
A simple normative network approximates local non-Hebbian learning in the cortex

Siavash Golkar¹ David Lipshutz¹ Yanis Bahroun¹
Anirvan M. Sengupta^{1,2} Dmitri B. Chklovskii^{1,3}

¹ Center for Computational Neuroscience, Flatiron Institute

² Department of Physics and Astronomy, Rutgers University

³ Neuroscience Institute, NYU Medical Center

{sgolkar, dlipshutz, ybahroun, mitya}@flatironinstitute.org
anirvans.physics@gmail.com

Abstract

To guide behavior, the brain extracts relevant features from high-dimensional data streamed by sensory organs. Neuroscience experiments demonstrate that the processing of sensory inputs by cortical neurons is modulated by instructive signals which provide context and task-relevant information. Here, adopting a normative approach, we model these instructive signals as supervisory inputs guiding the projection of the feedforward data. Mathematically, we start with a family of Reduced-Rank Regression (RRR) objective functions which include Reduced Rank (minimum) Mean Square Error (RRMSE) and Canonical Correlation Analysis (CCA), and derive novel offline and online optimization algorithms, which we call Bio-RRR. The online algorithms can be implemented by neural networks whose synaptic learning rules resemble calcium plateau potential dependent plasticity observed in the cortex. We detail how, in our model, the calcium plateau potential can be interpreted as a backpropagating error signal. We demonstrate that, despite relying exclusively on biologically plausible local learning rules, our algorithms perform competitively with existing implementations of RRMSE and CCA.

1 Introduction

In the brain, extraction of behaviorally-relevant features from high-dimensional data streamed by sensory organs occurs in multiple stages. Early stages of sensory processing, e.g., the retina, lack feedback and are naturally modeled by unsupervised learning algorithms [1]. In contrast, subsequent processing by cortical circuits is modulated by instructive signals from other cortical areas [2], which provide context and task-related information [3], thus calling for supervised learning models.

Unsupervised models of early sensory processing, despite employing many simplifying assumptions, have successfully bridged the salient features of biological neural networks, such as the architecture, synaptic learning rules and receptive field structure, with computational tasks such as dimensionality reduction, decorrelation, and whitening [4, 5, 6, 7, 8]. The success of such models was driven by two major factors. First, following a normative framework, their synaptic learning rules, network architecture and activity dynamics were derived by optimizing a principled objective, leading to an analytic understanding of the circuit computation without the need for numerical simulation [9]. Second, these models went beyond purely theoretical explorations by appealing to and explaining various experimental observations of early sensory organs available at the time [5, 8, 9].

In contrast to early sensory processing, subsequent processing in the cortex (both neocortex [10, 2, 11, 12, 13] and hippocampus [14, 15, 16, 17]) is guided by supervisory signals. In particular, in cortical pyramidal neurons, proximal dendrites receive and integrate feedforward inputs leading to the generation of action potentials (i.e., the output of the neuron). The distal dendrites of the apical tuft, in contrast, receive and integrate instructive signals resulting in local depolarization. When the local depolarization is large relative to inhibitory currents, this generates a calcium plateau potential that propagates throughout the entire neuron. If the calcium plateau coincides with feedforward input, it strengthens corresponding proximal synapses, thereby providing an instructive signal in these circuits [12, 11, 15, 16].

In this work, we model cortical processing as a projection of feedforward sensory input that is modulated by instructive signals from other cortical areas. Inspired by the success of the normative approach in early sensory processing, we adopt it here. Mathematically, the projections of sensory input can be learned by minimizing the prediction error or maximizing the correlation of the projected input with the instructive signal. These correspond to two instances of the Reduced-Rank Regression (RRR) objectives: Reduced-Rank (minimum) Mean Square Error (RRMSE) [18] and Canonical Correlation Analysis (CCA) [19].

To serve as a viable model of brain function, an algorithm must satisfy at least the following two criteria [9]. First, because sensory inputs are streamed to the brain and require real-time processing, it must be modeled by an online learning algorithm that does not store any significant fraction of the data. To satisfy this requirement, unlike standard offline formulations, which output projection matrices, at each time step, the algorithm must compute the projection from the input of that time step. The projection matrices are updated at each time step and can be represented in synaptic weights. Second, a neural network implementation of such an algorithm must rely exclusively on local synaptic learning rules. Here, locality means that the plasticity rules depend exclusively on the variables available to the biological synapse, i.e., the physicochemical activities of the pre- and post-synaptic neurons in the synaptic neighborhood. The Hebbian update rule is an example of local learning, where the change of synaptic weight is proportional to the correlation between the output activities of the pre- and post-synaptic neurons [20].

Contributions

- We derive novel algorithms for a family of RRR problems, which include RRMSE and CCA, and implement them in biologically plausible neural networks that resemble cortical micro-circuits.
- We demonstrate within the confines of our model how the calcium plateau potential in cortical microcircuits encodes a backpropagating error signal.
- We show numerically on a real-world dataset that our algorithms perform competitively compared with current state-of-the-art algorithms.

2 Related works

Our contributions are related to several lines of computational and theoretical research. One of the earliest normative models of cortical computation is based on the predictive coding framework where the feedback attempts to predict the feedforward input. When trained on natural images, this approach can explain extra-classical response properties observed in the visual cortex [21, 22]. The predictive coding framework has recently been used for the supervised training of deep networks with Hebbian learning rules [23]. However, these models have not been mapped onto the anatomy and physiology, especially the non-Hebbian synaptic plasticity, of cortical microcircuits [15, 16].

A prescient paper [24] proposed that supervised learning in the cortex can be implemented by multi-compartmental pyramidal neurons with non-Hebbian learning rules driven by calcium plateau potentials. Building on this proposal, [25, 26, 27] demonstrated possible biological implementations of backpropagation in deep networks. Neuroscience experiments have motivated the development of several biologically realistic models of microcircuits with multi-compartmental neurons and non-Hebbian learning rules [28, 29, 30]. Specifically, [29, 30] showed that calcium plateau potentials, generated in the apical tuft, can modulate the efficacy of proximal synapses. These demonstrations, however, are limited in that they were shown analytically in a small region of parameter space or they rely entirely on numerical simulations.

In the context of statistical learning, multiple RRMSE [31, 32, 33, 34, 35] and CCA [36, 37, 38, 39, 40, 41] algorithms have been developed. Of these algorithms, none satisfy the minimal criteria for biological plausibility. Biologically plausible formulations of CCA, as an unsupervised data integration algorithm following the normative approach, were proposed using deflation [42] and fully online in [43].

3 An objective function for reduced-rank regression problems

In this section, we review the Reduced-Rank Regression (RRR) problem which encompasses Canonical Correlation Analysis (CCA) and Reduced Rank (minimum) Mean Square Error (RRMSE) as special cases.

Notation. For positive integers m, n , let \mathbb{R}^m denote m -dimensional Euclidean space, and let $\mathbb{R}^{m \times n}$ denote the set of $m \times n$ real-valued matrices. We use boldface lower-case letters (e.g., \mathbf{v}) to denote vectors and boldface upper-case letters (e.g., \mathbf{M}) to denote matrices. Let \mathbf{I}_m denote the $m \times m$ identity matrix.

