(m)} - \vec{\phi}(y^{(m)}) \right) \quad (6)$$

where m stands for the m th instance, and $B = |\mathcal{D}_{\text{sample}}|$ denotes the sample size. Remark that $\vec{x}(t)$ represents the time evolution of the trained model, for any given item supplied as an input.

To enhance the complexity of the simple datasets used in our preliminary evaluation (see below for a discussion on this issue) and to mimic real-world scenarios where data may be subject to noise, we introduce uniform noise to individual pixels, chosen randomly. This alteration effectively *dirty* the image to be processed, offering a more realistic portrayal of what might be encountered in practical applications. By adjusting the number of pixels to be corrupted with $\epsilon \in [0, 1]$, the level of noise can be controlled, thereby allowing for systematic studies of noise resilience in the proposed model. In other words, the number of pixels corrupted, ζ , is $\zeta = \epsilon N$, where N is the size of the system, i.e. the number of total pixels of an image.

This deliberate induction of noise will be inserted into the training set and/or the test set. In such cases, we use ϵ_{train} to define the level of noise injected into the training set, which means that the model is trained with images that are corrupted with a particular level of noise, equal to ϵ_{train} . Similarly, we use ϵ_{test} to define the level of noise injected into the test set, once a particular model has been trained with its own level of noise.

4. Letter dataset

In this section, we present a synthetic dataset constructed through binarized ASCII characters, specifically tailored to our experimental requirements. Specifically, the chosen set involves five letters, from 'A' to 'E', each formulated as 7×7 in grey scale, see figure 2. This synthetic dataset provides a valuable resource for validating our algorithm's performance under various noise conditions, introduced as follows the procedure highlighted above. At this stage the objective is not to benchmark the performance of our learning model against conventional metrics of accuracy and efficiency, but rather demonstrating the functionality of the continuous dynamical model as a viable image classifier. In figure 2 the bottom row display the target pattern associated to each letter. Following the dictates of the proposed framework, the (arbitrary) patterns that define the attractors are composed by assigning to each pixel value $\pm a$, these latter identifying the location of the two wells of the non linear reaction term. The middle row of figure 2 presents the noisy version of the letters.

As an initial analysis, we focused on the scenario where the model was trained on noise-free images ($\epsilon_{\text{train}} = 0.0$). It should be remarked that this defines a rather academic condition, where train and test sets are made of identical reservoirs, populated with multiple copies of the same (no noise) letters. As such, and for this preliminary case study, we are not really probing the ability of the system to generalize beyond the limited set of images processed under training. The detailed outcomes of this investigation are delineated in figure 3, blue line. The overlap between the final state and the target, i.e. the metric m_f , reaches an impeccable score of 1.0. However, a discernible decrease in performance is observed as the noise level in the test images incrementally rises. The downgrade in the recorded performance is monotonous with the increase of ϵ_{test} and yields an an overlap of zero at $\epsilon_{\text{test}} = 1.0$.

We then transitioned our focus to assess the model's prowess when trained with images bearing a noise level of $\epsilon_{\text{train}} \neq 0$. The figure 3 offers a multidimensional perspective on the model's performance, by comparing various levels of noise during both training and testing. With a modest training noise of $\epsilon_{\text{train}} = 0.1$, the model still maintains near-ideal precision for low-noise testing, but the decline in performance manifests more gradually when compared to previous observations. As we escalate the training noise, these trends persist, up to a given threshold in ϵ_{train} , above which the images are too corrupted to be correctly classified.

In figure 4 Panel (a), we report the trend displayed by the MSE over time t for the model with $\epsilon_{\text{train}} = 0.3$. The MSE is evaluated on various test sets, identified by the level of noise introduced in the test set, namely ϵ_{test} . As a first remarkable observation, we notice that the MSE reaches a stable asymptotic plateau which indirectly points to the convergence of the underlying dynamical model. The ability of the system to classify a given datum supplied as an input, is stably maintained across time and not just referred to the limited horizon of the training time. Further, for $\epsilon_{\text{test}} \leq 0.3$, the MSE decreases approaching zero exponentially fast.




Figure 2. Representation and corresponding targets (asymptotic attractors) of the letters ‘A’, ‘B’, ‘C’, ‘D’, and ‘E’ as 7×7 grey scale. The top row illustrates the representation of each letter, enclosed within rectangular borders. The middle row displays the same letters but with a noise factor of $\epsilon = 0.2$ applied, introducing slight distortions. The bottom row presents the target pattern for each letter, with the black line of the target for ‘B’ starting immediately after the line of the target for ‘A’, the target for ‘C’ starting after ‘B’, ‘D’ after ‘C’, and ‘E’ after ‘D’, each surrounded by a sea of white pixels. This configuration allows for precise mapping to each corresponding letter, serving as the planted attractors within the matrix Φ , as illustrated in the main body of the paper. Recall in particular that the attractors are shaped by employing the two entry values $\pm a$. Here, $-a$ refers to pixel colored in white, black pixels are associated to a .


