A biologically plausible neural network for multi-channel Canonical Correlation Analysis

David Lipshutz[*,1], Yanis Bahroun[*,1], Siavash Golkar[*,1],
Anirvan M. Sengupta[1,2], and 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

March 29, 2021

Abstract

Cortical pyramidal neurons receive inputs from multiple distinct neural populations and integrate these inputs in separate dendritic compartments. We explore the possibility that cortical microcircuits implement Canonical Correlation Analysis (CCA), an unsupervised learning method that projects the inputs onto a common subspace so as to maximize the correlations between the projections. To this end, we seek a multi-channel CCA algorithm that can be implemented in a biologically plausible neural network. For biological plausibility, we require that the network operates in the online setting and its synaptic update rules are local. Starting from a novel CCA objective function, we derive an online optimization algorithm whose optimization steps can be implemented in a single-layer neural network with multi-compartmental neurons and local non-Hebbian learning rules. We also derive an extension of our online CCA algorithm with adaptive output rank and output whitening. Interestingly, the extension maps onto a neural network whose neural architecture and synaptic updates resemble neural circuitry and non-Hebbian plasticity observed in the cortex.

---
[*]Equal contribution

Contents

1   Introduction

Our brains can effortlessly extract a latent source contributing to two synchronous data streams, often from different sensory modalities. Consider, for example, following an actor while watching a movie with a soundtrack. We can easily pay attention to the actor's gesticulation and voice while filtering out irrelevant visual and auditory signals. How can biological neurons accomplish such multi-sensory integration?

In this paper, we explore an algorithm for solving a linear version of this problem known as Canonical Correlation Analysis (CCA) [19]. In CCA, the two synchronous datasets, known as views, are projected onto a common lower-dimensional subspace so that the projections are maximally correlated. For simple generative models, the sum of these projections yields an optimal estimate of the latent source [3]. CCA is a popular method because it has a closed form exact solution in terms of the Singular Value Decomposition (SVD) of the correlation matrix. Therefore, the projections can be computed using fast and well understood spectral numerical methods.

To serve as a viable model of a neuronal circuit, the CCA algorithm must map onto a neural network consistent with basic biological facts. For our purposes, we say that a network is "biologically plausible" if it satisfies the following two minimal requirements: (i) the network operates in the online setting, i.e., upon receiving an input, it computes the corresponding output without relying on the storage of any significant fraction of the full dataset, and (ii) the learning rules are local in the sense that each synaptic update depends only on the variables that are available as biophysical quantities represented in the pre- or post-synaptic neurons.

There are a number of neural network implementations of CCA [22, 38, 15, 49]; however, most of these networks use non-local learning rules and are therefore not biologically plausible. One exception is the normative neural network model derived by Pehlevan et al. [37]. They start with formulating an objective for single-(output) channel CCA and derive an online optimization algorithm (previously proposed in [22]) that maps onto a pyramidal neuron with three electrotonic compartments: soma, as well as apical and basal dendrites. The apical and basal synaptic inputs represent the two views, the two dendritic compartments extract highly correlated CCA projections of the inputs and the soma computes the sum of projections and outputs it downstream as action potentials. The communication between the compartments is implemented by calcium plateaus that also mediate non-Hebbian but local synaptic plasticity.

Whereas Pehlevan et al. [37] also propose circuits of pyramidal neurons for multi-channel CCA their implementations lack biological plausibility. In one implementation, they resort to deflation where the circuit sequentially finds projections of the two views. Implementing this algorithm in a neural network requires a centralized mechanism to facilitate the sequential updates, and there is no experimental evidence of such a biological mechanism. In another implementation that does not
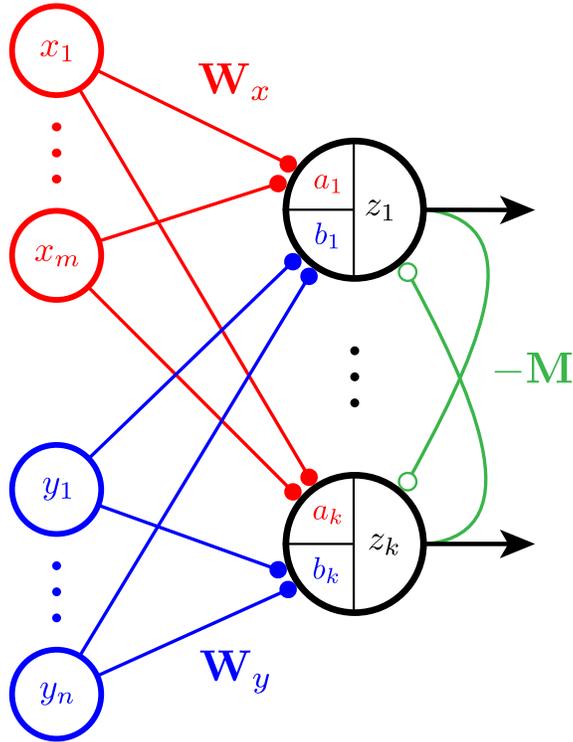
Figure 1: Single-layer network architecture with k multi-compartmental neurons for outputting the sum of the canonical correlation subspace projections (CCSPs) $\mathbf{z} = (z_1, \ldots, z_k)$, see Algorithm 2. Here $\mathbf{a} = \mathbf{W}_x\mathbf{x}$ and $\mathbf{b} = \mathbf{W}_y\mathbf{y}$ are projections of the views $\mathbf{x} = (x_1, \ldots, x_m)$ and $\mathbf{y} = (y_1, \ldots, y_n)$ onto a common k-dimensional subspace. The output, $\mathbf{z} = \mathbf{M}^{-1}(\mathbf{a} + \mathbf{b})$, is the sum of the CCSPs and is computed using recurrent lateral connections. The components of $\mathbf{a}$, $\mathbf{b}$ and $\mathbf{z}$ are represented in three separate compartments of the neurons. Filled circles denote non-Hebbian synapses and empty circles denote anti-Hebbian synapses. Importantly, each synaptic update depends only on variables represented locally.

require a centralized mechanism, the neural network has asymmetric lateral connections among pyramidal neurons. However, that algorithm is not derived from a principled objective for CCA and the network architecture does not match the neuronal circuitry observed in cortical microcircuits.

In this work, starting with a novel similarity-based CCA objective function, we derive a novel offline CCA algorithm (Algorithm 1) and an online multi-channel CCA algorithm (Algorithm 2), which can be implemented in a single-layer network composed of three-compartment neurons and with local non-Hebbian synaptic update rules, Figure 1. While our neural network implementation of CCA captures salient features of cortical microcircuits, the network includes direct lateral connections between the principal neurons (see Figure 1), which is in contrast to

cortical microcircuits where lateral influence between cortical pyramidal neurons is often indirect and mediated by interneurons. With this in mind, we derive an extension of our on-line CCA algorithm (Algorithm 3), which adaptively chooses the rank of the output based on the level of correlation captured, and also whitens the output. This extension is especially relevant for online unsupervised learning algorithms which are often confronted with the challenge of adapting to non-stationary input streams. In addition, the algorithm naturally maps onto a neural network with multi-compartmental principal neurons and without direct lateral connections between the principal neurons (see Figure 3 of Section 5). Interestingly, both the neural architecture and local, non-Hebbian plasticity resemble neural circuitry and synaptic plasticity in cortical microcircuits.

There are a number of existing consequential models of cortical microcircuits with multi-compartmental neurons and non-Hebbian plasticity [21, 47, 16, 42, 17, 41, 31]. These models provide mechanistic descriptions of the neural dynamics and synaptic plasticity and account for many experimental observations, including the nonlinearity of neural outputs and the layered organization of the cortex. While our neural network model is single-layered and linear, it is derived from a principled CCA objective function, which has several advantages. First, since biological neural networks evolved to adaptively perform behaviorally relevant computations, it is natural to view them as optimizing a relevant objective function. Second, our approach clarifies which features of the network (e.g., multi-compartmental neurons and non-Hebbian synaptic updates) are central to computing correlations. Finally, since the optimization algorithm is derived from a CCA objective that can be solved offline, the neural activities and synaptic weights can be analytically predicted for any input without resorting to numerical simulation. In this way, our neural network model is interpretable and analytically tractable, and provides a useful complement to nonlinear, layered neural network models.

Organization. The remainder of this work is organized as follows. We state the CCA problem in Section 2. In Section 3, we introduce a novel objective for the CCA problem and derive offline and online CCA algorithms. In Section 4, we derive an extension of our CCA algorithm, and in Section 5, we map the extension onto a simplified cortical microcircuit. We provide results of numerical simulations in Section 6.

Notation. For positive integers p, q, let $\mathbb{R}^p$ denote p-dimensional Euclidean space, and let $\mathbb{R}^{p \times q}$ denote the set of $p \times q$ real-valued matrices equipped with the Frobenius norm $\|\cdot\|_\mathrm{F}$. 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. We let $O(p)$ denote the set of $p \times p$ orthogonal matrices and $\mathcal{S}_{++}^p$ denote the set of $p \times p$ positive definite matrices. We let $\mathbf{I}_p$ denote the $p \times p$ identity matrix.
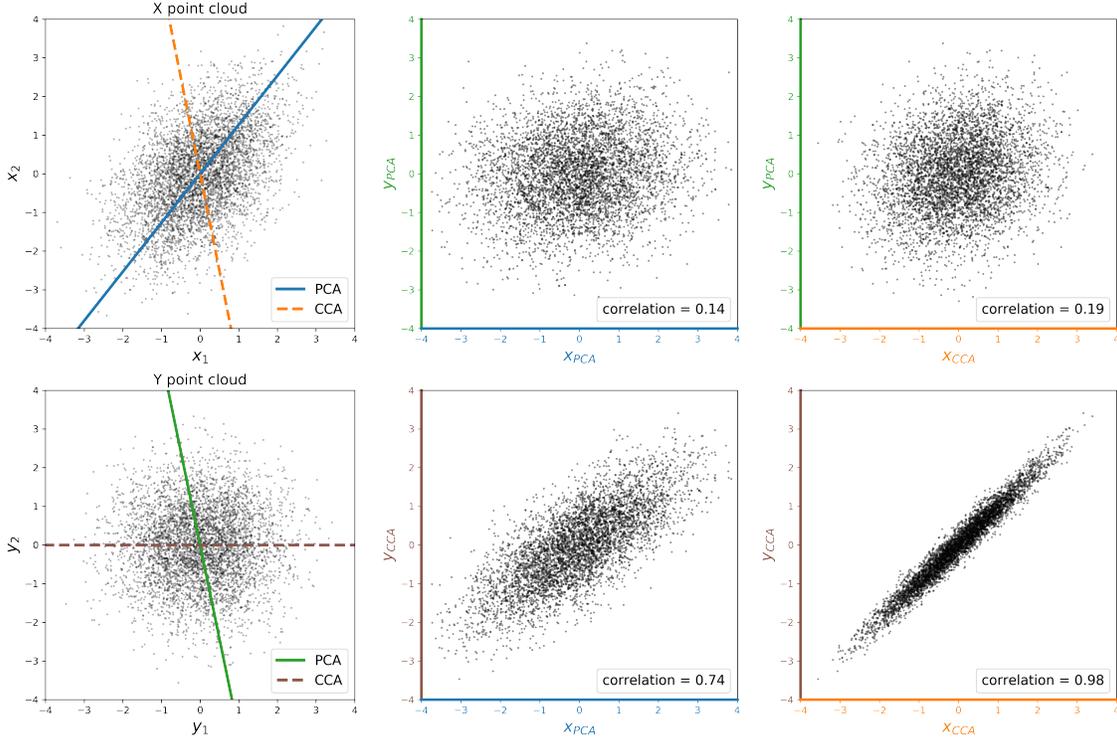
5

Figure 2: Illustration of CCA. The left column depicts point clouds of 2-dimensional views $\mathbf{X}$ and $\mathbf{Y}$ with lines indicating the span of the top principal component and top canonical correlation basis vector for each view (labeled "PCA" and "CCA", respectively). The center and right columns depict point clouds of joint 1-dimensional projections of $\mathbf{X}$ and $\mathbf{Y}$ onto their top principal component or top canonical correlation basis vector, with the correlations between the 2 projections listed. As illustrated in the lower right plot, the correlation between the projected views is maximized when each view is projected onto its top canonical correlation basis vector.

2 Canonical Correlation Analysis

Given T pairs of full-rank, centered input data samples $(\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_T, \mathbf{y}_T) \in \mathbb{R}^m \times \mathbb{R}^n$ and $k \leq \min(m, n)$, the CCA problem is to find k-dimensional linear projections of the views $\mathbf{x}_1, \ldots, \mathbf{x}_T$ and $\mathbf{y}_1, \ldots, \mathbf{y}_T$ that are maximally correlated, see Figure 2. To be precise, consider the CCA objective

$$\underset{\mathbf{V}_x \in \mathbb{R}^{m \times k}, \mathbf{V}_y \in \mathbb{R}^{n \times k}}{\arg\max} \quad \mathrm{Tr}\left(\mathbf{V}_x^\top \mathbf{C}_{xy} \mathbf{V}_y\right) \tag{1}$$

6

subject to the whitening constraint[1]

$$\mathbf{V}_x^\top \mathbf{C}_{xx} \mathbf{V}_x + \mathbf{V}_y^\top \mathbf{C}_{yy} \mathbf{V}_y = \mathbf{I}_k, \tag{2}$$

where we have defined the sample covariance matrices

$$\mathbf{C}_{xx} := \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t \mathbf{x}_t^\top, \qquad \mathbf{C}_{xy} := \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t \mathbf{y}_t^\top, \qquad \mathbf{C}_{yy} := \frac{1}{T} \sum_{t=1}^T \mathbf{y}_t \mathbf{y}_t^\top. \tag{3}$$

To compute the solution of the CCA objective (1)–(2), define the $m \times n$ *correlation matrix*

$$\mathbf{R}_{xy} := \mathbf{C}_{xx}^{-1/2} \mathbf{C}_{xy} \mathbf{C}_{yy}^{-1/2}. \tag{4}$$

Let $\rho_1 \geq \cdots \geq \rho_{\min(m,n)}$ denote the singular values, and let $\mathbf{U}_x \in O(m)$ and $\mathbf{U}_y \in O(n)$ denote the matrices whose column vectors are respectively the left- and right-singular vectors of the correlation matrix. The i^{th} singular value ρ_i is referred to as the i^{th} *canonical correlation*, and the i^{th} column vectors of $\mathbf{C}_{xx}^{-1/2} \mathbf{U}_x$ and $\mathbf{C}_{yy}^{-1/2} \mathbf{U}_y$ are jointly referred to as the i^{th} pair of *canonical correlation basis vectors*, for $i = 1, \ldots, \min(m,n)$. The maximal value of the trace in Equation (1) is the normalized sum of canonical correlations: $(\rho_1 + \cdots + \rho_k)/2$. For simplicity, we assume $\rho_k > \rho_{k+1}$ so the subspace spanned by the first k canonical correlation basis vectors is unique. In this case, every solution of the CCA objective (1)–(2), denoted $(\overline{\mathbf{V}}_x, \overline{\mathbf{V}}_y)$, is of the form

$$\overline{\mathbf{V}}_x = \mathbf{C}_{xx}^{-1/2} \mathbf{U}_x^{(k)} \mathbf{Q}, \qquad\qquad \overline{\mathbf{V}}_y = \mathbf{C}_{yy}^{-1/2} \mathbf{U}_y^{(k)} \mathbf{Q}, \tag{5}$$

where $\mathbf{U}_x^{(k)}$ (resp. $\mathbf{U}_y^{(k)}$) is the $m \times k$ (resp. $n \times k$) matrix whose i^{th} column vector is equal to the i^{th} column vector of $\mathbf{U}_x$ (resp. $\mathbf{U}_y$) for $i = 1, \ldots, k$, and $\mathbf{Q} \in O(k)$ is any orthogonal matrix. Since the column vectors of any solution $(\overline{\mathbf{V}}_x, \overline{\mathbf{V}}_y)$ span the same subspaces as the first k pairs of canonical correlation basis vectors, we refer to the column vectors of $\overline{\mathbf{V}}_x$ and $\overline{\mathbf{V}}_y$ as *basis vectors*. We refer to the k-dimensional projections $\overline{\mathbf{V}}_x^\top \mathbf{x}_t$ and $\overline{\mathbf{V}}_y^\top \mathbf{y}_t$ as *canonical correlation subspace projections (CCSPs)*.

The focus of this work is to derive a single-layer biologically plausible network whose input at each time t is the pair $(\mathbf{x}_t, \mathbf{y}_t)$ and the output is the following sum of the CCSPs:

$$\mathbf{z}_t := \overline{\mathbf{V}}_x^\top \mathbf{x}_t + \overline{\mathbf{V}}_y^\top \mathbf{y}_t, \tag{6}$$

which, as mentioned in the introduction, is a highly relevant statistic (see, also, Section 6.1). This is in contrast to many existing CCA networks which output one

---

[1]This constraint differs slightly from the usual CCA whitening constraint $\mathbf{V}_x^\top \mathbf{C}_{xx} \mathbf{V}_x = \mathbf{V}_y^\top \mathbf{C}_{yy} \mathbf{V}_y = \mathbf{I}_k$; however, the constraints are equivalent up to a scaling factor of 2.

or both CCSPs [22, 38, 15, 49, 14]. The components of the two input vectors $\mathbf{x}_t$ and $\mathbf{y}_t$ are represented by the activity of upstream neurons belonging to two different populations, which are integrated in separate compartments of the principal neurons in our network. The components of the output vector $\mathbf{z}_t$ are represented by the activity of the principal neurons in our network, see Figure 1.

While CCA is typically viewed as an *unsupervised* learning method, it can also be interpreted as a special case of the *supervised* learning method Reduced-Rank Regression, in which case one input is the feature vector and the other input is the label (see, e.g., page 38 of [48]). With this supervised learning view of CCA, the natural output of a CCA network is the CCSP of the feature vector. In separate work [14], we derive an algorithm for the general Reduced-Rank Regression problem, which includes CCA as a special case, for outputting the projection of the feature vector. The algorithm derived in [14] resembles the adaptive CCA with output whitening algorithm that we derive in Section 4 of this work (see Algorithm 3 as well as Appendix B.2 for a detailed comparison of the two algorithms); however, there are significant advantages to the algorithm derived here. First, our network outputs the (whitened) sum of the CCSPs, which, as discussed above, is a relevant statistic in applications. The algorithm in [14] only outputs the CCSP of the feature vector, which is natural when viewing CCA as a supervised learning method, but not when viewing CCA as an unsupervised learning method for integrating multi-view inputs. Second, in contrast to the algorithm derived in [14], our adaptive CCA with output whitening algorithm allows for adaptive output rank. This is particularly important for analyzing non-stationary input streams, a challenge that brains regularly face.

