

Asymmetric canonical correlation analysis of Riemannian and high-dimensional data

James Buenfil*¹ and Eardi Lila†²

¹Department of Statistics, University of Washington

²Department of Biostatistics, University of Washington

Abstract

In this paper, we introduce a novel statistical model for the integrative analysis of Riemannian-valued functional data and high-dimensional data. We apply this model to explore the dependence structure between each subject’s dynamic functional connectivity – represented by a temporally indexed collection of positive definite covariance matrices – and high-dimensional data representing lifestyle, demographic, and psychometric measures. Specifically, we employ a reformulation of canonical correlation analysis that enables efficient control of the complexity of the functional canonical directions using tangent space sieve approximations. Additionally, we enforce an interpretable group structure on the high-dimensional canonical directions via a sparsity-promoting penalty. The proposed method shows improved empirical performance over alternative approaches and comes with theoretical guarantees. Its application to data from the Human Connectome Project reveals a dominant mode of covariation between dynamic functional connectivity and lifestyle, demographic, and psychometric measures. This mode aligns with results from static connectivity studies but reveals a unique temporal non-stationary pattern that such studies fail to capture.

1 Introduction

One of the primary goals of large-scale neuroimaging studies, such as the Human Connectome Project, ABCD, and the UK Biobank, is to understand the relationship between complex neuroimaging traits and non-imaging high-dimensional variables, including cognitive abilities, neurodegenerative conditions, mental health disorders, psychometric test scores, and other external factors (Zhu et al. 2023). In the context of functional connectivity studies, such complex imaging data are typically networks that are derived from fMRI data and are characterized by a single covariance matrix that captures the temporal correlation between the fMRI signals of different brain regions. For instance, Xia et al. (2018) study correlation patterns between functional connectivity and psychiatric symptoms. Other studies, such as Smith et al. (2015) and Liu et al. (2022), investigate the relationship between functional connectivity and behavioral and demographic measures.

Traditional analyses often view brain functional networks as static. Yet, there is growing evidence that these networks are inherently dynamic and exhibit significant temporal fluctuations (Hutchison et al. 2013), which appear to be linked to various aspects of human behavior (Liégeois et al. 2019). Therefore, they are best represented by a time-indexed collection of covariances, that is, a Riemannian manifold-valued function where the manifold consists of the space of symmetric positive definite (SPD) matrices.

This work seeks to identify joint variation between these functional dynamic networks and multivariate variables, such as lifestyle, demographic, and psychometric measures. To this purpose, we

*buenfil@uw.edu

†elila@uw.edu

develop a novel asymmetric canonical correlation analysis model that allows us to explore the underlying relationships between Riemannian manifold-valued functional data and high-dimensional variables. While our motivation stems from dynamic functional connectivity, the proposed method is general and can be applied to a variety of other settings.

Numerous models have been developed to model manifold-valued functional data, (see, e.g., Dai and Müller 2018; Lin and Yao 2019; Dubey and Müller 2020; Zhang et al. 2020; Dubey and Müller 2021; Zhou and Müller 2021; Bhattacharjee and Müller 2021; Stöcker and Greven 2021), which can be more broadly viewed as object data (Marron and Dryden 2021) – a generalization of functional data (Ramsay and Silverman 2015; Hsing and Eubank 2013; Kokoszka and Reimherr 2017). Regression models for manifold-valued data with low-dimensional predictors have been proposed in Petersen and Müller (2019), Zhao et al. (2020), and Zhao et al. (2021). See also Petersen et al. (2022) for a review. Nonetheless, models that facilitate the integration of manifold-valued functional data with high-dimensional variables have not been extensively explored.