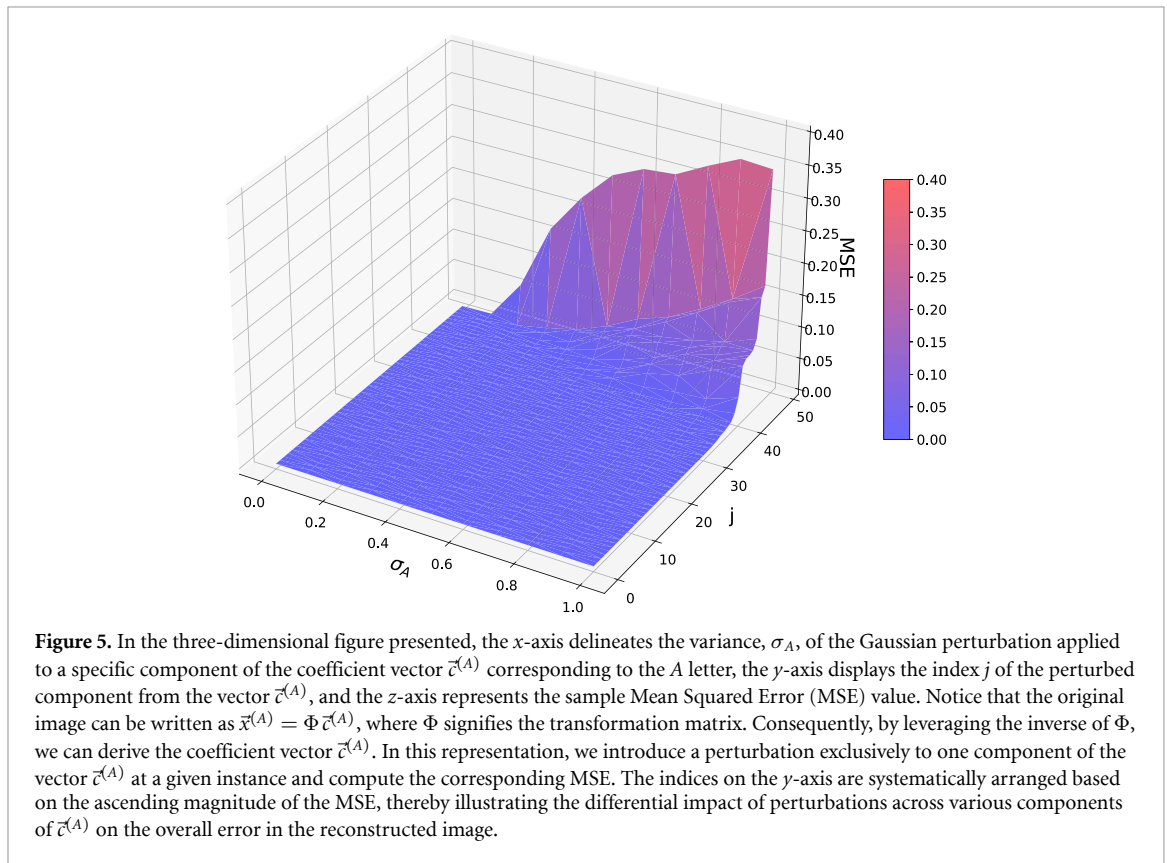
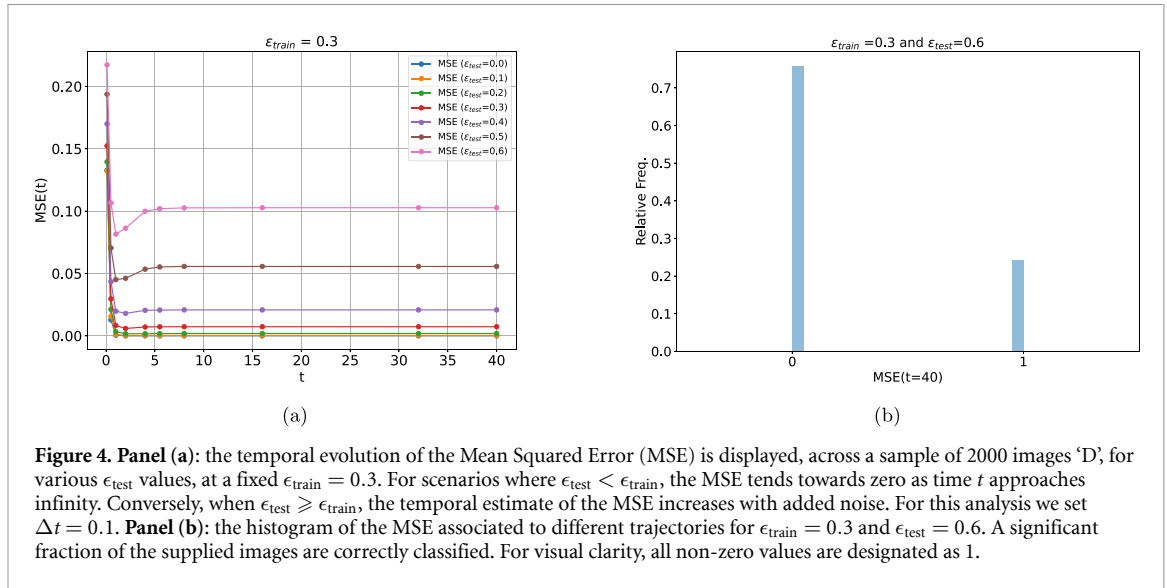


Figure 3. The figure displays the relationship between m_f and the testing noise ϵ_{test} for various levels of training noise, for 5 classes. Each line represents a different training noise level, ranging from 0.0 to 1.0, with a corresponding color code. The horizontal line illustrates a reference point. For this analysis we set $T = 40.0$, $\Delta t = 0.1$.

Conversely, for $\epsilon_{test} \geq 0.4$, the MSE stabilizes at non-zero values: trajectories within the phase space do not align with the anticipated attractors. A detailed analysis reveals instead that, for $\epsilon_{test} \geq 0.4$, the recorded MSE displays a bimodal distribution, see 4 Panel (b), with a significant fraction of images that are indeed correctly classified, also for noise amount considerably larger than the one imposed during training. From a more fundamental point of view, the displayed bimodal distribution is reminiscent of a metastable phase associated to a first order phase transition, that putatively separates the regime where SA-nODE works properly from that where predictions are found to be inaccurate.



Upon training, one obtains a basis of the scanned multi-dimensional space which is tailored to the problem at hand. The basis is formed by the column vectors of matrix Φ . Every supplied item \vec{x} can be decomposed by using the aforementioned basis. The set of obtained coefficients, $\vec{c} = \Phi^{(-1)}\vec{x}$ returns a complete and equivalent representation of the analysed datum \vec{x} . This framework can be used to probe the model's robustness and vulnerability to external source of disturbance from a different perspective, alternative to perturbing each individual pixel (which amounts to operate with the canonical basis viewpoint). More specifically, we can proceed by perturbing individual eigen-directions—or punctually alter each coefficient \vec{c} of the above expansion. In doing so, several pixels get simultaneously modulated by the external noise source, following a pattern of correlated activation that indirectly stems from the accomplished training. In doing so we empirically, senses the size of the basins of attraction, post training, by using a correlated version of noise that is self-consistently shaped by the learning protocol.