Let $\{(\mathbf{x}_t, \mathbf{y}_t)\}_{t=1}^T$ be a sequence of pairs of data points with $\mathbf{x}_t \in \mathbb{R}^m$, $\mathbf{y}_t \in \mathbb{R}^n$. We refer to \mathbf{x}_t as the predictor variable and \mathbf{y}_t as the response variable. Define the data matrices $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_T] \in \mathbb{R}^{m \times T}$ and $\mathbf{Y} := [\mathbf{y}_1, \dots, \mathbf{y}_T] \in \mathbb{R}^{n \times T}$. Let $\mathbf{C}_{xx} := \frac{1}{T} \mathbf{X} \mathbf{X}^\top$, $\mathbf{C}_{yy} := \frac{1}{T} \mathbf{Y} \mathbf{Y}^\top$, and $\mathbf{C}_{xy} := \frac{1}{T} \mathbf{X} \mathbf{Y}^\top$ be the empirical covariance matrices. Throughout this paper, we assume that \mathbf{X} and \mathbf{Y} are centered and full rank.

3.1 Problem formulation

The goal of RRR is to find a low-rank projection matrix $\mathbf{P} \in \mathbb{R}^{n \times m}$ that minimizes the error between $\mathbf{P} \mathbf{X}$ and \mathbf{Y} . The low-rank constraint favors the extraction of features that are most predictive of the response variables, thus preventing over-fitting [18]. We can formalize this as follows:

$$\arg \min_{\mathbf{P} \in \mathbb{R}^{n \times m}} \frac{1}{T} \|\mathbf{Y} - \mathbf{P} \mathbf{X}\|_{\Sigma}^2 \quad \text{subject to} \quad \text{rank}(\mathbf{P}) \leq k, \quad (1)$$

where $k \leq \min(m, n)$ determines the rank of the problem, $\Sigma \in \mathbb{R}^{n \times n}$ is a positive definite matrix, and $\|\cdot\|_{\Sigma}$ is the Σ -norm defined by $\|\mathbf{A}\|_{\Sigma}^2 := \text{Tr} \mathbf{A}^\top \Sigma \mathbf{A}$ for $\mathbf{A} \in \mathbb{R}^{n \times T}$. Intuitively, the Σ -norm is a generalized norm that can take into account the noise statistics of the samples [44]. Two common choices for Σ are $\Sigma = \mathbf{I}_n$ and $\Sigma = \mathbf{C}_{yy}^{-1}$. When $\Sigma = \mathbf{I}_n$, the RRR problem reduces to minimizing the mean square error (MSE) with a low-rank constraint. We refer to this objective as Reduced Rank (minimum) Mean Square Error (RRMSE) [18].¹ For $\Sigma = \mathbf{C}_{yy}^{-1}$, the objective in Eq. (1) is equivalent to Canonical Correlation Analysis (CCA) (see Sec. A of the supplementary materials).

3.2 Parametrizing the projection matrix

The low-rank constraint, $\text{rank}(\mathbf{P}) \leq k$, in Eq. (1) can be enforced by expressing $\mathbf{P} = \Sigma^{-1} \mathbf{V}_y \mathbf{V}_x^\top$, where $\mathbf{V}_x \in \mathbb{R}^{m \times k}$ and $\mathbf{V}_y \in \mathbb{R}^{n \times k}$ (the inclusion of Σ^{-1} here is for convenience in the derivation below). The matrix \mathbf{V}_x^\top projects the inputs \mathbf{x}_t onto a k -dimensional subspace and the column vectors of $\Sigma^{-1} \mathbf{V}_y$ span the range of the projection matrix \mathbf{P} . Plugging into Eq. (1), we have

$$\min_{\mathbf{V}_x \in \mathbb{R}^{m \times k}} \min_{\mathbf{V}_y \in \mathbb{R}^{n \times k}} \frac{1}{T} \|\mathbf{Y} - \Sigma^{-1} \mathbf{V}_y \mathbf{V}_x^\top \mathbf{X}\|_{\Sigma}^2. \quad (2)$$

The minimum of this objective is not unique: given a solution $(\mathbf{V}_x, \mathbf{V}_y)$ and any invertible matrix $\mathbf{M} \in \mathbb{R}^{k \times k}$, $(\mathbf{V}_x \mathbf{M}^\top, \mathbf{V}_y \mathbf{M}^{-1})$ is also a solution. To constrain the solution set, we impose the whitening constraint $\mathbf{V}_x^\top \mathbf{C}_{xx} \mathbf{V}_x = \mathbf{I}_k$. Expanding the quadratic in (2), dropping terms that do not depend on \mathbf{V}_x or \mathbf{V}_y , and using the whitening constraint, we arrive at

$$\min_{\mathbf{V}_x \in \mathbb{R}^{m \times k}} \min_{\mathbf{V}_y \in \mathbb{R}^{n \times k}} \text{Tr}(\mathbf{V}_y^\top \Sigma^{-1} \mathbf{V}_y - 2 \mathbf{V}_x^\top \mathbf{C}_{xy} \mathbf{V}_y) \quad \text{subject to} \quad \mathbf{V}_x^\top \mathbf{C}_{xx} \mathbf{V}_x = \mathbf{I}_k. \quad (3)$$

¹Also referred to as reduced rank Wiener filter or simply reduced rank regression.

The output of our algorithms will be the low-rank projection of \mathbf{X} , which we call $\mathbf{Z} := \mathbf{V}_x^\top \mathbf{X}$. Intuitively, for RRMSE ($\Sigma = \mathbf{I}_n$), optimization of this objective would find \mathbf{Z} which is most informative, in terms of MSE loss, of the response variable \mathbf{Y} . For CCA ($\Sigma = \mathbf{C}_{yy}^{-1}$), optimization of this objective finds the projection \mathbf{Z} which has the highest correlation with the response variable \mathbf{Y} .

We parametrize the normalizing matrix by its inverse as $\Sigma^{-1} = \Sigma_s^{-1} := s \mathbf{C}_{yy} + (1-s) \mathbf{I}_n$ with $0 \leq s \leq 1$. RRR with this normalizing matrix corresponds to a family of objectives which interpolate between RRMSE at $s = 0$ and CCA at $s = 1$.

4 Algorithm derivation

In this section, starting from Eq. (3), we derive offline and online algorithms for the family of RRR objectives parametrized by s .

4.1 Offline algorithms

Noting that imposing the constraint $\mathbf{V}_x^\top \mathbf{C}_{xx} \mathbf{V}_x = \mathbf{I}_k$ via a Lagrange multiplier leads to non-local update rules (see Sec. B of the supplementary materials), following [45] we impose the weaker inequality constraint $\mathbf{V}_x^\top \mathbf{C}_{xx} \mathbf{V}_x \preceq \mathbf{I}_k$ by introducing the matrix $\mathbf{Q} \in \mathbb{R}^{k \times k}$

$$\min_{\mathbf{V}_x \in \mathbb{R}^{m \times k}} \min_{\mathbf{V}_y \in \mathbb{R}^{n \times k}} \max_{\mathbf{Q} \in \mathbb{R}^{k \times k}} \text{Tr} \mathbf{V}_y^\top \Sigma_s^{-1} \mathbf{V}_y - 2 \mathbf{V}_x^\top \mathbf{C}_{xy} \mathbf{V}_y + \mathbf{Q} \mathbf{Q}^\top (\mathbf{V}_x^\top \mathbf{C}_{xx} \mathbf{V}_x - \mathbf{I}_k), \quad (4)$$

