


**Training of sparse and dense deep neural networks: Fewer parameters, same performance**Lorenzo Chicchi,<sup>1</sup> Lorenzo Giambagli<sup>1</sup>,<sup>✉</sup> Lorenzo Buffoni<sup>1</sup>,<sup>✉</sup> Timoteo Carletti<sup>2</sup>,<sup>✉</sup> Marco Ciavarella,<sup>1</sup> and Duccio Fanelli<sup>1</sup><sup>1</sup>*Dipartimento di Fisica e Astronomia, Università di Firenze, INFN and CSDC, Via Sansone 1, 50019 Sesto Fiorentino, Florence, Italy*<sup>2</sup>*naXys, Namur Institute for Complex Systems, University of Namur, Namur B5000, Belgium* (Received 14 June 2021; revised 6 October 2021; accepted 26 October 2021; published 29 November 2021)

Deep neural networks can be trained in reciprocal space by acting on the eigenvalues and eigenvectors of suitable transfer operators in direct space. Adjusting the eigenvalues while freezing the eigenvectors yields a substantial compression of the parameter space. This latter scales by definition with the number of computing neurons. The classification scores as measured by the displayed accuracy are, however, inferior to those attained when the learning is carried in direct space for an identical architecture and by employing the full set of trainable parameters (with a quadratic dependence on the size of neighbor layers). In this paper, we propose a variant of the spectral learning method as in Giambagli *et al.* [*Nat. Commun.* **12**, 1330 (2021)], which leverages on two sets of eigenvalues for each mapping between adjacent layers. The eigenvalues act as veritable knobs which can be freely tuned so as to (1) enhance, or alternatively silence, the contribution of the input nodes and (2) modulate the excitability of the receiving nodes with a mechanism which we interpret as the artificial analog of the homeostatic plasticity. The number of trainable parameters is still a linear function of the network size, but the performance of the trained device gets much closer to those obtained via conventional algorithms, these latter requiring, however, a considerably heavier computational cost. The residual gap between conventional and spectral trainings can be eventually filled by employing a suitable decomposition for the nontrivial block of the eigenvectors matrix. Each spectral parameter reflects back on the whole set of internode weights, an attribute which we effectively exploit to yield sparse networks with stunning classification abilities as compared to their homologs trained with conventional means.

DOI: [10.1103/PhysRevE.104.054312](https://doi.org/10.1103/PhysRevE.104.054312)**I. INTRODUCTION**

Automated learning from data via deep neural networks [1–4] is becoming popular in an ever-increasing number of applications [5–8]. Systems can learn from data by identifying distinctive features which form the basis of decision making with minimal human intervention. The weights, which link adjacent nodes across feedforward architectures, follow the optimization algorithm and store the information needed for the trained network to perform the assigned tasks with unprecedented fidelity [9–11]. A radically new approach to the training of a deep neural network has been recently proposed which anchors the process to reciprocal space rather than to the space of the nodes [12]. Reformulating the learning in reciprocal space enables one to shape key collective modes, the eigenvectors, which are implicated in the process of progressive embedding, from the input layer to the detection point. Even more interestingly, one can assume the eigenmodes of the interlayer transfer operator to align along suitable random directions and identify the associated eigenvalues as a target for the learning scheme. This results in a dramatic compression of the training parameter space, yielding accuracies which are superior to those attained with conventional methods restricted to operate with an identical number of tunable parameters. Nonetheless, neural networks trained in the space of nodes with no restrictions on the set of adjusted weights achieve better classification scores as compared to their

spectral homologs with quenched eigendirections. In the former case, the number of free parameters grows as the product of the sizes of adjacent layer pairs, thus quadratically in terms of hosted neurons. In the latter, the number of free parameters increases linearly with the size of the layers (hence with the number of neurons) when the eigenvalues are solely allowed to change. Also training the eigenvectors amounts to dealing with a set of free parameters equivalent to that employed when the learning is carried out in direct space: in this case, the two methods yield performances which are therefore comparable.

Starting from this setting, we begin by discussing a straightforward generalization of the spectral learning scheme presented in [12], which proves, however, effective in securing a significant improvement on the recorded classification scores while still optimizing a number of parameters which scales linearly with the size of the network. The proposed generalization paves the way to a biomimetic interpretation of the spectral training scheme. The eigenvalues can be tuned so as to magnify or damp the contribution associated with the input nodes. At the same time, they modulate the excitability of the receiving nodes as a function of the local field. This latter effect is reminiscent of the homeostatic plasticity [13] as displayed by living neurons. Further, we will show that the residual gap between conventional and spectral training methods can be eventually filled by resorting to suitable decompositions of the nontrivial block of the eigenvectors

matrix, which place the emphasis on a limited set of collective variables.

Finally, we will prove that working in reciprocal space turns out to be by far more performant when aiming at training sparse neural networks. Because of the improvement in terms of computational load, and due to the advantage of operating with collective target variables as we will make clear in the following, it is surmised that modified spectral learning of the type here discussed should be considered as a viable standard for deep neural networks training in artificial intelligence applications. Stated differently, the results reported in this work provide evidence that neural networks can be efficiently trained with substantially lower computational cost while maintaining comparable accuracy. The quest for innovative neural network schemes beyond state-of-the-art technology constitutes a rather fertile field of investigation which can be tackled via diverse strategies [14–16]. Compression and acceleration techniques are routinely employed to suit the scope. These are customarily divided into four distinct categories (parameter pruning and quantization, low-rank factorization, transferred or compact convolutional filters, and knowledge distillation) as thoroughly reviewed in the comprehensive survey [17]. We here face the problem from a different angle by aiming at reducing the number of trainable parameters rather than compactifying the underlying network as a whole. This is also the spirit of the methods put forward in [18,19].