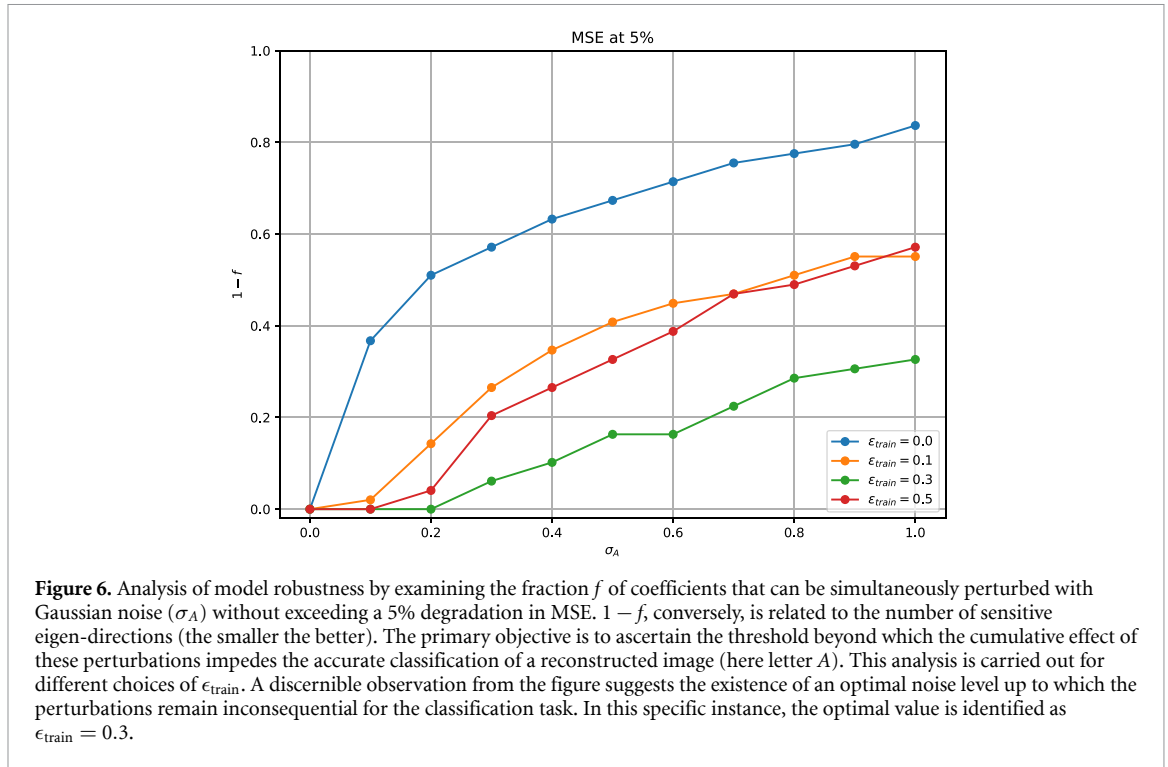


Figure 5 reports on this analysis, where the robustness of each eigen-direction to a Gaussian perturbation $\mathcal{N}(0, \sigma)$ is probed. The MSE, calculated in a steady state, serves as a benchmark metric. We observe that, for a model trained with $\epsilon_{\text{train}} = 0.3$, many coefficients are resilient to individual perturbations. Stated differently, only a few directions can trigger the system to become unstable. In fact, the vast majority of imposed collective eigen-perturbations do not affect the ability of the system to carry out the assigned classification task. In a world where images are frequently subjected to noise and interference, understanding noise resistance through coefficient perturbation can guide the implementation of resilient designs and/or information compression strategies, ensuring that essential information is preserved despite reductions.

Figure 6 shows how the addition of noise in the training process results in a marked increase in the model's robustness, which can be assessed in reciprocal space. Specifically, we reason as follows: let us focus on the letter A. The coefficients in the reciprocal space of this data can be perturbed with a Gaussian distribution of variance σ_A , causing a degradation of the MSE. As shown in figure 5, several modes are, in principle, robust to perturbations applied to individual coefficients. However, from this analysis it is not obvious to anticipate what could happen when a fraction f of the coefficients gets simultaneously perturbed with Gaussian random noise of variance σ_A and how this collective perturbation eventually reverberates on the recorded MSE.

To understand the effect of training with noise ϵ_{train} , we fix the degradation of the MSE at 5% and the perturbation amplitude σ_A . A natural question hence arises: how many coefficients can be perturbed without exceeding the 5% degradation threshold in the MSE? Alternatively, what fraction ($1-f$) of the eigenvectors must not be perturbed? This latter fraction is, in our view, a measure of the network's robustness. Indeed, if it is low, only a few coefficients need to remain unperturbed, while the remaining f can be perturbed simultaneously without causing a degradation greater than 5% of the MSE. The result of this analysis is reported in figure 6 against different training-noise amplitudes. Such representation shows the existence of an optimal training-noise ($\epsilon_{\text{train}} = 0.3$) level. With this latter choice, almost 70% of the eigenvectors can be perturbed with a maximal strength of the imposed noise ($\sigma_A = 1.0$) without observing a relative difference in the MSE larger than 5%. The fraction of relevant eigenvectors is, moreover, consistently lower with respect to the other training-noises.

Upon completing the analysis of the synthetic dataset, we can now shift our focus to assess the robustness of Sa-nODE in comparison to the most commonly cited benchmarks in literature: MNIST and Fashion MNIST (In appendix B we present the robustness of our model under noise condition, for both datasets).

5. Applying SA-nODE to MNIST and Fashion MNIST

In the preceding section we introduced the concept of SA-nODE and applied it to the analysis of a synthetic dataset, with different grades of imposed noise. To further test the adequacy of the proposed scheme, we turn to considering the so called MNIST and Fashion MNIST datasets.

The Modified National Institute of Standards and Technology (MNIST) dataset [37] is one of the most iconic and widely utilized datasets in the field of machine learning, particularly in the domain of pattern recognition and computer vision. Stemming from a rich history, it has played a pivotal role as a benchmark dataset, testing a plethora of algorithms and machine learning techniques over the years.

The MNIST dataset consists of a collection of 70 000 handwritten digit images. These grey scale images are uniformly sized at 28×28 pixels (i.e. $N = 784$), with digits centered in the frame. The dataset is subdivided into a training set of 60 000 images and a test set comprising 10 000 images. Each image is labelled with the corresponding digit it represents, ranging from 0 to 9.

The simplicity of the dataset, combined with its relatively small size, makes it suitable for beginners to explore the intricacies of various machine learning algorithms without the need for substantial computational resources. Furthermore, the dataset's balanced composition, with an almost equal number of samples for each digit, ensures that models trained on it are not biased towards any particular digit.

The accuracy of the trained SA-nODE algorithm can be computed from evaluation of m_f , as introduced above. More specifically, we assumed that an image is correctly classified only if the similarity

$m_f^{(j)} = \frac{1}{N} \sum_{i=1}^N \text{sign}(x_i^{(j)}) \text{sign}(\bar{\phi}_i^{(j)})$ is greater than or equal to 95%.

