Neurocomputing 455 (2021) 368–378

Contents lists available at ScienceDirect

Neurocomputing

journal homepage: www.elsevier.com/locate/neucom

Neural network learning of improved compressive sensing sampling and receptive field structure

Victor J. Barranca [a,*]

[a] Department of Mathematics and Statistics, Swarthmore College, 500 College Avenue, Swarthmore, PA 19081, USA

ARTICLE INFO

Article history:
Received 7 December 2020
Revised 5 April 2021
Accepted 16 May 2021
Available online 20 May 2021
Communicated by Zidong Wang

Keywords:
Signal processing
Mathematical neuroscience
Neural networks
Receptive fields

ABSTRACT

While the theory of compressive sensing (CS) in modern signal processing typically indicates that uniformly random sampling facilitates the efficient recovery of sparse signals, such measurements are infeasible in many engineering applications and are not well reflected by the constraints of natural systems, including neuronal networks in the brain. Uniformly random sampling also does not leverage the underlying structure of many classes of signals, and may therefore be suboptimal in these cases. We address these issues by formulating a novel neural network framework for learning improved CS sampling based on the intrinsic structure present in classes of training signals. Beyond sparsity in an appropriate domain, this approach does not assume knowledge of any specific signal statistics and is purely data-driven. The learning methodology is biologically realistic in that it utilizes (1) asymmetric feedback and feedforward connections in the neural network and (2) only information from adjacent layers in training the CS measurement matrix. Observing a broad spectrum of learned sampling paradigms that improve CS signal reconstructions relative to uniformly random sampling, our learned sampling is widely applicable across logistical constraints. Motivated by the receptive field structure of sensory systems, we specifically analyze natural scene inputs and demonstrate improved CS reconstruction as a result of training across several choices of penalization schemes on the sampling weights. Considering this learning is effective even under sparse and spatially localized constraints, as commonly observed in the brain, we hypothesize that neuronal connectivity may have manifested with the aim of providing a compressive encoding of data by leveraging its sparse structure, thereby achieving efficient signal transmission.

1. Introduction

The ubiquity of *compressive sensing* (CS) in modern signal processing theory marks a significant shift in the notion of efficient sampling, with potentially important implications in understanding information encoding in the brain. While signal acquisition prior to CS theory had largely utilized uniform sampling with rate determined by the Shannon-Nyquist theorem [1], compressive sensing instead facilitated the successful reconstruction of signals using far fewer *non-uniform* samples provided the signals were *sparse* in an appropriate domain [2,3]. Leveraging the widespread sparsity of technological and biological data, CS has furnished numerous interdisciplinary scientific advances, taking root in applications such as tomography, medical imaging, radar, microscopy, neuroscience, and genetics [4–10].

The current mathematical framework for the CS reconstruction of sparse data is primarily made rigorous when sampling is *uniformly random*, typically determined by Gaussian or Bernoulli distributions [11–13]. However, in various practical applications, completely random sampling is potentially infeasible or suboptimal. In many cases, sensors may be highly limited in number or expensive [14] and their completely random arrangement may not satisfy logistical constraints, which, for example, may be limited by the size of gaps between sensors [15–17]. Uniformly random sampling also does not take into account the structural characteristics often known for classes of signals in applications [18,19], and, particularly when the number of available measurements is significantly limited, the question of how to improve CS data acquisition in these common cases still remains an active area of investigation. For this reason, various *structured* sampling and recovery methods amenable to CS theory have been proposed for specific classes of signals with known structural features [20–22].

Considering that the unifying characteristics of data are often unknown in many applications and since an optimal CS sampling may be unclear even with partial knowledge of signal structure, several application-driven methodologies have recently been proposed to develop CS sampling protocols using mutual coherence

* Corresponding author.
E-mail address: vbarran1@swarthmore.edu

https://doi.org/10.1016/j.neucom.2021.05.061
0925-2312/

minimization to address these issues [23]. Other recent frameworks instead use deep learning to entirely circumvent CS theory while still making use of sparsity, inferring statistical dependencies among signals and leveraging them to compute signal reconstructions upon sufficient training [24]; however, such approaches still do not address the issue of improving compressive sensing sampling in particular.

In this work, we formulate a novel *neural network-based* and *data-driven* framework for learning improved compressive sensing sampling paradigms in a biologically motivated fashion. Under the observation that signals of interest commonly share certain unifying properties in applications, we iteratively refine the uniformly random sampling typical in CS theory by learning additional structure to incorporate into sampling based on CS reconstruction errors driven by the class of signals composing the training data. Since our method is largely motivated by the fact that learning in the brain is based on similar predictive errors in perception [25–27], we further show how realistic neuronal connectivity may manifest using analogous principles.

The use of neural networks in signal processing has a long history [28–31], and deep learning has achieved profound success in identifying as well as utilizing the structure common in large data sets [32–34]. With its widespread applicability and generalizability, the back-propagation method is the cornerstone of much of modern supervised deep learning [35]. However, the biological plausibility of most methods of deep learning, particularly back-propagation, is still questionable and how the brain may implement neural-network type learning remains unknown. In conflict with the requirements of conventional back-propagation, connectivity in the brain is generally asymmetric and upstream neurons do not necessarily possess knowledge of all downstream synapses [36–39,16]. One recently developed approach for achieving a biologically feasible mechanism for deep learning is *feedback alignment*, which effectively back-propagates error signals via fixed randomly generated feedback weights distinct from the feedforward connectivity adapted via standard error-driven updates [40,41]. Making a natural parallel with sensory processing in the brain and addressing the issue of reconciling back-propagation with the highly nonlinear and generally iterative CS reconstruction algorithms, our methodology for learning improved CS sampling utilizes a modified feedback alignment framework for adapting measurements to input signal structure.

Neural network architecture in deep learning was largely inspired by the hierarchical receptive field structure in the cat primary visual cortex [42]. Motivated by information processing in mammalian visual systems, we focus on CS sampling in the context of natural scenes, though our framework applies to data with sparse structure in general. Amenable to CS theory, typical visual stimuli are sparse in various frequency-based domains [43], and odors as well as soundwaves are similarly known to possess sparse representations [44,45]. Beyond sparsity, natural scenes generally demonstrate strong spatial correlations among nearby pixels and are composed of constituent modes with amplitudes that decay with frequency [19]; we investigate if such additional structure may be further leveraged in improving CS sampling. Analogous to the underdetermined nature of CS signal processing, there are many regions of the brain, such as the early visual system, in which upstream neurons far outnumber downstream neurons [46–48]. Despite these networks of highly dissimilar sizes, information must be preserved across such layers for successful sensation. We hypothesize that evolution has optimized sensory systems to efficiently encode sparse stimuli and explore the extent to which neuronal networks can *learn* optimal sampling of stimuli based on their sparsity and intrinsic properties.

Beyond sensory inputs, network connectivity among neurons in the brain is generally sparse yet often demonstrates spatial localization and modularity [39,49–51]. The receptive field architecture, prominent in the visual, somatosensory, auditory, and olfactory systems, is localized in the sense that specific neurons are most stimulated by a particular range of similar stimuli [52–58]. In the early visual system, for example, this translates to a given downstream ganglion cell sampling spatially clustered image properties prescribed by the activity of upstream photoreceptors, with the response depending on where the light falls in the receptive field, i.e., in the center or surround area, in addition to the light intensity distribution. Taking into account this sparse and localized architecture in the brain, we further explore how sampling, or receptive field structure, may potentially evolve via CS learning with constraints on the density and locality of the connections as observed in vivo.

The organization of this work is as follows. We first provide background by reviewing the primary facets of compressive sensing theory in Section 2.1. We then develop our methodology for learning CS sampling in Section 2.2. Next, we investigate the performance of several variants of our reconstruction framework on natural scene data in Section 3.1, analyzing how and why the sampling successfully adjusts upon learning in Section 3.2. In Section 3.3, we visually investigate the learned sampling paradigms in the spatial and frequency domains, exploring potential parallels with sensory systems. Finally, in Section 4, we examine the ramifications of this work and possible avenues for future investigation.

2. Methods

2.1. Compressive sensing theory

In the context of compressive sensing theory, signal acquisition is performed via linear measurements. For an unknown n-component signal, $x \in R^n$, each measurement of x is obtained by taking a weighted sum of its elements. The set of m measurements (samples) can be generated by a sequence of dot products between the measurement weighting vectors A_i and the signal x for $i = 1, \ldots, m$. Aggregating all such measurement vectors as individual rows in a measurement matrix A, sampling may thus be expressed compactly as

$$Ax = b, \tag{1}$$

where $b \in R^m$ contains the collected measurements. Reconstructing unknown signal x is then equivalent to solving linear system $Ax = b$. However, in the case of highly efficient sampling, when $m \ll n$, this system is generally underdetermined and possess an infinite number of solutions. Depending on the structure of x and the design of the measurement matrix A, CS theory shows an accurate reconstruction is nevertheless possible in this case. A schematic of this sampling and reconstruction paradigm is given by the top of Fig. 1(a).