## II. SPECTRAL TRAINING

To test the effectiveness of the proposed method we will consider classification tasks operated on three distinct database of images. The first is the celebrated MNIST database of handwritten digits [20], the second is Fashion-MNIST (F-MNIST), a data set of Zalando’s article images, and the third is CIFAR-10 a collection of images from different classes (airplanes, cars, birds, cats, deer, dogs, frogs, horses, ships, and trucks). In all considered cases, use can be made of a deep neural network to perform the sought classification, namely, to automatically assign the image supplied as an input to the class to which it belongs. The neural network is customarily trained via standard backpropagation algorithms to tune the weights that connect consecutive stacks of the multilayered architecture. The assigned weights, the target of the optimization procedure, bear the information needed to allocate the examined images to their reference category.

Consider a deep feedforward network made of  $\ell$  distinct layers and label each layer with the progressive index  $i$  ( $= 1, \dots, \ell$ ). Denote by  $N_i$  the number of computing units, the neurons, that belong to layer  $i$ . The total number of parameters that one seeks to optimize in a dense neural network setting (all neurons of any given layer with  $i < \ell - 1$  are linked to every neurons of the adjacent layer) equals  $\sum_{i=1}^{\ell-1} N_i N_{i+1}$ , when omitting additional bias. As we shall prove in the following, impressive performance can be also achieved by pursuing a markedly different procedure, which requires acting on just  $N_1 + N_\ell + 2 \sum_{i=2}^{\ell-1} N_i$  free parameters (not including bias). To this end, let us begin by reviewing the essence of spectral learning method as set forth in [12].

Introduce  $N = \sum_{i=1}^{\ell} N_i$  and create a column vector  $\vec{n}_1$ , of size  $N$ , whose first  $N_1$  entries are the intensities (from the top left to the bottom right, moving horizontally) as displayed on the pixels of the input image. All other entries of  $\vec{n}_1$  are set to zero. The ultimate goal is to transform  $\vec{n}_1$  into an output vector  $\vec{n}_\ell$ , of size  $N$ , whose last  $N_\ell$  elements reflect the intensities of the output nodes where reading takes eventually place. This is achieved with a nested sequence of linear transformations, as exemplified in the following. Let us focus on the generic  $\vec{n}_k$ , with  $k = 1, \dots, \ell - 1$  (as obtained after  $k$  applications of the procedure outlined below with reference to the specific transition from  $k$  to  $k + 1$ ). This latter vector undergoes a linear transformation to yield  $\vec{n}_{k+1} = \mathbf{A}^{(k)} \vec{n}_k$ . Further,  $\vec{n}_{k+1}$  is processed via a suitably defined nonlinear function, denoted by  $f(\cdot, \beta_k)$ , where  $\beta_k$  stands for an optional bias. Focus now on  $\mathbf{A}^{(k)}$ , a  $N \times N$  matrix with a rather specific structure, as we will highlight below. Posit  $\mathbf{A}^{(k)} = \Phi^{(k)} \Lambda^{(k)} (\Phi^{(k)})^{-1}$ , by invoking a spectral decomposition.  $\Lambda^{(k)}$  is the diagonal matrix of the eigenvalues of  $\mathbf{A}^{(k)}$ . By construction we impose  $(\Lambda^{(k)})_{jj} = 1$  for  $j < \sum_{i=1}^k N_i$  and  $j > \sum_{i=1}^{k+1} N_i$ . The remaining  $N_k$  elements are initially set to random numbers, e.g., extracted from a uniform distribution, and define the target of the learning scheme [21]. Returning to the spectral decomposition of  $\mathbf{A}^{(k)}$ ,  $\Phi^{(k)}$  is assumed to be the identity matrix  $\mathbb{1}_{N \times N}$ , with the inclusion of a subdiagonal rectangular block  $\phi^{(k)}$  of size  $N_{k+1} \times N_k$  (the block starts at line  $N_{k+1}$  and column  $N_k$ ; see [12]). This choice corresponds to dealing with a feedforward arrangement of nested layers. A straightforward calculation returns  $(\Phi^{(k)})^{-1} = 2\mathbb{1}_{N \times N} - \Phi^{(k)}$ , which readily yields  $\mathbf{A}^{(k)} = \Phi^{(k)} \Lambda^{(k)} (2\mathbb{1}_{N \times N} - \Phi^{(k)})$ . In the simplest setting that we shall inspect in the following, the off-diagonal elements of matrix  $\Phi^{(k)}$  are frozen to nominal values, selected at random from a given distribution. In this minimal version, the spectral decomposition of the transfer operators  $\mathbf{A}^{(k)}$  enables one to isolate a total of  $N = \sum_{i=1}^{\ell} N_i$  adjustable parameters, the full collection of nontrivial eigenvalues, which can be self-consistently trained. To implement the learning scheme on these premises, we consider  $\vec{n}_\ell$ , the image on the output layer of the input vector  $\vec{n}_1$ ,