In formulae, the accuracy (\mathcal{H}) of our algorithm is formulated as the average over the test dataset ($\mathcal{D}_{\text{test}}$):

$$\mathcal{H} = \frac{1}{|\mathcal{D}_{\text{test}}|} \sum_{j=1}^{|\mathcal{D}_{\text{test}}|} \delta(m_f^{(j)}), \quad (7)$$

where

$$\delta(m_f^{(j)}) = \begin{cases} 1 & \text{if } m_f^{(j)} \geq 0.95 \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

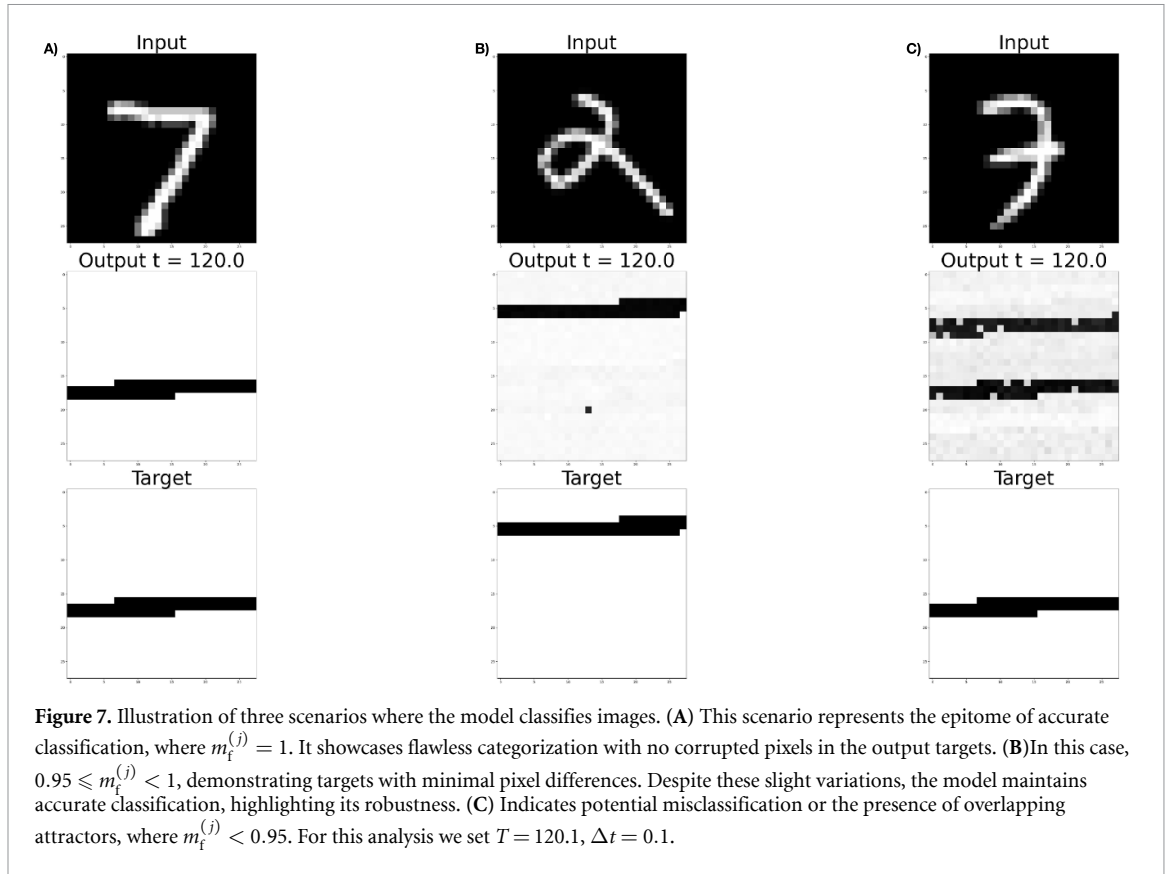
This allows us to account for images that exhibit minimal pixel differences from the correct attractors, while excluding those with more than 5% of corrupted pixels. This refined metric provides a comprehensive evaluation of our algorithm's classification capabilities, in terms of reported accuracy.

In such conditions, the accuracy achieved by SA-nODE is 98.06%. Among the 10^4 test images, 9776, an overwhelming majority, are perfectly classified with $m_f^{(j)} = 1$, indicating that no pixels are corrupted in the output images. In other words, the 9776 initial conditions have converged to the attractors within a specified maximum time frame and the state achieved corresponds accurately with the classification requirements. Furthermore, 30 images exhibit $0.95 \leq m_f^{(j)} < 1$, and the remaining images, though a small fraction, have $m_f^{(j)} < 0.95$.

To provide a comprehensive visual insight into the varying degrees of pixel-wise similarity, three cases typical examples are provided in Figure 7: Case A shows the perfect classification, thus $m_f^{(j)} = 1$, with no corrupted pixels. Case B displays an instance in the range $0.95 \leq m_f^{(j)} < 1$, showcasing output images with minimal pixel differences, while maintaining accurate classification. Case C indicates potential misclassification, without necessarily having corrupted pixels or overlapping attractors. In this particular instance, we can observe the algorithm's uncertainty as it produces an output that is a combination of the targets for both a three and a seven. We present a confusion matrix, other analyses and examples of misclassified images in appendix C.

As a second application, we consider Fashion MNIST dataset. This is specifically designed as a drop-in replacement for the classical MNIST dataset, albeit with a contemporary twist. Instead of handwritten digits, Fashion MNIST encompasses a variety of images pertaining to clothing and fashion articles. Introduced by Zalando, an e-commerce company, [38] the dataset responds to the criticism that MNIST was considered too easy, or even overused, in the deep learning community. While the images in the Fashion MNIST might appear simple, the subtle nuances between certain clothing categories make classification a more challenging task than distinguishing between handwritten digits.

Fashion MNIST consists of 70 000 images, representative of 10 distinct fashion categories, such as t-shirts, trousers, pullovers, dresses, coats, sandals, and more. Similar to the MNIST dataset, these grey scale images are standardized to a size of 28×28 pixels. The dataset is neatly split into 60 000 training images and 10 000 test images. Each image is associated with a label that denotes one of the ten categories.



SA-nODE achieves an accuracy of 88.21%. More precisely, among the 10^4 test images, 8740, an overwhelming majority, are perfectly classified with $m_f^{(j)} = 1$, indicating that no pixels are corrupted in the output images. Furthermore, 81 images exhibit $0.95 \leq m_f^{(j)} < 1$, and the remaining images, though a small fraction, have $m_f^{(j)} < 0.95$. We present a confusion matrix, other analyses and examples of misclassified images in appendix C.