A key facet of compressive sensing is the observation that many signals demonstrate sparse structure, namely x possesses a representation which contains only a small number of non-zero components in some domain. A signal with n components is defined as k-sparse when there exists a representation that contains at most k components whose magnitude exceeds a small threshold with $k \ll n$. In the case of natural scenes, for example, frequency-based transformations, such as a two-dimensional discrete Fourier, cosine, and wavelet transforms, generally yield highly sparse representations, as demonstrated in Fig. 1(b) for a representative natural scene [43,19].

For sparse signals, it is intuitive to conjecture that the number of samples necessary for an accurate reconstruction should be determined by the sparsity of the signal. CS theory therefore provides a framework for the measurement and subsequent

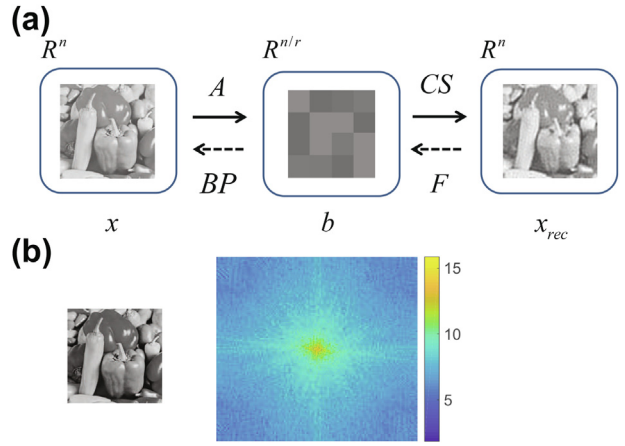

Fig. 1. (a) Schematic model of the learning process. A relatively large n-component signal x is sampled by measurement matrix A, producing m-component measurement vector b. A reduction factor of r is assumed in the sampling, such that $m = n/r$. Next, in accordance with compressive sensing theory, the optimization problem given by Eq. (2) is solved to produce reconstruction x_{rec} (denoted by CS). The reconstruction error is back-propagated (denoted by BP) via the learning rule in Eq. (7) and fixed feedback connectivity matrix F, driving adjustments in the measurement matrix A. This process in repeated across an ensemble of training inputs to produce an improved CS measurement matrix. (b) Example natural scene and its frequency domain representation via the two-dimensional discrete Fourier transform. The natural logarithm of the absolute value is depicted, accentuating differences between lower amplitude frequencies.

reconstruction of signals under the assumption of sparsity, selecting the sparsest reconstruction agreeing with the set of linear measurements [11]. This is equivalent to minimizing the ℓ_0 norm of x, which gives the number of non-zero entries of the signal. Such a minimization problem is computationally intractable in most applications [13], generally requiring a search over all possible subsets of nonzero elements in x consistent with Eq. (1). To address this issue, CS theory further demonstrates that by carefully designing the measurement matrix A, a computationally feasible surrogate ℓ_1 minimization problem yields an identical solution with high probability [13,2]. The resultant optimization problem is

$$\arg\min_{x\in R^n} |x|_{\ell_1} \text{ subject to } Ax = b, \tag{2}$$

where $|x|_{\ell_1} = \sum_{i=1}^{n} |x_i|$. This ℓ_1 minimization problem is solvable in polynomial time using multiple fast algorithms, such as the orthogonal matching pursuit (OMP), the least angle regression (LARS), and the least absolute shrinkage and selection operator (LASSO) methods [59,60], making the reconstruction now computationally tractable. Note that if the measured signal is not sparse in the sampled domain, but is instead sparse under a transform, T^{-1}, then linear system $M\hat{x} = b$, where $M = AT$ and $\hat{x} = T^{-1}x$, may be analogously considered. Once $\hat{x}$ is determined, the solution in the non-sparse domain, $x = T\hat{x}$, can be computed subsequently.

With the reconstruction methodology formulated, the remaining key idea of CS theory addresses how to select a measurement matrix such that solving the ℓ_1 minimization problem (2) indeed yields a viable x with as few measurements as possible. Intuitively, the measurements should be as distinct as possible from the basis used for the sparse domain, with each measurement collecting a large number of nonzero basis element contributions, yielding measurements that give distinct information about the components of x in the sparse domain and thereby revealing a high degree of structure in the measured signal. If instead, for example, each measurement only generated data regarding a single component of x in the sparse domain, then generally the full set of n measurements would be necessary to capture all such components and thus recover x. In CS theory, the maximum correlation (coherence) between the weighting vectors in the measurement matrix A and the column vectors in transformation matrix T is often gauged by the following *maximum mutual coherence* measure for the matrix product $M = AT$,

$$\mu_{max}(M) = \max_{\substack{1\le i,j\le n \\ i\neq j}} \frac{|M_i \cdot M_j|}{\|M_i\|\|M_j\|}, \tag{3}$$

which generally becomes larger with increasing correlation and facilitates rigorous estimates on the number of necessary measurements based on signal sparsity.

Crucially, random weighting vectors in measurements will typically be uncorrelated with any fixed set of basis vectors. Consequently, random measurement matrices are particularly viable for CS recovery. For measurement matrices with independent identically distributed Gaussian or Bernoulli entries, it can be proven that with high probability a successful CS reconstruction is achievable given sufficiently many measurements, as determined by the sparsity of x [11–13]. It is important to emphasize that such randomness in measurements is sufficient but not necessary for successful CS reconstructions, and that certain structured measurement matrices in fact yield improved reconstructions, depending on the class of signals analyzed [20–22,61]. In the subsequent sections, we will demonstrate that it is possible to refine the measurement matrix via learning, iteratively introducing structure based on the inherent properties of the class of signals to be reconstructed and thereby improving CS reconstruction quality.

We remark that for the signal class of natural scenes considered in detail in this work, the reconstruction step in solving Eq. (2) is performed in the sparse domain via the vectorization of the two-dimensional discrete cosine transform of the input image pixel matrix, $\hat{x} = (\hat{x}_1, \ldots, \hat{x}_n) = (D \otimes D)x$. In this case, x is vectorized, $\otimes$ denotes the $n \times n$ Kronecker product

$$D \otimes D = \begin{bmatrix} D_{11}D & \cdots & D_{1\sqrt{n}}D \\ \vdots & \ddots & \vdots \\ D_{\sqrt{n}1}D & \cdots & D_{\sqrt{n}\sqrt{n}}D \end{bmatrix},$$

and D is the $\sqrt{n} \times \sqrt{n}$ one-dimensional discrete cosine transform matrix with entries

$$D_{ij} = \left(D^{-1}\right)^T_{ij} = \omega(i) \cos\left(\frac{(i-1)(2j-1)\pi}{2\sqrt{n}}\right),$$

such that $\omega(1) = (1/n)^{1/4}$ and $\omega(i \neq 1) = (4/n)^{1/4}$.

Since the two-dimensional discrete cosine transform of the stimulus, $\hat{x}$, is sparse, recovering $\hat{x}$ is thereby reduced to Eq. (2) where $x = \hat{x}$ and the effective left-hand side in the constraint is $A(D \otimes D)^{-1}$. We solve this optimization problem using a greedy algorithm known as the Orthogonal Matching Pursuit [59], though a host of alternative algorithms may be similarly applied instead. Once $\hat{x}$ is recovered, we invert the two-dimensional discrete cosine transform and the vectorization to obtain the reconstructed signal x_{rec}.

2.2. Neural network training methodology for optimal CS sampling

From a learning perspective, compressive sensing sampling and signal reconstruction may be viewed as a two-step neural network operation. In the first step, the signal x is sampled by measurement matrix A to produce measurements b as in Eq. (1). Next, the ℓ_1 optimization problem given by Eq. (2) is solved to produce a reconstruction x_{rec}. It is important to underline that only the first operation is linear, since known CS reconstruction algorithms are

generally iterative and nonlinear. The input layer is given by x, the middle layer by b, and the output layer by x_{rec}.