3   A biologically plausible CCA algorithm

To derive a network that computes the sums of CCSPs for arbitrary input datasets, we adopt a normative approach in which we identify an appropriate cost function whose optimization leads to an online algorithm that can be implemented by a network with local learning rules. Previously, such approach was taken to derive a biologically plausible PCA network from a similarity matching objective function [35]. We leverage this work by reformulating a CCA problem in terms of PCA of a modified dataset and then solving it using similarity matching.

3.1   A similarity matching objective

First, we note that the sums of CCSPs $\mathbf{z}_1, \ldots, \mathbf{z}_T$ are equal to the principal subspace projections of the data $\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_T$, where $\boldsymbol{\xi}_t$ is the following d-dimensional vector of

concatenated whitened inputs (recall $d := m + n$):

$$\boldsymbol{\xi}_t := \begin{bmatrix} \mathbf{C}_{xx}^{-1/2} \mathbf{x}_t \\ \mathbf{C}_{yy}^{-1/2} \mathbf{y}_t \end{bmatrix}. \tag{7}$$

(See Appendix A for a detailed justification.) Next, we use the fact that the principal subspace projections can be expressed in terms of solutions of similarity matching objectives. To this end, we define the matrices $\boldsymbol{\Xi} := [\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_T] \in \mathbb{R}^{d \times T}$ and $\mathbf{Z} := [\mathbf{z}_1, \ldots, \mathbf{z}_T] \in \mathbb{R}^{k \times T}$, so that $\mathbf{Z}$ is a linear projection of $\boldsymbol{\Xi}$ onto its k-dimensional principal subspace. As shown in [9, 51], the principal subspace projection $\mathbf{Z}$ is a solution of following similarity matching objective:

$$\underset{\mathbf{Z} \in \mathbb{R}^{k \times T}}{\arg \min} \frac{1}{2T^2} \| \mathbf{Z}^\top \mathbf{Z} - \boldsymbol{\Xi}^\top \boldsymbol{\Xi} \|_{\mathrm{F}}^2. \tag{8}$$

The objective (8), which comes from classical multidimensional scaling [9], minimizes the difference between the similarity of output pairs, $\mathbf{z}_t^\top \mathbf{z}_{t'}$, and the similarity of input pairs, $\boldsymbol{\xi}_t^\top \boldsymbol{\xi}_{t'}$, where similarity is measured in terms of inner products. Finally, defining $\mathbf{X} := [\mathbf{x}_1, \ldots, \mathbf{x}_T] \in \mathbb{R}^{m \times T}$ and $\mathbf{Y} := [\mathbf{y}_1, \ldots, \mathbf{y}_T] \in \mathbb{R}^{n \times T}$, we use the definition of $\boldsymbol{\xi}_t$ in Equation (7) to rewrite the objective (8) as follows:

$$\underset{\mathbf{Z} \in \mathbb{R}^{k \times T}}{\arg \min} \frac{1}{2T^2} \| \mathbf{Z}^\top \mathbf{Z} - \mathbf{X}^\top \mathbf{C}_{xx}^{-1} \mathbf{X} - \mathbf{Y}^\top \mathbf{C}_{yy}^{-1} \mathbf{Y} \|_{\mathrm{F}}^2. \tag{9}$$

In the remainder of this section, we derive our online CCA algorithm. Then, in Sections 4 and 5, we derive an extension of our CCA algorithm and map it onto the cortical microcircuit. The reader who is primarily interested in the derivation of the extension and its relation to the cortical microcircuit can safely skip to Section 4.

3.2 A min-max objective

While the similarity matching objective (9) can be minimized by taking gradient descent steps with respect to $\mathbf{Z}$, this would not lead to an online algorithm because such computation requires combining data from different time steps. Rather, we introduce auxiliary matrix variables, which store sufficient statistics allowing for the CCA computation using solely contemporary inputs and will correspond to synaptic weights in the network implementation, and rewrite the minimization problem (9) as a min-max problem.

Expanding the square in Equation (9) and dropping terms that do not depend on $\mathbf{Z}$ yields the minimization problem

$$\underset{\mathbf{Z} \in \mathbb{R}^{k \times T}}{\min} -\frac{1}{T^2} \operatorname{Tr} \left(\mathbf{Z}^\top \mathbf{Z} \mathbf{X}^\top \mathbf{C}_{xx}^{-1} \mathbf{X} \right) - \frac{1}{T^2} \operatorname{Tr} \left(\mathbf{Z}^\top \mathbf{Z} \mathbf{Y}^\top \mathbf{C}_{yy}^{-1} \mathbf{Y} \right) + \frac{1}{2T^2} \operatorname{Tr} \left(\mathbf{Z}^\top \mathbf{Z} \mathbf{Z}^\top \mathbf{Z} \right).$$

Next, we introduce dynamic matrix variables $\mathbf{W}_x$, $\mathbf{W}_y$ and $\mathbf{M}$ in place of the matrices $\frac{1}{T}\mathbf{Z}\mathbf{X}^\top \mathbf{C}_{xx}^{-1}$, $\frac{1}{T}\mathbf{Z}\mathbf{Y}^\top \mathbf{C}_{yy}^{-1}$ and $\frac{1}{T}\mathbf{Z}\mathbf{Z}^\top$, respectively, and rewrite the minimization problem as a min-max problem:

$$\min_{\mathbf{Z}\in\mathbb{R}^{k\times T}} \min_{\mathbf{W}_x\in\mathbb{R}^{k\times m}} \min_{\mathbf{W}_y\in\mathbb{R}^{k\times n}} \max_{\mathbf{M}\in\mathcal{S}_{++}^k} L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{Z})$$

where

$$L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{Z}) := \frac{1}{T}\operatorname{Tr}\left(-2\mathbf{Z}^\top\mathbf{W}_x\mathbf{X} - 2\mathbf{Z}^\top\mathbf{W}_y\mathbf{Y} + \mathbf{Z}^\top\mathbf{M}\mathbf{Z}\right)$$
$$+ \operatorname{Tr}\left(\mathbf{W}_x\mathbf{C}_{xx}\mathbf{W}_x^\top + \mathbf{W}_y\mathbf{C}_{yy}\mathbf{W}_y^\top - \frac{1}{2}\mathbf{M}^2\right). \quad (10)$$

To verify the above substitutions are valid, it suffices to optimize over the matrices $\mathbf{W}_x$, $\mathbf{W}_y$ and $\mathbf{M}$; e.g., by differentiating $L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{Z})$ with respect to $\mathbf{W}_x$, $\mathbf{W}_y$ or $\mathbf{M}$, setting the derivative equal to zero, and solving for $\mathbf{W}_x$, $\mathbf{W}_y$ or $\mathbf{M}$. Finally, we interchange the order of minimization with respect to $\mathbf{Z}$ and $(\mathbf{W}_x, \mathbf{W}_y)$, as well as the order of minimization with respect to $\mathbf{Z}$ and maximization with respect to $\mathbf{M}$:

$$\min_{\mathbf{W}_x\in\mathbb{R}^{k\times m}} \min_{\mathbf{W}_y\in\mathbb{R}^{k\times n}} \max_{\mathbf{M}\in\mathcal{S}_{++}^k} \min_{\mathbf{Z}\in\mathbb{R}^{k\times T}} L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{Z}). \quad (11)$$

The second interchange is justified by the fact that $L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{Z})$ satisfies the saddle point property with respect to $\mathbf{Z}$ and $\mathbf{M}$, which follows from its strict convexity in $\mathbf{Z}$ (since $\mathbf{M}$ is positive definite) and strict concavity in $\mathbf{M}$.

Given an optimal quadruple of the min-max problem (11), we can compute the basis vectors, as follows. First, minimizing the objective $L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{Z})$ over $\mathbf{Z}$ yields the relation

$$\overline{\mathbf{Z}} := \arg\min_{\mathbf{Z}\in\mathbb{R}^{k\times T}} L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{Z}) = \mathbf{M}^{-1}\mathbf{W}_x\mathbf{X} + \mathbf{M}^{-1}\mathbf{W}_y\mathbf{Y}. \quad (12)$$

Therefore, if $(\overline{\mathbf{W}}_x, \overline{\mathbf{W}}_y, \overline{\mathbf{M}}, \overline{\mathbf{Z}})$ is an optimal quadruple of the min-max problem (11), it follows from Equation (6) that the corresponding basis vectors satisfy

$$\overline{\mathbf{V}}_x^\top = \overline{\mathbf{M}}^{-1}\overline{\mathbf{W}}_x \qquad \text{and} \qquad \overline{\mathbf{V}}_y^\top = \overline{\mathbf{M}}^{-1}\overline{\mathbf{W}}_y. \quad (13)$$

3.3 An offline CCA algorithm

Before deriving our online CCA algorithm, we first demonstrate how the objective (11) can be optimized in the offline setting, where one has access to the data matrices $\mathbf{X}$ and $\mathbf{Y}$ in their entirety. In this case, the algorithm solves the min-max problem

([11](#)) by alternating minimization and maximization steps. First, for fixed $\mathbf{W}_x$, $\mathbf{W}_y$ and $\mathbf{M}$, we minimize the objective function $L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{Z})$ over $\mathbf{Z}$ to obtain the minimum $\overline{\mathbf{Z}}$ defined in Equation ([12](#)). Then, with $\overline{\mathbf{Z}}$ fixed, we perform a gradient descent-ascent step with respect to $(\mathbf{W}_x, \mathbf{W}_y)$ and $\mathbf{M}$:

$$\mathbf{W}_x \leftarrow \mathbf{W}_x - \eta \frac{\partial L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \overline{\mathbf{Z}})}{\partial \mathbf{W}_x}$$

$$\mathbf{W}_x \leftarrow \mathbf{W}_y - \eta \frac{\partial L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \overline{\mathbf{Z}})}{\partial \mathbf{W}_y}$$

$$\mathbf{M} \leftarrow \mathbf{M} + \frac{\eta}{\tau} \frac{\partial L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \overline{\mathbf{Z}})}{\partial \mathbf{M}}.$$

Here $\eta > 0$ is the learning rate for $\mathbf{W}_x$ and $\mathbf{W}_y$, which may depend on the iteration, and $\tau > 0$ is the ratio of the learning rates for $\mathbf{W}_x$ (or $\mathbf{W}_y$) and $\mathbf{M}$. Substituting in the explicit expressions for the partial derivatives of $L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \overline{\mathbf{Z}})$ yields our offline CCA algorithm (Algorithm [1](#)), which we refer to as Offline-CCA.

Algorithm 1: Offline-CCA

    input: data matrices $\mathbf{X}$, $\mathbf{Y}$; dimension k
    initialize: matrices $\mathbf{W}_x$, $\mathbf{W}_y$ and positive definite matrix $\mathbf{M}$
    $\mathbf{C}_{xx} \leftarrow \frac{1}{T}\mathbf{X}\mathbf{X}^\top$; $\mathbf{C}_{yy} \leftarrow \frac{1}{T}\mathbf{Y}\mathbf{Y}^\top$ ▷ covariance matrices
    repeat
      $\mathbf{Z} \leftarrow \mathbf{M}^{-1}\mathbf{W}_x\mathbf{X} + \mathbf{M}^{-1}\mathbf{W}_x\mathbf{Y}$ ▷ optimize over output
      $\mathbf{W}_x \leftarrow \mathbf{W}_x + 2\eta \left(\frac{1}{T}\mathbf{Z}\mathbf{X}^\top - \mathbf{W}_x\mathbf{C}_{xx}\right)$ ▷ gradient descent-ascent steps
      $\mathbf{W}_y \leftarrow \mathbf{W}_y + 2\eta \left(\frac{1}{T}\mathbf{Z}\mathbf{Y}^\top - \mathbf{W}_y\mathbf{C}_{yy}\right)$
      $\mathbf{M} \leftarrow \mathbf{M} + \frac{\eta}{\tau} \left(\frac{1}{T}\mathbf{Z}\mathbf{Z}^\top - \mathbf{M}\right)$
    until convergence

Recall that $\mathbf{M}$ is optimized over the set of positive definite matrices $\mathcal{S}_{++}^k$. To ensure that $\mathbf{M}$ remains positive definite after each update, note that the update rule for $\mathbf{M}$ can be rewritten as the following convex combination (provided $\eta \leq \tau$): $\mathbf{M} \leftarrow (1 - \frac{\eta}{\tau})\mathbf{M} + \frac{\eta}{\tau}(\frac{1}{T}\mathbf{Z}\mathbf{Z}^\top)$. Since $\frac{1}{T}\mathbf{Z}\mathbf{Z}^\top$ is positive semidefinite, to guarantee that $\mathbf{M}$ remains positive definite given a positive definite initialization, it suffices to assume that $\eta < \tau$.

3.4 An online CCA algorithm

In the online setting, the input data $(\mathbf{x}_t, \mathbf{y}_t)$ are streamed one at a time and the algorithm must compute its output $\mathbf{z}_t$ without accessing any significant fraction of $\mathbf{X}$ and $\mathbf{Y}$. To derive an online algorithm, it is useful to write the cost function as

an average over time-separable terms:

$$L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{Z}) = \frac{1}{T} \sum_{t=1}^{T} l_t(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{z}_t),$$

where

$$
\begin{aligned}
l_t(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{z}_t) := {}& -2\mathbf{z}_t^\top \mathbf{W}_x \mathbf{x}_t - 2\mathbf{z}_t^\top \mathbf{W}_y \mathbf{y}_t + \mathbf{z}_t^\top \mathbf{M} \mathbf{z}_t \\
& + \mathrm{Tr}\left(\mathbf{W}_x \mathbf{x}_t \mathbf{x}_t^\top \mathbf{W}_x^\top + \mathbf{W}_y \mathbf{y}_t \mathbf{y}_t^\top \mathbf{W}_y^\top - \frac{1}{2}\mathbf{M}^2 \right). \quad (14)
\end{aligned}
$$

At iteration t, to compute the output $\mathbf{z}_t$, we minimize the cost function $l_t(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{z}_t)$ with respect to $\mathbf{z}_t$ by running the following gradient descent dynamics to equilibrium:

$$\frac{d\mathbf{z}_t(\gamma)}{d\gamma} = \mathbf{a}_t + \mathbf{b}_t - \mathbf{M}\mathbf{z}_t(\gamma), \tag{15}$$

where we have defined the following k-dimensional projections of the inputs: $\mathbf{a}_t :=$ $\mathbf{W}_x \mathbf{x}_t$ and $\mathbf{b}_t := \mathbf{W}_y \mathbf{y}_t$. These dynamics, which will correspond to recurrent neural dynamics in our network implementation, are assumed to occur on a fast timescale, allowing $\mathbf{z}_t(\gamma)$ to equilibrate at $\bar{\mathbf{z}}_t := \mathbf{M}^{-1}(\mathbf{a}_t + \mathbf{b}_t)$ before the algorithm outputs its value. After $\mathbf{z}_t(\gamma)$ equilibrates, we update the matrices $(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M})$ by taking a stochastic gradient descent-ascent step of the cost function $l_t(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \bar{\mathbf{z}}_t)$ with respect to $(\mathbf{W}_x, \mathbf{W}_y)$ and $\mathbf{M}$:

$$
\begin{aligned}
\mathbf{W}_x &\leftarrow \mathbf{W}_x - \eta \frac{\partial l_t(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \bar{\mathbf{z}}_t)}{\partial \mathbf{W}_x} \\
\mathbf{W}_x &\leftarrow \mathbf{W}_y - \eta \frac{\partial l_t(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \bar{\mathbf{z}}_t)}{\partial \mathbf{W}_y} \\
\mathbf{M} &\leftarrow \mathbf{M} + \frac{\eta}{\tau} \frac{\partial l_t(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \bar{\mathbf{z}}_t)}{\partial \mathbf{M}}.
\end{aligned}
$$

Substituting in the explicit expressions for the partial derivatives of $l_t(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \bar{\mathbf{z}}_t)$ yields our online CCA algorithm (Algorithm 2), which we refer to as Bio-CCA.

Algorithm 2 can be implemented in a biologically plausible single-layer network with k neurons that each consist of three separate compartments, Figure 1. At each time step, the inputs $\mathbf{x}_t$ and $\mathbf{y}_t$ are multiplied by the respective feedforward synapses $\mathbf{W}_x$ and $\mathbf{W}_y$ to yield the k-dimensional vectors $\mathbf{a}_t$ and $\mathbf{b}_t$, which are represented in the first two compartments of the k neurons. Lateral synapses, $-\mathbf{M}$, connect the k neurons. The vector of neuronal outputs, $\mathbf{z}_t$, equals the normalized sum of the CCSPs, and is computed locally using recurrent dynamics in Equation (15). The

---
Algorithm 2: Bio-CCA
---
    input data $\{(\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_T, \mathbf{y}_T)\}$; dimension k
    initialize matrices $\mathbf{W}_x$, $\mathbf{W}_y$, and positive definite matrix $\mathbf{M}$.
    for $t = 1, 2, \ldots, T$ **do**
        $\mathbf{a}_t \leftarrow \mathbf{W}_x\mathbf{x}_t$; $\mathbf{b}_t \leftarrow \mathbf{W}_y\mathbf{y}_t$ ▷ projection of inputs
        run
            $\frac{d\mathbf{z}_t(\gamma)}{d\gamma} = \mathbf{a}_t + \mathbf{b}_t - \mathbf{M}\mathbf{z}_t(\gamma)$ ▷ neural dynamics
        until convergence
        $\mathbf{W}_x \leftarrow \mathbf{W}_x + 2\eta(\mathbf{z}_t - \mathbf{a}_t)\mathbf{x}_t^\top$ ▷ synaptic updates
        $\mathbf{W}_y \leftarrow \mathbf{W}_y + 2\eta(\mathbf{z}_t - \mathbf{b}_t)\mathbf{y}_t^\top$
        $\mathbf{M} \leftarrow \mathbf{M} + \frac{\eta}{\tau}(\mathbf{z}_t\mathbf{z}_t^\top - \mathbf{M})$
    end for
---

synaptic updates can be written elementwise, as follows:

$$
\begin{aligned}
W_{x,ij} &\leftarrow W_{x,ij} + \eta(z_{t,i} - a_{t,i})x_{t,j}, && 1 \le i \le k,\ 1 \le j \le m, \\
W_{y,ij} &\leftarrow W_{y,ij} + \eta(z_{t,i} - b_{t,i})y_{t,j}, && 1 \le i \le k,\ 1 \le j \le n, \\
M_{ij} &\leftarrow M_{ij} + \frac{\eta}{\tau}(z_{t,i}z_{t,j} - M_{ij}), && 1 \le i, j \le k.
\end{aligned}
$$

As shown above, the update to synapse $W_{x,ij}$ (resp. $W_{y,ij}$), which connects the j^{th} input $x_{t,j}$ (resp. $y_{t,j}$) to the i^{th} output neuron, depends only on the quantities $z_{t,i}$, $a_{t,i}$ (resp. $b_{t,i}$), and $x_{t,j}$ (resp. $y_{t,j}$), which are represented in the pre- and post-synaptic neurons, so the updates are local, but non-Hebbian due to the contribution from the $a_{t,i}$ (resp. $b_{t,i}$) term. Similarly, the update to synapse $-M_{ij}$, which connects the j^{th} output neuron to the i^{th} output neuron, is inversely proportional to $z_{t,i}z_{t,j}$, the product of the outputs of the pre- and post-synaptic neurons, so the updates are local and anti-Hebbian.