where $\mathbf{Q} \mathbf{Q}^\top$ is the positive semi-definite Lagrange multiplier enforcing the inequality. As in [45], the dynamics of the optimization enforce that the inequality constraint is saturated, i.e., $\mathbf{V}_x^\top \mathbf{C}_{xx} \mathbf{V}_x = \mathbf{I}_k$ is satisfied at the optimum of the objective (for a different proof see Sec. C). In the offline setting, objective (4) can be optimized using gradient descent-ascent dynamics derived by taking partial derivatives:

$$\mathbf{V}_x^\top \leftarrow \mathbf{V}_x^\top + \eta (\mathbf{V}_y^\top \mathbf{C}_{yx} - \mathbf{Q} \mathbf{Q}^\top \mathbf{V}_x^\top \mathbf{C}_{xx}) \quad (5)$$

$$\mathbf{V}_y^\top \leftarrow \mathbf{V}_y^\top + \eta (\mathbf{V}_x^\top \mathbf{C}_{xy} - \mathbf{V}_y^\top \Sigma_s^{-1}) \quad (6)$$

$$\mathbf{Q} \leftarrow \mathbf{Q} + \frac{\eta}{\tau} (\mathbf{V}_x^\top \mathbf{C}_{xx} \mathbf{V}_x - \mathbf{I}_k) \mathbf{Q}, \quad (7)$$

where $\eta > 0$ is the learning rate for \mathbf{V}_x and \mathbf{V}_y , and $\tau > 0$ is a parameter controlling the ratio of the descent and ascent steps.

4.2 Online algorithms

In the online (or streaming) setting, the input is presented one sample at a time, and the algorithm must find the projection without storing any significant fraction of the dataset.

To derive an online algorithm, we rewrite the objective function (4) making the dependence of the objective on each individual sample manifest:

$$\min_{\mathbf{V}_x} \min_{\mathbf{V}_y} \max_{\mathbf{Q}} \frac{1}{T} \sum_{t=1}^T \mathbf{V}_y^\top (s \mathbf{y}_t \mathbf{y}_t^\top + (1-s) \mathbf{I}_n) \mathbf{V}_y - 2 \mathbf{V}_x^\top \mathbf{x}_t \mathbf{y}_t^\top \mathbf{V}_y + \mathbf{Q} \mathbf{Q}^\top (\mathbf{V}_x^\top \mathbf{x}_t \mathbf{x}_t^\top \mathbf{V}_x - \mathbf{I}_k). \quad (8)$$

If we now perform stochastic gradient descent/ascent [46], i.e., perform the gradient updates with respect to individual samples, we arrive at our online algorithm. Explicitly, at time t , we have:

$$\mathbf{V}_x^\top \leftarrow \mathbf{V}_x^\top + \eta (\mathbf{a}_t - \mathbf{Q} \mathbf{n}_t) \mathbf{x}_t^\top \quad (9)$$

$$\mathbf{V}_y^\top \leftarrow \mathbf{V}_y^\top + \eta (\mathbf{z}_t \mathbf{y}_t^\top - s \mathbf{a}_t \mathbf{y}_t^\top - (1-s) \mathbf{V}_y^\top) \quad (10)$$

$$\mathbf{Q} \leftarrow \mathbf{Q} + \frac{\eta}{\tau} (\mathbf{z}_t \mathbf{n}_t^\top - \mathbf{Q}). \quad (11)$$

where $\mathbf{z}_t := \mathbf{V}_x^\top \mathbf{x}_t$ is the output of the algorithm, $\mathbf{a}_t := \mathbf{V}_y^\top \mathbf{y}_t$ and $\mathbf{n}_t := \mathbf{Q}^\top \mathbf{z}_t$.

Our algorithms, which we call Bio-RRR, are summarized in Alg. 1 (offline) and Alg. 2 (online).

Algorithm 1: Offline Bio-RRR

input: $\mathbf{X} \in \mathbb{R}^{m \times T}$, $\mathbf{Y} \in \mathbb{R}^{n \times T}$
initialize \mathbf{V}_x , \mathbf{V}_y , and \mathbf{Q} .
 $\mathbf{C}_{xx} \leftarrow \mathbf{X}\mathbf{X}^\top / T$; $\mathbf{C}_{xy} \leftarrow \mathbf{X}\mathbf{Y}^\top / T$
 $\Sigma_s^{-1} \leftarrow s \mathbf{Y}\mathbf{Y}^\top / T + (1-s) \mathbf{I}_n$
repeat:
 $\mathbf{V}_x^\top \leftarrow \mathbf{V}_x^\top + \eta(\mathbf{V}_y^\top \mathbf{C}_{yx} - \mathbf{Q}\mathbf{Q}^\top \mathbf{V}_x^\top \mathbf{C}_{xx})$
 $\mathbf{V}_y^\top \leftarrow \mathbf{V}_y^\top + \eta(\mathbf{V}_x^\top \mathbf{C}_{xy} - \mathbf{V}_y^\top \Sigma_s^{-1})$
 $\mathbf{Q} \leftarrow \mathbf{Q} + \frac{\eta}{\tau}(\mathbf{V}_x^\top \mathbf{C}_{xx} \mathbf{V}_x - \mathbf{I}_k) \mathbf{Q}$
until convergence
output: $\mathbf{Z} = \mathbf{V}_x^\top \mathbf{X}$ \triangleright projected predictor

Algorithm 2: Online Bio-RRR

input: $\mathbf{x}_t \in \mathbb{R}^m$, $\mathbf{y}_t \in \mathbb{R}^n$ \triangleright new sample
 $\mathbf{V}_x, \mathbf{V}_y, \mathbf{Q}$ \triangleright previous matrices
 $\mathbf{z}_t \leftarrow \mathbf{V}_x^\top \mathbf{x}_t$; $\mathbf{n}_t \leftarrow \mathbf{Q}^\top \mathbf{z}_t$; $\mathbf{a}_t \leftarrow \mathbf{V}_y^\top \mathbf{y}_t$
 $\mathbf{V}_x^\top \leftarrow \mathbf{V}_x^\top + \eta(\mathbf{a}_t - \mathbf{Q}\mathbf{n}_t) \mathbf{x}_t^\top$
 $\mathbf{V}_y^\top \leftarrow \mathbf{V}_y^\top + \eta(\mathbf{z}_t \mathbf{y}_t^\top - s \mathbf{a}_t \mathbf{y}_t^\top - (1-s) \mathbf{V}_y^\top)$
 $\mathbf{Q} \leftarrow \mathbf{Q} + \frac{\eta}{\tau}(\mathbf{z}_t \mathbf{n}_t^\top - \mathbf{Q})$
output: \mathbf{z}_t \triangleright projected sample
 $\mathbf{V}_x, \mathbf{V}_y, \mathbf{Q}$ \triangleright updated matrices

5 Biological implementation and comparison with experiment

In this section, we introduce a biological neural circuit that implements the online RRR algorithm and demonstrate that the details of this circuit resemble neurophysiological properties of pyramidal cells in the neocortex and the hippocampus.

5.1 Neural circuit

The algorithm for online RRR summarized by the update rules in Eqs. (9)–(11) can be implemented in a neural circuit with schematic shown in Fig. 1. In this circuit, the individual components of the output of Bio-RRR, z_1, \dots, z_k , are represented as the outputs of k neurons. The matrices \mathbf{V}_x and \mathbf{V}_y are encoded as the weights of synapses between the output neurons and the inputs of the network (blue and pink nodes in Fig. 1). Explicitly the element V_x^{ij} (resp. V_y^{ij}) is the efficacy of the synapse connecting x_i (resp. y_i) to the j^{th} output neuron z_j . Because of the disjoint nature of the two inputs, we model these as synapsing respectively onto the distal (apical tuft) and proximal (mostly basal) dendrites of the output neurons, Fig. 1. The quantities $\mathbf{z}_t = \mathbf{V}_x^\top \mathbf{x}_t$ and $\mathbf{a}_t = \mathbf{V}_y^\top \mathbf{y}_t$ are then the integrated dendritic currents in each dendritic compartment.