6. Conclusion

We have here proposed a variant of the Neural ODEs architecture, which accounts for a closed analytical expression of the imposed reaction term. Individual nodes are in particular assumed to be subject to a local reaction which follows from a prototypical double well potential. As such, the uncoupled dynamics can accommodate for two symmetric equilibria. These latter solutions define a minimal alphabet that we use to construct a set of arbitrarily crafted attractors of the coupled (spatially extended) system dynamics. The coupling among computing nodes, each associated to a pixel of the image to be classified, is linear and topologically encoded in an adjacency weighted matrix, whose elements represent the target of the training process. This latter matrix is formulated in reciprocal domain and the target attractors are assigned to belong to its kernel. Stationary stability is *a priori* granted by forcing the trained eigenvalues to populate a given domain, as stipulated by a linear stability analysis. Classifying amounts to shaping the basin of the planted attractors, so that different items will be eventually directed towards distinct targets, depending on their specific category of pertinence. In essence, SA-nODE, as we chose to denote the algorithm, is made of a finite set of particles or spins, each of which jumping between two allowed states, identified by the positions of the symmetric potential wells. Because of mutual (linear) interactions, and following spectral training, the particles are eventually frozen in a heterogeneous (in the space of the pixel) and stationary stable pattern, yielding an image that is uniquely representative of the processed element. The proposed classification strategy has been successfully tested against mock and reference datasets, as reported in the main body of the paper (we recall that SA-nODE achieved an accuracy of 98.06% on the MNIST test set and 88.21% on the Fashion MNIST test set). The interest of SA-nODE resides in having constructed a simple, dynamically sound model, which yield competitive classification scores. The model takes a closed, compact and intuitive

form, which can be further dissected by leveraging on the vast arsenal of non-linear dynamics tools, to shed light, from a different angle, on the inherent ability of neural networks to effectively learn and decide. We anticipate that other models, including several of biomimetic inspiration, thus relevant to computational neuroscience, could be proposed, which yield dynamically assisted classification via stable attractors. We will report in future works on these possible extensions.

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files).

Acknowledgments

This work is supported by #NEXTGENERATIONEU (NGEU) and funded by the Ministry of University and Research (MUR), National Recovery and Resilience Plan (NRRP), Project MNESYS (PE0000006) "A Multiscale integrated approach to the study of the nervous system in health and disease" (DR. 1553 11.10.2022).

Appendix A. Converge to the continuous model

To challenge numerical convergence, we varied the integration timesteps, Δt , and recovered the results reported in figure 8. We can hence utterly conclude that the reported analysis provides a faithful representation of the underlying continuum dynamical system. These numerical integrations were carried out using the fourth-order Runge–Kutta method [39]. This method is employed for its precision, having a truncation error of $O(\Delta t^4)$, where Δt is the time step of the integration.

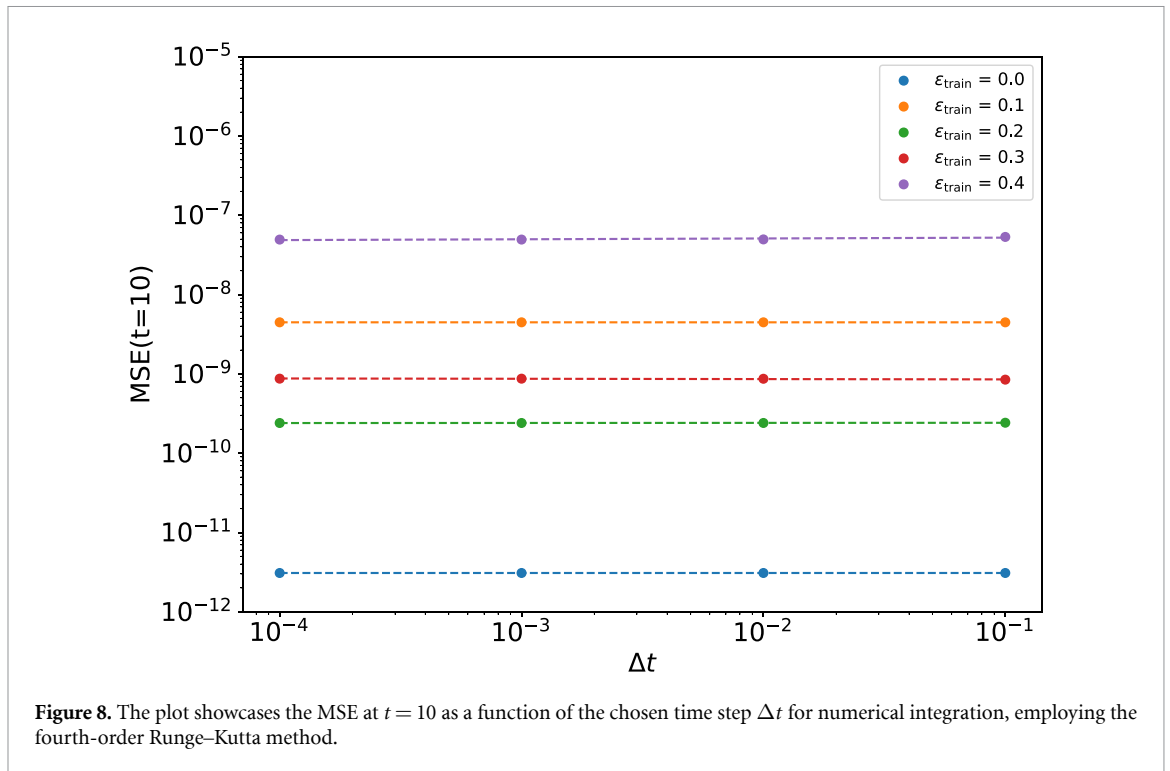


Table B1. This table delineates the performance accuracy of two distinct models, SA-nODE and MLP, when trained under varying levels of noise (ϵ_{train}) and tested in the case $\epsilon_{\text{test}} = 0.0$. The accuracy metrics are separately catalogued for two datasets: MNIST and Fashion MNIST. Each cell indicates the model's accuracy, i.e. equation (7), facilitating a comparative assessment of the models' robustness and adaptability in the presence of noise-induced perturbations in the input data during training process.

ϵ_{train}	MNIST		Fashion MNIST	
	$\mathcal{H}_{\text{SA-nODE}}$	\mathcal{H}_{MLP}	$\mathcal{H}_{\text{SA-nODE}}$	\mathcal{H}_{MLP}
0.0	0.9806 ± 0.0014	0.9825 ± 0.0013	0.8821 ± 0.0032	0.8955 ± 0.0030
0.1	0.9766 ± 0.0015	0.9730 ± 0.0016	0.8788 ± 0.0033	0.8825 ± 0.0032
0.2	0.9692 ± 0.0017	0.9641 ± 0.0018	0.8685 ± 0.0035	0.8745 ± 0.0033

Appendix B. MNIST and Fashion MNIST trained with noises images

We explore the influence of introducing pixel errors into images on the accuracy of our classification algorithm. Specifically, our focus lies in investigating the potential expansion of the basin of attraction MNIST and Fashion MNIST. In table B1, we delineate the behaviour of classification accuracy in the MNIST and Fashion MNIST datasets when trained under varying noise conditions—specifically, with $\epsilon_{\text{train}} = 0$ (without noise) and $\epsilon_{\text{train}} \neq 0$ (with uniform random noise). Interestingly, SA-nODE model performs slightly better than MLP in the presence of noise only for MNIST dataset. For the MNIST dataset, the SA-nODE's robustness is evidenced by its competitive accuracy \mathcal{H} relative to the MLP at almost all noise levels. The MLP has one hidden layer, 512 neurons and ReLU activation function, and an output layer with softmax activation and 10 neuron. In this case we used a categorical cross-entropy loss function.