4 Online adaptive CCA with output whitening

We now introduce an extension of Bio-CCA which addresses two biologically relevant issues. First, Bio-CCA a priori sets the output rank at k; however, it may be advantageous for a neural circuit to instead adaptively set the output rank depending on the level of correlation captured. In particular, this can be achieved by projecting each view onto the subspace spanned by the canonical correlation basis vectors which correspond to canonical correlations that exceed a threshold. Second, it is useful from an information theoretic perspective for neural circuits to whiten their outputs [39], and there is experimental evidence that neural outputs in the cortex are decorrelated [10, 32]. Both adaptive output rank and output whitening modifications were implemented for a PCA network by Pehlevan and Chklovskii [34], and can be

adapted to the CCA setting. Here we present the modifications without providing detailed proofs, which can be found in the supplement of [34].

In order to implement these extensions, we need to appropriately modify the similarity matching objective function (9). First, to adaptively choose the output rank, we add a quadratic penalty $\mathrm{Tr}(\mathbf{Z}^\top\mathbf{Z})$ to the objective function (9):

$$\underset{\mathbf{Z}\in\mathbb{R}^{k\times T}}{\arg\min}\,\frac{1}{2T^2}\|\mathbf{Z}^\top\mathbf{Z}-\mathbf{X}^\top\mathbf{C}_{xx}^{-1}\mathbf{X}-\mathbf{Y}^\top\mathbf{C}_{yy}^{-1}\mathbf{Y}\|_{\mathrm{F}}^2+\frac{\alpha}{T}\,\mathrm{Tr}\left(\mathbf{Z}^\top\mathbf{Z}\right). \qquad (16)$$

The effect of the quadratic penalty is to rank constrain the output, with $\alpha\geq 0$ acting as a threshold parameter on the eigenvalues values of the output covariance.

Next, to whiten the output, we expand the square in Equation (16) and replace the quartic term $\mathrm{Tr}(\mathbf{Z}^\top\mathbf{Z}\mathbf{Z}^\top\mathbf{Z})$ by a Lagrange constraint enforcing $\mathbf{Z}^\top\mathbf{Z}\preceq T\mathbf{I}_T$ (i.e., $T\mathbf{I}_T-\mathbf{Z}^\top\mathbf{Z}$ is positive semi-definite):

$$\underset{\mathbf{Z}\in\mathbb{R}^{k\times T}}{\arg\min}\,\underset{\mathbf{N}\in\mathbb{R}^{k\times T}}{\max}\,\frac{1}{T^2}\,\mathrm{Tr}(-\mathbf{Z}^\top\mathbf{Z}\mathbf{X}^\top\mathbf{C}_{xx}^{-1}\mathbf{X}-\mathbf{Z}^\top\mathbf{Z}\mathbf{Y}^\top\mathbf{C}_{yy}^{-1}\mathbf{Y}+\alpha T\mathbf{Z}^\top\mathbf{Z}) \qquad (17)$$

$$+\frac{1}{T^2}\,\mathrm{Tr}[\mathbf{N}^\top\mathbf{N}(\mathbf{Z}^\top\mathbf{Z}-T\mathbf{I}_T)].$$

The effect of the Lagrange constraint in Equation (17) is to enforce that all non-zero eigenvalues of the output covariance are set to one.

Solutions of the objective (17) can be expressed in terms of the eigendecomposition of the Gram matrix $\mathbf{\Xi}^\top\mathbf{\Xi}=T\mathbf{U}_\xi\mathbf{\Lambda}_\xi\mathbf{U}_\xi^\top$, where $\mathbf{U}_\xi\in O(T)$ is a matrix of eigenvectors and $\mathbf{\Lambda}_\xi=\mathrm{diag}(\lambda_1,\ldots,\lambda_d,0,\ldots,0)$ is the $T\times T$ diagonal matrix whose non-zero entries $\lambda_1\geq\cdots\geq\lambda_d>0$ are the eigenvalues of the $d\times d$ covariance matrix

$$\mathbf{C}_{\xi\xi}:=\frac{1}{T}\sum_{t=1}^T\boldsymbol{\xi}_t\boldsymbol{\xi}_t^\top=\begin{bmatrix}\mathbf{I}_m & \mathbf{R}_{xy}\\ \mathbf{R}_{xy}^\top & \mathbf{I}_n\end{bmatrix}. \qquad (18)$$

Assume, for technical purposes, that $\alpha\notin\{\lambda_1,\ldots,\lambda_d\}$. Then, as shown in [35, Theorem 3], every solution $\widehat{\mathbf{Z}}$ of objective (17) is of the form

$$\widehat{\mathbf{Z}}=\mathbf{Q}\sqrt{T\widehat{\mathbf{\Lambda}}_\xi^{(k)}}\mathbf{U}_\xi^{(k)\top},\qquad\widehat{\mathbf{\Lambda}}_\xi^{(k)}=\mathrm{diag}(H(\lambda_1-\alpha),\ldots,H(\lambda_k-\alpha)),$$

where $\mathbf{Q}\in O(k)$ is any orthogonal matrix, $\mathbf{U}_\xi^{(k)}\in\mathbb{R}^{T\times k}$ is the $T\times k$ matrix whose i^{th} column vector is equal to the i^{th} column vector of $\mathbf{U}_\xi$, for $i=1,\ldots,k$, and H is the Heaviside step function defined by $H(r)=1$ if $r>0$ and $H(r)=0$ otherwise. Finally, we note that, in view of Equation (18) and the SVD of $\mathbf{R}_{xy}$, the top $\min(m,n)$ eigenvalues of $\mathbf{C}_{\xi\xi}$ satisfy

$$\lambda_i=1+\rho_i,\qquad i=1,\ldots,\min(m,n),$$

14

where we recall that $\rho_1, \ldots, \rho_{\min(m,n)}$ are the canonical correlations. Thus, $H(\lambda_i - \alpha) = H(\rho_i - (\alpha - 1))$, for $i = 1, \ldots, k$. In other words, the objective (17) outputs the sum of the projections of the inputs $\mathbf{x}_t$ and $\mathbf{y}_t$ onto the canonical correlation subspace spanned by the (at most k) pairs of canonical correlation basis vectors associated with canonical correlations exceeding the threshold $\max(\alpha - 1, 0)$, and sets the non-zero output covariance eigenvalues to one, thus implementing both the adaptive output rank and output whitening modifications.

With the modified objective (17) in hand, the next step is to derive an online algorithm. Similar to Section 3.2, we introduce dynamic matrix variables $\mathbf{W}_x$, $\mathbf{W}_y$ and $\mathbf{P}$ in place of $\frac{1}{T}\mathbf{Z}\mathbf{X}^\top\mathbf{C}_{xx}^{-1}$, $\frac{1}{T}\mathbf{Z}\mathbf{Y}^\top\mathbf{C}_{yy}^{-1}$ and $\frac{1}{T}\mathbf{Z}\mathbf{N}^\top$ to rewrite the objective (17) as follows:

$$\arg\min_{\mathbf{Z}\in\mathbb{R}^{k\times T}} \max_{\mathbf{N}\in\mathbb{R}^{k\times T}} \min_{\mathbf{W}_x\in\mathbb{R}^{k\times m}} \min_{\mathbf{W}_y\in\mathbb{R}^{k\times n}} \max_{\mathbf{P}\in\mathbb{R}^{k\times k}} \widetilde{L}(\mathbf{W}_x, \mathbf{W}_y, \mathbf{P}, \mathbf{Z}, \mathbf{N}), \tag{19}$$

where

$$\begin{aligned}
\widetilde{L}(\mathbf{W}_x, \mathbf{W}_y, \mathbf{P}, \mathbf{Z}, \mathbf{N}) :={}& \frac{1}{T}\operatorname{Tr}\left(-2\mathbf{Z}^\top\mathbf{W}_x\mathbf{X} - 2\mathbf{Z}^\top\mathbf{W}_y\mathbf{Y} + \alpha\mathbf{Z}^\top\mathbf{Z}\right) \\
&+ \frac{1}{T}\operatorname{Tr}\left(2\mathbf{N}^\top\mathbf{P}^\top\mathbf{Z} - \mathbf{N}^\top\mathbf{N}\right) \\
&+ \operatorname{Tr}\left(\mathbf{W}_x\mathbf{C}_{xx}\mathbf{W}_x^\top + \mathbf{W}_y\mathbf{C}_{yy}\mathbf{W}_y^\top - \mathbf{P}\mathbf{P}^\top\right).
\end{aligned}$$

After interchanging the order of optimization, we solve the min-max optimization problem by taking online gradient descent-ascent steps, with descent step size η and ascent step size $\frac{\eta}{\tau}$. Since the remaining steps are similar to those taken in Section 3 to derive Bio-CCA, we defer the details to Appendix B.1 and simply state the online algorithm (Algorithm 3), which we refer to as Adaptive Bio-CCA with output whitening.

5   Relation to cortical microcircuits

We now show that Adaptive Bio-CCA with output whitening (Algorithm 3) maps onto a neural network with local, non-Hebbian synaptic update rules that emulate salient aspects of synaptic plasticity found experimentally in cortical microcircuits (both in the neocortex and the hippocampus).

Cortical microcircuits contain two classes of neurons: excitatory pyramidal neurons and inhibitory interneurons. Pyramidal neurons receive excitatory synaptic inputs from two distinct sources via their apical and basal dendrites. The apical dendrites are all oriented in a single direction and the basal dendrites branch from the cell body in the opposite direction [45, 23], Figure 3. The excitatory synaptic currents in the apical and basal dendrites are first integrated separately in their respective compartments [45, 23]. If the integrated excitatory current in the apical

---

Algorithm 3: Adaptive Bio-CCA with output whitening

---

    input data $\{(\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_T, \mathbf{y}_T)\}$; max output-dimension k; threshold α

    initialize weight matrices $\mathbf{W}_x$, $\mathbf{W}_y$, and $\mathbf{P}$.

    for $t = 1, 2, \ldots, T$ **do**

        $\mathbf{a}_t \leftarrow \mathbf{W}_x \mathbf{x}_t$; $\mathbf{b}_t \leftarrow \mathbf{W}_y \mathbf{y}_t$ $\triangleright$ projection of inputs

        run

            $\frac{d\mathbf{z}_t(\gamma)}{d\gamma} = \mathbf{a}_t + \mathbf{b}_t - \mathbf{P}\mathbf{n}_t(\gamma) - \alpha\mathbf{z}_t(\gamma)$ $\triangleright$ neural dynamics

            $\frac{d\mathbf{n}_t(\gamma)}{d\gamma} = \mathbf{P}^\top \mathbf{z}_t(\gamma) - \mathbf{n}_t(\gamma)$

        until convergence

        $\mathbf{W}_x \leftarrow \mathbf{W}_x + \eta(\mathbf{z}_t - \mathbf{a}_t)\mathbf{x}_t^\top$ $\triangleright$ synaptic updates

        $\mathbf{W}_y \leftarrow \mathbf{W}_y + \eta(\mathbf{z}_t - \mathbf{b}_t)\mathbf{y}_t^\top$

        $\mathbf{P} \leftarrow \mathbf{P} + \frac{\eta}{\tau}(\mathbf{z}_t \mathbf{n}_t^\top - \mathbf{P})$

    end for

---

compartment exceeds the corresponding inhibitory input (the source of which is explained below) it produces a calcium plateau potential that propagates through the basal dendrites, driving plasticity [45, 23, 5]. When the apical calcium plateau potential and basal dendritic current coincidentally arrive in the soma, they generate a burst in spiking output [24, 23, 5]. Inhibitory interneurons integrate pyramidal outputs and reciprocally inhibit the apical dendrites of pyramidal neurons, thus closing the loop.

We propose that a network of k pyramidal neurons implements CCA on the inputs received by apical and basal dendrites and outputs the whitened sum of CCSPs (Algorithm 3). In our model, each pyramidal neuron has three compartments — two compartments for the apical and basal dendritic currents, and one compartment for the somatic output. The two datasets $\mathbf{X}$ and $\mathbf{Y}$ are represented as activity vectors $\mathbf{x}_t$ and $\mathbf{y}_t$ streamed onto the apical and basal dendrites respectively, Figure 3. At each time step, the activity vectors are multiplied by the corresponding synaptic weights to yield localized apical and basal dendritic currents, $\mathbf{a}_t = \mathbf{W}_x\mathbf{x}_t$ and $\mathbf{b}_t = \mathbf{W}_y\mathbf{y}_t$, thus implementing projection onto the common subspace. This is followed by the following linear recurrent neural dynamics:

$$\frac{d\mathbf{z}_t(\gamma)}{d\gamma} = \mathbf{a}_t + \mathbf{b}_t - \mathbf{P}\mathbf{n}_t(\gamma) - \alpha\mathbf{z}_t(\gamma) \tag{20}$$

$$\frac{d\mathbf{n}_t(\gamma)}{d\gamma} = \mathbf{P}^\top\mathbf{z}_t(\gamma) - \mathbf{n}_t(\gamma), \tag{21}$$

where the components of $\mathbf{z}_t$ are represented by the spiking activity of pyramidal neurons, the components of $\mathbf{n}_t$ are represented by the activity of inhibitory interneurons, the components of $\mathbf{P}$ are represented by the synaptic weights from the interneurons to the pyramidal neurons, the components of $\mathbf{P}^\top$ are represented by the synaptic
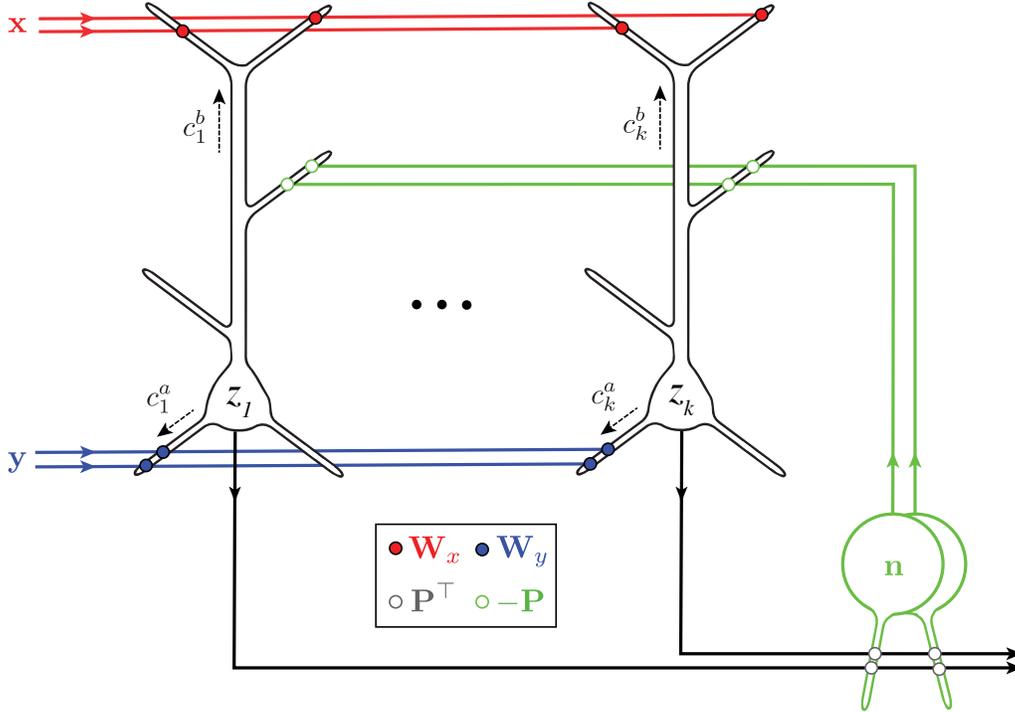
Figure 3: Cortical microcircuit implementation of Adaptive Bio-CCA with output whitening (Algorithm 3). The black cell bodies denote pyramidal neurons, with the apical tufts pointing upwards. The red and blue lines denote the axons respectively transmitting the apical input $\mathbf{x}$ and basal input $\mathbf{y}$. The black lines originating from the bases of the pyramidal neurons are their axons, which transmit their output $\mathbf{z}$. The green cell bodies denote the interneurons and the green lines are their axons, which transmit their output $\mathbf{n}$. Filled circles denote non-Hebbian synapses whose updates are proportional to the input (i.e., $\mathbf{x}$ or $\mathbf{y}$) and the weighted sum of the calcium plateau potential plus backpropagating somatic output [i.e., $\mathbf{c}^b + (1 - \alpha)\mathbf{z}$ or $\mathbf{c}^a + (1 - \alpha)\mathbf{z}$]. The directions of travel of these weighted sums are depicted using dashed lines with arrows. Empty circles denote Hebbian or anti-Hebbian synapses whose updates are proportional or inversely proportional to the pre- and post-synaptic activities.

17

weights from the pyramidal neurons to the interneurons, and α is the threshold parameter of the adaptive algorithm. These dynamics equilibrate at $\mathbf{n}_t = \mathbf{P}^\top \mathbf{z}_t$ and

$$\mathbf{z}_t = (\mathbf{P}\mathbf{P}^\top + \alpha \mathbf{I}_k)^{-1}(\mathbf{a}_t + \mathbf{b}_t). \tag{22}$$

Provided $\alpha > 0$, we can rearrange Equation (22) to write the output as

$$\mathbf{z}_t = \alpha^{-1}(\mathbf{b}_t + \mathbf{c}_t^a),$$

where the components of $\mathbf{c}_t^a := \mathbf{a}_t - \mathbf{P}\mathbf{n}_t$ are represented by the apical calcium plateau potentials within each pyramidal neuron. In other words, the output is proportional to the sum of the basal dendritic current and the apical calcium plateau potential, which is consistent with experimental evidence showing that the output depends on both the basal inputs and apical calcium plateau potential [5, 6, 28].