Canonical correlation analysis (CCA) is one of the principal tools for data integration (Hotelling 1936; Urtio et al. 2018; Zhuang et al. 2020; Yang et al. 2021) and can be used to identify shared structure between two low-dimensional sets of variables by seeking linear combinations of these sets – with weights referred to as canonical vectors – that exhibit maximum correlation. Extensions of CCA to high-dimensional data have been proposed, for instance, in Witten et al. (2009), Lin et al. (2013), Chen et al. (2013), Gaynanova et al. (2016), Gao et al. (2017), Yoon et al. (2020), and Wang and Zhou (2021). The setting of functional data has been considered in He et al. (2010), Shin and Lee (2015), and Huang and Renaut (2015) and that of more complex imaging data in Cho et al. (2021) and Liu et al. (2021). Inferential aspects have been explored in Yang and Pan (2015), McKeague and Zhang (2022), and Kessler and Levina (2023). Methods that estimate both shared and individual structure have been proposed in Lock et al. (2013), Feng et al. (2018), Carmichael (2020), Shu et al. (2020), and Yuan and Gaynanova (2021), and their connection to CCA has been studied in Murden et al. (2022).

Yet, despite the large body of literature on CCA and its extensions, existing approaches are not able to effectively estimate common structure between Riemannian manifold-valued functional data and high-dimensional multivariate variables, and more broadly, between imaging and high-dimensional data. To bridge this gap, we propose a model that leverages a regression-based characterization of CCA which allows us to incorporate appropriate notions of complexity for the functional and high-dimensional canonical vectors. Specifically, our approach takes advantage of the inherent smoothness and geometric nature of the functional data, employing tangent space approximations based on a data-driven function basis computed using the Riemannian Functional Principal Components Analysis (RFPCA) framework (Dai and Müller 2018; Lin and Yao 2019; Shao et al. 2022). Moreover, it tackles the high dimensionality of the multivariate data by imposing sparsity. It therefore performs feature selection, resulting in models that are more interpretable and mitigate overfitting issues. In the motivating application, this will result in the identification of a small and interpretable set of multivariate variables linked to specific functional dynamic connectivity patterns.

The asymmetric setting considered in this work is of interest not just for its potential applications but also methodologically, as it has some distinct features that are not found in the purely sparse or functional settings. Specifically, we show that if the functional data can be efficiently represented using a finite subspace, the proposed method can consistently estimate the high-dimensional canonical vectors without requiring the direct estimation of the precision matrix of the high-dimensional data – a notoriously difficult problem in high-dimensions and typically solvable only under specific structural assumptions (Cai et al. 2016). This feature renders the proposed methodology novel even in the simpler setting of classical functional and high-dimensional data integration.

In addition to accommodating manifold-valued functional data and high-dimensional data, our proposed method has several other desirable properties in comparison to existing CCA models, which we highlight below:

1. It can estimate multiple canonical directions simultaneously, without requiring iterative deflation strategies and leveraging shared sparsity structure across canonical vectors.

2. It is computationally efficient, with its complexity essentially reducing to solving a regularized multivariate linear regression problem.
3. It does not require a consistent estimator for the precision matrix of the high-dimensional data.
4. The canonical vectors satisfy the correct orthogonality conditions, ensuring that the proposed approach is invariant to data rescaling, while simultaneously maintaining an interpretable sparsity structure on the high-dimensional canonical vectors.
5. When the number of observations is larger than the dimension of the high-dimensional data and the hyper-parameter controlling sparsity is set to 0, our approach reduces to classical multivariate CCA.

The rest of the paper is organized as follows. In Section 2, we introduce the proposed asymmetric CCA model. In Section 3, we introduce the associated estimator and in Section 4, we explore its theoretical properties. In Section 5, we apply our method to data from the Human Connectome Project to study dynamic functional connectivity, and in Section 6, we study its empirical performance by means of simulation studies. Proofs and more technical details are left to the appendices.