As a whole, this network may be described as an *autoencoder*, where the input and output layers both possess n components and the goal of the network operation is to minimize the distance between the input and output signals. With this aim, the ubiquitous back-propagation method, for example, distributes error signals across the network following a gradient descent approach, typically requiring layer operations to be both known and differentiable to send feedback to upstream layers [35]. For a three-layer network, as in the CS context, the error may be gauged by

$$e = \frac{1}{2}\sum_{i=1}^{n}\left(x_{rec_i} - x_i\right)^2. \tag{4}$$

In general, back-propagation assumes that matrix A determines the connectivity between the input and middle layers and that matrix C determines the connectivity between the middle and output layers. Supposing the inputs into the middle and output layers are potentially integrated by nonlinear and differentiable activation functions, f and g, respectively, the total inputs into arbitrary component i of the middle and output layers, respectively, are prescribed by

$$b_i = f\left(h_i^{(1)}\right) = f\left(\sum_{j=1}^{n} A_{ij}x_j\right) \tag{5a}$$

$$x_{rec_i} = g\left(h_i^{(2)}\right) = g\left(\sum_{j=1}^{m} C_{ij}b_j\right). \tag{5b}$$

We remark that we will adapt these broad assumptions to the context of CS subsequently, but we first briefly summarize the main theory of back-propagation to provide context for our CS learning framework. With the aim of minimizing the reconstruction error, gradient descent and the chain rule ultimately yield the following adjustments to the two connectivity matrices

$$\Delta C_{ij} = -\gamma\frac{\partial e}{\partial C_{ij}} = \gamma\left(x_i - x_{rec_i}\right)g'\left(h_i^{(2)}\right)b_j \tag{6a}$$

$$\Delta A_{jk} = -\gamma\frac{\partial e}{\partial A_{jk}} = \gamma\sum_{i=1}^{n}\left(x_i - x_{rec_i}\right)g'\left(h_i^{(2)}\right)C_{ij}f'\left(h_j^{(1)}\right)x_k, \tag{6b}$$

where γ is the learning rate controlling the magnitude of the connectivity adjustments, ΔC_{ij} and ΔA_{jk}, that occur on each iteration of the learning algorithm.

In the compressive sensing neural network setting, particularly when the number of components, n, in the input layer is far larger than the number of components in the middle layer, m, a closed form expression for the second, typically iterative, nonlinear operation is unknown and thus standard back-propagation is infeasible. Back-propagation also makes the assumption of symmetric connectivity between layers, i.e., reciprocal connections, and shared knowledge of connectivity between all layers, which are both in conflict with known physiological neuronal connectivity [36–39,16].

To address the theoretical and biological challenges imposed by learning CS sampling via back-propagation, we adapt a novel learning methodology known as feedback alignment. In its original context, feedback alignment was developed with the aim of extending the framework of back-propagation to improve the biological plausibility of neural network learning [40]. In rewiring the first layer connectivity matrix, A, as in Eq. (6b), feedback alignment replaces the previously required knowledge of the second layer connectivity matrix, C, with a fixed randomly generated feedback matrix F. In this case, the feedforward connections may adapt with learning while the feedback connections remain constant. Hence, the feedback alignment learning rule, which may be extended similarly to networks with additional layers, simultaneously eliminates the assumptions of symmetric connectivity between layers and shared knowledge of the connectivity structure among layers. Intuitively, the network learns how to make the static feedback useful, such that the feedforward connectivity adjustments eventually facilitate an alignment between the feedforward and feedback weights that allows error signals to be effectively transmitted between layers.

In the specific case of our CS neural network architecture, while the connectivity between the first two layers is indeed linear, namely with activation function $f(x) = x$ and feedforward connectivity prescribed by measurement matrix A, the reconstruction operation performed between the middle and output layers is nonlinear and not prescribed by a closed form expression. Hence, motivated by feedback alignment, we fix the nature of the CS reconstruction via solution to Eq. (2) and back-propagate error signals based on this operation via fixed feedback connectivity matrix F in order to adjust the measurement matrix A. Hence, the CS sampling learning rule is

$$\Delta A_{jk} = -\gamma\frac{\partial e}{\partial A_{jk}} = \gamma\sum_{i=1}^{n}\left(x_i - x_{rec_i}\right)F_{ji}x_k. \tag{7}$$

Learning is carried out iteratively via Eq. (7) based on the CS reconstruction error over an ensemble of signals in a common class that composes the training set. A schematic model of the full recovery process is given by the complete flow diagram in Fig. 1(a). Note that while both feedforward connectivity matrices are trained in conventional feedback alignment, the nonlinear, iterative second feedforward layer operation remains fixed here and nonetheless training only the first layer feedforward connectivity matrix is sufficient for successful learning. It is important to emphasize that this framework is relatively intuitive, biologically feasible, and generalizable to diverse classes of signals with underlying sparse structure because it is data-driven.

3. Results

3.1. Performance analysis

We utilize the neural network methodology for learning enhanced CS sampling discussed in Section 2.2 via stochastic gradient descent on a training set of 10000 inputs in the form of 50×50 pixel patches of natural scenes extracted from the Computational Visual Cognition Laboratory of MIT database [62]. A validation set of 2000 patches drawn from the same database distinct from the training set is utilized to gauge the reconstruction performance. To quantitatively assess the accuracy of each CS reconstruction, we measure the relative reconstruction error defined by

$$\text{Error} = \|x - x_{\text{rec}}\|/\|x\|, \tag{8}$$

where the Euclidean norm $\|x\| = \sqrt{\sum_{i=1}^{n} x_i^2}$ is utilized. For concreteness, the signals are sampled with reduction factor $r = 5$, such that $n = 2500$ and $m = n/r = 500$. Alternative reduction factors may be utilized analogously, yielding potentially distinct learned measurement matrices, though the details of the training process remain identical. A learning rate of $\gamma = 10^{-7}$ is shown to yield relatively fast convergence to a local minimum in error. The measurement matrix A is initialized with independent identically Gaussian distributed elements, as typical in CS theory, with mean $\mu = 0.02$ and standard deviation $\sigma = 0.2$, and the fixed feedback connectivity matrix F is analogously generated with entries given by new realizations of the same Gaussian distributed random variable.

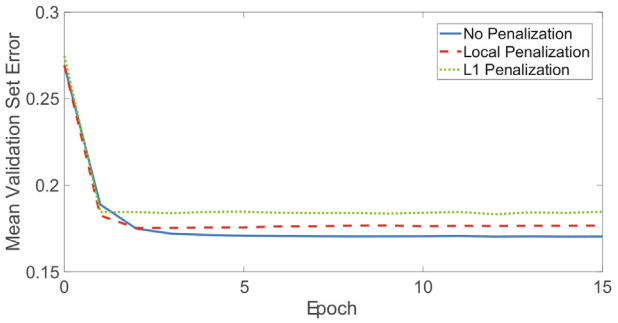

Fig. 2. Mean relative reconstruction error for the CS image reconstructions averaged over the validation set of natural scene patches across epochs. The training parameters are described in detail in Section 3.1. Each input is reconstructed by solving Eq. (2) and the measurement matrix is adjusted according to learning rule (7) for each training input. The error using no penalization is plotted in solid blue, using localized penalization in dashed red, and using ℓ_1 penalization in dotted green.

In Fig. 2, we plot the mean CS relative reconstruction error over the validation set across 15 epochs. Inputs are reconstructed by solving Eq. (2) and the measurement matrix is adjusted according to the learning rule in Eq. (7) for each training image. We observe an initial rapid decrease in error across the first three epochs, subsequently leveling off upon processing further training data. The CS reconstruction error is reduced by over one third as a result of learning, demonstrating marked improvements in the resultant sampling scheme beyond the conventional uniformly random sampling with which the network was initialized. Considering that the error remains approximately constant upon sufficient training, we see that there exists a relatively broad range of low error local minima that corresponds to a widespread space of sampling schemes over which improved compressive sensing is achievable. Hence, as we will affirm in our later analysis in Section 3.2, once a certain degree of structure is present in the sampling, based on the class of training inputs considered and number of samples utilized, signal information may be captured with high accuracy by leveraging the distribution of dominant components in the sparse domain.

For comparison, we also investigate the result of learning under several penalization schemes on the sampling structure with the goal of utilizing a small number of weights in each measurement and later to better reflect biological realism. The learning rule given by Eq. (7) adapted to account for a general penalization term is

$$\Delta A_{jk} = -\gamma \frac{\partial}{\partial A_{jk}}\left(e + \lambda R(A_{jk})\right) = \gamma\left(\sum_{i=1}^{n}(x_i - x_{rec_i})F_{ji}x_k - \lambda R'(A_{jk})\right),$$

where cost function $R(\cdot)$ penalizes non-zero entries in measurement matrix A and parameter λ determines the penalization weight.