Similarly, the auxiliary variable \mathbf{n} is represented by the activity of k interneurons with \mathbf{Q} encoded in the weights of synapses connecting \mathbf{n} to \mathbf{z} (purple nodes on the upper dendritic branch of \mathbf{z}) and \mathbf{Q}^\top encoded in the weights of synapses from \mathbf{z} to \mathbf{n} (gray nodes). In a biological setting, the implied equality of weights of synapses from \mathbf{z} to \mathbf{n} and the transpose of those from \mathbf{n} to \mathbf{z} can be guaranteed approximately by application of the same Hebbian learning rule (see supplementary materials Sec. D).

The proximal synaptic weights, given by the elements of \mathbf{V}_x , are updated by the product of two factors represented in the corresponding post- and pre-synaptic neurons (Eq. 9).

$$\delta \mathbf{V}_x^\top \propto (\mathbf{a}_t - \mathbf{Q}\mathbf{n}_t) \mathbf{x}_t^\top$$

The first factor $(\mathbf{a}_t - \mathbf{Q}\mathbf{n}_t)$, is the difference between the excitatory synaptic current in the apical tuft ($\mathbf{a}_t = \mathbf{V}_y^\top \mathbf{y}_t$) and the inhibitory current induced by interneurons synapsing onto the distal compartment ($\mathbf{Q}\mathbf{n}_t$). Biologically, this factor can be approximated by the calcium plateau potential traveling down the apical shaft. The second factor is the input \mathbf{x}_t to the proximal dendrites. Therefore, the synaptic weight update is proportional to quantities that are available to the synapse locally.

The synaptic learning rule for \mathbf{V}_y (Eq. 10) also involves the products of pre- and post-synaptic variables but weighted by the parameter s ,

$$\delta \mathbf{V}_y^\top \propto \left[\mathbf{z}_t \mathbf{y}_t^\top - (1-s) \mathbf{V}_y^\top - s \mathbf{a}_t \mathbf{y}_t^\top \right]$$

In the case of RRMSE ($s = 0$), the update is Hebbian ($\mathbf{z}_t \mathbf{y}_t^\top$) with a homeostasis decay term ($-\mathbf{V}_y^\top$). In the case of CCA ($s = 1$), the synaptic weight update is proportional to $(\mathbf{z}_t - \mathbf{a}_t) \mathbf{y}_t^\top$, where the difference between the (dendritically backpropagated) output activity of the pyramidal neuron (\mathbf{z}_t) and the total synaptic input to the distal compartment (\mathbf{a}_t) can be computed in the corresponding

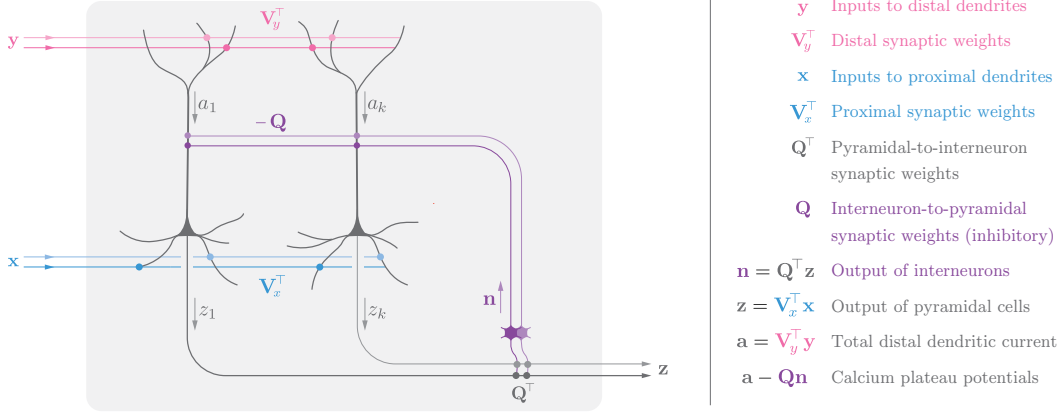


Figure 1: Cortical microcircuit for Bio-RRR. Pyramidal neurons (black) receive inputs x onto the dendrites proximal to the cell bodies (black triangles) weighted by V_x^T , and inputs y onto the distal dendrites weighted by V_y^T . The calcium plateau potential is the difference between the total distal dendritic current for each pyramidal neuron, $\mathbf{a} = (a_1, \dots, a_k)$, and the corresponding component of the inhibitory input, $-Q\mathbf{n}$. Output activity of pyramidal neurons, $\mathbf{z} = (z_1, \dots, z_k)$, is fed back via inhibitory interneurons (purple). The equivalence of the pyramidal-to-interneuron weight matrix, Q^T , and the transpose of the interneuron-to-pyramidal weight matrix, Q , follows from the operation of the local learning rules, see Sec. D of supplementary materials.

post-synaptic neuron (cf. [28]). In the intermediate cases of $0 < s < 1$, the update rule for V_y linearly interpolates between these two cases and remains local.

Finally, this circuit has the advantage of being purely feedforward in the sense that the output computation does not require equilibration of recurrent activity in lateral connections as was the case in e.g. [9]. This is due to the segregation between the proximal compartment that computes the output of the neuron and the distal compartment which receives the inhibitory lateral feedback.

5.2 Comparison with neuroscience experiments

The Bio-RRR circuit derived above has many features in common with cortical microcircuits but also deviates from them in a number of ways. Microcircuits in the cortex contain two classes of neurons: excitatory pyramidal neurons and inhibitory interneurons.² The pyramidal neurons can be considered the output neurons as their axon projections leave the local circuit. Similar to the output neurons of our circuit in Fig. 1, pyramidal neurons have two integration sites, the proximal compartment comprised of the basal and proximal apical dendrites providing inputs to the soma, and the distal compartment comprised of the apical dendritic tuft [49, 3]. These two compartments receive excitatory inputs from two separate sources [50, 2].

The inputs onto the two compartments are processed differently [2, 51, 52, 49, 3]. The proximal inputs directly drive the pyramidal neuron output by generating action potentials. If the distal inputs are stronger than the inhibitory post-synaptic currents driven by the interneurons, they generate a calcium plateau potential, which can also cause action potentials in the pyramidal neurons [2]. This is in contrast to our RRR algorithms, where only the proximal input contributes to the output, $\mathbf{z} = V_x^T \mathbf{x}$. Neglecting the contribution of the apical inputs to the action potential generation can be justified by the temporal sparsity of calcium plateau potentials. The situation where both proximal and distal inputs contribute significantly to the generation of action potentials can be modeled by an alternative biologically plausible implementation of CCA [43].

The calcium plateau potentials generated by the apical tuft inputs drive the plasticity of proximal synapses [12, 14, 15, 16]. Because this update is not purely dependent on the action potentials of the pre- and post-synaptic neurons, such plasticity is called non-Hebbian [16]. This resembles the synaptic updates of V_x in Eq. (9). However, while the teaching signal for the proximal synapses in Bio-RRR (i.e., $\mathbf{a}_t - Q\mathbf{n}_t$) is signed and graded, in the cortex, these signals are generally believed to be stereotypical [2]. Graded calcium mediated signals were recently observed in [29].