2 Model

2.1 Elements of Riemannian geometry

Let \mathcal{M} be an M -dimensional Riemannian manifold and let $T_x\mathcal{M}$ denote the tangent space at a point $x \in \mathcal{M}$ equipped with Riemannian metric $\langle \cdot, \cdot \rangle_x$. Moreover, for any $x \in \mathcal{M}$, denote the exponential map by $\text{Exp}_x : U \rightarrow \mathcal{M}$, where $U \subset T_x\mathcal{M}$ is an open set containing the origin that guarantees that this map is a bijection onto its range $\text{Im}(\text{Exp}_x)$. The logarithmic map at x , denoted by $\text{Log}_x : \text{Im}(\text{Exp}_x) \rightarrow \mathcal{M}$ is the inverse of the exponential map Exp_x . We denote by $d_{\mathcal{M}}(\cdot, \cdot)$ the Riemannian distance function on \mathcal{M} , which generalizes the Euclidean distance to manifolds. We refer to Lee (2012) and Lee (2018) for an introduction to the differential geometric concepts used in this work.

In our final application, \mathcal{M} will represent the non-Euclidean manifold of SPD matrices equipped with the affine-invariant metric (see, e.g., Fletcher and Joshi 2007; Pennec et al. 2019)). The affine-invariant metric at $P \in \mathcal{M}$ between $W, Z \in T_P\mathcal{M}$ is defined as $\langle W, Z \rangle_{\mathcal{M}} = \text{tr}(P^{-1}WP^{-1}Z)$. Let \exp and \log denote the matrix exponential and logarithm, defined here on the sets of symmetric matrices and positive definite matrices, respectively, and let $\|\cdot\|_F$ denote the Frobenius norm of a matrix. Then, the affine-invariant Riemannian distance is defined as $d_{\mathcal{M}}(P, Q) = \|\log(P^{-1/2}QP^{-1/2})\|_F$. The logarithmic map $\text{Log}_P(Q) = P^{1/2} \log(P^{-1/2}QP^{-1/2}) P^{1/2}$ will allow us to compute unconstrained tangent space representations of our data, i.e., symmetric matrices. Roughly speaking, the tangent space representations allow us to apply simple Euclidean mathematical operations without breaking the geometry of the space of SPD matrices and the exponential map $\text{Exp}_P(W) = P^{1/2} \exp(P^{-1/2}WP^{-1/2}) P^{1/2}$ will allow us to map tangent space elements back to the manifold \mathcal{M} . In this case, the exponential and logarithmic maps Exp and Log are global bijections between \mathcal{M} and the space of symmetric matrices.

Next, we present the mathematical tools necessary to model Riemannian-valued functions. Let \mathcal{T} be a compact subset of \mathbb{R} and let $\mu : \mathcal{T} \rightarrow \mathcal{M}$ be a sufficiently smooth curve on \mathcal{M} . A vector field V along μ is a map from \mathcal{T} to the tangent bundle $T\mathcal{M}$ such that $V(t) \in T_{\mu(t)}\mathcal{M}$ for all $t \in \mathcal{T}$. The collection of vector fields V along μ defines a vector space. Define $L^2(T\mu)$ to be the space of square integrable vector fields V along μ equipped with inner product $\langle\langle U, V \rangle\rangle_{\mu} := \int_{\mathcal{T}} \langle V(t), U(t) \rangle_{\mu(t)} dt$ and induced norm defined by $\|\cdot\|_{\mu}^2 = \langle\langle \cdot, \cdot \rangle\rangle_{\mu}$, where U and V are both vector fields along μ . Then, $L^2(T\mu)$ is a separable Hilbert space (Lin and Yao 2019).

For a curve μ and Riemannian-valued function $y : \mathcal{T} \rightarrow \mathcal{M}$, we denote as $\text{Log}_{\mu} y$ the function $t \mapsto \text{Log}_{\mu(t)} y(t)$. Similarly, for a vector field V along μ , we denote as $\text{Exp}_{\mu} V$ the function $t \mapsto \text{Exp}_{\mu(t)} V(t)$. In our setting, y will be random, and μ will represent the mean of y . Under appropriate assumptions,

the vector field $\text{Log}_\mu y$ along μ will be a random element of $L^2(T\mu)$, which intuitively represents a linearized and centered version of y . Indeed, if $\mathcal{M} = \mathbb{R}^d$, then $\text{Log}_{\mu(t)} y(t) = y(t) - \mu(t)$ for every $t \in \mathcal{T}$.