$$\vec{n}_\ell = f(\mathbf{A}^{(\ell-1)} \dots f(\mathbf{A}^{(2)} f(\mathbf{A}^{(1)} \vec{n}_1, \beta_1), \beta_2), \beta_{\ell-1}), \quad (1)$$

and calculate  $\vec{z} = \sigma(\vec{n}_\ell)$  where  $\sigma(\cdot)$  stands for the softmax operation. We then introduce the categorical cross-entropy loss function  $CE(l(\vec{n}_1), \vec{z})$  where the quantity  $l(\vec{n}_1)$  identifies the label attached to  $\vec{n}_1$  reflecting the category to which it belongs via one-hot encoding [22]. More specifically, the  $k$ th element of vector  $l(\vec{n}_1)$  is equal to unit (the other entries being identically equal to zero) if the image supplied as an input is associated to the class of items grouped under label  $k$ .

The loss function can be minimized by acting on a limited set of free parameters, the collection of  $N$  nontrivial eigenvalues of matrices  $\mathbf{A}^{(k)}$  (i.e.,  $N_1 + N_2$  eigenvalues of  $\mathbf{A}^{(1)}$ ,  $N_3$  eigenvalues of  $\mathbf{A}^{(2)}$ , ...,  $N_\ell$  eigenvalues of  $\mathbf{A}^{(\ell-1)}$ ). In principle, the subdiagonal blocks  $\phi^{(k)}$  (the nonorthogonal entries of the basis that diagonalises  $\mathbf{A}^{(k)}$ ) can be optimized in parallel, but this choice nullifies the gain in terms of parameter containment, as achieved via spectral decomposition, when the eigenvalues get solely modulated. The remaining part of this paper is dedicated to overcoming this limitation while

securing the decisive enhancement of the neural network's performance.

### III. IMPROVING THE LEARNING SCHEME AND TESTING ITS PERFORMANCE

The first idea, as effective as it is simple, is to extend the set of trainable eigenvalues. When mapping layer  $k$  into layer  $k + 1$ , we can in principle act on  $N_k + N_{k+1}$  eigenvalues without restricting the training to the  $N_{k+1}$  elements, which were identified as the sole target of the spectral method in its original conception (except for the first mapping, from the input layer to its adjacent counterpart). As we shall clarify in the following, the eigenvalues can be trained twice, depending on whether they originate from incoming or outgoing nodes, along the successive arrangement of nested layers. The global number of trainable parameters is hence  $N_1 + N_\ell + 2 \sum_{i=2}^{\ell-1} N_i$ , as anticipated above (in Appendix A we show that this addition yields a substantial improvement of the displayed accuracy, for the case of the F-MNIST database). A straightforward calculation, carried out in Appendix B, returns a closed analytical expression for  $w_{ij}^{(k)}$ , the weights of the edges linking nodes  $i$  and  $j$  in direct space, as a function of the underlying spectral quantities. This can be written as

$$w_{ij}^{(k)} = (\lambda_{m(j)}^{(k)} - \lambda_{l(i)}^{(k)}) \Phi_{l(i),m(j)}^{(k)}, \quad (2)$$

where  $l(i) = \sum_{s=1}^k N_s + i$  and  $m(j) = \sum_{s=1}^{k-1} N_s + j$ , with  $i \in (1, \dots, N_{k+1})$  and  $j \in (1, \dots, N_k)$ . More specifically,  $j$  runs on the nodes at the departure layer ( $k$ ), whereas  $i$  identifies those sitting at destination (layer  $k + 1$ ). In the above expression,  $\lambda_{m(j)}^{(k)}$  stand for the first (nontrivial)  $N_k$  eigenvalues of  $\mathbf{\Lambda}^{(k)}$ . The remaining (nontrivial)  $N_{k+1}$  eigenvalues are labeled  $\lambda_{l(i)}^{(k)}$ . To help comprehension, denote by  $x_j^{(k)}$  the activity (i.e., the value of the state vector) on nodes  $j$ . Then

$$x_i^{(k+1)} = \sum_{j=1}^{N_k} (\lambda_{m(j)}^{(k)} \Phi_{l(i),m(j)}^{(k)} x_j^{(k)}) - \lambda_{l(i)}^{(k)} \sum_{j=1}^{N_k} (\Phi_{l(i),m(j)}^{(k)} x_j^{(k)}). \quad (3)$$

The eigenvalues  $\lambda_{m(j)}^{(k)}$  modulate the density at the origin, while  $\lambda_{l(i)}^{(k)}$  appears to regulate the local node's excitability relative to the network activity in its neighborhood. This is the artificial analog of the *homeostatic plasticity*, the strategy implemented by living neurons to maintain the synaptic basis for learning, respiration, and locomotion [13].