²There are multiple types of interneurons targeting pyramidal cells [47, 48]. The interneurons of Bio-RRR most closely resemble the somatostatin-expressing interneurons, which preferentially inhibit the apical dendrites.

Whereas the pyramidal neurons of the cortex fire all-or-nothing action potentials, Bio-RRR neurons are analog and linear as in firing-rate models. Furthermore, the goal of the RRR objectives is to reduce the dimensionality of the feedforward input, whereas sensory cortical processing is thought to expand dimensionality [53, 54, 55, 56]. These two disparities between our networks and realistic circuits are closely linked in that it is impossible to perform meaningful dimensionality expansion with linear neurons. However, due to the analytical tractability of simplified linear models, they provide insights that are difficult to obtain in more realistic models amenable only to numerical simulations.

The above comparisons of our algorithm with experiment apply equally to Bio-RRR with any $0 \leq s \leq 1$. The property which distinguishes different members of this family of algorithms is the update rule associated with the synapses of the distal compartment \mathbf{V}_y given in Eq. (10). There is conflicting experimental evidence regarding the plasticity of the distal apical dendrites in different areas of the brain. In the neocortex, the plasticity is thought to be Hebbian [11, 57], whereas in the hippocampus, experimental evidence points to non-Hebbian plasticity [12]. As discussed in the previous section, our online RRMSE and CCA algorithms require that distal synapses follow Hebbian and non-Hebbian plasticity rules, respectively. For a given cortical circuit, determining whether CCA or RRMSE or some intermediate value of s provides the best fit would require a close examination of the plasticity rules of the distal compartment.

6 Interpretation of calcium plateau potential in Bio-RRR

Experimentally, the calcium plateau potentials act as instructive signals in cortical pyramidal neurons by driving plasticity in the proximal dendrites [15, 16, 30]. Several prior works [24, 25, 26, 27] have suggested that the calcium plateau potential carries the backpropagation error. Here, we show that the calcium plateau potential plays a similar role in Bio-RRR provided the network is close to the optimum of the objective. In the process, we will also show how Bio-RRR avoids the weight transport problem of Artificial Neural Networks (ANNs) trained with the backpropagation algorithm (backprop).

We first describe how a two-layer ANN trained with backprop would implement RRR. We then compare the Bio-RRR learning rule for \mathbf{V}_x^\top , which approximates the calcium plateau potential, with that of the first layer weights of this ANN. For simplicity, we focus on the RRMSE case ($s = 0$), but the interpretation of the role of the calcium plateau potential in the \mathbf{V}_x^\top learning rule holds for any s .

The RRMSE objective given by

$$\min_{\mathbf{V}_x \in \mathbb{R}^{m \times k}} \min_{\mathbf{V}_y \in \mathbb{R}^{n \times k}} \frac{1}{T} \|\mathbf{Y} - \mathbf{V}_y \mathbf{V}_x^\top \mathbf{X}\|^2, \quad (12)$$

can be implemented as a two-layer linear ANN, where \mathbf{V}_x^\top and \mathbf{V}_y are the weights of the first and second layer of the network. We define $\hat{\mathbf{y}}_t = \mathbf{V}_y \mathbf{V}_x^\top \mathbf{x}_t$ as the network’s prediction of the label \mathbf{y}_t given input \mathbf{x}_t . When trained by backprop, the weight updates of this network are given by taking derivatives of the loss with respect to the weights [46]. Specifically, the learning rule for the weights of the first layer is given by:

$$\delta \mathbf{V}_x^\top \propto (\mathbf{V}_y^\top \boldsymbol{\epsilon}_t) \mathbf{x}_t^\top, \quad \boldsymbol{\epsilon}_t = (\mathbf{y}_t - \hat{\mathbf{y}}_t), \quad (13)$$

where we have defined $\boldsymbol{\epsilon}_t$ as the prediction error for the sample at time t . The update for \mathbf{V}_x^\top , the weights of the first layer of the ANN, requires the computation and backpropagation of the error signal $\boldsymbol{\epsilon}_t$. A cartoon of this process is given in Fig. 2a, where the forward and backward passes are respectively denoted in blue and red. Here, the weights \mathbf{V}_y are used both in the forward pass when computing the error $\boldsymbol{\epsilon}_t = \mathbf{y}_t - \mathbf{V}_y \mathbf{V}_x^\top \mathbf{x}_t$, and also their transpose in the backward pass when propagating the error back to the first layer (13). This symmetry between the forward and backward weights is a general property of SGD in deep networks but is not biologically realistic and is referred to as the “weight transport problem” [58, 59, 60]. Several solutions exist to facilitate the backpropagation of the computed error in a biologically plausible manner [61, 62, 63, 64].

Next, we show how Bio-RRR circumvents the weight transport problem. Comparing the above procedure for computing the \mathbf{V}_x weight updates to that of Bio-RRR given by:

$$\delta \mathbf{V}_x^\top \propto (\mathbf{a}_t - \mathbf{Q} \mathbf{n}_t) \mathbf{x}_t^\top \quad (14)$$

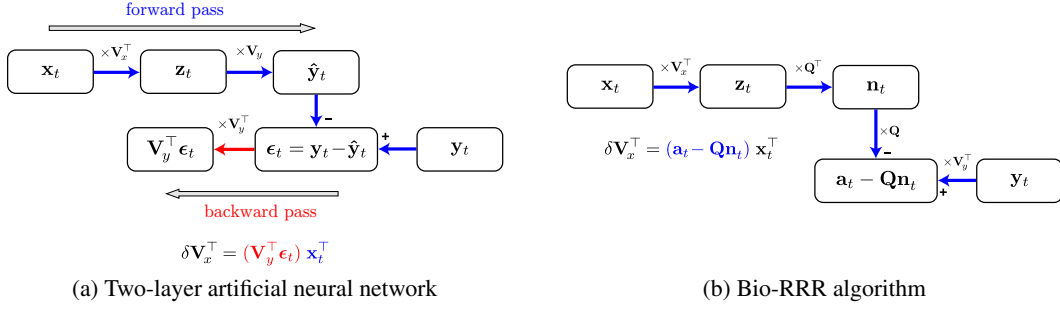


Figure 2: Schematic of a two-layer ANN implementation of RRR (left) and Bio-RRR (right), demonstrating the computation of the learning rule for \mathbf{V}_x^\top (for $\boldsymbol{\Sigma} = \mathbf{I}_n$). The blue and red arrows respectively denote the forward and backward passes. In Bio-RRR, $\mathbf{a}_t - \mathbf{Q}\mathbf{n}_t$ (encoded in the calcium plateau potential) replaces the backpropagated error $\mathbf{V}_y^\top \boldsymbol{\epsilon}_t$ in the \mathbf{V}_x^\top learning rule.

we see that the backpropagated error term in Eq. (13) is now replaced by the term $(\mathbf{a}_t - \mathbf{Q}\mathbf{n}_t)$ which emulates the calcium plateau potential. A diagram showing how this quantity is computed is given in Fig. 2b. We see that, unlike in the backprop computation depicted in Fig. 2a, in Bio-RRR no weights are reused and therefore weight transport problem is circumvented. This is because the Bio-RRR algorithm does not require the computation of the inferred value $\hat{\mathbf{y}}_t$ and the error signal $\boldsymbol{\epsilon}_t = \mathbf{y}_t - \hat{\mathbf{y}}_t$.