Later, we will need to compare vector fields along different curves μ and $\hat{\mu}$. To this purpose, following Lin and Yao (2019), we introduce the parallel transport operator. We denote the parallel transport operator on \mathcal{M} along geodesics as $\mathcal{P}_{x,p} : T_x\mathcal{M} \rightarrow T_p\mathcal{M}$. A fundamental property of this operator is that it preserves inner products of tangent vectors, i.e., for any $u, v \in T_x\mathcal{M}$, $\langle u, v \rangle_x = \langle \mathcal{P}_{x,p}u, \mathcal{P}_{x,p}v \rangle_p$. We can then define parallel transport for vector fields U, V along curves $f, h : \mathcal{T} \rightarrow \mathcal{M}$. Specifically, given $U \in L^2(Tf)$ and $V \in L^2(Th)$, we define $\Gamma_{f,h}U \in L^2(Th)$ as the map $t \mapsto \mathcal{P}_{f(t),h(t)}U(t)$. Therefore, $\Gamma_{f,h}$ can be viewed as a map from $L^2(Tf)$ to $L^2(Th)$. Therefore, while U and V cannot be ‘compared’ directly since for every t , $U(t)$ and $V(t)$ may belong to different tangent spaces, we can compare $\Gamma_{f,h}U$ and V , since they are both elements of $L^2(Th)$. In particular, $\|\Gamma_{f,h}U - V\|_h$ quantitatively describes the difference between U and V . We refer to Proposition 2 of Lin and Yao (2019) for additional properties of the parallel transport operator.

2.2 Modeling Riemannian-valued data

Let (y, X) be a pair of random variables, where $X \in \mathbb{R}^p$ is a zero-mean random vector, with covariance $\Sigma_X \in \mathbb{R}^{d \times d}$, representing the high-dimensional multivariate variables, and the process y is a Riemannian-valued random process with continuous sample paths. We assume that $\forall x \in \mathcal{M}, \forall t \in \mathcal{T}$ we have $\mathbb{E}[d_{\mathcal{M}}^2(y(t), x)] < \infty$.

Next, we define the Fréchet mean of the process y on \mathcal{M} as

$$\mu(t) = \arg \min_{x \in \mathcal{M}} \mathbb{E}[d_{\mathcal{M}}^2(y(t), x)].$$

We assume that the Fréchet mean $\mu(t)$ exists and is unique for every $t \in \mathcal{T}$, and μ is a continuous function. For more details on the Fréchet mean, see Bhattacharya and Patrangenaru (2003). Moreover, we assume

$$\Pr\{\text{For all } t \in \mathcal{T} : y(t) \in \text{Im}(\text{Exp}_{\mu(t)})\} = 1,$$

which ensures that $\text{Log}_{\mu(t)}y(t)$ is defined almost surely for all $t \in \mathcal{T}$.

Let the tensor product $U \otimes V : L^2(T\mu) \rightarrow L^2(T\mu)$, between $U, V \in L^2(T\mu)$, be defined as $(U \otimes V)(W) = \langle U, W \rangle_\mu V$ for all $W \in L^2(T\mu)$. If $\mathbb{E}[\|\text{Log}_\mu y\|_\mu^2] < \infty$, then the covariance function \mathcal{C} of $\text{Log}_\mu y$ is defined as $\mathcal{C} = \mathbb{E}[\text{Log}_\mu y \otimes \text{Log}_\mu y]$ and is nonnegative and trace class. Therefore, it admits the eigendecomposition

$$\mathcal{C} = \sum_{j=1}^{\infty} \omega_j \phi_j \otimes \phi_j, \tag{1}$$

with ω_j a sequence of real numbers converging to 0, and $\phi_j \in L^2(T\mu)$ satisfying $\langle \phi_j, \phi_k \rangle_\mu = \delta_{jk}$, where $\delta_{jk} = 1$ if $j = k$, and 0 otherwise. The functions $\{\phi_j\}$ are called the population loading functions, or population principal components, of $\text{Log}_\mu y$. Moreover, with probability one, we have that the process $\text{Log}_\mu y$ admits a Principal Component expansion

$$\text{Log}_\mu y = \sum_{j=1}^{\infty} Y_j \phi_j,$$

where $Y_j = \langle \phi_j, \text{Log}_\mu y \rangle_\mu$ are pairwise uncorrelated random variables, and satisfy $\mathbb{E}[Y_j] = 0$ and $\text{Var}(Y_j) = \omega_j$. The variables Y_j are called the population principal scores. For further details on the principal component basis and eigendecomposition of the \mathcal{C} , see Lemma A.1 in the appendices.