To incorporate sparsity into the sampling paradigm, we first utilize an ℓ_1 penalization on the connection weights in measurement matrix A, such that $R(A_{jk}) = \mid A_{jk} \mid$ and therefore $R'(A_{jk}) = sign(A_{jk})$. In Fig. 2, we observe CS reconstruction results using ℓ_1 penalization analogous to those in the absence of penalization, indicating that sparsity may be enforced in the weighting of signal components while still yielding improved CS reconstructions over a broad space of sampling paradigms. We remark that since neuronal network connectivity in the brain is generally sparse [39,49], the sampling reflected under ℓ_1 penalization is more physiological and consistent with the notion that the brain may achieve efficient processing of sparse stimuli via CS principles. As seen in our learning framework, for sensory systems, such sparse feedforward connectivity may yield a similarly accurate compressive encoding as dense connectivity while additionally satisfying physical constraints on energy expenditure and the speed of information transfer [63,64].

A yet closer step towards biological realism is to utilize spatial localization in sampling akin to receptive fields [52–58]. To incorporate spatial structure into the sampling, each pixel in the sampled image is assigned a unique (x, y) location with integer coordinates on a $[1, \sqrt{n}] \times [1, \sqrt{n}]$ Cartesian grid reflecting all available pixel locations on n-component image x. Each row of the measurement matrix A is associated with a distinct random location (x_j, y_j) on this Cartesian grid and entry A_{jk} is penalized based on the normalized Euclidean distance, $d(A_{jk})$, from the row location to location (x_k, y_k) of pixel k. We note that this distance measure normalizes the raw Euclidean distance by the maximum possible Euclidean distance between grid points to produce $O(1)$ quantities, and it is held constant throughout the learning process for a given entry of A. Based on this distance measure, we now impose the cost function $R(A_{jk}) = d(A_{jk}) \cdot \mid A_{jk} \mid$ with derivative $R'(A_{jk}) = d(A_{jk}) \cdot sign(A_{jk})$. Hence, sparsity is encouraged by the ℓ_1 norm with increased penalization for more distant connections, thereby yielding distance-dependent sparse sampling. Using the localized penalization, we again observe a rapid initial decrease in CS reconstruction error with learning and diminishing gains in accuracy upon sufficient training. The accuracy achieved is comparable to that of the other two methods, demonstrating a minor improvement from the non-local ℓ_1 penalization.

While post-processing steps, such as applying a rectified linear function to the CS reconstructions, may improve results further by about 20%, we focus on the effect of learned sampling in this work rather than other means of improved reconstruction to limit confounding factors. Applying our framework to alternative signal databases, such as MNIST or sinusoidal functions modeling sound waves, produces similar results, but we investigate natural scenes here for concreteness and to more closely draw parallels with the visual system.

3.2. Sampling analysis

To elucidate the changes in sampling structure that manifest with training for each learning method, we first analyze the spatial distribution of the dominant connections in the measurement matrix A. Eliminating the near-zero connections that are pruned from learning, we limit our considerations to all connections that are at least two standard deviations from the mean connection strength across their respective rows. We note that similar results are observed if instead connections only at least one standard deviation from the row mean are included. Across the remaining dominant connections, we generate for each row a point cloud of Cartesian grid locations corresponding to the sampled pixels. On a row-by-row basis, the centroid of the (x, y) point cloud of sampled pixels is computed and the mean Euclidean distance from the centroid is determined. Finally, taking the average of the mean distance from the centroid across all rows yields an average distance measuring the spatial dispersion of the sampling for the full measurement matrix.

Agreeing with the intuition for the connection penalization schemes discussed in the previous section, we observe in Fig. 3 that only the learning method with localized penalization demonstrates any significant change in the average distance of measurements upon training. Via localized penalization, we observe an initial sharp decrease in distance with learning that saturates upon sufficient training, showing the same qualitative structure as the CS reconstruction error and minimizing distance when high reconstruction accuracy is achieved. We conclude that by adding some degree of spatial localization into the CS sampling, as demonstrated by realistic neuronal networks, improved reconstructions

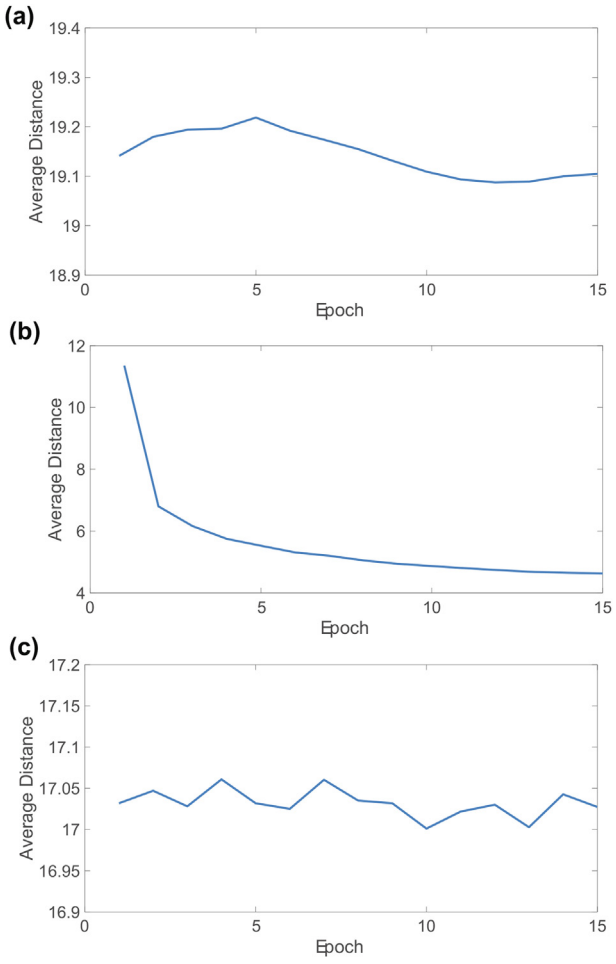

Fig. 3. Mean sampling distance averaged across the rows of A, computed upon the completion of each epoch using: (a) no penalization, (b) localized penalization, and (c) ℓ_1 penalization. The distance measure utilized is discussed in detail in Section 3.2.

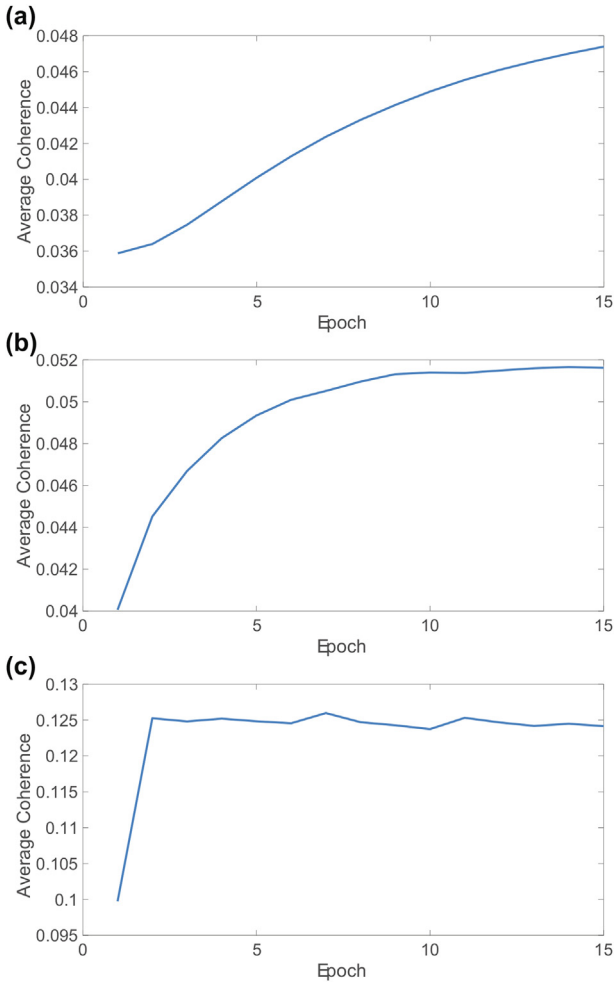

Fig. 4. Average mutual coherence in sampling, computed upon the completion of each epoch using: (a) no penalization, (b) localized penalization, and (c) ℓ_1 penalization.

are indeed obtainable. Since nearly comparable accuracy is yielded via training using the other penalization methods, we also note that alternative types of sampling structure may yield improved CS recovery, but likely the physical constraints of sensory systems encourage the implementation of localized sampling and engineered sampling devices may similarly be subject to such spatial limitations [15–17].