Next, we compare the synaptic update rules with experimental evidence. Rearranging Equation (22) and substituting into the synaptic update rules in Algorithm 3, we can rewrite the synaptic updates as follows:

$$\mathbf{W}_x \leftarrow \mathbf{W}_x + \eta \left(\mathbf{c}_t^b + (1-\alpha)\mathbf{z}_t \right) \mathbf{x}_t^\top$$
$$\mathbf{W}_y \leftarrow \mathbf{W}_y + \eta \left(\mathbf{c}_t^a + (1-\alpha)\mathbf{z}_t \right) \mathbf{y}_t^\top$$
$$\mathbf{P} \leftarrow \mathbf{P} + \frac{\eta}{\tau}(\mathbf{z}_t \mathbf{n}_t^\top - \mathbf{P}),$$

where the components of $\mathbf{c}_t^b := \mathbf{b}_t - \mathbf{P}\mathbf{n}_t$ are represented by basal calcium plateau potentials within each pyramidal neuron. The learning signal for the basal (resp. apical) synaptic updates of this circuit is the correlation between the sum of the apical (resp. basal) calcium plateau potentials plus the scaled spiking activity of the pyramidal neurons, $\mathbf{c}_t^b + (1-\alpha)\mathbf{z}_t$ [resp. $\mathbf{c}_t^a + (1-\alpha)\mathbf{z}_t$], and the synaptic inputs to the basal (resp. apical) dendrites, $\mathbf{x}_t$ (resp. $\mathbf{y}_t$). When $\alpha = 1$ the spiking (action potentials) of the post-synaptic neuron is not required for synaptic plasticity, whereas when $\alpha \neq 1$ the spiking of the post-synaptic neuron affects synaptic plasticity along with the calcium plateau. Therefore, our model can account for a range of experimental observations which have demonstrated that spiking in the post-synaptic neuron contributes to plasticity in some contexts [13, 11, 6, 28], but does not appear to affect plasticity in other contexts [46, 43]. Unlike Hebbian learning rules which depend only on the correlation of the spiking output of the post-synaptic neuron with the pre-synaptic spiking, the mechanisms involving the calcium plateau potential represented internally in a neuron are called non-Hebbian. Because synapses have access to both the corresponding presynaptic activity and to the calcium plateau potential, the learning rule remains local.

Note that the update rule for the synapses in the apical dendrites, $\mathbf{W}_x$, depend on the basal calcium plateau potentials $\mathbf{c}_t^b$. Experimental evidence is focused on apical calcium plateau potentials and it is not clear whether differences between

basal inputs and inhibitory signals generate calcium signals for driving plasticity in the apical dendrites. Alternatively, the learning rule for $\mathbf{W}_x$ coincides with the learning rule for the apical dendrites in [14], where a biological implementation in terms of local depolarization and backpropagating spikes was proposed. Due to the inconclusive evidence pertaining to plasticity in the apical tuft, we find it useful to put forth both interpretations.

Multi-compartmental models of pyramidal neurons have been invoked previously in the context of biological implementation of the backpropagation algorithm [21, 47, 16, 17, 42, 41]. Under this interpretation, the apical compartment represents the target output, the basal compartment represents the algorithm prediction and calcium plateau potentials communicate the error from the apical to the basal compartment, which is used for synaptic weight updates. The difference between these models and ours is that we use a normative approach to derive not only the learning rules but also the neural dynamics of the CCA algorithm ensuring that the output of the network is known for any input. On the other hand, the linearity of neural dynamics in our network means that stacking our networks will not lead to any nontrivial results expected of a deep learning architecture. We leave introducing nonlinearities into neural dynamics and stacking our network to future work.

We conclude this section with comments on the interneuron-to-pyramidal neuron synaptic weight matrix $\mathbf{P}$ and pyramidal neuron-to-interneuron synaptic weight matrix $\mathbf{P}^\top$, as well as the computational role of the interneurons in this network. First, the algorithm appears to require a weight sharing mechanism between the two sets of synapses to ensure the symmetry between the weight matrices, which is biologically unrealistic and commonly referred to as the weight transport problem. However, even without any initial symmetry between these feedforward and feedback synaptic weights, because of the symmetry of the local learning rules, the difference between the two will decay exponentially without requiring weight transport (see Appendix B.3). Second, in Equations (20)–(21), the interneuron-to-pyramidal neuron synaptic weight matrix $\mathbf{P}$ is preceded by a negative sign and the pyramidal neuron-to-interneuron synaptic weight matrix $\mathbf{P}^\top$ is preceded by a positive sign, which is consistent with the fact that, in simplified cortical microcircuits, interneuron-to-pyramidal neuron synapses are inhibitory whereas the pyramidal neuron-to-interneuron synapses are excitatory. That being said, this interpretation is superficial because the weight matrices are not constrained to be non-negative, which is due to the fact that we are implementing a linear statistical method. Imposing non-negativity constraints on the weights $\mathbf{P}$ and $\mathbf{P}^\top$ may be useful for implementing nonlinear statistical methods; however, this requires further investigation. Finally, the activities of the interneurons $\mathbf{N}$ were introduced in Equation (17) to decorrelate the output. This is consistent with previous models of the cortex (e.g., [20, 50]), which have introduced inhibitory interneurons to decorrelate excitatory outputs; however, in contrast to the current work, the models proposed in [20, 50] are not normative.

6 Numerical experiments

We now evaluate the performance of the online algorithms, Bio-CCA and Adaptive Bio-CCA with output whitening. In each plot, the lines and shaded regions respectively denote the means and 90% confidence intervals over 5 runs. Detailed descriptions of the implementations are given in Appendix C.1. All experiments were performed in Python on an iMac Pro equipped with a 3.2 GHz 8-Core Intel Xeon W CPU. The evaluation code is available at https://github.com/flatironinstitute/bio-cca.

6.1 Datasets

We first describe the evaluation datasets.

Synthetic. We generated a synthetic dataset with $T = 100,000$ samples according to the probabilistic model for CCA introduced by Bach and Jordan [3]. In particular, let $\mathbf{s}_1, \ldots, \mathbf{s}_T$ be i.i.d. 8-dimensional latent mean-zero Gaussian vectors with identity covariance. Let $\mathbf{T}_x \in \mathbb{R}^{50 \times 8}$, $\mathbf{T}_y \in \mathbb{R}^{30 \times 8}$, $\mathbf{\Psi}_x \in \mathcal{S}_{++}^{50}$ and $\mathbf{\Psi}_y \in \mathcal{S}_{++}^{30}$ be randomly generated matrices and define the 50-dimensional observations $\mathbf{x}_1, \ldots, \mathbf{x}_T$ and 30-dimensional observations $\mathbf{y}_1, \ldots, \mathbf{y}_T$ by

$$\mathbf{x}_t := \mathbf{T}_x \mathbf{s}_t + \boldsymbol{\phi}_t, \qquad \mathbf{y}_t := \mathbf{T}_y \mathbf{s}_t + \boldsymbol{\psi}_t, \qquad t = 1, \ldots, T,$$

where $\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_T$ (resp. $\boldsymbol{\psi}_1, \ldots, \boldsymbol{\psi}_T$) are i.i.d. 50-dimensional (resp. 30-dimensional) mean-zero Gaussian vectors with covariance $\mathbf{\Psi}_x$ (resp. $\mathbf{\Psi}_y$). Thus, conditioned on the latent random variable $\mathbf{s}$, the observation $\mathbf{x}$ (resp. $\mathbf{y}$) has a Gaussian distribution with mean $\mathbf{T}_x \mathbf{s}$ (resp. $\mathbf{T}_y \mathbf{s}$) with covariance $\mathbf{\Psi}_x$ (resp. $\mathbf{\Psi}_y$), i.e.,

$$\mathbf{x}_t | \mathbf{s}_t \sim \mathcal{N}(\mathbf{T}_x \mathbf{s}_t, \mathbf{\Psi}_x), \qquad \mathbf{y}_t | \mathbf{s}_t \sim \mathcal{N}(\mathbf{T}_y \mathbf{s}_t, \mathbf{\Psi}_y).$$

For this generative model, Bach and Jordan [3] showed that the posterior expectation of the latent vector $\mathbf{s}_t$ given the observation $(\mathbf{x}_t, \mathbf{y}_t)$ is a linear transformation of the sum of the 8-dimensional CCSPs $\mathbf{z}_t$; that is, $\mathbb{E}[\mathbf{s}_t | (\mathbf{x}_t, \mathbf{y}_t)] = \mathbf{L} \mathbf{z}_t$ for some 8×8 matrix $\mathbf{L}$. (To see this, set $M_1 = M_2 = P_d^{1/2}$ in the paragraph following [3, Theorem 2].) The first 10 canonical correlations are plotted in Figure 4 (left). Observe that the first 8 canonical correlations are close to 1 and the remaining canonical correlations are approximately 0. This sharp drop in the canonical correlations is a consequence of the linear generative model and is generally not the case in real data (see, e.g., the right panel in Figure 4). Still, we find it useful to test our algorithms on this synthetic dataset since the generative model is well studied and relevant to CCA [3].
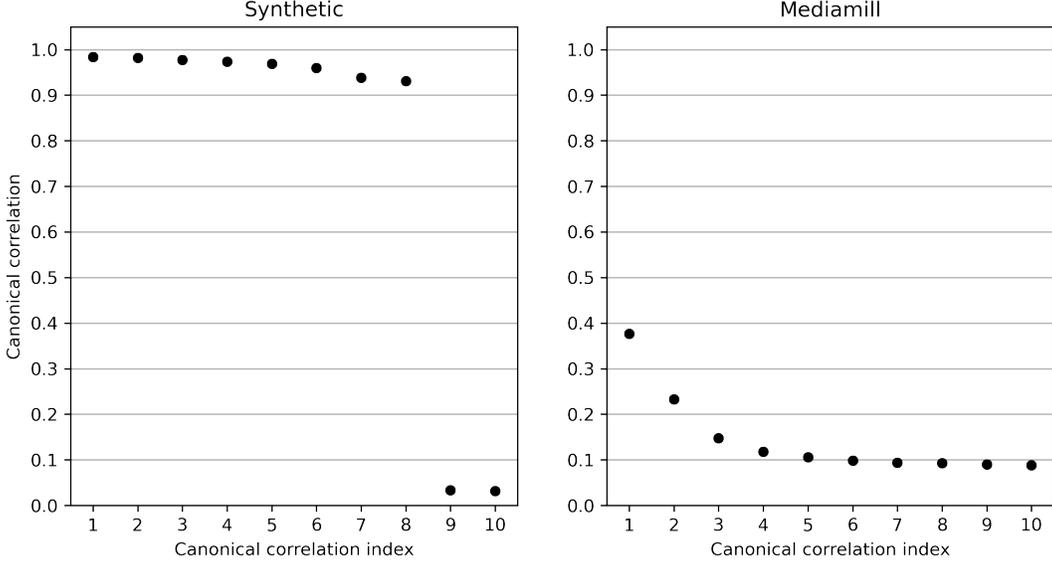
Figure 4: Top 10 canonical correlations $\rho_1, \dots, \rho_{10}$ of the synthetic dataset (left) and `Mediamill` (right).

Mediamill. The dataset `Mediamill` [44] consists of $T = 43,907$ samples (including training and testing sets) of video data and text annotations, and has been previously used to evaluate CCA algorithms [1, 37]. The first view consists of 120-dimensional visual features extracted from representative video frames. The second view consists of 101-dimensional vectors whose components correspond to manually labeled semantic concepts associated with the video frames (e.g., "basketball" or "tree"). To ensure that the problem is well-conditioned, we add Gaussian noise with covariance matrix $\varepsilon \mathbf{I}_{120}$ (resp. $\varepsilon \mathbf{I}_{101}$), for $\varepsilon = 0.1$, to the first (resp. second) view to generate the data matrix $\mathbf{X}$ (resp. $\mathbf{Y}$). The first 10 canonical correlations are plotted in Figure 4 (right).

Non-stationary. To evaluate Adaptive Bio-CCA with output whitening, we generated a non-stationary synthetic dataset with $T = 300,000$ samples, which are streamed from 3 distinct distributions that are generated according to the probabilistic model in [3]. In this case, the first $N = 100,000$ samples are generated from a 4-dimensional latent source, the second N samples are generated from an 8-dimensional latent source, and the final N samples are generated from a 1-dimensional latent source.

Specifically, we let $\mathbf{s}_1, \dots, \mathbf{s}_N$ (resp. $\mathbf{s}_{N+1}, \dots, \mathbf{s}_{2N}$ and $\mathbf{s}_{2N+1}, \dots, \mathbf{s}_T$) be i.i.d. 4-dimensional (resp. 8-dimensional and 1-dimensional) mean-zero Gaussian vectors with identity covariance. We then let $\mathbf{T}_x^{(1)} \in \mathbb{R}^{50 \times 4}$, $\mathbf{T}_x^{(2)} \in \mathbb{R}^{50 \times 8}$, $\mathbf{T}_x^{(3)} \in \mathbb{R}^{50 \times 1}$, $\mathbf{T}_y^{(1)} \in \mathbb{R}^{30 \times 4}$, $\mathbf{T}_y^{(2)} \in \mathbb{R}^{30 \times 8}$, $\mathbf{T}_y^{(3)} \in \mathbb{R}^{30 \times 1}$, $\mathbf{\Psi}_x \in \mathcal{S}_{++}^{50}$ and $\mathbf{\Psi}_y \in \mathcal{S}_{++}^{30}$ be randomly

21

generated matrices and define the 50-dimensional observations $\mathbf{x}_1, \ldots, \mathbf{x}_T$ and 30-dimensional observations $\mathbf{y}_1, \ldots, \mathbf{y}_T$ by

$$
\begin{aligned}
\mathbf{x}_t &:= \mathbf{T}_x^{(1)}\mathbf{s}_t + \boldsymbol{\phi}_t, & \mathbf{y}_t &:= \mathbf{T}_y^{(1)}\mathbf{s}_t + \boldsymbol{\psi}_t, & t &= 1, \ldots, N, \\
\mathbf{x}_t &:= \mathbf{T}_x^{(2)}\mathbf{s}_t + \boldsymbol{\phi}_t, & \mathbf{y}_t &:= \mathbf{T}_y^{(2)}\mathbf{s}_t + \boldsymbol{\psi}_t, & t &= N+1, \ldots, 2N, \\
\mathbf{x}_t &:= \mathbf{T}_x^{(3)}\mathbf{s}_t + \boldsymbol{\phi}_t, & \mathbf{y}_t &:= \mathbf{T}_y^{(3)}\mathbf{s}_t + \boldsymbol{\psi}_t, & t &= 2N+1, \ldots, T,
\end{aligned}
$$

where, as before, $\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_T$ (resp. $\boldsymbol{\psi}_1, \ldots, \boldsymbol{\psi}_T$) are i.i.d. 50-dimensional (resp. 30-dimensional) mean-zero Gaussian vectors with covariance $\boldsymbol{\Psi}_x$ (resp. $\boldsymbol{\Psi}_y$). In Figure 5, we plot the first 10 canonical correlations for each of the 3 distributions.
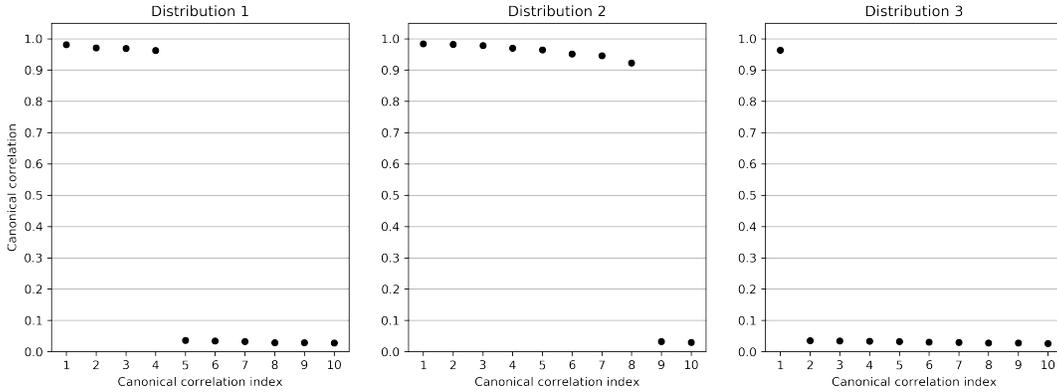

Figure 5: Top 10 canonical correlations $\rho_1, \ldots, \rho_{10}$ of the 3 distributions that contribute to the non-stationary synthetic dataset.

6.2 Bio-CCA

We now evaluate the performance of Bio-CCA (Algorithm 2) on the synthetic dataset and `Mediamill`.

Competing algorithms. We compare the performance of Bio-CCA with the following state-of-the-art online CCA algorithms:

- A two-time-scale algorithm for computing the top canonical correlation basis vectors (i.e., $k = 1$) introduced by Bhatia et al. [4]. The algorithm is abbreviated "Gen-Oja" due to its resemblance to Oja's method [33].

- An inexact Matrix Stochastic Gradient method for solving CCA, abbreviated "MSG-CCA", which was derived by Arora et al. [1].

- The asymmetric neural network proposed by Pehlevan et al. [37], which we abbreviate as "Asym-NN".

22

- The biologically plausible Reduced-Rank Regression algorithm derived by Golkar et al. [14], abbreviated "Bio-RRR", which implements a supervised version of CCA when $s = 1$ (see Algorithm 4 in Appendix B.2).

Detailed descriptions of the implementations of each algorithm are provided in Appendix C.1.

Performance metrics. To evaluate the performance of Bio-CCA, we use 2 performance metrics. The first performance metric is the following $[0, 2]$-valued normalized objective error function:

$$\text{Normalized objective error}(t) := \frac{\rho_{\max} - \text{Tr}(\mathbf{V}_{x,t}^\top \mathbf{C}_{xy} \mathbf{V}_{y,t})}{\rho_{\max}}. \qquad (23)$$

Here $\rho_{\max} := (\rho_1 + \cdots + \rho_k)/2$ is the optimal value of the CCA objective (1)–(2), and $(\mathbf{V}_{x,t}, \mathbf{V}_{y,t})$ are the basis vectors reported by the respective algorithm after iteration t, normalized to ensure they satisfy the orthonormality constraint (2). (We do not evaluate Bio-RRR using this metric because the algorithm only outputs one set of basis vectors.)

The second performance metric is the $(x-)$subspace error function is defined by

$$\text{Subspace error}(t) := \|\mathbf{V}_{x,t}(\mathbf{V}_{x,t}^\top \mathbf{V}_{x,t})^{-1}\mathbf{V}_{x,t}^\top - \overline{\mathbf{V}}_x(\overline{\mathbf{V}}_x^\top \overline{\mathbf{V}}_x)^{-1}\overline{\mathbf{V}}_x^\top\|_{\mathrm{F}}^2, \qquad (24)$$

where $\overline{\mathbf{V}}_x$ is the matrix of optimal basis vectors defined as in Equation (5). (We do not evaluate MSG-CCA using this metric because the algorithm outputs the product $\mathbf{V}_{x,t}\mathbf{V}_{y,t}^\top$ rather than outputting the basis vectors $\mathbf{V}_{x,t}$ and $\mathbf{V}_{y,t}$ separately.)