To illustrate the effectiveness of the proposed methodology we reference Fig. 1, which summarizes a first set of results obtained for MNIST. To keep the analysis as simple as possible here we have chosen to deal with  $\ell = 3$ . The sizes of the input ( $N_1$ ) and output ( $N_3$ ) layers are set by the specificity of the considered data set. Conversely, the size of the intermediate layer ( $N_2$ ) can be changed at will. We then monitor the relative accuracy, i.e., the accuracy displayed by the deep neural networks trained according to different strategies, normalized to the accuracy achieved with an identical network trained with conventional methods. In the upper panel of Fig. 1, the performance of the neural networks trained via the modified spectral strategy (referred to as to *Spectral*) is displayed in blue (triangles). The recorded accuracy is satisfactory (about




FIG. 1. The case of MNIST. Upper panel: the accuracy of the different learning strategies, normalized to the accuracy obtained for an identical deep neural network trained in direct space, as a function of the size of the intermediate layer,  $N_2$ . Triangles stand for the relative accuracy obtained when employing the spectral method (*Spectral*). Pentagons refer to the setting which extends the training to the eigenvectors' blocks via a SVD decomposition. Specifically, matrices  $\mathbf{U}_k$  and  $\mathbf{V}_k$  are randomly generated (with a uniform distribution of the entries) and stay unchanged during optimization. The singular values are instead adjusted together with the eigenvalues which stem from the spectral method (this configuration is labeled *S-SVD*). Diamonds are when the eigenvalues and the elements of the triangular matrix  $\mathbf{R}$  (as follows a QR decomposition of the eigenvectors' blocks) are simultaneously adjusted (*S-QR*). Here  $\mathbf{Q}$  is not taking part in the optimization process (its entries are random number extracted from a uniform distribution and fixed throughout the process). Errors are computed after 10 independent realizations of the respective procedures. Lower panel:  $\rho$  the ratio of the number of tuned parameters (modified spectral, *S-SVD*, and *S-QR* methods vs the conventional one) is plotted against  $N_2$ . In calculating  $\rho$  the contribution of the bias is properly acknowledged. As a reference, the best accuracy obtained over the explored range for the deep network trained with conventional means is 98%.

90% of that obtained with usual means and a few percent more than that obtained with the spectral method of original conception [12]), despite the modest number of trained parameters. To exemplify this, in the bottom panel of Fig. 1 we plot the relative ratio of the number of tuned parameters (*Spectral* vs the conventional one) against  $N_2$  (blue triangles): the reduction in the number of parameters as follows the modified spectral method is staggering. Working with the other employed data set, respectively, F-MNIST and CIFAR-10, yields analogous conclusions (see Appendix C). In Appendix E we also consider the case  $\ell = 4$ , for F-MNIST

One further improvement can be achieved by replacing  $\phi^{(k)}$  with its equivalent singular value decomposition (SVD), a factorization that generalizes the eigendecomposition to rectangular (in this framework,  $N_{k+1} \times N_k$ ) matrices (see [23] for an application to neural networks). In formulas, this amounts to postulate  $\phi^{(k)} = \mathbf{U}_k \mathbf{\Sigma}_k \mathbf{V}_k^T$  where  $\mathbf{V}_k$  and  $\mathbf{U}_k$  are, respectively,  $N_k \times N_k$  and  $N_{k+1} \times N_{k+1}$  real orthogonal matrices. On

the other hand,  $\Sigma_k$  is a  $N_{k+1} \times N_k$  rectangular diagonal matrix, with non-negative real numbers on the diagonal. The diagonal entries of  $\Sigma_k$  are the singular values of  $\phi_k$ . The symbol  $(\cdot)^T$  stands for the transpose operation. The learning scheme hence can be reformulated as follows. For each  $k$ , generate two orthogonal random matrices  $\mathbf{U}_k$  and  $\mathbf{V}_k$ . These latter are not updated during the successive stages of the learning process. At variance, the  $M_{k+1} = \min(N_k, N_{k+1})$  nontrivial elements of  $\Sigma_k$  take active part to the optimization process. For each  $k$ ,  $M_{k+1} + N_k + N_{k+1}$  parameters can be thus modulated to optimize the information transfer, from layer  $k$  to layer  $k + 1$ . Stated differently,  $M_{k+1}$  free parameters add up to the  $N_k + N_{k+1}$  eigenvalues that get modulated under the original spectral approach. One can hence count on a larger set of parameters as compared to that made available via the spectral method, restricted to operate with the eigenvalues. Nonetheless, the total number of parameters still scales with the linear size  $N$  of the deep neural network, and not quadratically, as for a standard training carried out in direct space. This addition (referred to as the *S-SVD* scheme) yields an increase of the recorded classification score, as compared to the setting where the *Spectral* method is solely employed, which is, however, not sufficient to fill the gap with conventional schemes (see Fig. 1). Similar scenarios are found for F-MNIST and CIFAR-10 (see Appendix C), with a varying degree of improvement, which reflects the specificity of the considered data set.