With the aim of investigating the general structure introduced in CS sampling with learning beyond distance-dependence, we next analyze correlations between columns in the effective measurement matrix $M = AT$. To quantify the mutual coherence between the measurement matrix A and transformation matrix $T = (D \otimes D)^{-1}$ while not over-weighting a single outlier as the network learns, we utilize the following *average mutual coherence* measure

$$\mu_{avg}(M) = \frac{1}{n(n-1)} \sum_{\substack{1 \leqslant i,j \leqslant n \\ i \neq j}} \frac{|M_i \cdot M_j|}{\|M_i\| \|M_j\|},$$

which gives a more global analogue to the maximum mutual coherence measure μ_{max}, discussed in Section 2.1, typically utilized in CS theory because it lends itself to theoretical results [11–13].

We depict in Fig. 4, the average mutual coherence with training, which initially increases and then levels off for each learning method. This suggests that increasing the average mutual coherence up to a sufficient level, based on the training input signal structure, improves CS reconstruction quality. The average mutual coherence reaches a clear plateau since further increases in coherence may prove counterproductive to the demands of CS theory and would instead cause the reconstruction error to increase.

It is important to remark that neural network learning in our CS framework indeed coincides with the mechanisms observed in explaining the success of feedback alignment. Given the feedback connectivity matrix F remains fixed throughout learning whereas the feedforward measurement matrix A is adjusted, we demonstrate that the two align with learning as the network is trained to make the feedback information useful.

To gauge the correlation between the feedforward and feedback signals, we compute the angle between them with training, which was shown to diminish with training in the original formulation of feedback alignment [40]. Specifically, we use the angle measure

$$\theta(v, w) = \cos^{-1}\left(\frac{v \cdot w}{\|v\| \|w\|}\right)$$

for $v = F(x_{rec} - x)$ and $w = A(x_{rec} - x)$.

Before training, the angle between the feedforward and feedback signals is near $\pi/2$ radians, and thus the signals are unaligned. However, as shown in Fig. 5, with each penalization method, this angle precipitously decreases before ultimately saturating once the signals are well aligned upon sufficient training. Thus, fixed

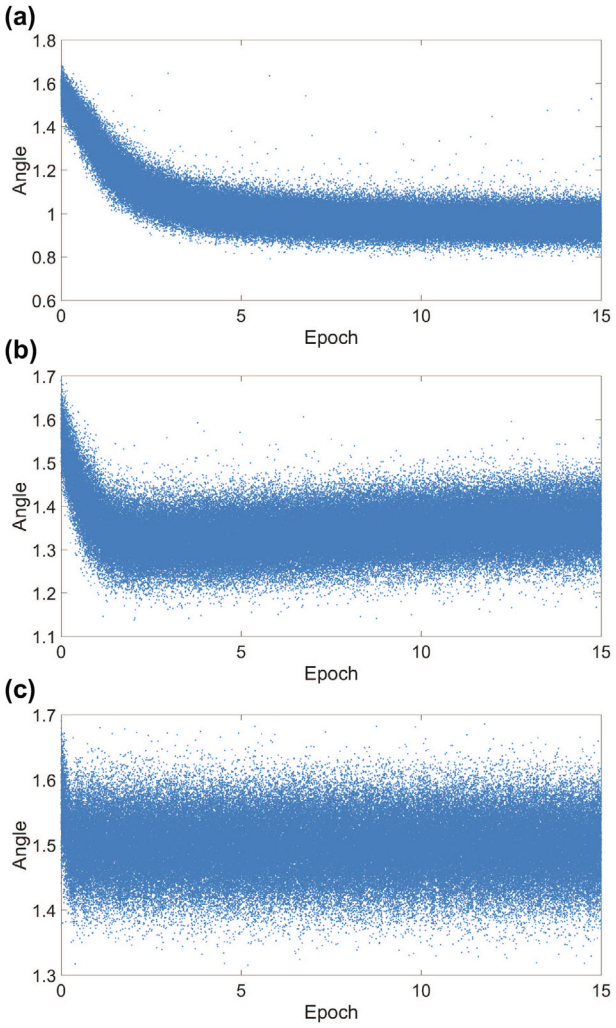

Fig. 5. Angle between feedforward and feedback signals using: (a) no penalization, (b) localized penalization, and (c) ℓ_1 penalization.

random feedback connections, which are not reciprocal with the feedforward connections as observed for physiological neuronal networks [36–39,16], send instructive training signals for rewiring the feedforward measurement matrix and thereby improve CS reconstructions in a biologically plausible manner. We posit that connectivity between layers of disparate sizes in physiological neuronal networks may develop according to similar learning principles in order to facilitate an efficient compressive coding of information.

3.3. Sampling structure

For the signal class of natural scenes analyzed in this work, we may visually inspect the sampling structure that arises from our learning framework. For each penalization type, we observe that the sampling shifts from the independent identically distributed Gaussian entries conventional in CS theory with which the network was initialized to a more structured though still relatively incoherent sampling scheme.

In Fig. 6, the spatial domain representation of the sampling across 400 representative rows of the measurement matrix yielding minimal mean reconstruction error over the validation set is depicted for each penalization method. Note that the spatial domain representation for a given row of size n is computed by reshaping the weighting vector into the form of a $\sqrt{n} \times \sqrt{n}$ image

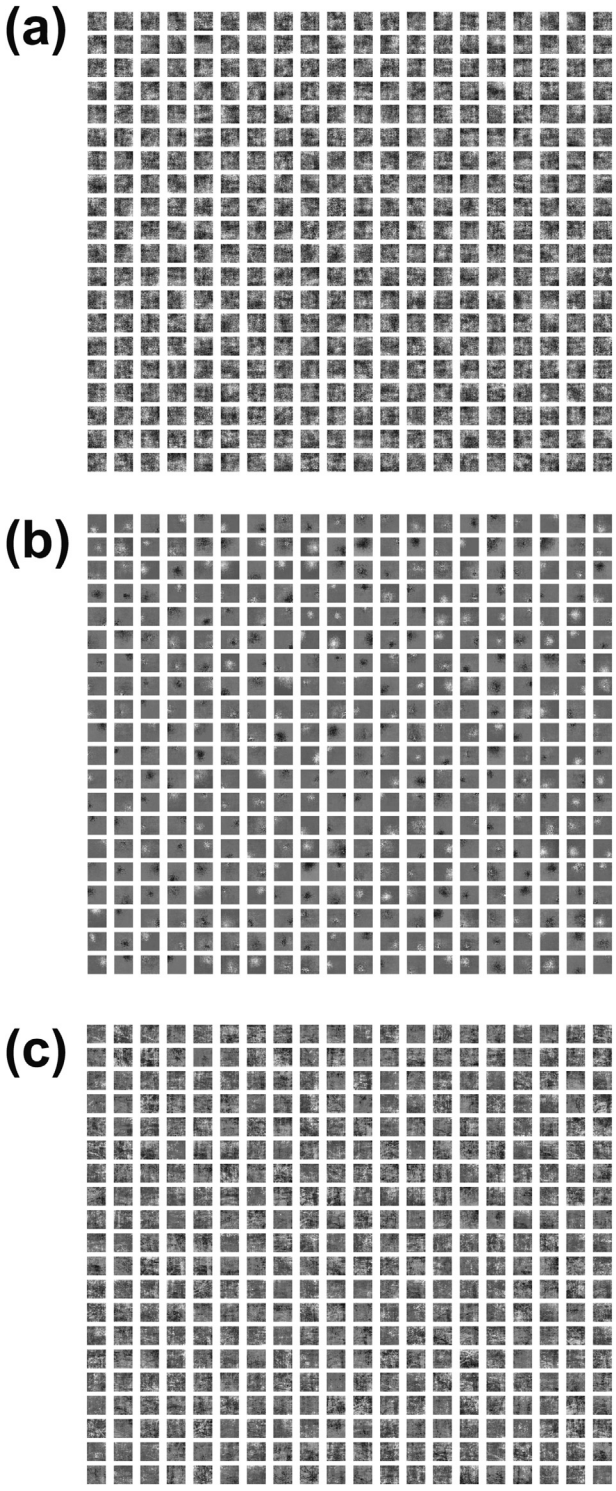

Fig. 6. Spatial domain representation of each row of A yielding minimal mean validation error using (a) no penalization, (b) localized penalization, and (c) ℓ_1 penalization. Positive entries are scaled in white and negative entries are scaled in black.

matrix with entries corresponding to the weight given to each respective spatial location sampled.