Although Bio-RRR does not explicitly compute prediction error, the update for \mathbf{V}_x^\top can still be interpreted in the context of error backpropagation. To this end, we look at the optimum of the objective where, from Eq. (5), we have

$$\mathbf{Q}\mathbf{Q}^\top \mathbf{V}_x^\top = \mathbf{V}_y^\top \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \Rightarrow \mathbf{Q}\mathbf{n}_t = \mathbf{Q}\mathbf{Q}^\top \mathbf{V}_x^\top \mathbf{x}_t = \mathbf{V}_y^\top \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \mathbf{x}_t = \mathbf{V}_y^\top \tilde{\mathbf{y}}_t,$$

where we have used $\mathbf{n}_t = \mathbf{Q}^\top \mathbf{z}_t$ and $\mathbf{z}_t = \mathbf{V}_x^\top \mathbf{x}_t$, and we have defined $\tilde{\mathbf{y}}_t := \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \mathbf{x}_t$. As $\mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} = \arg \min_{\mathbf{W}} \|\mathbf{Y} - \mathbf{W}\mathbf{X}\|_{\boldsymbol{\Sigma}}^2$ is the optimum of the rank-unconstrained regression objective, $\tilde{\mathbf{y}}_t$ is the best estimate of \mathbf{y}_t given the samples received thus far. Using these quantities and the definition of $\mathbf{a}_t = \mathbf{V}_y^\top \mathbf{y}_t$, we can rewrite the quantity $\mathbf{a}_t - \mathbf{Q}\mathbf{n}_t$ and the \mathbf{V}_x^\top update in Eq. (9) as

$$\mathbf{a}_t - \mathbf{Q}\mathbf{n}_t = \mathbf{V}_y^\top (\mathbf{y}_t - \tilde{\mathbf{y}}_t) \Rightarrow \mathbf{V}_x^\top \leftarrow \mathbf{V}_x^\top + \eta \left[\underbrace{\mathbf{V}_y^\top (\mathbf{y}_t - \tilde{\mathbf{y}}_t)}_{\text{prediction error}} \right] \mathbf{x}_t^\top. \quad (15)$$

Therefore, while the error term $\mathbf{y}_t - \tilde{\mathbf{y}}_t$ and backpropagation are not present explicitly in Bio-RRR, at the optimum, the calcium plateau potential is equal to a backpropagated error signal, and the update of \mathbf{V}_x^\top is proportional to the covariance of this backpropagated error signal and the input \mathbf{x}_t^\top .

7 Numerical experiments

In this section, we report the results of numerical simulations for our algorithms with $s = 0$ denoted as Bio-RRMSE and $s = 1$ denoted as Bio-CCA, and compare with current non-biologically plausible algorithms. For our experiments, we use the MediaMill dataset [65], a commonly used real-world benchmark consisting of $T = 2 \times 10^4$ samples of video data and text annotations. For our experiments, we take the predictor variables \mathbf{X} to be the 100-dimensional textual features and the response variable to be the 120-dimensional visual features extracted from representative video frames.

RRMSE. The performance of our RRMSE algorithm on MediaMill is given in Fig. 3a in terms of the objective function in Eq. (3) with $\boldsymbol{\Sigma} = \mathbf{I}_n$. For ranks $k \in \{1, 2, 4\}$, we plot this both as a function of iteration (top) and as a function of the CPU runtime (bottom). Here, the black dashed line denotes the value of the objective at its global minimum. For comparison, we provide the performance of the iterative quadratic minimum distance (IQMD) algorithm [32] and the 2-layer ANN discussed in Sec. 6. We see that IQMD is the most sample efficient, and ANN and Bio-RRMSE are within variance of each other and match the performance of IQMD in runtime. For plots of these algorithms in the offline (batch) setting, see Sec. E.

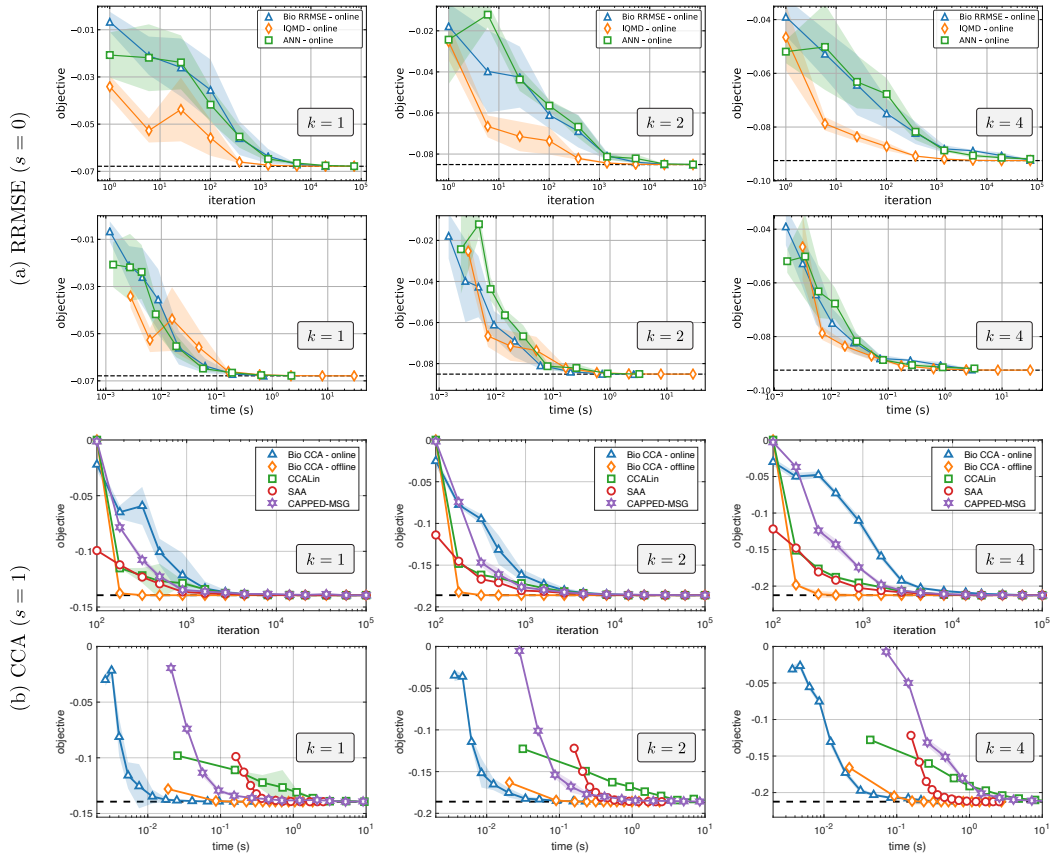


Figure 3: Comparisons of RRR algorithms for (a) RRMSE ($s = 0$) and (b) CCA ($s = 1$) in terms of the objective value Eq. (3) vs. iteration and runtime. Mean \pm standard deviation over 5 runs of the experiment.

CCA. We evaluate the performance of our CCA algorithm on the MediaMill dataset and compare with CCALin [66] and SAA [67], which are offline algorithms, and MSG-CCA [36] which is online. In Fig. 3b, we plot the performance of the different algorithms for CCA projection dimensions $k \in \{1, 2, 4\}$ in terms of the objective function given in Eq. (3). We see that on this dataset, our offline algorithm (Bio-CCA - offline) is the most sample efficient and our online algorithm (Bio-CCA - online) is fastest in terms of CPU runtime.