A decisive leap forward is, however, accomplished by employing a QR factorization of matrix  $\phi^{(k)}$ . For  $N_{k+1} > N_k$ , this corresponds to writing the  $N_{k+1} \times N_k$  matrix  $\phi_k$  as the product of an orthogonal  $N_{k+1} \times N_k$  matrix  $\mathbf{Q}_k$  and an upper triangular  $N_k \times N_k$  matrix  $\mathbf{R}_k$ . Conversely, when  $N_{k+1} < N_k$ , we factorize  $\phi_k^T$ , in such a way that the square matrix  $\mathbf{R}_k$  has linear dimension  $N_{k+1}$ . In both cases, matrix  $\mathbf{Q}_k$  is randomly generated and stays frozen during gradient descent optimization. The  $M_{k+1}(M_{k+1} + 1)/2$  entries of the  $M_{k+1} \times M_{k+1}$  matrix  $\mathbf{R}_k$  can be adjusted so as to improve the classification ability of the trained network (this strategy of training, integrated to the *Spectral* method, is termed *S-QR*). Results are depicted in Fig. 1 with (red) diamonds. The achieved performance is practically equivalent to that obtained with a conventional approach to learning. Also in this case  $\rho < 1$ , the gain in parameter reduction being noticeable when  $N_1$  is substantially different (smaller or larger) than  $N_2$  for the case at hand. Interestingly enough, for a chief improvement of the performance, over the SVD reference case, it is sufficient to train a portion of the off-diagonal elements of  $\mathbf{R}$ . In Appendix D we report the recorded accuracy against  $p$ , the probability to train the entries that populate the non-null triangular part of  $\mathbf{R}_k$ . The value of the accuracy attained with conventional strategies to the training is already approached at values of  $p$  which are significantly different from unit [24].

The quest for a limited subset of key parameters which define the target of a global approach to the training is also important for its indirect implications, besides the obvious reduction in terms of algorithmic complexity. As a key application to exemplify this point, we shall consider the problem of performing the classification tasks considered above by training a neural network with a prescribed degree of imposed sparsity. This can be achieved by applying a nonlinear filter on each individual weight  $w_{ij}$ . The nonlinear mask is devised

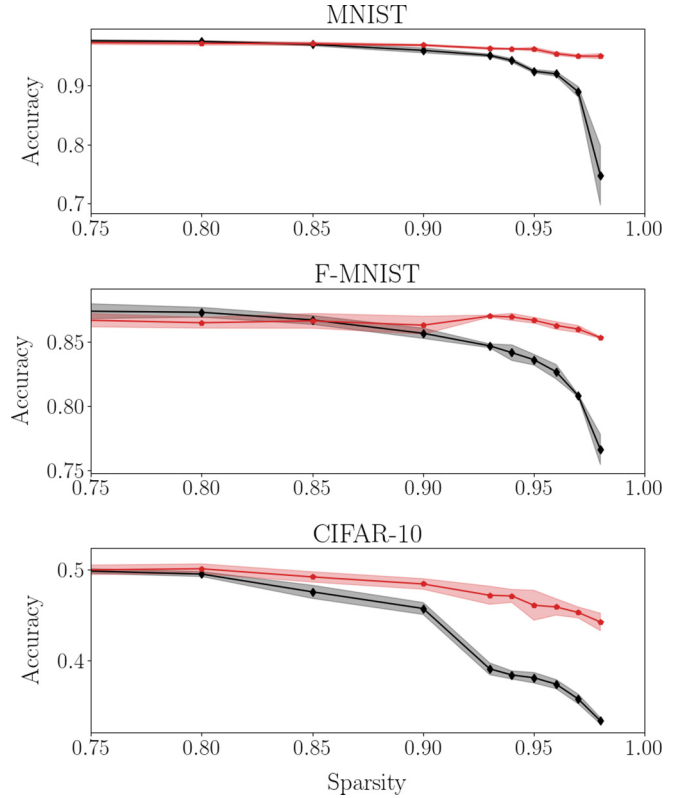


FIG. 2. Training sparse networks. The accuracy of the trained network against the degree of imposed sparsity. Black diamonds refer to the usual training in direct space, while red pentagons refer to the *S-QR* method. From top to bottom: results are reported for MNIST, F-MNIST, and CIFAR-10, respectively. In all cases,  $\ell = 3$ .

so as to return zero (no link present) when  $|w_{ij}| < C$ . Here  $C$  is an adaptive cutoff which can be freely adjusted to allow for the trained network to match the requested amount of sparsity. This latter is measured by a scalar quantity, spanning the interval  $[0,1]$ : when the degree of sparsity is set to zero, the network is dense. At the opposite limit, when the sparsity equals one, the nodes of the network are uncoupled and the information cannot be transported across layers. Working with the usual approach to the training, which seeks to modulate individual weights in direct space, one has to face an obvious problem. When the weight of a given link is turned into zero, then it gets excluded by the subsequent stages of the optimization process. Consequently, a weight that has been silenced cannot regain an active role in the classification handling. This is not the case when operating under the spectral approach to learning, or when complemented by the supplemental features tested above. The target of the optimization, the spectral attributes of the transfer operators, are not biased by any filtering masks: as a consequence, acting on them, one can rescue from oblivion weights that are deemed useless at a given iteration (and, as such, silenced), but which might prove of help at later stages of the training. In Fig. 2 the effect of the imposed sparsity on the classification accuracy is represented for conventional vs the *S-QR* method. The latter is definitely more performant in terms of displayed accuracy when the degree of sparsity gets more pronounced. The drop

in accuracy as exhibited by the sparse network trained with the  $S$ - $QR$  modality is clearly less pronounced than that reported for an equivalent network optimized in direct space. Deviations between the two proposed methodologies indeed become appreciable in the very sparse limit, i.e., when the residual active links are too few for a proper functioning of the direct scheme. In fact, edges which could prove central to the classification, but that are set silent at the beginning, cannot come back to active. At variance, the method anchored to reciprocal space can identify an optimal pool of links (still constraint to the total allowed for) reversing to the active state those that were initially set to null. Interestingly, it can be shown that a few hubs emerge in the intermediate layer, which collect and process the information delivered from the input stack.