Likewise, we similarly plot in Fig. 7 the frequency domain representation of the sampling, which corresponds to the two-dimensional discrete Fourier transform of the spatial domain representation for each constituent row. Initially, the uniformly random structure of the measurement matrix A weights all

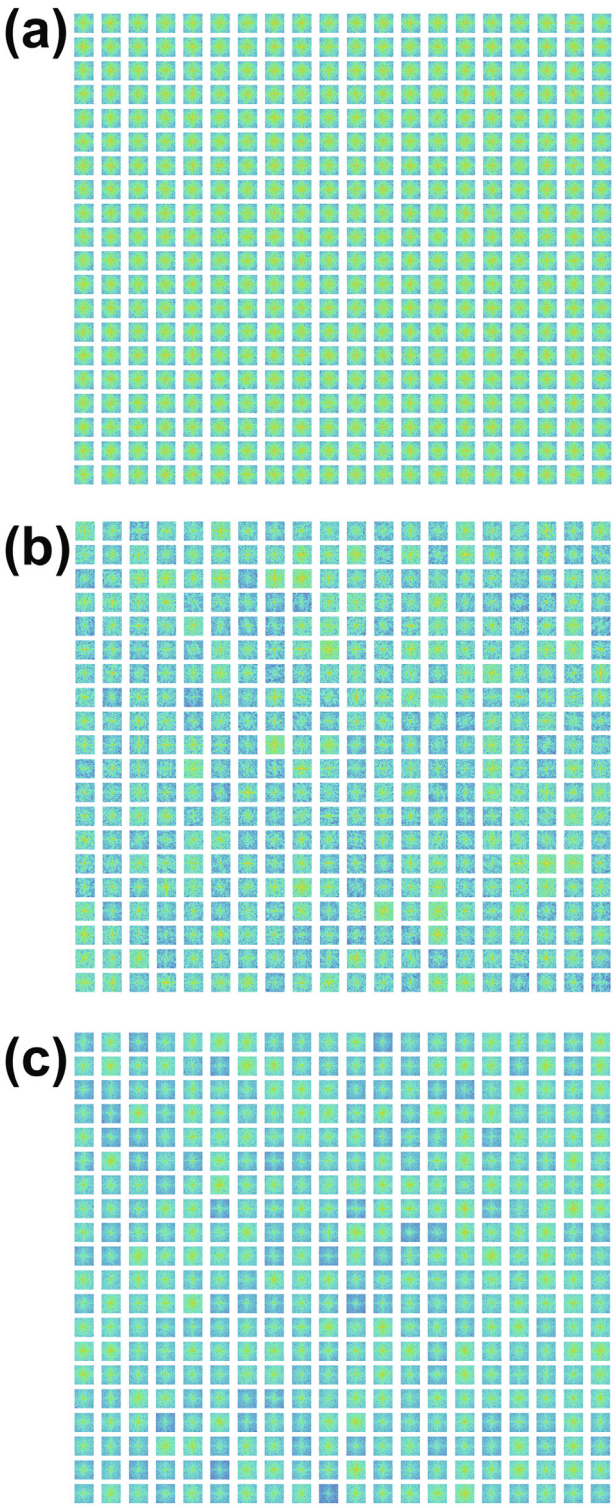

Fig. 7. Frequency domain representation of each row of A yielding minimal mean validation error using (a) no penalization, (b) localized penalization, and (c) ℓ_1 penalization. Entries distant from the center correspond to higher frequency components, and brighter (warmer) coloring indicates higher amplitude.

constituent frequency components in the sparse domain approximately equally for each row measurement. As a result of learning, the dominant low frequency components composing natural scenes are given more weighting.

Particularly when a large reduction factor is imposed, corresponding to a relatively small number of measurements, higher CS reconstruction quality may be obtained from sampling that captures components carrying the most signal information even if some lower amplitude component information is lost as a result. While additional spatial structure clearly manifests for optimal CS sampling, it is generally not smooth and demonstrates some degree of randomness so as to maintain sufficient incoherence with the sparse domain basis and thereby produce samples that yield distinct information.

We remark that while sampling structure more heavily weighting dominant components in the sparse domain is introduced from learning even with no penalization or ℓ_1 penalization, the localized sampling scheme visually demonstrates the clearest clustering in the spatial domain. In the sparse domain, this is indicated in the localized penalization sampling by a particularly rapid decrease in amplitude with increasing frequency.

We characterize the distribution of amplitudes in the sampled frequency components for a given measurement matrix row, corresponding to one set of weighted measurements, using the frequency entropy

$$H = -\sum_i p_i \log_2 p_i,$$

where p_i denotes the probability of observing the i^{th} frequency-component amplitude, binned over the amplitudes corresponding to all frequency components composing the sparse domain representation [65]. An average frequency entropy over the set of all measurement matrix rows, $\overline{H}$, can then be used to quantify the overall spread of frequency amplitudes corresponding to all constituent weighted measurements. In Fig. 8, we plot the measurement matrix average frequency entropy after each epoch for each penalization type. We observe that while the average frequency entropy ultimately saturates upon sufficient training regardless of the penalization scheme, only in the case of localized penalization does the average frequency entropy increase initially with learning. For the alternative penalizations, the average frequency entropy either decreases initially or remains approximately the same, and thus localized sampling ultimately achieves the greatest frequency entropy after training. Note that while uniformly random sampling, as in the case of the untrained measurement matrices, generally yields high $\overline{H}$ because of the unstructured diversity of sampled frequencies, it lacks concentration of amplitudes about the lower frequency components and therefore should be distinguished from the frequency entropies obtained after training. Hence, the abrupt transition in amplitudes from the heavily-weighted low frequency components to the lightly-weighted moderate frequency components that manifests through training with localized penalization produces sampling with both structure and relatively large frequency entropy.

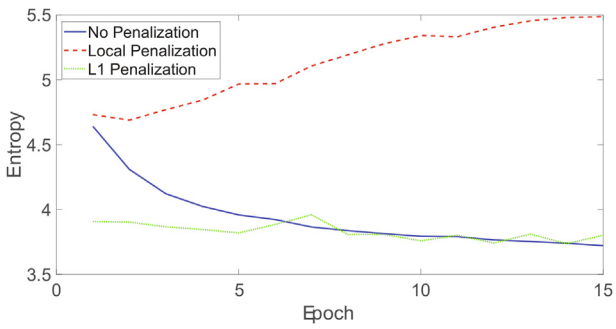

Fig. 8. Average frequency entropy of A after each epoch. The average frequency entropy using no penalization is plotted in solid blue, using localized penalization in dashed red, and using ℓ_1 penalization in dotted green.

The key differences in structural organization between the trained measurement matrices using the ℓ_1 and localized sampling penalizations support the notion that beyond energetic constraints on the number of synapses, as enforced by the ℓ_1 penalty, evolutionary pressure towards spatial locality was a potentially necessary ingredient in forming the largely distance-dependent and modular neuronal connectivity observed in vivo. The localized connections learned were often all of the same sign for a given row of A, with some demonstrating a sharp transition between excitatory and inhibitory weights in the spatial domain as observed in center-surround or simple cell receptive fields. We note that altering the penalization scheme and reduction factor may change the spatial structure of samples, therefore producing different types as well as distributions of receptive fields. However, the broad region of locally minimal CS reconstruction error upon sufficient training depicted in Fig. 2 across penalization schemes indicates that there are numerous means of achieving a relatively accurate compressive encoding of signals, and thus numerous specific system constraints are likely still amenable to improved reconstructions by introducing some degree of structure in sampling. This speaks to the robustness of CS across applications as well as a possible explanation for the variety of receptive field types found downstream in many sensory systems whose constituent layers may be optimized for the compressive encoding of specific signal features via particular sparse representations [66–68].

4. Discussion

Formulating a neural network framework for the compressive sensing sampling and reconstruction of sparse signals, we develop a new data-driven methodology for learning more effective CS measurement protocols. Our training method significantly improves signal reconstruction quality across several connection weight penalization schemes and classes of signals. The methodology demonstrates both logistical pragmatism and biological plausibility, utilizing asymmetric feedforward and feedback connectivity as well as only requiring neuronal knowledge of local connectivity information. Learning results in more structured sampling, driven by the potentially unknown structural characteristics of the class of signals analyzed in the sparse domain, while still maintaining a sufficient degree of incoherence as typically suggested by conventional CS theory. When utilizing spatially localized penalization of measurements in particular, we demonstrate that learning yields sparse and distance-dependent sampling weights, as commonly observed in the feedforward connectivity of physiological neuronal networks, suggesting that the brain may learn to encode information across neuronal layers of diverse sizes in agreement with CS principles.