For further details, including the choice of hyperparameters and plots of convergence of the RRR constraint, see supplementary materials Sec. E. For experiments comparing the performance of RRMSE and backprop on a number of standard image classification datasets, see Sec. F.

8 Conclusion

Employing a normative approach, we derived new offline and online algorithms for a family of optimization objectives, which include CCA and RRMSE as special cases. We implemented these algorithms in biologically plausible neural networks and discussed how they resemble recent experimentally observed plasticity rules in the hippocampus and the neocortex. We elaborated on how this algorithm circumvents the weight transport problem of backprop and how the teaching signal is encoded in a quantity that resembles the calcium plateau potential. Determining which algorithm, CCA or RRMSE, more closely resembles cortical processing would require a careful examination of synaptic plasticity in the distal compartment of pyramidal neurons.

Broader impact

Understanding the inner workings of the brain has the potential of having a tremendous impact on society. On the one hand, this can lead to better performing machine learning algorithms and better artificial intelligent agents. On the other, understanding how the brain works can pave the way for

better treatments of psychological and neurological disorders. While this paper does not tackle these lofty broad societal goals directly, it is a small step in clarifying how information is processed in the brain.

Acknowledgments and Disclosure of Funding

We are grateful to Jeffrey Magee and Jason Moore for insightful discussions related to this work. We further thank Nicholas Chua, Shiva Farashahi, Johannes Friedrich, Alexander Genkin, Tiberiu Tesileanu, and Charlie Windolf for providing feedback on the manuscript.

The authors did not receive any third party funding for the completion of this project.

References

- [1] Eero P Simoncelli. Vision and the statistics of the visual environment. *Current Opinion in Neurobiology*, 13(2):144–149, 2003.
- [2] Matthew Larkum. A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends in Neurosciences*, 36(3):141–151, 2013.
- [3] Guy Major, Matthew E. Larkum, and Jackie Schiller. Active properties of neocortical pyramidal neuron dendrites. *Annual Review of Neuroscience*, 36(1):1–24, 2013. PMID: 23841837.
- [4] Erkki Oja. Simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, 15(3):267–273, 1982.
- [5] Mandyam V Srinivasan, Simon B Laughlin, and Andreas Dubs. Predictive coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 216(1205):427–459, 1982.
- [6] Adrian A Wanner and Rainer W Friedrich. Whitening of odor representations by the wiring diagram of the olfactory bulb. *Nature Neuroscience*, 23(3):433–442, 2020.
- [7] Yang Dan, Joseph J Atick, and R Clay Reid. Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *Journal of Neuroscience*, 16(10):3351–3362, 1996.
- [8] Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.
- [9] Cengiz Pehlevan and Dmitri B Chklovskii. Neuroscience-inspired online unsupervised learning algorithms: Artificial neural networks. *IEEE Signal Processing Magazine*, 36(6):88–96, 2019.
- [10] Nelson Spruston. Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neuroscience*, 9(3):206–221, 2008.
- [11] Per Jesper Sjöström and Michael Häusser. A cooperative switch determines the sign of synaptic plasticity in distal dendrites of neocortical pyramidal neurons. *Neuron*, 51(2):227–238, 2006.
- [12] Nace L Golding, Nathan P Staff, and Nelson Spruston. Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature*, 418(6895):326–331, 2002.
- [13] Frédéric Gambino, Stéphane Pagès, Vassilis Kehayas, Daniela Baptista, Roberta Tatti, Alan Carleton, and Anthony Holtmaat. Sensory-evoked LTP driven by dendritic plateau potentials in vivo. *Nature*, 515(7525):116–119, 2014.
- [14] Katie C Bittner, Christine Grienberger, Sachin P Vaidya, Aaron D Milstein, John J Macklin, Junghyup Suh, Susumu Tonegawa, and Jeffrey C Magee. Conjunctive input processing drives feature selectivity in hippocampal ca1 neurons. *Nature neuroscience*, 18(8):1133, 2015.
- [15] Katie C Bittner, Aaron D Milstein, Christine Grienberger, Sandro Romani, and Jeffrey C Magee. Behavioral time scale synaptic plasticity underlies ca1 place fields. *Science*, 357(6355):1033–1036, 2017.
- [16] Jeffrey C. Magee and Christine Grienberger. Synaptic plasticity forms and functions. *Annual Review of Neuroscience*, 43, 2020.
- [17] Jason Hardie and Nelson Spruston. Synaptic depolarization is more effective than back-propagating action potentials during induction of associative long-term potentiation in hippocampal pyramidal neurons. *Journal of Neuroscience*, 29(10):3233–3241, 2009.

- [18] Raja Velu and Gregory C Reinsel. *Multivariate Reduced-Rank Regression: Theory and Applications*, volume 136. Springer Science & Business Media, 2013.
- [19] Alan J Izenman. Reduced-rank regression for the multivariate linear model. *Journal of Multivariate Analysis*, 5(2):248–264, 1975.
- [20] Donald Olding Hebb. *The Organization of Behavior: A Neuropsychological Theory*. Psychology Press, 2005.
- [21] Rajesh P N Rao and Dana H Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87, 1999.
- [22] Rajesh PN Rao and Dana H Ballard. Probabilistic models of attention based on iconic representations and predictive coding. In *Neurobiology of Attention*, pages 553–561. Elsevier, 2005.
- [23] James C.R. Whittington and Rafal Bogacz. An approximation of the error backpropagation algorithm in a predictive coding network with local hebbian synaptic plasticity. *Neural Computation*, 29(5):1229–1262, may 2017.
- [24] Konrad P Körding and Peter König. Supervised and unsupervised learning with two sites of synaptic integration. *Journal of Computational Neuroscience*, 11(3):207–215, 2001.
- [25] Jordan Guerguiev, Timothy P Lillicrap, and Blake A Richards. Deep learning with segregated dendrites. *arXiv preprint arXiv:1610.00161*, 2016.
- [26] João Sacramento, Rui Ponte Costa, Yoshua Bengio, and Walter Senn. Dendritic cortical microcircuits approximate the backpropagation algorithm. *Advances in Neural Information Processing Systems*, 2018-Decem(Nips):8721–8732, oct 2018.
- [27] Alexandre Payeur, Jordan Guerguiev, Friedemann Zenke, Blake A Richards, and Richard Naud. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *bioRxiv*, 2020.
- [28] Robert Urbanczik and Walter Senn. Learning by the Dendritic Prediction of Somatic Spiking. *Neuron*, 81(3):521–528, 2014.
- [29] Albert Gidon, Timothy Adam Zolnik, Pawel Fidzinski, Felix Bolduan, Athanasia Papoutsis, Panayiota Poirazi, Martin Holtkamp, Imre Vida, and Matthew Evan Larkum. Dendritic action potentials and computation in human layer 2/3 cortical neurons. *Science*, 367(6473):83–87, jan 2020.
- [30] Aaron D. Milstein, Yiding Li, Katie C. Bittner, Christine Grienberger, Ivan Soltesz, Jeffrey C. Magee, and Sandro Romani. Bidirectional synaptic plasticity rapidly modifies hippocampal representations independent of correlated activity. *bioRxiv*, 2020.
- [31] Søren Johansen. Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models. *Econometrica*, 59(6):1551–1580, 1991.
- [32] Yingbo Hua and Maziar Nikpour. Computing the reduced rank Wiener filter by IQMD. *IEEE Signal Processing Letters*, 6(9):240–242, 1999.
- [33] Yingbo Hua, Maziar Nikpour, and Petre Stoica. Optimal reduced-rank estimation and filtering. *IEEE Transactions on Signal Processing*, 49(3):457–469, 2001.
- [34] Ziping Zhao and Daniel P Palomar. Robust maximum likelihood estimation of sparse vector error correction model. In *2017 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 913–917. IEEE, 2017.
- [35] Yangzhuoran Yang and Ziping Zhao. RRRR: Online robust reduced-rank regression estimation. R package version 1.0.0. <https://pkg.yangzhuoranyang.com/rrrr/>. 2020.
- [36] Raman Arora, Teodor Vanislavov Marinov, Poorya Mianjy, and Nati Srebro. Stochastic approximation for canonical correlation analysis. In *Advances in Neural Information Processing Systems*, pages 4775–4784, 2017.
- [37] Pei Ling Lai and Colin Fyfe. A neural implementation of canonical correlation analysis. *Neural Networks*, 12(10):1391–1397, 1999.
- [38] Ali Pezeshki, Mahmood R Azimi-Sadjadi, and Louis L Scharf. A network for recursive extraction of canonical coordinates. *Neural Networks*, 16(5-6):801–808, 2003.