Appendix C. Confusion matrix, precision, recall, F1 score and misclassified images

The confusion matrix for the MNIST test set using SA-nODE is shown. Each output was compared with all possible attractors, and the matrix was filled based on the highest overlap for each case, in contrast to the metric described in the main text, i.e. equation (7).

C.1. MNIST

The confusion matrix for MNIST reads:

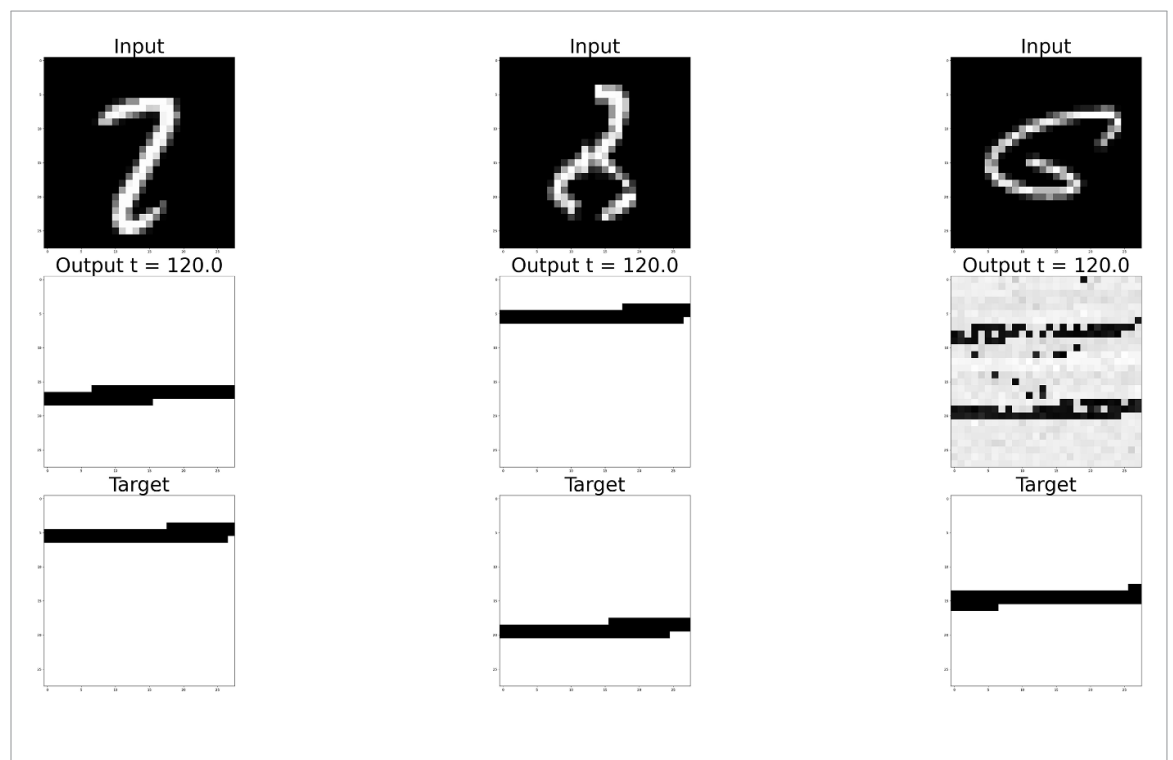
MNIST	Class 0	Class 1	Class 2	Class 3	Class 4	Class 5	Class 6	Class 7	Class 8	Class 9
Class 0	968	0	2	0	1	3	1	1	2	2
Class 1	0	1124	0	1	0	0	2	1	2	3
Class 2	1	2	1012	3	4	0	1	7	4	1
Class 3	0	3	4	992	2	15	0	4	1	4
Class 4	0	0	2	0	964	0	4	0	1	5
Class 5	1	1	0	5	0	864	1	1	3	3
Class 6	4	2	1	0	0	3	945	0	1	2
Class 7	3	0	8	4	2	1	0	1012	3	5
Class 8	1	3	3	3	1	5	3	0	955	2
Class 9	2	0	0	2	8	1	1	2	2	982

In this case, the overall accuracy (the trace of the matrix normalized to the test set dimensionality) is 98.18%, which is higher than the 98.06% accuracy mentioned in the main text. This difference arises because our definition of accuracy, i.e. (7), was specifically designed to demonstrate that the stable attractor was reached (or almost reached) within the fixed integration time.

Given the confusion matrix, it is possible to compute for each class the precision, the recall and the F1 Score metrics. We report them in the matrix below.

MNIST classes	Precision	Recall	F1 score
Class 0	0.9878	0.9878	0.9878
Class 1	0.9903	0.9921	0.9912
Class 2	0.9806	0.9778	0.9792
Class 3	0.9822	0.9678	0.9749
Class 4	0.9817	0.9877	0.9847
Class 5	0.9686	0.9829	0.9757
Class 6	0.9864	0.9864	0.9864
Class 7	0.9844	0.9750	0.9797
Class 8	0.9805	0.9785	0.9795
Class 9	0.9732	0.9820	0.9776

We present below a simple example of misclassified images for the MNIST dataset:



C.2. Fashion MNIST

The confusion matrix for Fashion MNIST reads:


Fashion MNIST	Class 0	Class 1	Class 2	Class 3	Class 4	Class 5	Class 6	Class 7	Class 8	Class 9
Class 0	835	2	25	27	6	0	118	0	2	0
Class 1	1	983	2	9	3	0	4	0	0	0
Class 2	15	0	792	14	106	0	74	1	4	0
Class 3	23	7	8	879	23	1	28	1	5	0
Class 4	4	3	95	36	792	1	67	0	3	0
Class 5	1	0	0	1	2	976	2	12	2	10
Class 6	113	2	73	29	58	0	693	0	12	0
Class 7	0	0	0	0	1	13	2	946	5	21
Class 8	8	3	4	5	8	3	12	0	965	2
Class 9	0	0	1	0	1	6	0	40	2	967

Again, the overall accuracy (the trace of the matrix normalized to the test set dimensionality) is 88.28%, which is higher than the 88.21% accuracy mentioned in the main text. This difference arises because our definition of accuracy, i.e. (7), was specifically designed to demonstrate that the stable attractor was reached (or almost reached) within the fixed integration time.