We conjecture that the numerous convergent brain pathways in which the number of neurons encoding information is significantly reduced downstream may serve to facilitate efficient processing in low dimensional spaces while preserving information via compressive coding so that it is available across subsequent expansive pathways [46–48]. This is made possible by the ubiquity of sparse representations in high dimensional spaces, which is likely leveraged in the brain through evolutionary selection. Note that while solving the minimization problem given by Eq. (2) facilitates our CS reconstruction of signals in silico, whether and how the brain implements ℓ_1 minimization is still an area of investigation. One implementation argued to be potentially biologically feasible is a locally competitive algorithm with thresholding, which takes the form of a nonlinear system of ordinary differential equations mimicking more detailed computations across a neuronal network with inhibitory interactions [69].

Considering prior theoretical work demonstrates that the firing rate dynamics in model neuronal networks can provide a successful compressive encoding of sensory information [70,71], investigating if it is possible for such systems to learn improved CS sampling based on their nonlinear network dynamics, analogous to the method in this work, marks an interesting area for future analysis. Through the conduit of learning receptive field structure in a biologically plausible way that is motivated by compressive coding, it would be informative to study if differences in receptive fields across species, such as the salt-and-pepper organization of the rodent primary visual cortex and the pinwheel orientation columns of many higher mammals, may manifest from learned sampling based on distinctions in encountered stimuli and resultant training data [72,73].

Though here we primarily investigated learned CS sampling of natural scenes, with particularly well known structure [19], our data-driven methodology may be analogously applied to alternative classes of sparse signals in engineering applications or non-visual stimuli in the context of sensory system models. For example, we have applied our framework to both sound data as well as handwritten text data and obtained similar improvements in CS reconstructions, but other reduction factors or training data may yield alternative sampling characteristics furnishing new application-specific insights. In contrast to modern deep learning, which leverages many layers to help improve performance, we focus on a three-layer network for interpretability. This allows us to consider the natural case in which only a single measurement matrix is utilized and trained, with adaptability to any nonlinear or iterative optimization scheme that produces the final layer CS reconstruction output. In engineered systems, one measurement matrix is typically utilized for practical purposes and our work gives insight into amenable measurement matrix structures beyond those considered in classical CS theory. However, in the realm on neuroscience, the primary visual cortex is composed of many layers that are often optimized for processing different stimulus features [74,75], and hence investigating information encoding across multiple compressive-and-expansive pathways provides a natural avenue for further biologically-motivated studies.

Since using uniformly random samples is costly or infeasible due to the spatial scale between sampling locations in most engineered systems across the natural sciences [14–17], it is promising that the space of measurement matrices over which learned CS sampling improves reconstructions is demonstrated to be extensive. This grants much needed flexibility in effective sampling designs. Considering improved CS reconstructions are learned even under highly limited, distance-dependent measurement constraints, it is likely possible to use spatially localized random clusters of measurements, which are often more feasible to implement in practice, and still yield particularly accurate CS reconstructions.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by NSF grant DMS-1812478 and a Swarthmore Faculty Research Support Grant.

References

[1] C.E. Shannon, Communication in the presence of noise, Proc. IRE 37 (1) (1949) 10–21.