- [39] Javier Vía, Ignacio Santamaría, and Jesús Pérez. A learning algorithm for adaptive canonical correlation analysis of several data sets. *Neural Networks*, 20(1):139–152, 2007.
- [40] Tatsuya Haga and Tomoki Fukai. Dendritic processing of spontaneous neuronal sequences for one-shot learning. *bioRxiv*, page 165613, 2017.
- [41] Kush Bhatia, Aldo Pacchiano, Nicolas Flammarion, Peter L. Bartlett, and Michael I Jordan. Gen-Oja: Simple & efficient algorithm for streaming generalized eigenvector computation. In *Advances in Neural Information Processing Systems*, pages 7016–7025, 2018.
- [42] Cengiz Pehlevan, Xinyuan Zhao, Anirvan M Sengupta, and Dmitri Chklovskii. Neurons as canonical correlation analyzers. *Frontiers in computational neuroscience*, 14:55, 2020.
- [43] David Lipshutz, Yanis Bahroun, Siavash Golkar, Anirvan M. Sengupta, and Dmitri B. Chklovskii. A biologically plausible neural network for multi-channel canonical correlation analysis. *arXiv preprint arXiv:2010.00525*, 2020.
- [44] Prasanta Chandra Mahalanobis. On the generalized distance in statistics. *Proceedings of the National Institute of Sciences (Calcutta)*, 2:49–55, 1936.
- [45] Cengiz Pehlevan and Dmitri Chklovskii. A normative theory of adaptive dimensionality reduction in neural networks. In *NeurIPS*, pages 2269–2277, 2015.
- [46] Léon Bottou. Online algorithms and stochastic approximations. In David Saad, editor, *Online Learning and Neural Networks*. Cambridge University Press, Cambridge, UK, 1998. revised, oct 2012.
- [47] Thomas Klausberger, Peter J Magill, László F Márton, J David B Roberts, Philip M Cobden, György Buzsáki, and Peter Somogyi. Brain-state-and cell-type-specific firing of hippocampal interneurons in vivo. *Nature*, 421(6925):844–848, 2003.
- [48] Therese Riedemann. Diversity and function of somatostatin-expressing interneurons in the cerebral cortex. *International Journal of Molecular Sciences*, 20(12):2952, 2019.
- [49] Matthew E. Larkum, Thomas Nevian, Maya Sandler, Alon Polsky, and Jackie Schiller. Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: A new unifying principle. *Science*, 325(5941):756–760, 2009.
- [50] Hiroto Takahashi and Jeffrey C Magee. Pathway interactions and synaptic plasticity in the dendritic tuft regions of CA1 pyramidal neurons. *Neuron*, 62(1):102–111, 2009.
- [51] Charles D Gilbert and Wu Li. Top-down influences on visual processing. *Nature Reviews Neuroscience*, 14(5):350–363, 2013.
- [52] Georg B Keller and Thomas D Mrsic-Flogel. Predictive processing: a canonical cortical computation. *Neuron*, 100(2):424–435, 2018.
- [53] Michael Brecht and Bert Sakmann. Dynamic representation of whisker deflection by synaptic potentials in spiny stellate and pyramidal cells in the barrels and septa of layer 4 rat somatosensory cortex. *The Journal of Physiology*, 543(1):49–70, 2002.
- [54] Michael R. DeWeese, Michael Wehr, and Anthony M. Zador. Binary spiking in auditory cortex. *Journal of Neuroscience*, 23(21):7940–7949, 2003.
- [55] Bruno A Olshausen and David J Field. Sparse coding of sensory inputs. *Current Opinion in Neurobiology*, 14(4):481 – 487, 2004.
- [56] Maurice J Chacron, André Longtin, and Leonard Maler. Efficient computation via sparse coding in electrosensory neural networks. *Current Opinion in Neurobiology*, 21(5):752 – 760, 2011. Networks, circuits and computation.
- [57] Björn M. Kampa, Johannes J. Letzkus, and Greg J. Stuart. Dendritic mechanisms controlling spike-timing-dependent synaptic plasticity. *Trends in Neurosciences*, 30(9):456–463, 2007.
- [58] Stephen Grossberg. Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science*, 1987.
- [59] Francis Crick. The recent excitement about neural networks, 1989.
- [60] D Zipser and D E Rumelhart. The neurobiological significance of the new learning models. In *Computational Neuroscience*. 1990.
- [61] Kendra S. Burbank and Gabriel Kreiman. Depression-biased reverse plasticity rule is required for stable learning at top-down connections. *PLoS Computational Biology*, 2012.

- [62] Kendra S. Burbank. Mirrored STDP Implements Autoencoder Learning in a Network of Spiking Neurons. *PLoS Computational Biology*, 2015.
- [63] Timothy P. Lillicrap, Daniel Cownden, Douglas B. Tweed, and Colin J. Akerman. Random synaptic feedback weights support error backpropagation for deep learning. *Nature Communications*, 7:1–10, 2016.
- [64] Mohamed Akrouf, Collin Wilson, Peter C. Humphreys, Timothy Lillicrap, and Douglas Tweed. Using Weight Mirrors to Improve Feedback Alignment. *arXiv*, 2019.
- [65] Cees GM Snoek, Marcel Worring, Jan C Van Gemert, Jan-Mark Geusebroek, and Arnold WM Smeulders. The challenge problem for automated detection of 101 semantic concepts in multimedia. In *Proceedings of the 14th ACM international conference on Multimedia*, pages 421–430, 2006.
- [66] Rong Ge, Chi Jin, Praneeth Netrapalli, and Aaron Sidford. Efficient algorithms for large-scale generalized eigenvector computation and canonical correlation analysis. In *International Conference on Machine Learning*, pages 2741–2750, 2016.
- [67] Chao Gao, Dan Garber, Nathan Srebro, Jialei Wan, and Weiran Wang. Stochastic canonical correlation analysis. *Journal of Machine Learning Research*, 20:1–46, 2019.
- [68] Lecun Yann, Cortes Corinna, and Burges Christopher. The MNIST database of Handwritten Digits. *The Courant Institute of Mathematical Sciences*, pages 1–10, 1998.
- [69] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms, 2017.
- [70] A Krizhevsky, V Nair, and G Hinton. CIFAR-10 and CIFAR-100 datasets, 2009.