**IV. CONCLUSIONS**

Taken altogether, it should be unequivocally concluded that a large body of free parameters, usually trained in machine learning applications, is *de facto* unessential. The spectral learning scheme, supplemented with a QR training of the non-trivial portion of eigenvectors' matrix, enabled us to identify a limited subset of key parameters which prove central to the learning procedure and reflect back with a global impact on the computed weights in direct space. This observation could materialize in a drastic simplification of current machine learning technologies, a challenge at reach via algorithmic optimization carried out in dual space. Quite remarkably, working in reciprocal space yields trained networks with better classification scores when operating at a given degree of imposed sparsity. This finding suggests that shifting the training to the spectral domain might prove beneficial for a wide gallery of deep neural networks applications. Additional tests on more complex data sets than those examined here are clearly needed before concluding on the widespread applicability of the proposed technique.

The codes employed in this paper, as well as a notebook to reproduce our results, can be found in the public repository of this project [25].

**APPENDIX A: IMPROVED SPECTRAL METHOD**

In this Appendix we report on the performance of the improved spectral method. As discussed in the main body of the paper we extend the set of trainable eigenvalues as compared to the method in its original conception. More specifically, when mapping layer  $k$  into layer  $k + 1$ , we can now act on  $N_k + N_{k+1}$  eigenvalues without restricting the training to the  $N_{k+1}$  elements (for  $k \geq 1$ ), as originally done in [12]. Results displayed in Fig. 3 suggest that extending the pool of trainable eigenvalues yields a substantial improvement of the classification scores. Tests refer to the F-MNIST database.

**APPENDIX B: ANALYTICAL CHARACTERISATION OF INTERNODE WEIGHTS IN DIRECT SPACE**

In the following, we will derive Eq. (2) as reported in the main body of the paper. Recall that  $\mathbf{A}^{(k)}$  is a  $N \times N$  matrix. From  $\mathbf{A}^{(k)}$  extract a square block of size  $(N_k +$




FIG. 3. The modified spectral method F-MNIST The accuracy of the different learning strategies, normalized to the accuracy obtained for an identical deep neural network trained in direct space, as a function of the size of the intermediate layer,  $N_2$ . Triangles stand for the relative accuracy obtained when employing the modified spectral method ( $N_k + N_{k+1}$  trainable parameters for each transition from layer  $k$  to layer  $k + 1$ ). Pentagons refer to the spectral method in its original version (when just  $N_{k+1}$  elements, for  $k \geq 1$ , are trained).




FIG. 4. The case of F-MNIST. Upper panel: the accuracy of the different learning strategies, normalized to the accuracy obtained for an identical deep neural network trained in direct space, as a function of the size of the intermediate layer,  $N_2$ . Triangles stand for the relative accuracy obtained when employing the spectral method (*Spectral*). Pentagons refer to the setting which extends the training to the eigenvectors' blocks via a SVD decomposition. Specifically, matrices  $\mathbf{U}_k$  and  $\mathbf{V}_k$  are randomly generated (with a uniform distribution of the entries) and stay unchanged during optimization. The singular values are instead adjusted together with the eigenvalues which originate from the spectral method (*S-SVD*). Diamonds represent when the eigenvalues and the elements of matrix  $\mathbf{R}$  (in a QR decomposition of the eigenvectors' blocks) are simultaneously adjusted (*S-QR*). Errors are computed after 10 independent realizations of the respective procedures. Lower panel: the ratio of the number of tuned parameters (*Spectral*, *S-SVD*, and *S-QR* methods) is plotted against  $N_2$ . As a reference, the best accuracy obtained over the explored range for the deep network trained with conventional means is 90%.




FIG. 5. The case of CIFAR-10. Same as in Fig. 4. The best accuracy obtained over the explored range for the deep network trained with conventional means is 52%.

$N_{k+1} \times (N_k + N_{k+1})$ . This is formed by the set of elements  $\mathbf{A}_{i',j'}^{(k)}$  with  $i' = \sum_{s=1}^{k-1} N_s + i$  and  $j' = \sum_{s=1}^{k-1} N_s + j$ , with  $i = 1, \dots, N_k + N_{k+1}$ ,  $j = 1, \dots, N_k + N_{k+1}$ . For the sake of simplicity, we use  $\mathbf{A}^{(k)}$  to identify the obtained matrix. We proceed in analogy for  $\mathbf{\Lambda}^{(k)}$  and  $\mathbf{\Phi}^{(k)}$ . Then

$$\begin{aligned} A_{ij}^{(k)} &= [\mathbf{\Phi}^{(k)} \mathbf{\Lambda}^{(k)} (2\mathbb{1} - \mathbf{\Phi}^{(k)})]_{ij} \\ &= [2\mathbf{\Phi}^{(k)} \mathbf{\Lambda}^{(k)}]_{ij} - [\mathbf{\Phi}^{(k)} \mathbf{\Lambda}^{(k)} \mathbf{\Phi}^{(k)}]_{ij} \\ &= \alpha_{ij}^{(k)} - \beta_{ij}^{(k)}. \end{aligned} \quad (\text{B1})$$




FIG. 6. The case of MNIST. The (relative) classification accuracy is plotted (red, diamonds, and solid line) against  $p$ , the probability to train the entries that populate the non-null triangular part of  $R$ . The corresponding value of the relative accuracy as computed via the  $S$ -SVD is also reported (green, pentagons, and solid lines). Here  $\ell = 3$ , with  $N_2 = 500$ .