[2] E. Candes, J. Romberg, T. Tao, Stable signal recovery from incomplete and inaccurate measurements, Commun. Pur. Appl. Math. 59 (8) (2006) 1207–1223.
[3] D. Donoho, Compressed sensing, IEEE Trans. Inform. Theory 52 (2006) 1289–1306.
[4] M. Lustig, D. Donoho, J.M. Pauly, Sparse MRI: The application of compressed sensing for rapid MR imaging, Magn. Reson. Med. 58 (6) (2007) 1182–1195.
[5] M.A. Herman, T. Strohmer, High-resolution radar via compressed sensing, Trans. Sig. Proc. 57 (6) (2009) 2275–2284.
[6] J.H. Ender, On compressive sensing applied to radar, Signal Process. 90 (5) (2010) 1402–1414.
[7] D. Hayden, Y.H. Chang, J. Goncalves, C.J. Tomlin, Sparse network identifiability via compressed sensing, Automatica 68 (2016) 9–17.
[8] J.B. Lee, A. Yonar, T. Hallacy, C.-H. Shen, J. Milloz, J. Srinivasan, A. Kocabas, S. Ramanathan, A compressed sensing framework for efficient dissection of neural circuits, Nat. Methods 16 (1) (2019) 126.
[9] Y. Mishchenko, L. Paninski, A Bayesian compressed-sensing approach for reconstructing neural connectivity from subsampled anatomical data, J. Comput. Neurosci. 33 (2) (2012) 371–388.
[10] W. Dai, M.A. Sheikh, O. Milenkovic, R.G. Baraniuk, Compressive sensing DNA microarrays, J. Bioinf. Syst. Biol. (2009) 162824.
[11] E.J. Candes, M.B. Wakin, An Introduction To Compressive Sampling, Signal Process, Mag., IEEE 25 (2) (2008) 21–30.
[12] R. Baraniuk, Compressive sensing, IEEE Signal Process. Mag. (2007) 118–120.
[13] A. Bruckstein, D. Donoho, M. Elad, From sparse solutions of systems of equations to sparse modeling of signals and images, SIAM Rev. 51 (1) (2009) 34–81.
[14] M. Rani, S. Dhok, R. Deshmukh, A systematic review of compressive sensing: Concepts, implementations and applications, IEEE Access 6 (2018) 4875–4894.
[15] G. Hennenfent, F.J. Herrmann, Simply denoise: Wavefield reconstruction via jittered undersampling, Geophysics 73 (3) (2008) V19–V28.
[16] N.T. Markov, M. Ercsey-Ravasz, D.C. Van Essen, K. Knoblauch, Z. Toroczkai, H. Kennedy, Cortical high-density counterstream architectures, Science 342 (6158) (2013) 1238406.
[17] I.H. Stevenson, J.M. Rebesco, L.E. Miller, K.P. Kording, Inferring functional connections between neurons, Curr. Opin. Neurobiol. 18 (6) (2008) 582–588.
[18] M.L. Malloy, R.D. Nowak, Near-optimal adaptive compressed sensing, IEEE Trans. Inf. Theory 60 (7) (2014) 4001–4012.
[19] E.P. Simoncelli, B.A. Olshausen, Natural image statistics and neural representation, Annu. Rev. Neurosci. 24 (1) (2001) 1193–1216.
[20] B. Adcock, A.C. Hansen, C. Poon, B. Roman, Breaking the coherence barrier: A new theory for compressed sensing, in: Forum of Mathematics, Sigma, vol. 5, Cambridge University Press, 2017. .
[21] M. Elad, Optimized projections for compressed sensing, IEEE Trans. Signal Process. 55 (12) (2007) 5695–5702.
[22] V.J. Barranca, G. Kovačič, D. Zhou, D. Cai, Improved compressive sensing of natural scenes using localized random sampling, Sci. Rep. 6 (2016) 31976.
[23] R. Obermeier, J.A. Martinez-Lorenzo, Sensing matrix design via mutual coherence minimization for electromagnetic compressive imaging applications, IEEE Trans. Comput. Imag. 3 (2) (2017) 217–229.
[24] A. Mousavi, A.B. Patel, R.G. Baraniuk, A deep learning approach to structured signal recovery, in: 2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton), IEEE, 2015, pp. 1336–1343. .
[25] W. Schultz, A. Dickinson, Neuronal coding of prediction errors, Annu. Rev. Neurosci. 23 (1) (2000) 473–500.
[26] G.B. Keller, T. Bonhoeffer, M. Hübener, Sensorimotor mismatch signals in primary visual cortex of the behaving mouse, Neuron 74 (5) (2012) 809–815.
[27] A.M. Bastos, W.M. Usrey, R.A. Adams, G.R. Mangun, P. Fries, K.J. Friston, Canonical microcircuits for predictive coding, Neuron 76 (4) (2012) 695–711.
[28] A. Cochocki, R. Unbehauen, Neural Networks for Optimization and Signal Processing, John Wiley & Sons Inc, 1993.
[29] H. Greenspan, B. Van Ginneken, R.M. Summers, Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique, IEEE Trans. Med. Imag. 35 (5) (2016) 1153–1159.
[30] T.-H. Chan, K. Jia, S. Gao, J. Lu, Z. Zeng, Y. Ma, Pcanet: A simple deep learning baseline for image classification?, IEEE Trans Image Process. 24 (12) (2015) 5017–5032.
[31] N. Wang, D.-Y. Yeung, Learning a deep compact image representation for visual tracking, in: Advances in Neural Information Processing Systems, 2013, pp. 809–817. .
[32] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, Nature 521 (7553) (2015) 436. .
[33] L. Deng, D. Yu, et al., Deep learning: methods and applications, Found. Trends Signal Process. 7 (3–4) (2014) 197–387.
[34] J. Schmidhuber, Deep learning in neural networks: An overview, Neural Networks 61 (2015) 85–117.
[35] D.E. Rumelhart, G.E. Hinton, R.J. Williams, et al., Learning representations by back-propagating errors, Cognit. Model. 5 (3) (1988) 1.
[36] F. Crick, The recent excitement about neural networks, Nature 337 (6203) (1989) 129–132.
[37] S. Grossberg, Competitive learning: From interactive activation to adaptive resonance, Cognit. Sci. 11 (1) (1987) 23–63.
[38] A.H. Marblestone, G. Wayne, K.P. Kording, Toward an integration of deep learning and neuroscience, Front. Comput. Neurosci. 10 (2016) 94.
[39] E. Ganmor, R. Segev, E. Schneidman, The architecture of functional interaction networks in the retina, J. Neurosci. 31 (8) (2011) 3044–3054.
[40] T.P. Lillicrap, D. Cownden, D.B. Tweed, C.J. Akerman, Random synaptic feedback weights support error backpropagation for deep learning, Nat. Commun. 7 (2016) 13276.
[41] A. Nøkland, Direct feedback alignment provides learning in deep neural networks, in: Advances in Neural Information Processing Systems, 2016, pp. 1037–1045. .
[42] D. Hubel, T. Wiesel, Receptive fields, binocular interaction and functional architecture of the cat's visual cortex, J. Physiol. (Lond.) 160 (1962) 106–154.
[43] D. Field, What is the goal of sensory coding?, Neural Comput 6 (4) (1994) 559–601, https://doi.org/10.1162/neco.1994.6.4.559.
[44] C. Poo, J.S. Isaacson, Odor representations in olfactory cortex: "sparse" coding, global inhibition, and oscillations, Neuron 62 (6) (2009) 850–861.
[45] T. Hromádka, M.R. DeWeese, A.M. Zador, Sparse representation of sounds in the unanesthetized auditory cortex, PLoS Biol. 6 (1) (2008) e16.
[46] H.B. Barlow, The ferrier lecture, 1980. critical limiting factors in the design of the eye and visual cortex., Proc. R. Soc. Lond. B Biol. Sci. 212 (1186) (1981) 1–34. .
[47] L. Buck, Information coding in the vertebrate olfactory system, Ann. Rev. Neurosci. 19 (1996) 517–544, https://doi.org/10.1146/annurev.ne.19.030196.002505.
[48] S. Dasgupta, C.F. Stevens, S. Navlakha, A neural algorithm for a fundamental computing problem, Science 358 (6364) (2017) 793–796.
[49] Y. He, Z.J. Chen, A.C. Evans, Small-world anatomical networks in the human brain revealed by cortical thickness from MRI, Cereb. Cortex 17 (10) (2007) 2407–2419.
[50] O. Sporns, The human connectome: a complex network, Ann. N.Y. Acad. Sci. 1224 (2011) 109–125.
[51] A. Barabasi, R. Albert, Emergence of scaling in random networks, Science 286 (5439) (1999) 509–512.
[52] T.N. Wiesel, Receptive fields of ganglion cells in the cat's retina, J. Physiol. 153 (1960) 583–594.
[53] D.H. Hubel, T.N. Wiesel, Receptive fields of optic nerve fibres in the spider monkey, J. Physiol. 154 (1960) 572–580.
[54] M.S. Graziano, C.G. Gross, A bimodal map of space: somatosensory receptive fields in the macaque putamen with corresponding visual receptive fields, Exp. Brain Res. 97 (1) (1993) 96–109.
[55] D.A. Wilson, Receptive fields in the rat piriform cortex, Chem. Senses 26 (5) (2001) 577–584.
[56] C. Welker, Receptive fields of barrels in the somatosensory neocortex of the rat, J. Comp. Neurol. 166 (2) (1976) 173–189.
[57] K. Mori, H. Nagao, Y. Yoshihara, The olfactory bulb: coding and processing of odor molecule information, Science 286 (5440) (1999) 711–715.
[58] E.I. Knudsen, M. Konishi, Center-surround organization of auditory receptive fields in the owl, Science 202 (4369) (1978) 778–780.
[59] J.A. Tropp, A.C. Gilbert, Signal Recovery From Random Measurements Via Orthogonal Matching Pursuit, IEEE Trans. Inform. Theory 53 (12) (2007) 4655–4666.
[60] D. Donoho, Y. Tsaig, Fast solution of l_1-norm minimization problems when the solution may be sparse, IEEE Trans. Inform. Theory 54 (11) (2008) 4789–4812.
[61] H. Monajemi, S. Jafarpour, M. Gavish, D.L. Donoho, Deterministic matrices matching the compressed sensing phase transitions of gaussian random matrices, Proc. Nat. Acad. Sci. 110 (4) (2013) 1181–1186, https://doi.org/10.1073/pnas.1219540110.
[62] A. Oliva, A. Torralba, Modeling the shape of the scene: A holistic representation of the spatial envelope, Int. J. Comput. Vision 42 (3) (2001) 145–175, https://doi.org/10.1023/A:1011139631724.
[63] R. Bogacz, M. Usher, J. Zhang, J.L. McClelland, Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice, Philos. Trans. R. Soc. Lond., B, Biol. Sci. 362 (1485) (2007) 1655–1670.
[64] R.J. Douglas, K.A. Martin, Recurrent neuronal circuits in the neocortex, Curr. Biol. 17 (13) (2007) 496–500.
[65] F. Rieke, D. Warland, R.R. van Steveninck, W. Bialek, Spikes: Exploring the Neural Code, MIT Press, Cambridge, 1996.
[66] D. Hubel, Eye, Brain, and Vision, Scientific American Library Series, Henry Holt and Company, New York, 1995.
[67] M.P. Sceniak, D.L. Ringach, M.J. Hawken, R. Shapley, Contrast's effect on spatial summation by macaque V1 neurons, Nat. Neurosci. 2 (8) (1999) 733–739.
[68] R. Desimone, T.D. Albright, C.G. Gross, C. Bruce, Stimulus-selective properties of inferior temporal neurons in the macaque, J. Neurosci. 4 (8) (1984) 2051–2062.
[69] C.J. Rozell, D.H. Johnson, R.G. Baraniuk, B.A. Olshausen, Sparse coding via thresholding and local competition in neural circuits, Neural Comput. 20 (10) (2008) 2526–2563.
[70] V.J. Barranca, G. Kovačič, D. Zhou, D. Cai, Sparsity and compressed coding in sensory systems, PLoS Comput. Biol. 10 (8) (2014) e1003793.

[71] V.J. Barranca, X.G. Zhu, A computational study of the role of spatial receptive field structure in processing natural and non-natural scenes, J. Theor. Biol. 454 (2018) 268–277.
[72] K. Ohki, R.C. Reid, Specificity and randomness in the visual cortex, Curr. Opin. Neurobiol. 17 (4) (2007) 401–407.
[73] S.D. Van Hooser, J.A.F. Heimel, S. Chung, S.B. Nelson, L.J. Toth, Orientation selectivity without orientation maps in visual cortex of a highly visual mammal, J. Neurosci. 25 (1) (2005) 19–28, https://doi.org/10.1523/JNEUROSCI.4042-04.2005.
[74] D.H. Hubel, T.N. Wiesel, Receptive fields of single neurones in the cat's striate cortex, J. Physiol. 148 (1959) 574–591.
[75] C.M. Niell, Cell types, circuits, and receptive fields in the mouse visual cortex, Annu. Rev. Neurosci. 38 (2015) 413–431.

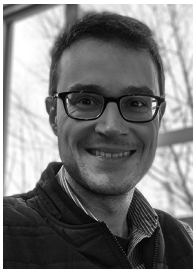

Victor J. Barranca is currently an associate professor of mathematics at Swarthmore College. He received his Ph. D. degree in mathematics from Rensselaer Polytechnic Institute in 2013. From 2013 to 2015, he was a post-doctoral research associate at New York University. He also conducted research at Shanghai Jiao Tong University in summer 2014. He is a recipient of the Joaquin B. Diaz Prize, the Ralph Ernest Huston Prize, National Science Foundation funding, and a James A. Michener Fellowship. His research interests include mathematical neuroscience, signal processing, and network dynamics.