Given the confusion matrix, it is possible to compute for each class the precision, the recall and the F1 Score metrics. We report them in the matrix below.

Fashion MNIST classes	Precision	Recall	F1 score
Class 0	0.8350	0.8227	0.8288
Class 1	0.9830	0.9810	0.9820
Class 2	0.7920	0.7873	0.7896
Class 3	0.8790	0.9015	0.8901
Class 4	0.7920	0.7912	0.7916
Class 5	0.9760	0.9702	0.9731
Class 6	0.6930	0.7071	0.7000
Class 7	0.9460	0.9575	0.9517
Class 8	0.9650	0.9554	0.9602
Class 9	0.9670	0.9508	0.9588

We present below a simple example of misclassified images for the Fashion MNIST dataset:



Appendix D. Invertibility of the SA-nODE algorithm

Invertible machine learning models can be used for data compression (or related problems [40]), enabling high-fidelity data decompression or synthetic data generation. Invertibility also aids in model interpretability by maintaining information throughout the network’s layers, which can illuminate the model’s decision processes. With these applications in mind we set we begin with an image, represented as $\vec{x}(t = 0)$, serving as the initial condition of our dynamics. This image is allowed to evolve in accordance with its inherent dynamical laws. Post-training, it is constructively known that the image concludes its dynamical trajectory at an attractor.

A pivotal question arises at this juncture: Having the output of our classification algorithm ($x(T)$), is it possible to reconstruct the originating image ($x(0)$)?

Addressing this question, we find that SA-nODE possesses a mechanism for reversing its dynamics. By employing a straightforward transformation, $\tau = T - t$ [41, 42], which effectively reverses the dynamics of the system, we can integrate the dynamics backwards and recapture the initial image. The modified evolutionary law for this backward integration becomes $\dot{\vec{x}} = -\vec{F}(\vec{x})$, with the derivative now taken concerning the variable τ . It is noteworthy to mention that the choice of T in the transformation $\tau = T - t$ is crucial. T should not be excessively large, but large enough to ensure that the system has reached the attractor, as illustrated in figure 9.




Figure 9. The figure depicts the temporal evolution of the 784 pixels of an MNIST image corresponding to the digit 0. Notably, all curves converge to distinct clusters in the positive and negative planes once $t > T_{inv} = 3.2$, where T_{inv} stands for the time where the inversion takes place. For this analysis we set $\Delta t = 0.01$.

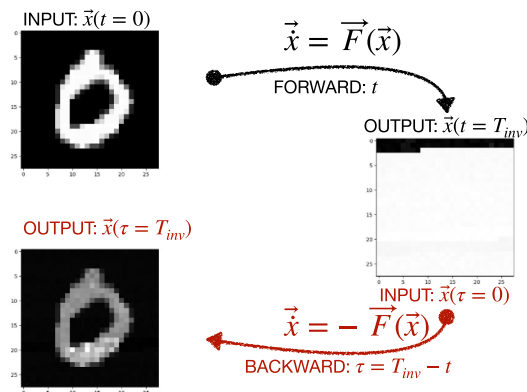


Figure 10. The figure illustrates the process of reversing the time arrow for an MNIST image corresponding to the digit 0. The classification can be easily verified by the image on the right, while the recovered image can be seen at the bottom on the left of the figure. For this analysis we set $\Delta t = 0.0001$, $T_{inv} = 3.2$.


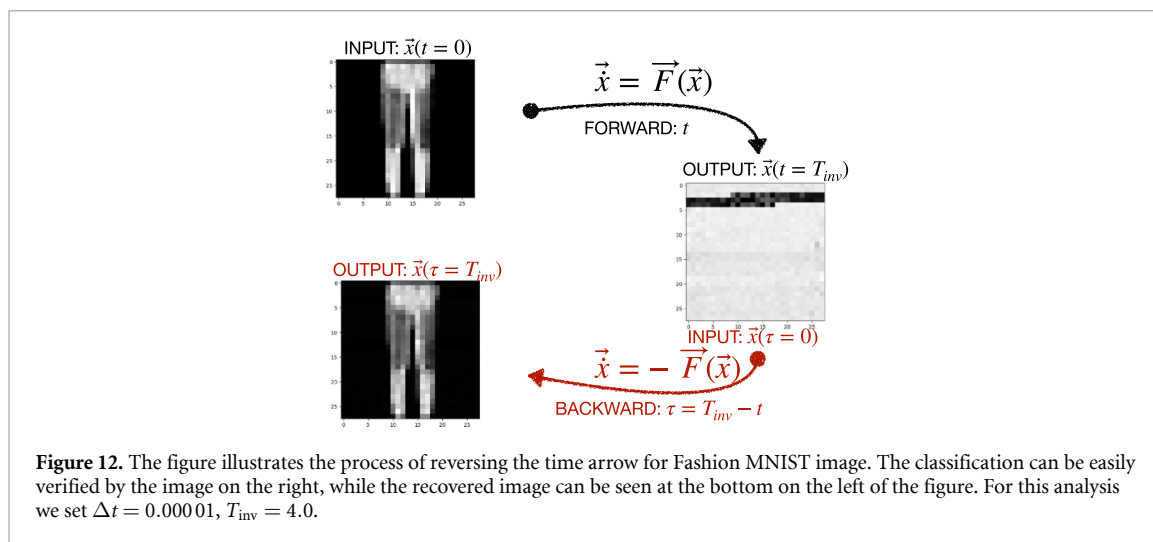


Figure 11. The figure depicts the temporal evolution of the 784 pixels of an Fashion MNIST image corresponding to trousers in figure 12. Notably, all curves converge to distinct clusters in the positive and negative planes once $t > T_{inv} = 4.0$. For this analysis we set $\Delta t = 0.01$.

Figure 9 and 10 (resp. 11 and 12) provide a graphical representation of the journey from a starting image to the recovered image, obtained via the backward dynamic integration. Notably, the figures obtained closely resembles the input image, with any discernible differences likely attributed to numerical errors inherent in the integration process.