FIG. 7. The case of F-MNIST. As in the caption of Fig. 6. Here  $N_2 = 500$ . The averages are carried out over 10 independent realizations.

Hereafter we shall omit the apex  $(k)$ . Let  $\lambda_1 \dots \lambda_{N_k + N_{k+1}}$  identify the eigenvalues of the transfer operator  $\mathbf{A}$ , namely the diagonal entries of  $\mathbf{\Lambda}$ . In formulas,  $\Lambda_{ij} = \sum_{j=1}^{N_k + N_{k+1}} \delta_{ij} \lambda_j$ .

The quantities  $\alpha_{ij}$  and  $\beta_{ij}$  can be cast in the form

$$\begin{aligned} \alpha_{ij} &= 2 \sum_{k=1}^{N_k + N_{k+1}} \Phi_{ik} \lambda_k \delta_{kj} = 2\Phi_{ij} \lambda_j, \\ \beta_{ij} &= \sum_{k,m=1}^{N_k + N_{k+1}} \Phi_{ik} \lambda_k \delta_{km} \Phi_{mj} = \sum_{m \in \mathcal{I} \cup \mathcal{J}} \Phi_{im} \lambda_m \Phi_{mj}, \end{aligned}$$

where  $j \in \mathcal{J} = (1, \dots, N_k)$  runs on the nodes at the departure layer  $(k)$ , whereas  $i \in \mathcal{I} = (N_k + 1, \dots, N_k + N_{k+1})$ . Hence,  $\mathcal{I} \cup \mathcal{J} = [1, \dots, N_k + N_{k+1}]$ . The above expression for  $\beta_{ij}$




FIG. 8. The case of CIFAR-10. As in the caption of Fig. 6. Here  $N_2 = 700$ . The averages are carried out over five independent realizations.




FIG. 9. The case of a multilayered architecture: the relative accuracy. The relative accuracy as obtained by training a four-layer network with  $N_2 = N_3$  via different strategies. The symbols are as specified in Fig. 4. The analysis refers to F-MNIST.

can be further processed to yield

$$\beta_{ij} = \sum_{m \in \mathcal{J}} \Phi_{im} \lambda_m \Phi_{mj} + \sum_{m \in \mathcal{I}} \Phi_{im} \lambda_m \Phi_{mj} = \Phi_{ij} \lambda_j + \lambda_i \Phi_{ij},$$

where, in the first sum,  $\Phi_{mj} \equiv \delta_{mj}$  for  $m \in \mathcal{J}$  (and because  $j \in \mathcal{J}$ ) while, in the second,  $\Phi_{im} \equiv \delta_{im}$  for  $m \in \mathcal{I}$  (and because  $i \in \mathcal{I}$ )

Finally we can express the difference in (B1) as

$$\alpha_{ij} - \beta_{ij} = 2\Phi_{ij}\lambda_j - \Phi_{ij}\lambda_j - \lambda_i\Phi_{ij} = (\lambda_j - \lambda_i)\Phi_{ij}. \quad (\text{B2})$$

From the above expression, one readily obtains the sought equation upon shifting the index  $i$  to have it spanning the interval  $[1, \dots, N_{k+1}]$ . Recall in fact that, by definition,  $\mathbf{w}$  (the matrix of the weights; see the main body of the paper) is a  $N_k \times N_{k+1}$  matrix.

#### APPENDIX C: TESTING THE *S*-SVD AND *S*-QR METHODS ON THE F-MNIST AND CIFAR-10 DATABASES

Figures 4 and 5 display the accuracy of the *S*-SVD and *S*-QR methods when applied to the case of F-MNIST and CIFAR-10. The analysis refers to a three-layer setting. The results displayed in the figures are in line with those discussed in the main body of the paper.

#### APPENDIX D: REDUCING THE NUMBER OF TRAINABLE PARAMETERS IN THE *S*-QR METHOD

Introduce  $p \in [0, 1]$ . When  $p = 0$ , the diagonal elements of  $\mathbf{R}$  in the *S*-QR method are solely trained. The off-diagonal




FIG. 10. The case of a multilayered architecture: training a sparse network. The accuracy of the trained network against the degree of imposed sparsity. Black diamonds refer to the usual training in direct space, while red pentagons refer to the *S*-QR method. The analysis is carried out for F-MNIST. Here  $N_2 = N_3 = 500$ .

elements are instead frozen to random values. In the opposite limit, when  $p = 1$  all elements of matrix  $\mathbf{R}$  are assumed to be trained. Intermediate values of  $p$  interpolate between the aforementioned limiting conditions. More specifically, the entries that undergo optimization are randomly chosen from the pool of the available ones, as reflecting the selected fraction. In Fig. 6 the relative accuracy for MNIST is plotted against  $p$ . Here the network is made of  $\ell = 3$  layers with  $N_2 = 500$ . A limited fraction of parameters is sufficient to approach the accuracy displayed by the network trained with conventional means. In Figs. 7 and 8 the results relative to F-MNIST and CIFAR-10 are respectively reported.