Evaluation on the synthetic dataset. In Figure 6 we plot the performance of Bio-CCA, in terms of both sample and runtime efficiency, against the competing algorithms for target dimensions $k = 1, 2, 4$ on the synthetic dataset, presented once in a randomly permuted order. For $k = 1$, Gen-Oja initially outperforms Bio-CCA in sample and runtime efficiency; however Bio-CCA eventually performs comparably with Gen-Oja when given sufficiently many samples (and outperforms Gen-Oja in terms of subspace error). For $k = 2, 4$, the sample and runtime efficiency of Bio-CCA outperforms the other competing algorithms. The MSG-CCA error does not begin to decay until the 10^3 iteration because the first 10^3 samples are used to obtain initial estimates of the covariance matrices $\mathbf{C}_{xx}$ and $\mathbf{C}_{yy}$. The poor performance of the asymmetric network [37] is due in part to the fact that the algorithm depends on the gaps between the canonical correlations, which are small for the synthetic dataset. In Figure 11 of Appendix C.2, we verify that the Bio-CCA basis vectors asymptotically satisfy the orthonormality constraint (2).
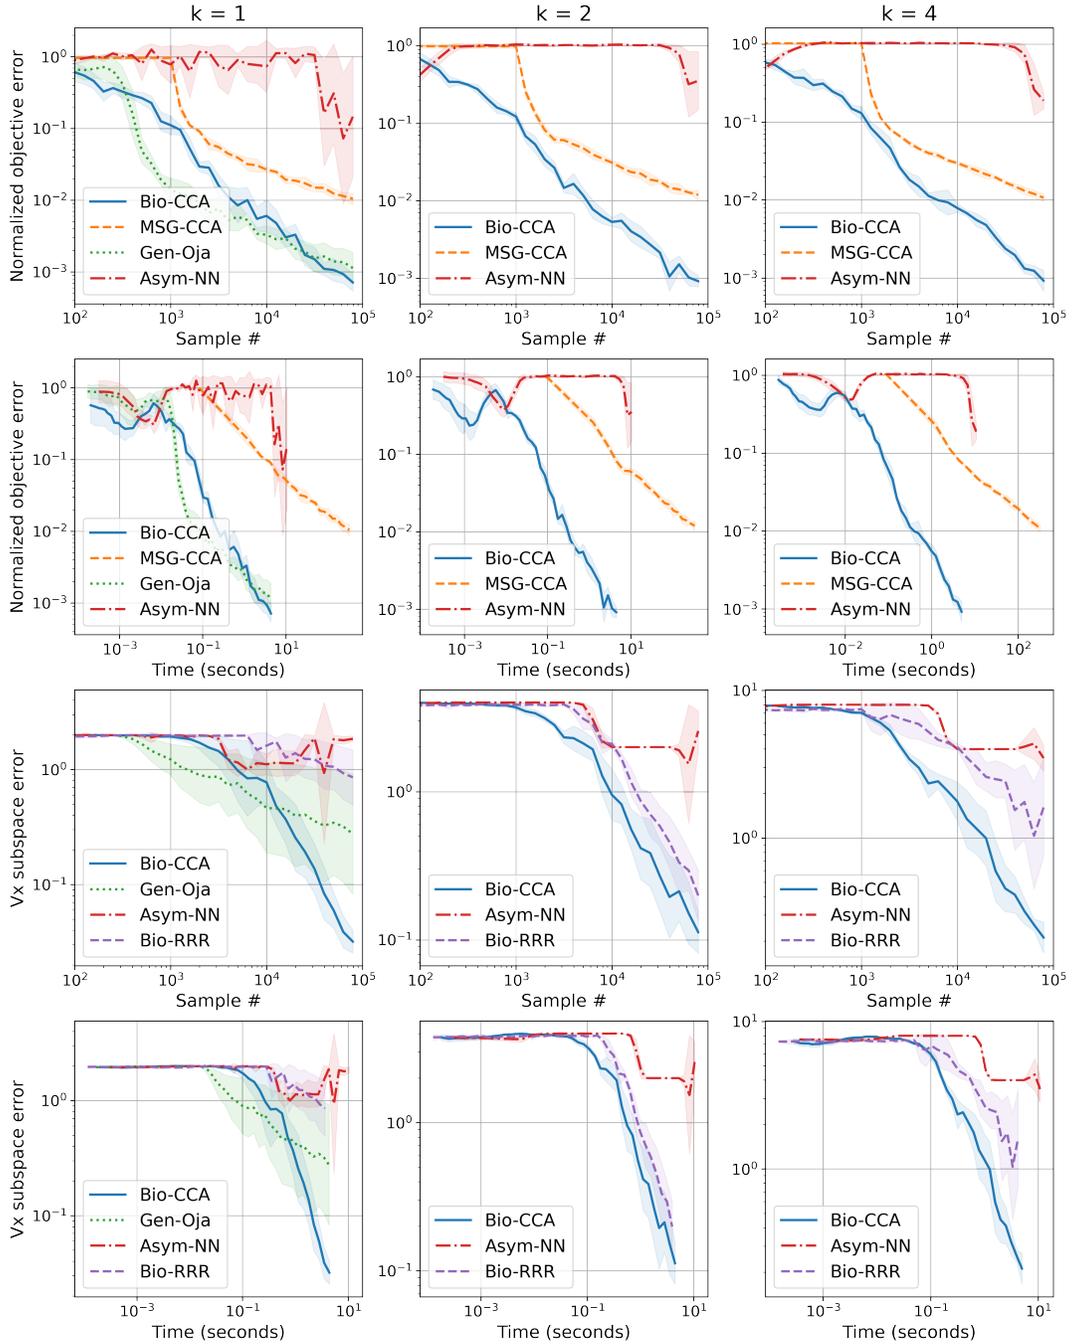
Figure 6: Comparisons of Bio-CCA (Algorithm 2) with the competing algorithms on the synthetic dataset, for $k = 1, 2, 4$, in terms of the normalized objective error defined in Equation (23) as a function of sample number and runtime (top two rows), and in terms of the subspace error defined in Equation (24) as a function of sample number and runtime (bottom two rows).

Evaluation on Mediamill. In Figure 7 we plot the performance of Bio-CCA, in terms of both sample and runtime efficiency, against the competing algorithms for target dimensions $k = 1, 2, 4$ on `Mediamill`, presented 3 times with a randomly permuted order in each presentation. When tested on `Mediamill`, the sample and runtime efficiency of Bio-CCA outperform the competing algorithms (for $k = 1$, Bio-RRR performs comparably with Bio-CCA). In Figure 12 of Appendix C.2, we verify that the Bio-CCA basis vectors asymptotically satisfy the orthonormality constraint (2).

6.3 Adaptive Bio-CCA with output whitening

Next, we evaluate the performance of Adaptive Bio-CCA with output whitening (Algorithm 3) on all 3 datasets. Since we are unaware of competing online algorithms for adaptive CCA with output whitening, to compare the performance of Algorithm 3 to existing methods, we also plot the performance of Bio-RRR [14] with respect to subspace error, where we a priori select the target dimension. We chose Bio-RRR because the algorithm also maps onto a neural network that resembles the cortical microcircuit and because it performs relatively well on the synthetic dataset and `Mediamill`.

Performance metric. To evaluate the performance of Adaptive Bio-CCA with output whitening, we use a subspace error metric. Recall that the target output rank of Algorithm 3, denoted r, is equal to the number of canonical correlations $\rho_1, \ldots, \rho_k$ that exceed $\max(\alpha - 1, 0)$, i.e.,

$$r := \max\{1 \leq i \leq k : \rho_i > \max(\alpha - 1, 0)\}. \tag{25}$$

Since the target dimension is often less than k, the subspace error defined in Equation (24) is not an appropriate metric because the projection matrices are rank k. Rather, we evaluate Adaptive Bio-CCA with output whitening using the *adaptive* subspace error function defined by

$$\text{Adaptive subspace error}(t) := \|\widetilde{\mathbf{U}}_{x,t}\widetilde{\mathbf{U}}_{x,t}^{\top} - \widetilde{\mathbf{V}}_x(\widetilde{\mathbf{V}}_x^{\top}\widetilde{\mathbf{V}}_x)^{-1}\widetilde{\mathbf{V}}_x^{\top}\|_{\mathrm{F}}^2, \tag{26}$$

where $\widetilde{\mathbf{U}}_{x,t}$ is the $m \times r$ matrix whose column vectors are the top r right-singular vectors of $\mathbf{W}_{x,t}$, and $\widetilde{\mathbf{V}}_x$ is the $m \times r$ matrix whose i^{th} column vector is equal to the i^{th} column vector of the matrix $\overline{\mathbf{V}}_x$ defined in Equation (5), for $i = 1, \ldots, r$. (The covariance matrices used in the definition for $\overline{\mathbf{V}}_x$ are those associated with the distribution that is being streamed at iteration t.) If $r = k$, then the error aligns with the subspace error defined in Equation (24). Since the output dimension of Bio-RRR is a priori set to r, the performance of Bio-RRR is evaluated using the subspace error defined in Equation (24).
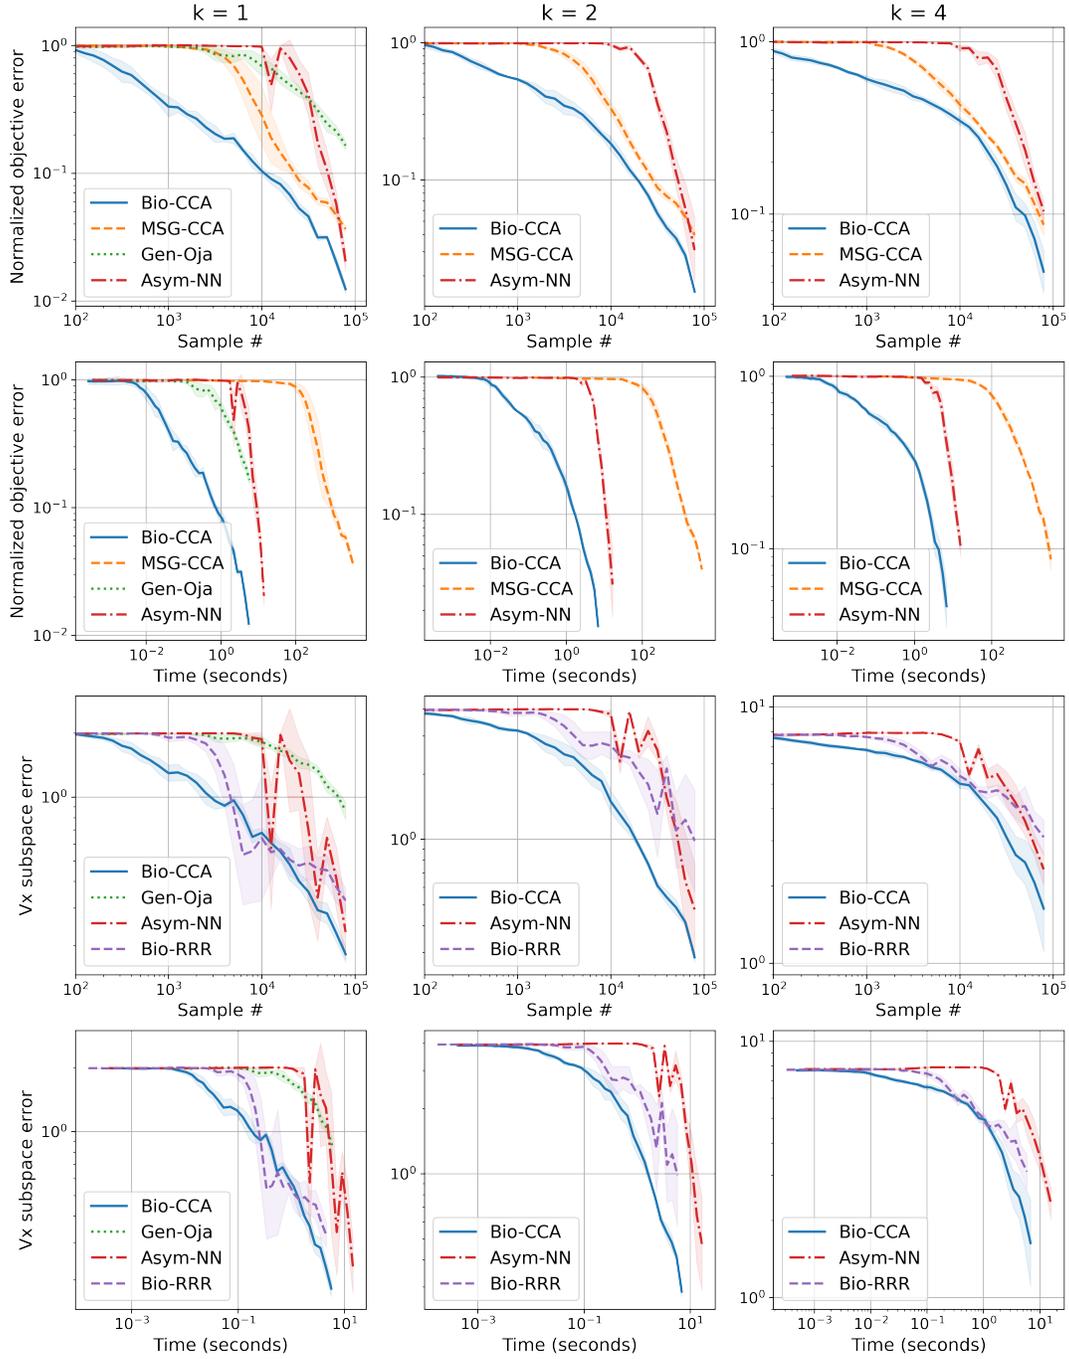
Figure 7: Comparisons of Bio-CCA (Algorithm 2) with the competing algorithms on `Mediamill`, for $k = 1, 2, 4$, in terms of the normalized objective error defined in Equation (23) as a function of sample number and runtime (top two rows), and in terms of the subspace error defined in Equation (24) as a function of sample number and runtime (bottom two rows).
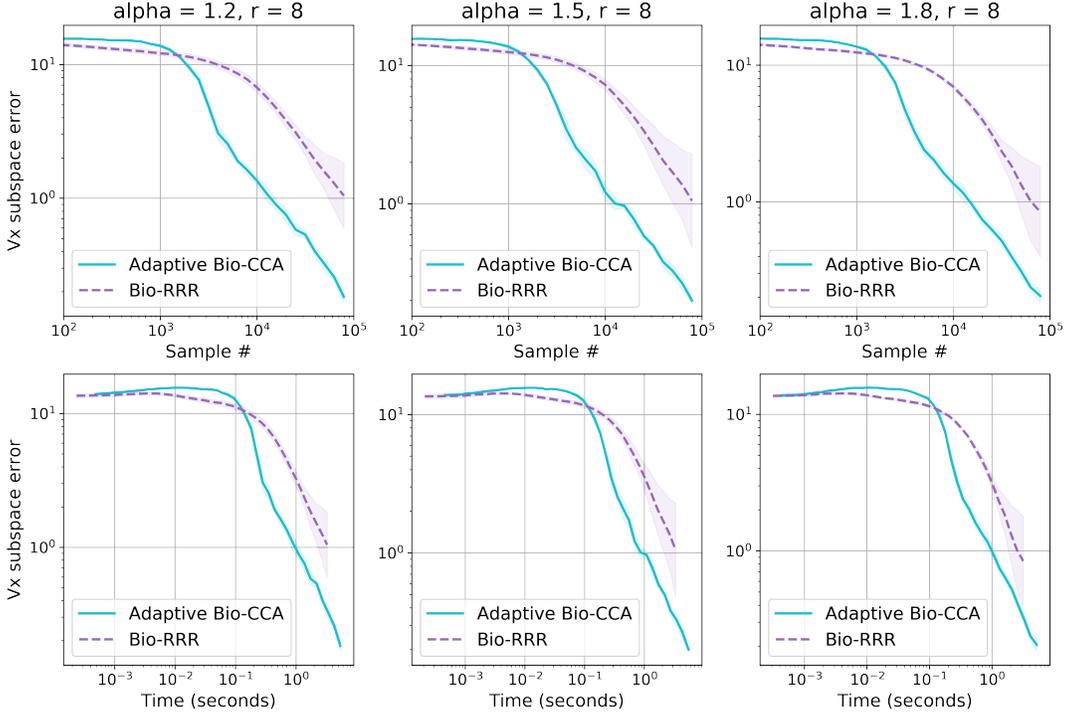
Figure 8: Comparisons of Adaptive Bio-CCA with output whitening (Algorithm 3) and Bio-RRR on the synthetic dataset, for $\alpha = 1.2, 1.5, 1.8$, in terms of the subspace error as a function of sample number (top row) and runtime (bottom row). For each α, r is the target output rank defined as in Equation (25).

Evaluation on the synthetic dataset. From Figure 4 (left) we see that the first 8 canonical correlations are close to 1 and the remaining canonical correlations are approximately 0. Therefore, for $k \geq 8$ and $\alpha \in (1.1, 1.9)$, Algorithm 3 should project the inputs $\mathbf{x}_t$ and $\mathbf{y}_t$ onto the 8-dimensional subspace spanned by the top 8 pairs of canonical correlation basis vectors, and set the non-zero output covariance eigenvalues to one. In Figure 8 we plot the performance of Algorithm 3 with $k = 10$ for $\alpha = 1.2, 1.5, 1.8$ on the synthetic dataset (presented once with a randomly permuted order). We see that Adaptive Bio-CCA with output whitening outperforms Bio-RRR, even though the target output dimension of Bio-RRR was set a priori. In Figure 13 of Appendix C.2, we verify that the Adaptive Bio-CCA with output whitening basis vectors asymptotically satisfy the whitenening constraint.