ORCID iDs

Raffaele Marino  <https://orcid.org/0000-0002-2311-4380>

Lorenzo Buffoni  <https://orcid.org/0000-0002-8191-3375>

Duccio Fanelli  <https://orcid.org/0000-0001-8545-9424>

References

- [1] Bishop C M 2007 *Pattern recognition and machine learning* 5th edn (Springer) (<https://link.springer.com/book/9780387310732>)
- [2] Shalev-Shwartz S and Ben-David S 2014 *Understanding Machine Learning: From Theory to Algorithms* (Cambridge University Press)
- [3] LeCun Y, Bengio Y and Hinton G 2015 Deep learning *Nature* **521** 436–44
- [4] Goodfellow I, Bengio Y and Courville A 2016 *Deep learning* (MIT press) (www.deeplearningbook.org)
- [5] Prince S J 2023 *Understanding Deep Learning* (MIT press) (<http://udlbook.com>)
- [6] Liu Y et al 2023 Summary of ChatGPT-related research and perspective towards the future of large language models *Meta Radiol.* **1** 100017
- [7] Esteva A et al 2019 A guide to deep learning in healthcare *Nat. Med.* **25** 24–29
- [8] Ching T et al 2018 Opportunities and obstacles for deep learning in biology and medicine *J. R. Soc. Interface* **15** 20170387
- [9] Marino R and Ricci-Tersenghi F 2024 Phase transitions in the mini-batch size for sparse and dense two-layer neural networks *Mach. Learn.: Sci. Technol.* **5** 015015
- [10] Razzak M I, Naz S and Zaib A 2018 Deep learning for medical image processing: overview, challenges and the future *Classification in Bioapps: Automation of Decision Making* ed N Dey, A S Ashour and S Borra (Springer) pp 323–50
- [11] Marino R and Macris N 2023 Solving non-linear Kolmogorov equations in large dimensions by using deep learning: a numerical comparison of discretization schemes *J. Sci. Comput.* **94** 8
- [12] Heaton J B, Polson N G and Witte J H 2017 Deep learning for finance: deep portfolios *Appl. Stoch. Models Bus. Ind.* **33** 3–12
- [13] Sezer O B, Gudelek M U and Ozbayoglu A M 2020 Financial time series forecasting with deep learning : a systematic literature review: 2005–2019 *Appl. Soft Comput.* **90** 106181
- [14] Marino R 2021 Learning from survey propagation: a neural network for MAX-E-3-SAT *Mach. Learn.: Sci. Technol.* **2** 035032
- [15] Rudin C 2019 Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead *Nat. Mach. Intell.* **1** 206–15
- [16] Giambagli L, Buffoni L, Carletti T, Nocentini W and Fanelli D 2021 Machine learning in spectral domain *Nat. Commun.* **12** 1330
- [17] Buffoni L, Civitelli E, Giambagli L, Chicchi L and Fanelli D 2022 Spectral pruning of fully connected layers *Sci. Rep.* **12** 11201
- [18] Chen R T Q, Rubanova Y, Bettencourt J and Duvenaud D 2018 Neural ordinary differential equations *Proc. 32nd Int. Conf. on Neural Information Processing Systems. NIPS'18* (Curran Associates Inc.) pp 6572–83
- [19] Bishop C M and Bishop H 2023 *Deep Learning* (Springer)
- [20] Pontryagin L S 2018 *Mathematical Theory of Optimal Processes* (Routledge)
- [21] He K, Zhang X, Ren S and Sun J 2016 Identity Mappings in Deep Residual Networks *Computer Vision – Eccv 2016* ed B Leibe, J Matas, N Sebe and M Welling (Springer) pp 630–45
- [22] He K, Zhang X, Ren S and Sun J 2016 Deep Residual Learning for Image Recognition *2016 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* pp 770–8
- [23] Hasani R, Lechner M, Amini A, Rus D and Grosu R 2021 Liquid Time-constant Networks *Proc. AAAI Conf. on Artificial Intelligence* vol 35 pp 7657–66
- [24] Funahashi K I and Nakamura Y 1993 Approximation of dynamical systems by continuous time recurrent neural networks *Neural Netw.* **6** 801–6
- [25] Beer R D 1995 On the dynamics of small continuous-time recurrent neural networks *Adapt. Behav.* **3** 469–509
- [26] Zhang H, Wang Z and Liu D 2014 A comprehensive review of stability analysis of continuous-time recurrent neural networks *IEEE Trans. on Neural Netw. Learn. Syst.* **25** 1229–62
- [27] Lecun Y, Bottou L, Bengio Y and Haffner P 1998 Gradient-based learning applied to document recognition *Proc. of the IEEE.* **86** 2278–324
- [28] Sherrington D and Kirkpatrick S 1975 Solvable model of a spin-glass *Phys. Rev. Lett.* **35** 1792–6

- [29] Panchenko D 2013 *The Sherrington-Kirkpatrick Model* (Springer)
- [30] Zankoc C, Biancalani T, Fanelli D and Livi R 2017 Diffusion approximation of the stochastic Wilson–Cowan model *Chaos, Solitons & Fractals* **103** 504–12
- [31] Chicchi L, Giambagli L, Buffoni L, Carletti T, Ciavarella M and Fanelli D 2021 Training of sparse and dense deep neural networks: fewer parameters, same performance *Phys. Rev. E* **104** 054312
- [32] Chicchi L, Fanelli D, Giambagli L, Buffoni L and Carletti T 2023 Recurrent spectral network (RSN): shaping a discrete map to reach automated classification *Chaos Solitons Fractals* **168** 113128
- [33] Chicchi L, Giambagli L, Buffoni L, Marino R and Fanelli D 2024 Complex recurrent spectral network *Chaos Solitons Fractals* **184** 114998
- [34] Kingma D and Ba J 2015 Adam: a method for stochastic optimization *Int. Conf. on Learning Representations (ICLR)*
- [35] Marino R 2024 SA-nODE (<https://github.com/RaffaeleMarino/SA-nODE>)
- [36] Mézard M, Parisi G and Virasoro M A 1987 *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications* vol 9 (World Scientific Publishing Company)
- [37] Deng L 2012 The MNIST database of handwritten digit images for machine learning research [Best of the Web] *IEEE Signal Process. Mag.* **29** 141–2
- [38] Xiao H, Rasul K, Vollgraf R Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms 2017
- [39] Butcher J C 2016 *Numerical Methods for Ordinary Differential Equations* (Wiley)
- [40] Marino R and Kirkpatrick S 2023 Hard optimization problems have soft edges *Sci. Rep.* **13** 3671
- [41] Han J, Jentzen A and Weinan E 2018 Solving high-dimensional partial differential equations using deep learning *Proc. Natl Acad. Sci.* **115** 8505–10
- [42] Marino R and Aurell E 2016 Advective-diffusive motion on large scales from small-scale dynamics with an internal symmetry *Phys. Rev. E* **93** 062147