#### APPENDIX E: TESTING THE PERFORMANCE OF THE INTRODUCED METHODS ON A MULTILAYERED ARCHITECTURE

In this Appendix we will test the setting of a multilayered architecture by generalising beyond the case study  $\ell = 3$  that we employed in the main body of the paper. More specifically, we have trained according to different modalities a four-layer ( $\ell = 4$ ) deep neural network, by modulating  $N_2 = N_3$  over a finite window. As usual, the size of the incoming and outgoing layers are set by the specificity of the examined data sets. The results reported in Fig. 9 refer to F-MNIST and confirm that the *S*-QR strategy yields performance that are comparable to those reached with conventional learning approaches, but relying on a much smaller set of trainable parameters. In Fig. 10 the effect of the imposed sparsity on the classification accuracy is displayed for both the conventional and *S*-QR method. Similar conclusions can be reached for MNIST and CIFAR-10.

- [1] C. M. Bishop, *Pattern Recognition and Machine Learning*, 2011 edition (Springer, New York, 2011).  
 [2] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, Wiley Series in Telecommunications (Wiley, New York, 1991).

- [3] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (Springer Science & Business Media, New York, 2009).  
 [4] A. Burkov, *The Hundred-Page Machine Learning Book*, <http://themlbook.com/> (2019).

- [5] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA, 2018).
- [6] A. Graves, A.-R. Mohamed, and G. Hinton, Speech recognition with deep recurrent neural networks, in *Proc. 2013 IEEE Intl. Conf. on Acoustics, Speech and Signal Processing* (IEEE, New York, 2013), pp. 6645–6649.
- [7] N. Sebe, I. Cohen, A. Garg, and T. S. Huang, *Machine Learning in Computer Vision*, Vol. 29 (Springer Science & Business Media, New York, 2005).
- [8] S. Grigorescu, B. Trasnea, T. Cocias, and G. Macesanu, A survey of deep learning techniques for autonomous driving, *J. Field Robotics* **37**, 362 (2020).
- [9] M. Chen, S. Mao, and Y. Liu, Big data: A survey, *Mobile Netw. Appl.* **19**, 171 (2014).
- [10] E. Meyers and L. Wolf, Using biologically inspired features for face processing, *Intl. J. Comput. Vision* **76**, 93 (2008).
- [11] L. Caponetti, C. A. Buscicchio, and G. Castellano, Biologically inspired emotion recognition from speech, *EURASIP J. Adv. Signal Proc.* **2011**, 24 (2011).
- [12] L. Giambagli, L. Buffoni, T. Carletti, W. Nocentini, and D. Fanelli, Machine learning in spectral domain, *Nat. Commun.* **12**, 1330 (2021).
- [13] F. R. Surmeier, D. A mechanism for homeostatic plasticity, *Nat. Neurosci.* **7**, 691 (2004).
- [14] I. V. Oseledets, Tensor-train decomposition, *SIAM J. Sci. Comput.* **33**, 2295 (2011).
- [15] M. Dusenberry, G. Jerfel, Y. Wen, Y. Ma, J. Snoek, K. Heller, B. Lakshminarayanan, and D. Tran, Efficient and scalable Bayesian neural nets with rank-1 factors, in *Proceedings of the 37th International Conference on Machine Learning* (PMLR, 2020), pp. 2782–2792.
- [16] L. Buffoni, E. Civitelli, L. Giambagli, L. Chicchi, and D. Fanelli, Spectral pruning of fully connected layers: Ranking the nodes based on the eigenvalues, [arXiv:2108.00940](https://arxiv.org/abs/2108.00940).
- [17] Y. Cheng, D. Wang, P. Zhou, and T. Zhang, A survey of model compression and acceleration for deep neural networks, [arXiv:1710.09282](https://arxiv.org/abs/1710.09282).
- [18] G. Hinton, O. Vinyals, and J. Dean, Distilling the knowledge in a neural network, [arXiv:1503.02531](https://arxiv.org/abs/1503.02531).
- [19] D. Zhang, H. Wang, M. Figueiredo, and L. Balzano, Learning to share: Simultaneous parameter tying and sparsification in deep learning, in *Proc. Intl. Conf. on Learning Representations* (2018).
- [20] Y. LeCun, C. Cortes, and C. J. C. Borges, The MNIST database of handwritten digits, <http://yann.lecun.com/exdb/mnist/> (1998).
- [21] The only noticeable exception is when  $k = 1$ . In this case, the first  $N_1$  diagonal elements of  $\Lambda^{(1)}$  take part in the training.
- [22] C. C. Aggarwal, *Neural Networks and Deep Learning* (Springer, New York, 2018).
- [23] M. Gabri e, A. Manoel, Cl ement Luneau, J. Barbier, N. Macris, F. Krzakala, and L. Zdeborova, Entropy and mutual information in models of deep neural networks, *J. Stat. Mech.: Theory Exp.* **124014** (2019).
- [24] The possibility of optimizing the initial choice for the spectral parameters to be trained remains to be fully investigated. Here we have performed an empiric search by scanning across different distributions and parameters range for the initial random matrix elements. We could not, however, identify any noticeable effects on the displayed classification accuracy, at least within the explored ranges.
- [25] [https://github.com/Jamba15/Spectral\\_parameters\\_reduction](https://github.com/Jamba15/Spectral_parameters_reduction).