Evaluation on Mediamill. From Figure 4 (right) we see that the canonical correlations of the dataset `Mediamill` exhibit a more gradual decay than the canonical correlations of the synthetic dataset. As we increase the threshold α in the interval $(1.1, 1.4)$, the rank of the output of Algorithm 3 decreases. In Figure 9 we plot
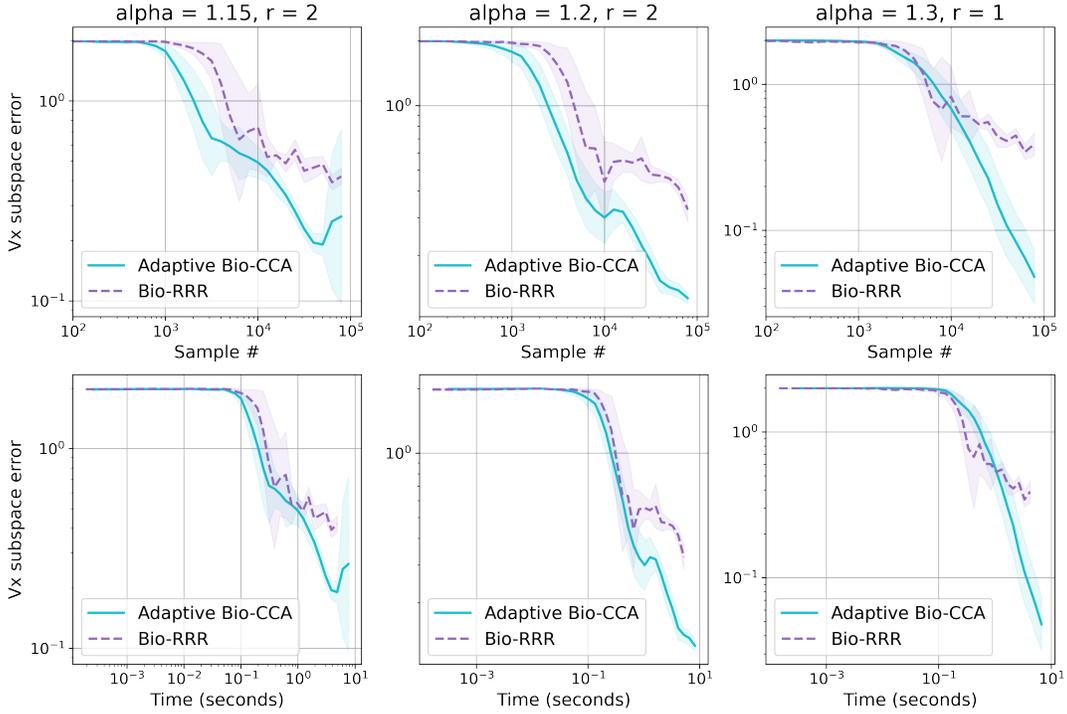
Figure 9: Comparisons of Adaptive Bio-CCA with output whitening (Algorithm 3) with Bio-RRR on the dataset `Mediamill`, for $\alpha = 1.2, 1.5, 1.8$, in terms of the subspace error as a function of sample number (top row) and runtime (bottom row). For each α, r is the target output rank defined as in Equation (25).

the performance of Algorithm 3 with $k = 10$ for $\alpha = 1.15, 1.2, 1.3$ on `Mediamill` (presented three times with a randomly permuted order in each presentation). As with the synthetic dataset, we find that Adaptive Bio-CCA with output whitening outperforms Bio-RRR. In Figure 14 of Appendix C.2, we verify that the Adaptive Bio-CCA with output whitening basis vectors asymptotically satisfy the whitening constraint.

Evaluation on the non-stationary dataset. From Figure 5, we see that the distributions that contribute to the non-stationary dataset all have canonical correlations that are close to 0 or 1. The first distribution has 4 canonical correlations close to 1, the second distribution has 8 canonical correlations close to 1 and the third distribution has 1 canonical correlation close to 1. Therefore, for $k \geq 8$ and $\alpha \in (1.1, 1.9)$, Algorithm 3 should initially (for the first 100,000 inputs) project the inputs $\mathbf{x}_t$ and $\mathbf{y}_t$ onto a 4-dimensional subspaces, then (for the second 100,000 inputs) project the inputs onto an 8-dimensional subspaces, and finally (for the final 100,000 inputs) project the inputs onto a 1-dimensional subspaces. In Figure 10 (left) we plot the performance of Algorithm 3 with $k = 10$ and $\alpha = 1.5$ on
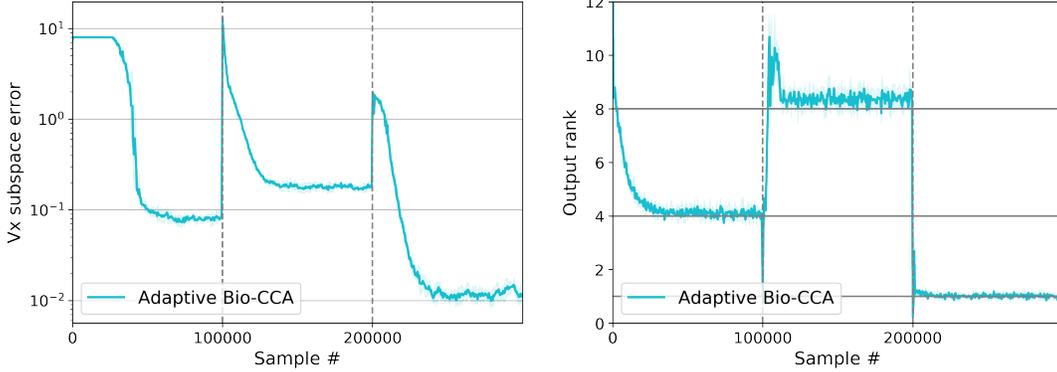
Figure 10: On the left we plot the performance of Adaptive Bio-CCA with output whitening (with $\alpha = 1.5$) on the non-stationary synthetic dataset in terms of the subspace error defined in Equation (26). On the right we plot the output rank defined in Equation (27).

the non-stationary dataset in terms of the adaptive subspace error.[2] Note that the adaptive subspace error spikes each time the algorithm is streamed inputs from a new distribution (these epochs are denoted by the gray vertical dashed lines) before decaying. In Figure 10 (right) we plot the output rank of Adaptive Bio-CCA with output whitening defined by

$$\text{Output rank}(t) = \text{Tr}(\mathbf{V}_{x,t}^{\top}\mathbf{C}_{xx}\mathbf{V}_{x,t} + \mathbf{V}_{x,t}^{\top}\mathbf{C}_{xy}\mathbf{V}_{y,t} + \mathbf{V}_{y,t}^{\top}\mathbf{C}_{xy}^{\top}\mathbf{V}_{x,t} + \mathbf{V}_{y,t}^{\top}\mathbf{C}_{yy}\mathbf{V}_{y,t}),$$
(27)

where, at each iteration t, the covariance matrices correspond to the distribution from which $(\mathbf{x}_t, \mathbf{y}_t)$ is streamed. Note that the output rank adapts each time the algorithm is streamed inputs from a new distribution to match the dimension of the latent variable generating the distribution. In particular, we see that the algorithm is able to quickly adapt its output rank to new distributions.

7 Discussion

In this work, we derived an online algorithm for CCA that can be implemented in a neural network with multi-compartmental neurons and local, non-Hebbian learning rules. We also derived an extension that adaptively chooses the output rank and whitens the output. Remarkably, the neural architecture and non-Hebbian learning rules of our extension resembled neural circuitry and non-Hebbian plasticity in cortical pyramidal neurons. Thus, our neural network model may be useful for

---

[2]Since the competing algorithms are not adaptive and need to have their output dimension set by hand, we do not include a competing algorithm for comparison.

understanding the computational role of multi-compartmental neurons with non-Hebbian plasticity.

While our neural network model captures salient features of cortical microcircuits, there are important biophysical properties that are not explained by our model. First, our model uses linear neurons to solve the linear CCA problem, which substantially limits its computational capabilities and is a major simplification of cortical pyramidal neurons which can perform nonlinear operations [12]. However, studying the analytically tractable and interpretable linear neural network model is useful for understanding more complex nonlinear models. Such an approach has proven successful for studying deep networks in the machine learning literature [2]. In future work, we plan to incorporate nonlinear neurons in our model.

Second, our neural network implementation requires the same number of interneurons and principal neurons, whereas in the cortex there are approximately 4 times more pyramidal neurons than interneurons [24]. In our model, the interneurons decorrelate the output and, in practice, the optimal fixed points of the algorithm can destabilize when there are fewer interneurons than principal neurons (see Remark 3A in the supplementary material of [34]). In biological circuits, these instabilities could be mitigated by other biophysical constraints; however, a theoretical justification would require additional work.

Third, the output of our neural network is the equally weighted sum of the basal and apical projections. However, experimental evidence suggests that the pyramidal neurons integrate their apical and basal inputs asymmetrically [25, 23, 29]. In addition, in our model, the apical learning rule is non-Hebbian and depends on a calcium plateau potential that travels from the basal dendrites to the apical tuft. Experimental evidence for calcium plateau potential dependent plasticity is focused on the basal dendrites, with inconclusive evidence on the plasticity rules for the apical dendrites [13, 43].

To provide an alternative explanation of cortical computation, in a separate work [14], we derive an online algorithm for the general *supervised* learning method Reduced-Rank Regression, which includes CCA as a special case (see Appendix B.2 for a detailed comparison of the two algorithms). The algorithm also maps onto a neural network with multi-compartmental neurons and non-Hebbian plasticity in the basal dendrites. Both models adopt a normative approach in which the algorithms are derived from principled objective functions. This approach is highly instructive as the differences between the models highlight which features of the network that are central to implementing an unsupervised learning method versus a supervised learning method.

There are three main differences between the biological interpretation of the two algorithms. First, the output of the network in [14] is the projection of the basal inputs, with no apical contribution. Second, the network in [14] allows for a range of apical synaptic update rules, including Hebbian updates. Third, the adaptive network derived here includes a threshold parameter α, which adaptively

sets the output dimension and is not included in [14]. In our model, this parameter corresponds the contribution of the somatic output to plasticity in the basal dendrites. These differences can be compared to experimental outcomes to provide evidence that cortical microcircuits implement unsupervised algorithms, supervised algorithms, or mixtures of both. Thus, we find it informative to put forth and contrast the two models.

Finally, we did not prove theoretical guarantees that our algorithms converge. As we show in Appendix D, Offline-CCA and Bio-CCA can be viewed as gradient descent-ascent and stochastic gradient descent-ascent algorithms for solving a nonconvex-concave min-max problem. While gradient descent-ascent algorithms are natural methods for solving such min-max problems, they are not always guaranteed to converge to a desired solution. In fact, when the gradient descent step size not sufficiently small relative to the gradient ascent step size (i.e., when τ is not sufficiently small), gradient descent-ascent algorithms for solving nonconvex-concave min-max problems can converge to limit cycles [18, 30]. Establishing local or global convergence, and convergence rate guarantees for general gradient descent-ascent algorithms is an active area of research, and even recent advances [26] impose assumptions that are not satisfied in our setting. In Appendix D, we discuss these challenges and place our algorithms within the broader context of gradient descent-ascent algorithms for solving nonconvex-concave min-max problems.

Acknowledgements

A Sums of CCSPs as principal subspace projections

In this section, we show that the sums of the CCPSs $\mathbf{z}_1, \dots, \mathbf{z}_T$ can be expressed as the principal subspace projection of the whitened concatenated data $\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_T$, defined by

$$\boldsymbol{\xi}_t := \begin{bmatrix} \mathbf{C}_{xx}^{-1/2} \mathbf{x}_t \\ \mathbf{C}_{yy}^{-1/2} \mathbf{y}_t \end{bmatrix}. \tag{28}$$

To this end, recall the CCA objective

$$\underset{\mathbf{V}_x \in \mathbb{R}^{m \times k}, \mathbf{V}_y \in \mathbb{R}^{n \times k}}{\arg\max} \operatorname{Tr}\left(\mathbf{V}_x^\top \mathbf{C}_{xy} \mathbf{V}_y\right)$$

subject to the whitening constraint $\mathbf{V}_x^\top \mathbf{C}_{xx} \mathbf{V}_x + \mathbf{V}_y^\top \mathbf{C}_{yy} \mathbf{V}_y = \mathbf{I}_k$. We can rewrite the CCA objective as a generalized eigenproblem:

$$\underset{\mathbf{V}_x \in \mathbb{R}^{m \times k}, \mathbf{V}_y \in \mathbb{R}^{n \times k}}{\arg\max} \operatorname{Tr} \begin{bmatrix} \mathbf{V}_x \\ \mathbf{V}_y \end{bmatrix}^\top \begin{bmatrix} & \mathbf{C}_{xy} \\ \mathbf{C}_{xy}^\top & \end{bmatrix} \begin{bmatrix} \mathbf{V}_x \\ \mathbf{V}_y \end{bmatrix}, \tag{29}$$

subject to the whitening constraint

$$\begin{bmatrix} \mathbf{V}_x \\ \mathbf{V}_y \end{bmatrix}^\top \begin{bmatrix} \mathbf{C}_{xx} & \\ & \mathbf{C}_{yy} \end{bmatrix} \begin{bmatrix} \mathbf{V}_x \\ \mathbf{V}_y \end{bmatrix} = \mathbf{I}_k. \tag{30}$$

The equivalence can be readily verified by expanding the matrix products in Equations (29)–(30).

Next, we transform this generalized eigenproblem formulation (29)–(30) into a standard eigenproblem formulation. Since the trace of the left hand side of Equation (30) is constrained to equal k, we can add it to the trace in Equation (29), without affecting the output of the argmax, to arrive at the CCA objective:

$$\underset{\mathbf{V}_x \in \mathbb{R}^{m \times k}, \mathbf{V}_y \in \mathbb{R}^{n \times k}}{\arg\max} \operatorname{Tr} \begin{bmatrix} \mathbf{V}_x \\ \mathbf{V}_y \end{bmatrix}^\top \begin{bmatrix} \mathbf{C}_{xx} & \mathbf{C}_{xy} \\ \mathbf{C}_{xy}^\top & \mathbf{C}_{yy} \end{bmatrix} \begin{bmatrix} \mathbf{V}_x \\ \mathbf{V}_y \end{bmatrix} \tag{31}$$

subject to the whitening constraint in Equation (30).

Our final step is a substitution of variables. Define the $d \times k$ matrix

$$\mathbf{V}_\xi := \begin{bmatrix} \mathbf{C}_{xx}^{1/2} \mathbf{V}_x \\ \mathbf{C}_{yy}^{1/2} \mathbf{V}_y \end{bmatrix}. \tag{32}$$

After substituting into Equations (31) and (30), we see that $(\overline{\mathbf{V}}_x, \overline{\mathbf{V}}_y)$ is a solution of the CCA objective if and only if $\overline{\mathbf{V}}_\xi$ is a solution of:

$$\underset{\mathbf{V}_\xi \in \mathbb{R}^{d \times k}}{\arg\max} \operatorname{Tr} \mathbf{V}_\xi^\top \mathbf{C}_{\xi\xi} \mathbf{V}_\xi \qquad \text{subject to} \qquad \mathbf{V}_\xi^\top \mathbf{V}_\xi = \mathbf{I}_k, \tag{33}$$

where we recall that $\mathbf{C}_{\xi\xi}$ is the covariance matrix defined by $\mathbf{C}_{\xi\xi} := \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{\xi}_t \boldsymbol{\xi}_t^\top$. Importantly, Equation (33) is the variance maximization objective for the standard PCA eigenproblem, which is optimized when the column vectors of $\mathbf{V}_\xi$ form an orthonormal basis spanning the k-dimensional principal subspace of the data $\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_T$. Therefore, by Equations (28) and (32), the projection of the vector $\boldsymbol{\xi}_t$ onto its principal subspace is precisely the desired sum of CCSPs:

$$\mathbf{z}_t := \overline{\mathbf{V}}_x^\top \mathbf{x}_t + \overline{\mathbf{V}}_y^\top \mathbf{y}_t = \overline{\mathbf{V}}_\xi^\top \boldsymbol{\xi}_t.$$

B Adaptive Bio-CCA with output whitening

B.1 Detailed derivation of Algorithm 3

Recall the min-max objective given in Equation (19):

$$\underset{\mathbf{Z}\in\mathbb{R}^{k\times T}}{\arg\min}\ \underset{\mathbf{N}\in\mathbb{R}^{k\times T}}{\max}\ \underset{\mathbf{W}_x\in\mathbb{R}^{k\times m}}{\min}\ \underset{\mathbf{W}_y\in\mathbb{R}^{k\times n}}{\min}\ \underset{\mathbf{P}\in\mathbb{R}^{k\times k}}{\max}\ \widetilde{L}(\mathbf{W}_x,\mathbf{W}_y,\mathbf{P},\mathbf{Z},\mathbf{N}), \qquad (34)$$

where

$$
\begin{aligned}
\widetilde{L}(\mathbf{W}_x,\mathbf{W}_y,\mathbf{P},\mathbf{Z},\mathbf{N}) := & \frac{1}{T}\operatorname{Tr}\left(-2\mathbf{Z}^\top\mathbf{W}_x\mathbf{X} - 2\mathbf{Z}^\top\mathbf{W}_y\mathbf{Y} + \alpha\mathbf{Z}^\top\mathbf{Z}\right) \\
& + \frac{1}{T}\operatorname{Tr}\left(2\mathbf{N}^\top\mathbf{P}^\top\mathbf{Z} - \mathbf{N}^\top\mathbf{N}\right) \\
& + \operatorname{Tr}\left(\mathbf{W}_x\mathbf{C}_{xx}\mathbf{W}_x^\top + \mathbf{W}_y\mathbf{C}_{yy}\mathbf{W}_y^\top - \mathbf{P}\mathbf{P}^\top\right).
\end{aligned}
$$

Since $\widetilde{L}(\mathbf{W}_x,\mathbf{W}_y,\mathbf{P},\mathbf{Z},\mathbf{N})$ is strictly convex in $\mathbf{W}_x$, $\mathbf{W}_y$ and $\mathbf{Z}$, and strictly concave in $\mathbf{P}$ and $\mathbf{N}$, we can interchange the order of optimization to obtain:

$$\underset{\mathbf{W}_x\in\mathbb{R}^{k\times m}}{\min}\ \underset{\mathbf{W}_y\in\mathbb{R}^{k\times n}}{\min}\ \underset{\mathbf{P}\in\mathbb{R}^{k\times k}}{\max}\ \underset{\mathbf{Z}\in\mathbb{R}^{k\times T}}{\min}\ \underset{\mathbf{N}\in\mathbb{R}^{k\times T}}{\max}\ \widetilde{L}(\mathbf{W}_x,\mathbf{W}_y,\mathbf{P},\mathbf{Z},\mathbf{N}). \qquad (35)$$

The interchange of the maximization with respect to $\mathbf{N}$ and the minimization with respect to $\mathbf{W}_x$ and $\mathbf{W}_y$ is justified by the fact that, for fixed $\mathbf{Z}$ and $\mathbf{P}$, $\widetilde{L}(\mathbf{W}_x,\mathbf{W}_y,\mathbf{P},\mathbf{Z},\mathbf{N})$ is strictly convex in $(\mathbf{W}_x,\mathbf{W}_y)$ and strictly concave in $\mathbf{N}$. Similarly, the interchange of the minimization with respect to $\mathbf{Z}$ and the maximization with respect to $\mathbf{P}$ is justified by the fact that, for fixed $\mathbf{N}$, $\mathbf{W}_x$ and $\mathbf{W}_y$, $\widetilde{L}(\mathbf{W}_x,\mathbf{W}_y,\mathbf{P},\mathbf{Z},\mathbf{N})$ is convex in $\mathbf{Z}$ and strictly concave in $\mathbf{P}$. In order to derive an online algorithm, we write the objective in terms of time-separable terms:

$$\widetilde{L}(\mathbf{W}_x,\mathbf{W}_y,\mathbf{P},\mathbf{Z},\mathbf{N}) = \frac{1}{T}\sum_{t=1}^{T}\widetilde{l}_t(\mathbf{W}_x,\mathbf{W}_y,\mathbf{P},\mathbf{z}_t,\mathbf{n}_t),$$

where

$$
\begin{aligned}
\widetilde{l}_t(\mathbf{W}_x,\mathbf{W}_y,\mathbf{P},\mathbf{z}_t,\mathbf{n}_t) := & -2\mathbf{z}_t^\top\mathbf{W}_x\mathbf{x}_t - 2\mathbf{z}_t^\top\mathbf{W}_y\mathbf{y}_t + \alpha\mathbf{z}_t^\top\mathbf{z}_t + 2\mathbf{n}_t^\top\mathbf{P}^\top\mathbf{z}_t - \mathbf{n}_t^\top\mathbf{n}_t \\
& + \operatorname{Tr}\left(\mathbf{W}_x\mathbf{x}_t\mathbf{x}_t^\top\mathbf{W}_x^\top + \mathbf{W}_y\mathbf{y}_t\mathbf{y}_t^\top\mathbf{W}_y^\top - \mathbf{P}\mathbf{P}^\top\right).
\end{aligned}
$$

At each time step t, for fixed $\mathbf{W}_x$, $\mathbf{W}_y$ and $\mathbf{M}$, we first simultaneously maximize the objective function $\widetilde{l}_t(\mathbf{W}_x,\mathbf{W}_y,\mathbf{P},\mathbf{z}_t,\mathbf{n}_t)$ over $\mathbf{n}_t$ and minimize over $\mathbf{z}_t$ by running the fast gradient descent-ascent dynamics in Eqs. (20)–(21) to equilibrium. Since $\widetilde{l}_t(\mathbf{W}_x,\mathbf{W}_y,\mathbf{P},\mathbf{z}_t,\mathbf{n}_t)$ is convex in $\mathbf{z}_t$ and concave in $\mathbf{n}_t$, these fast dynamics equilibriate at the saddle point where $\bar{\mathbf{z}}_t = (\mathbf{P}\mathbf{P}^\top + \alpha\mathbf{I}_k)^{-1}(\mathbf{a}_t + \mathbf{b}_t)$ and $\bar{\mathbf{n}}_t = \mathbf{P}^\top\mathbf{z}_t$.

Then, with $(\overline{\mathbf{n}}_t, \overline{\mathbf{z}}_t)$ fixed, we perform a gradient descent-ascent step of the objective function with respect to $(\mathbf{W}_x, \mathbf{W}_y)$ and $\mathbf{P}$:

$$\mathbf{W}_x \leftarrow \mathbf{W}_x - \frac{\eta}{2} \frac{\partial \widetilde{l}_t(\mathbf{W}_x, \mathbf{W}_y, \mathbf{P}, \overline{\mathbf{z}}_t, \overline{\mathbf{n}}_t)}{\partial \mathbf{W}_x}$$

$$\mathbf{W}_x \leftarrow \mathbf{W}_y - \frac{\eta}{2} \frac{\partial \widetilde{l}_t(\mathbf{W}_x, \mathbf{W}_y, \mathbf{P}, \overline{\mathbf{z}}_t, \overline{\mathbf{n}}_t)}{\partial \mathbf{W}_y}$$

$$\mathbf{P} \leftarrow \mathbf{P} + \frac{\eta}{2\tau} \frac{\partial \widetilde{l}_t(\mathbf{W}_x, \mathbf{W}_y, \mathbf{P}, \overline{\mathbf{z}}_t, \overline{\mathbf{n}}_t)}{\partial \mathbf{P}}.$$

Substituting in with the partial derivatives yields Algorithm 3.

B.2   Comparison with Bio-RRR

In this section, we compare Adaptive Bio-CCA with output whitening (Algorithm 3) and Bio-RRR ([14, Algorithm 2]). We first state the Bio-RRR algorithm.[3]

Algorithm 4: Bio-RRR

    input data $\{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_T, \mathbf{y}_T)\}$; max output-dimension k; weight $0 \le s \le 1$
    initialize weight matrices $\mathbf{W}_x$, $\mathbf{W}_y$, and $\mathbf{P}$.
    for $t = 1, 2, \dots, T$ **do**
        $\mathbf{a}_t \leftarrow \mathbf{W}_x \mathbf{x}_t$; $\quad \mathbf{z}_t \leftarrow \mathbf{W}_y \mathbf{y}_t$; $\quad \mathbf{n}_t \leftarrow \mathbf{P}^\top \mathbf{z}_t$
        $\mathbf{W}_x \leftarrow \mathbf{W}_x + \eta_x(\mathbf{z}_t \mathbf{x}_t^\top - s\mathbf{a}_t \mathbf{x}_t^\top - (1-s)\mathbf{W}_x)$
        $\mathbf{W}_y \leftarrow \mathbf{W}_y + \eta_y(\mathbf{a}_t - \mathbf{P}\mathbf{n}_t)\mathbf{y}_t^\top$
        $\mathbf{P} \leftarrow \mathbf{P} + \eta_p(\mathbf{z}_t \mathbf{n}_t^\top - \mathbf{P})$
    end for

Bio-RRR implements a supervised learning method for minimizing reconstruction error, with the parameter $0 \le s \le 1$ specifying the norm under which the error is measured (see [14, section 3] for details). Importantly, when $s = 1$, Algorithm 4 implements a supervised version of CCA. Setting $\alpha = 1$ in Algorithm 3 and $s = 1$ in Algorithm 4, the algorithms have identical network architectures and synaptic update rules, namely:

$$\mathbf{W}_x \leftarrow \mathbf{W}_x + \eta_x(\mathbf{z}_t - \mathbf{a}_t)\mathbf{x}_t^\top$$

$$\mathbf{W}_y \leftarrow \mathbf{W}_y + \eta_y \mathbf{c}_t^a \mathbf{y}_t^\top$$

$$\mathbf{P} \leftarrow \mathbf{P} + \eta_p(\mathbf{z}_t \mathbf{n}_t^\top - \mathbf{P}),$$

---

[3]In Golkar et al. [14] the inputs $\mathbf{x}_t$ are the basal inputs and the inputs $\mathbf{y}_t$ are the apical inputs. Here we switch the inputs to be consistent with Algorithm 3.

where we recall that $\mathbf{c}_t^a := \mathbf{a}_t - \mathbf{P}\mathbf{n}_t$. For this parameter choice, the main difference between the algorithms is that Adaptive Bio-CCA with output whitening is an unsupervised learning algorithm whereas Bio-RRR is a supervised learning algorithm, which is reflected in their outputs $\mathbf{z}_t$: the output of Algorithm 3 is the whitened sum of the basal dendritic currents and the apical calcium plateau potential, i.e., $\mathbf{z}_t = \mathbf{b}_t + \mathbf{c}_t^a$, whereas the output of [14, Algorithm 2] is the (whitened) CCSP of the basal inputs, i.e., $\mathbf{z}_t = \mathbf{b}_t$. In other words, the apical inputs do not directly contribute to the output of the network in [14]; only indirectly via plasticity in the basal dendrites. Experimental evidence suggests that apical calcium plateau potentials contribute significantly to the outputs of pyramidal cells, which supports the model derived here. Furthermore, the model in this work allows one to adjust the parameter α to adaptively set the output rank, which is important for analyzing non-stationary input streams. In Table 1 we summarize the differences between the two algorithms.

	Adaptive Bio-CCA	Bio-RRR
unsupervised/supervised	unsupervised	supervised
whitened outputs	$\checkmark$	$\checkmark$
adaptive output rank	$\checkmark$	$\times$

Table 1: Comparison of Adaptive Bio-CCA with output whitening and Bio-RRR.

B.3   Decoupling the interneuron synapses

The neural network for Adaptive Bio-CCA with output whitening derived in Section 4 requires the pyramidal neuron-to-interneuron synaptic weight matrix $\mathbf{P}^\top$ to be the the transpose of the interneuron-to-pyramidal neuron synaptic weight matrix $\mathbf{P}$. Enforcing this symmetry via a centralized mechanism is not biologically plausible. Rather, following [14, Appendix D], we show that the symmetry between these two sets of weights naturally follows from the local learning rules.

To begin, we replace the pyramidal neuron-to-interneuron weight matrix $\mathbf{P}^\top$ by a new weight matrix $\mathbf{R}$ with Hebbian learning rules:

$$\mathbf{R} \leftarrow \mathbf{R} + \frac{\eta}{\tau}(\mathbf{n_t}\mathbf{z_t}^\top - \mathbf{R}). \tag{36}$$

Let $\mathbf{P}_0$ and $\mathbf{R}_0$ denote the initial values of $\mathbf{P}$ and $\mathbf{R}$. Then, in view of the updates rule for $\mathbf{P}$ in and $\mathbf{R}$ in Algorithm 3 and Equation (36), respectively, the difference $\mathbf{P}^\top - \mathbf{R}$ after T updates is given by

$$\mathbf{P}^\top - \mathbf{R} = \left(1 - \frac{\eta}{\tau}\right)^T (\mathbf{P}_0^\top - \mathbf{R_0}).$$

In particular, the difference decays exponentially (recall that $\eta < \tau$ by assumption).

C Numerics

C.1 Experimental details

Bio-CCA: We implemented Algorithm 2. We initialized $\mathbf{W}_x$ (rexp. $\mathbf{W}_y$) to be a random matrix with i.i.d. mean-zero normal entries with variance $1/m$ (resp. $1/n$). We initialized $\mathbf{M}$ to be the indentity matrix $\mathbf{I}_k$. We used the time-dependent learning rate of the form $\eta_t = \eta_0/(1 + \gamma t)$. To find the optimal hyperparameters, we performed a grid search over $\eta_0 \in \{10^{-2}, 10^{-3}, 10^{-4}\}$, $\gamma \in \{10^{-3}, 10^{-4}, 10^{-5}\}$ and $\tau \in \{1, 0.5, 0.1\}$. The best performing parameters are reported in Table 2. Finally, to ensure the reported basis vectors satisfy the orthonormality constraints (2), we report the following normalized basis vectors:

$$\mathbf{V}_x^\top := \left\{\mathbf{M}^{-1}(\mathbf{W}_x\mathbf{C}_{xx}\mathbf{W}_x^\top + \mathbf{W}_y\mathbf{C}_{yy}\mathbf{W}_y^\top)\mathbf{M}^{-1}\right\}^{-1/2}\mathbf{M}^{-1}\mathbf{W}_x, \qquad (37)$$

$$\mathbf{V}_y^\top := \left\{\mathbf{M}^{-1}(\mathbf{W}_x\mathbf{C}_{xx}\mathbf{W}_x^\top + \mathbf{W}_y\mathbf{C}_{yy}\mathbf{W}_y^\top)\mathbf{M}^{-1}\right\}^{-1/2}\mathbf{M}^{-1}\mathbf{W}_y. \qquad (38)$$

	synthetic	Mediamill
Bio-CCA (η_0, γ, τ)	$(10^{-3}, 10^{-4}, 0.1)$	$(10^{-2}, 10^{-4}, 0.1)$
Gen-Oja (β_0, γ)	$(1, 10^{-2})$	$(10^{-2}, 10^{-3})$
Asym-NN (η_0, α)	$(2 \times 10^{-4}, 5 \times 10^{-6})$	$(2 \times 10^{-3}, 5 \times 10^{-6})$
Bio-RRR (η_0, γ, μ)	$(10^{-3}, 10^{-4}, 10)$	$(10^{-2}, 10^{-5}, 10)$
Adaptive Bio-CCA (η_0, γ, τ)	$(10^{-3}, 10^{-4}, 0.1)$	$(10^{-2}, 10^{-4}, 0.5)$

Table 2: Optimal time-dependent learning rates.

MSG-CCA: We implemented the online algorithm stated in [1]. MSG-CCA requires a training set to estimate the covariance matrices $\mathbf{C}_{xx}$ and $\mathbf{C}_{yy}$. We provided the algorithm with 1000 samples to initially estimate the covariance matrix. Following [1], we use the time-dependent learning rate $\eta_t = 0.1/\sqrt{t}$.

Gen-Oja: We implemented the online algorithm stated in [4]. The algorithm includes 2 learning rates: α_t and β_t. As stated in [4], the Gen-Oja's performance is robust to changes in the learning rate α_t, but sensitive to changes in the learning rate β_t. Following [4], we set α_t to be constant and equal to $1/R^2$ where $R^2 := \mathrm{Tr}(\mathbf{C}_{xx}) + \mathrm{Tr}(\mathbf{C}_{yy})$. To optimize over β_t, we used a time-dependent learning rate of the form $\beta_t = \beta_0/(1 + \gamma t)$ and performed a grid search over $\beta_0 \in \{1, 10^{-1}, 10^{-2}\}$ and $\gamma \in \{10^{-1}, 10^{-2}, 10^{-3}\}$. The best performing parameters are reported in Table 2.

Asymmetric CCA Network: We implemented the online multi-channel CCA algorithm derived in [37]. Following [37], we use the linearly decaying learning rate $\eta_t = \eta_0 \times \max(1 - \alpha t, 0.1)$. To optimize the performance of the algorithm we performed a grid search over $\eta_0 \in \{2 \times 10^{-2}, 10^{-2}, 2 \times 10^{-3}, 10^{-3}\}$ and $\alpha \in \{5 \times 10^{-5}, 10^{-5}, 5 \times 10^{-6}, 10^{-6}\}$. The best performing parameters are reported in Table 2.

Bio-RRR: We implemented the online CCA algorithm derived in [14] with $s = 1$ (see Algorithm 4). The algorithm includes learning rates η_x, η_y and η_p. Following [14], we set $\eta_y = \eta_t$ and $\eta_x = \eta_p = \eta_y/\mu$ and use the time-dependent learning rate of the form $\eta_t = \eta_0/(1 + \gamma t)$. We performed a grid search over $\eta_0 \in \{10^{-2}, 10^{-3}, 10^{-4}\}$, $\gamma \in \{10^{-3}, 10^{-4}, 10^{-5}\}$ and $\mu \in \{1, 10, 100\}$ and list best performing parameters in Table 2.

Adaptive Bio-CCA with output whitening: We implemented Algorithm 3. We initialized $\mathbf{W}_x$, $\mathbf{W}_y$ and $\mathbf{P}$ to be random matrices with i.i.d. standard normal entries. To find the optimal hyperparameters, we performed a grid search over $\eta_0 \in \{10^{-2}, 10^{-3}, 10^{-4}\}$, $\gamma \in \{10^{-3}, 10^{-4}, 10^{-5}\}$ and $\tau \in \{1, 0.5, 0.1\}$. The best performing parameters are reported in Table 2.

C.2 Orthonormality constraints

Bio-CCA. To verify that the Bio-CCA basis vectors asymptotically satisfy the orthonormality constraint (2), we use the following orthonormality error function:

$$\text{Orthonormality Error}(t) := \frac{\|\mathbf{M}_t^{-1}(\mathbf{W}_{x,t}\mathbf{C}_{xx}\mathbf{W}_{x,t}^\top + \mathbf{W}_{y,t}\mathbf{C}_{yy}\mathbf{W}_{y,t}^\top)\mathbf{M}_t^{-1} - \mathbf{I}_k\|_F^2}{k}. \tag{39}$$

In Figure 11 (resp. Figure 12) we demonstrate that Bio-CCA asymptotically satisfies the CCA orthonormality constraints (2) on the synthetic dataset (resp. `Mediamill`).

Adaptive Bio-CCA with output whitening. To verify that the top r eigenvalues of the output covariance asymptotically approach 1, we let $\lambda_1(t) \geq \cdots \geq \lambda_k(t)$ denote the ordered eigenvalues of the matrix

$$\mathbf{C}_{zz}(t) := \mathbf{V}_{x,t}^\top\mathbf{C}_{xx}\mathbf{V}_{x,t} + \mathbf{V}_{x,t}^\top\mathbf{C}_{xy}\mathbf{V}_{y,t} + \mathbf{V}_{y,t}^\top\mathbf{C}_{xy}^\top\mathbf{V}_{x,t} + \mathbf{V}_{y,t}^\top\mathbf{C}_{yy}\mathbf{V}_{y,t},$$

and define the whitening error by

$$\text{Whitening Error}(t) := \frac{\sum_{i=1}^{r} |\lambda_i(t) - 1|^2 + \sum_{i=r+1}^{k} |\lambda_i(t)|^2}{k}. \tag{40}$$

In Figure 13 (resp. Figure 14) we demonstrate that Bio-CCA asymptotically satisfies the CCA orthonormality constraints (2) on the synthetic dataset (resp. `Mediamill`).
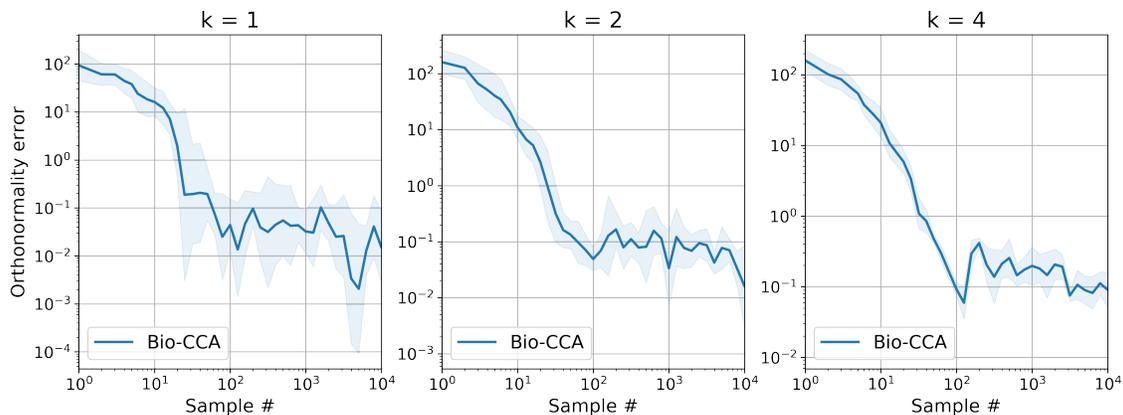
Figure 11: The deviation of the Bio-CCA solution from the CCA orthonormality constraint, in terms of the orthonormality error defined in Equation (39), on the synthetic dataset.
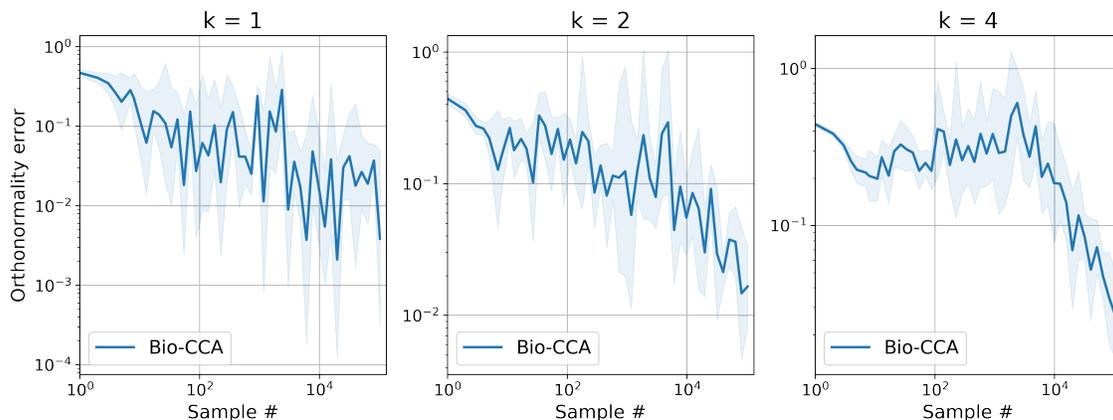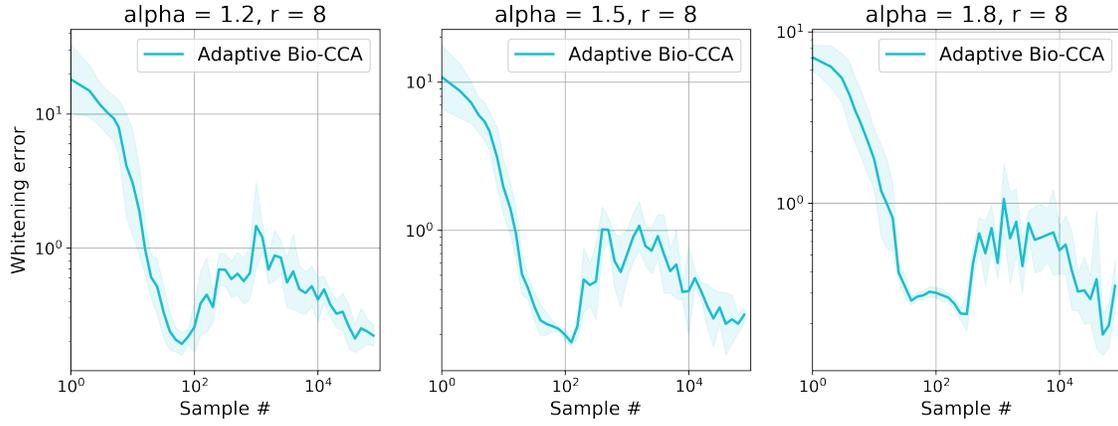

Figure 12: The deviation of the Bio-CCA solution from the CCA orthonormality constraint, in terms of the orthonormality error defined in Equation (39), on the dataset `Mediamill`.

Figure 13: The deviation of the Bio-CCA solution from the CCA orthonormality constraint, in terms of the orthonormality error defined in Equation (39), on the synthetic dataset.
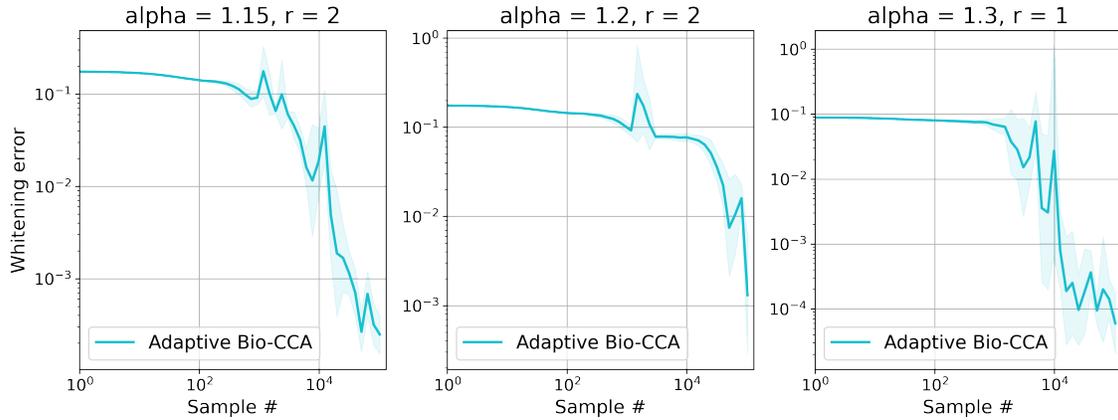

Figure 14: The deviation of the Adaptive Bio-CCA with output whitening solution from the whitening constraint, in terms of the whitening error defined in Equation (40), on the dataset `Mediamill`.

D   On convergence of the CCA algorithms

To interpret our offline and online CCA algorithms (Algorithms 1 & 2) mathematically, we first make the following observation. Both algorithms optimize the min-max objective (11), which includes optimization over the neural outputs $\mathbf{Z}$. In this way, the neural activities can be viewed as optimization steps, which is useful for a biological interpretation of the algorithms. However, since we assume a separation of time-scales in which the neural outputs equilibrate at their optimal values before the synaptic weight matrices are updated, the neural dynamics are superfluous when analyzing the algorithms from a mathematical perspective. Therefore, we set $\mathbf{Z}$ equal to its equilibrium value $\overline{\mathbf{Z}} = \mathbf{M}^{-1}(\mathbf{W}_x\mathbf{X} + \mathbf{W}_y\mathbf{Y})$ in the cost function $L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \mathbf{Z})$ to obtain a min-max problem in terms of the synaptic weights:

$$\min_{\mathbf{W}\in\mathbb{R}^{k\times d}} \max_{\mathbf{M}\in\mathcal{S}_{++}^k} F(\mathbf{W}, \mathbf{M}), \tag{41}$$

where $\mathbf{W} := [\mathbf{W}_x \ \mathbf{W}_y]$ is the matrix of concatenated weights and $F(\mathbf{W}, \mathbf{M}) := L(\mathbf{W}_x, \mathbf{W}_y, \mathbf{M}, \overline{\mathbf{Z}})$. After substitution, we see that $F : \mathbb{R}^{k\times d} \times \mathcal{S}_{++}^k \to \mathbb{R}$ is the nonconvex-concave function

$$F(\mathbf{W}, \mathbf{M}) = \mathrm{Tr}\left(-\mathbf{M}^{-1}\mathbf{W}\mathbf{A}\mathbf{W}^\top + \mathbf{W}\mathbf{B}\mathbf{W}^\top - \frac{1}{2}\mathbf{M}^2\right),$$

with partial derivatives

$$-\frac{\partial F(\mathbf{W}, \mathbf{M})}{\partial \mathbf{W}} = 2\mathbf{M}^{-1}\mathbf{W}\mathbf{A} - 2\mathbf{W}\mathbf{B}$$

$$\frac{\partial F(\mathbf{W}, \mathbf{M})}{\partial \mathbf{M}} = \mathbf{M}^{-1}\mathbf{W}\mathbf{A}\mathbf{W}^\top\mathbf{M}^{-1} - \mathbf{M}.$$

where we have defined

$$\mathbf{A} := \begin{bmatrix} \mathbf{C}_{xx} & \mathbf{C}_{xy} \\ \mathbf{C}_{xy}^\top & \mathbf{C}_{yy} \end{bmatrix}, \qquad\qquad \mathbf{B} := \begin{bmatrix} \mathbf{C}_{xx} & \\ & \mathbf{C}_{yy} \end{bmatrix}. \tag{42}$$

Similar objectives, with different values of $\mathbf{A}$ and $\mathbf{B}$, arise in the analysis of online principal subspace projection [36] and slow feature analysis [27] algorithms.

   The synaptic updates in both our offline and online algorithms can be viewed as (stochastic) gradient descent-ascent algorithms for solving the noncovex-concave min-max problem (41). To make the comparison with our offline algorithm, we substitute the optimal value $\overline{\mathbf{Z}}$ into the synaptic weight updates in Algorithm 1 to obtain:

$$\mathbf{W} \leftarrow \mathbf{W} + 2\eta(\mathbf{M}^{-1}\mathbf{W}\mathbf{A} - \mathbf{W}\mathbf{B}) \tag{43}$$

$$\mathbf{M} \leftarrow \mathbf{M} + \frac{\eta}{\tau}(\mathbf{M}^{-1}\mathbf{W}\mathbf{A}\mathbf{W}^\top\mathbf{M}^{-1} - \mathbf{M}), \tag{44}$$

Comparing these updates to the partial derivatives of F, we see that Offline-CCA is naturally interpreted as a gradient descent-ascent algorithm for solving the min-max problem (41), with descent step size η and ascent step size $\frac{\eta}{\tau}$. Similarly, to make the comparison with our online algorithm, we substitute the explicit expression for the equilibrium value $\bar{\mathbf{z}}_t = \mathbf{M}^{-1}(\mathbf{a}_t + \mathbf{b}_t)$, where $\mathbf{a}_t = \mathbf{W}_x\mathbf{x}_t$ and $\mathbf{b}_t = \mathbf{W}_y\mathbf{y}_t$, into the synaptic weight updates in Algorithm 2 to rewrite the updates:

$$\mathbf{W} \leftarrow \mathbf{W} + 2\eta(\mathbf{M}^{-1}\mathbf{W}\mathbf{A}_t - \mathbf{W}\mathbf{B}_t)$$
$$\mathbf{M} \leftarrow \mathbf{M} + \frac{\eta}{\tau}(\mathbf{M}^{-1}\mathbf{W}\mathbf{A}_t\mathbf{W}^\top\mathbf{M}^{-1} - \mathbf{M}),$$

where

$$\mathbf{A}_t := \begin{bmatrix} \mathbf{x}_t\mathbf{x}_t^\top & \mathbf{x}_t\mathbf{y}_t^\top \\ \mathbf{y}_t\mathbf{x}_t^\top & \mathbf{y}_t\mathbf{y}_t^\top \end{bmatrix}, \qquad\qquad \mathbf{B}_t := \begin{bmatrix} \mathbf{x}_t\mathbf{x}_t^\top & \\ & \mathbf{y}_t\mathbf{y}_t^\top \end{bmatrix}.$$

Comparing these updates to the partial derivatives of F, we see that our online algorithm is naturally interpreted as a stochastic gradient descent-ascent algorithm for solving the min-max problem (41), using the time t rank-1 approximations $\mathbf{A}_t$ and $\mathbf{B}_t$ in place of $\mathbf{A}$ and $\mathbf{B}$, respectively.

Establishing theoretical guarantees for solving nonconvex-concave min-max problems of the form (41) via stochastic gradient descent-ascent is an active area of research [40]. Borkar [7, 8] proved asymptotic convergence to the solution of the min-max problem for a two time-scale stochastic gradient descent-ascent algorithm, where the ratio between the learning rates for the minimization step and the maximization step, τ, depends on the iteration and converges to zero in the limit as the iteration number approaches infinity. Lin et al. [26] established convergence rate guarantees for a stochastic gradient descent-ascent algorithm to an equilibrium point (not necessarily a solution of the min-max problem). Both these results, however, impose assumptions that do not hold in our setting: the partial derivatives of F are Lipschitz continuous and $\mathbf{M}$ is restricted to a bounded convex set. Therefore, establishing global stability with convergence rate guarantees for our offline and online CCA algorithms requires new mathematical techniques that are beyond the scope of this work.

Even proving local convergence properties is non-trivial. In the special case that $\mathbf{B} = \mathbf{I}_d$, Pehlevan et al. [36] carefully analyzed the continuous dynamical system obtained by formally taking the step size η to zero in Equations (43)–(44). They computed an explicit value $\tau_0 \geq 1/2$, in terms the eigenvalues of $\mathbf{A}$ such that if $\tau \leq \tau_0$, then solutions of the min-max problem (41) are the only linearly stable fixed points of the continuous dynamics. The case that $\mathbf{B} \neq \mathbf{I}_d$ is more complicated, and the approach in [36] is not readily extended. In ongoing work, we take a step towards understanding the asymptotics of our algorithms by analyzing local stability properties for a general class of gradient descent-ascent algorithms, which includes Offline-CCA and Bio-CCA as special cases.

41

References

[1] 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.

[2] Sanjeev Arora, Nadav Cohen, Noah Golowich, and Wei Hu. A convergence analysis of gradient descent for deep linear neural networks. In International Conference on Learning Representations, 2019.

[3] Francis R Bach and Michael I Jordan. A probabilistic interpretation of canonical correlation analysis. Technical Report, 2005.

[4] 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.

[5] 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.

[6] 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.

[7] Vivek S Borkar. Stochastic approximation with two time scales. Systems & Control Letters, 29(5):291–294, 1997.

[8] Vivek S Borkar. Stochastic Approximation: A Dynamical Systems Viewpoint, volume 48. Springer, 2009.

[9] Trevor F Cox and Michael AA Cox. Multidimensional Scaling. Chapman and Hall/CRC, 2000.

[10] Alexander S Ecker, Philipp Berens, Georgios A Keliris, Matthias Bethge, Nikos K Logothetis, and Andreas S Tolias. Decorrelated neuronal firing in cortical microcircuits. Science, 327(5965):584–587, 2010.

[11] 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.

[12] Albert Gidon, Timothy Adam Zolnik, Pawel Fidzinski, Felix Bolduan, Athanasia Papoutsi, 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, 2020.

[13] 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.

[14] Siavash Golkar, David Lipshutz, Yanis Bahroun, Anirvan M Sengupta, and Dmitri B Chklovskii. A simple normative network approximates local non-hebbian learning in the cortex. In Advances in Neural Information Processing Systems, volume 33, pages 7283–7295, 2020.

[15] Zhenkun Gou and Colin Fyfe. A canonical correlation neural network for multicollinearity and functional data. Neural Networks, 17(2):285–293, 2004.

[16] Jordan Guerguiev, Timothy P Lillicrap, and Blake A Richards. Towards deep learning with segregated dendrites. ELife, 6:e22901, 2017.

[17] Tatsuya Haga and Tomoki Fukai. Dendritic processing of spontaneous neuronal sequences for single-trial learning. Scientific Reports, 8(1):1–22, 2018.

[18] Cars H Hommes and Marius I Ochea. Multiple equilibria and limit cycles in evolutionary games with logit dynamics. Games and Economic Behavior, 74 (1):434–441, 2012.

[19] Harold Hotelling. Relations between two sets of variates. Biometrika, 28(3-4): 321–377, 1936.

[20] Paul D King, Joel Zylberberg, and Michael R DeWeese. Inhibitory interneurons decorrelate excitatory cells to drive sparse code formation in a spiking model of V1. Journal of Neuroscience, 33(13):5475–5485, 2013.

[21] 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.

[22] Pei Ling Lai and Colin Fyfe. A neural implementation of canonical correlation analysis. Neural Networks, 12(10):1391–1397, 1999.

[23] Matthew Larkum. A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. Trends in Neurosciences, 36(3):141–151, 2013.

[24] Matthew E Larkum, J Julius Zhu, and Bert Sakmann. A new cellular mechanism for coupling inputs arriving at different cortical layers. Nature, 398(6725): 338–341, 1999.

[25] 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.

[26] Tianyi Lin, Chi Jin, and Michael I Jordan. On gradient descent ascent for nonconvex-concave minimax problems. International Conference on Machine Learning, 2020.

[27] David Lipshutz, Charles Windolf, Siavash Golkar, and Dmitri B Chklovskii. A biologically plausible neural network for slow feature analysis. In Advances in Neural Information Processing Systems, volume 33, pages 14986–14996, 2020.

[28] Jeffrey C Magee and Christine Grienberger. Synaptic plasticity forms and functions. Annual Review of Neuroscience, 43, 2020.

[29] Guy Major, Matthew E Larkum, and Jackie Schiller. Active properties of neocortical pyramidal neuron dendrites. Annual Review of Neuroscience, 36: 1–24, 2013.

[30] Panayotis Mertikopoulos, Christos Papadimitriou, and Georgios Piliouras. Cycles in adversarial regularized learning. In Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 2703–2717. SIAM, 2018.

[31] 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.

[32] Keiji Miura, Zachary F Mainen, and Naoshige Uchida. Odor representations in olfactory cortex: distributed rate coding and decorrelated population activity. Neuron, 74(6):1087–1098, 2012.

[33] Erkki Oja. Simplified neuron model as a principal component analyzer. Journal of Mathematical Biology, 15(3):267–273, 1982.

[34] Cengiz Pehlevan and Dmitri B Chklovskii. A normative theory of adaptive dimensionality reduction in neural networks. In Advances in Neural Information Processing Systems, pages 2269–2277, 2015.

[35] Cengiz Pehlevan, Tao Hu, and Dmitri B Chklovskii. A Hebbian/anti-Hebbian neural network for linear subspace learning: A derivation from multidimensional scaling of streaming data. Neural Computation, 27(7):1461–1495, 2015.

[36] Cengiz Pehlevan, Anirvan M Sengupta, and Dmitri B Chklovskii. Why do similarity matching objectives lead to Hebbian/anti-Hebbian networks? Neural Computation, 30(1):84–124, 2018.

[37] Cengiz Pehlevan, Xinyuan Zhao, Anirvan M Sengupta, and Dmitri B Chklovskii. Neurons as canonical correlation analyzers. Frontiers in Computational Neuroscience, 15(55), 2020.

[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] Mark D Plumbley. A Hebbian/anti-Hebbian network which optimizes information capacity by orthonormalizing the principal subspace. In 1993 Third International Conference on Artificial Neural Networks, pages 86–90. IET, 1993.

[40] Meisam Razaviyayn, Tianjian Huang, Songtao Lu, Maher Nouiehed, Maziar Sanjabi, and Mingyi Hong. Nonconvex min-max optimization: Applications, challenges, and recent theoretical advances. IEEE Signal Processing Magazine, 37(5):55–66, 2020.

[41] Blake A Richards and Timothy P Lillicrap. Dendritic solutions to the credit assignment problem. Current Opinion in Neurobiology, 54:28–36, 2019.

[42] João Sacramento, Rui Ponte Costa, Yoshua Bengio, and Walter Senn. Dendritic cortical microcircuits approximate the backpropagation algorithm. In Advances in Neural Information Processing Systems, pages 8721–8732, 2018.

[43] 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.

[44] 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.

[45] 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.

[46] Cezar M Tigaret, Valeria Olivo, Josef HLP Sadowski, Michael C Ashby, and Jack R Mellor. Coordinated activation of distinct Ca^{2+} sources and metabotropic glutamate receptors encodes Hebbian synaptic plasticity. Nature Communications, 7(1):1–14, 2016.

[47] Robert Urbanczik and Walter Senn. Learning by the dendritic prediction of somatic spiking. Neuron, 81(3):521–528, 2014.

[48] Raja Velu and Gregory C Reinsel. Multivariate Reduced-Rank Regression: Theory and Applications, volume 136. Springer Science & Business Media, 2013.

[49] 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.

[50] 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.

[51] Christopher KI Williams. On a connection between kernel PCA and metric multidimensional scaling. In Advances in Neural Information Processing Systems, pages 675–681, 2001.