2.3 Asymmetric Riemannian CCA

In this section, we introduce the asymmetric CCA model, which can be naturally formalized by mirroring the multivariate and functional versions (He et al. 2010) of the problem. We define the first

canonical direction pair (ψ_1, θ_1) as a solution, if one exists, to the following problem

$$\underset{\psi \in L^2(T\mu), \theta \in \mathbb{R}^p}{\text{maximize}} \quad \text{Corr}^2(\langle \langle \text{Log}_\mu y, \psi \rangle \rangle_\mu, X^\top \theta). \quad (2)$$

Analogously, we can define the subsequent pairs (ψ_k, θ_k) to maximize the same objective function, with the condition that each pair is orthogonal to the previous ones, namely, $\langle \langle \psi_k, \mathcal{C}\psi_{k'} \rangle \rangle_\mu = \delta_{kk'}$ and $\theta_k^\top \Sigma_X \theta_{k'} = \delta_{kk'}$. When they exist, we refer to ψ_k as the k th canonical function, and to θ_k as the k th canonical vector. Given the canonical function $\psi_k \in L^2(T\mu)$, we can map it back to the original space via the exponential map. This procedure is illustrated in Figure 1.

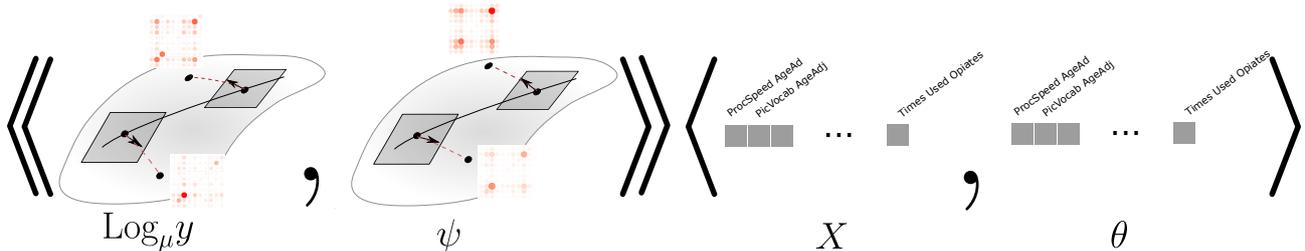


Figure 1: In this figure, we illustrate the process of projecting the Riemannian-valued functional data and the high-dimensional data to define maximally correlated variables. We leverage tools from differential geometry to compute linear tangent representations $\text{Log}_\mu y$ of the temporally-indexed Riemannian-valued data y , which are equipped with a notion of inner product $\langle \langle \cdot, \cdot \rangle \rangle_\mu$, that is, a projection operator. For the multivariate data, we use the conventional notion of projection, i.e., the Euclidean inner product. We therefore seek ψ and θ whose respective data projections define maximally correlated variables.

While equation (2) provides an intuitive formulation of the canonical correlation problem, it has been noted in Cupidon et al. (2008) that the maximum of this problem may not be attained by any $\psi \in L^2(T\mu)$, $\theta \in \mathbb{R}^p$. To address this issue, it is necessary to reformulate the problem with respect to the pair of canonical variables $(U, V) = (\langle \langle \text{Log}_\mu y, \psi \rangle \rangle_\mu, X^\top \theta)$, resulting in the following minimization problem for the first canonical pair:

$$\underset{U \in \bar{\mathcal{U}}, V \in \bar{\mathcal{V}}}{\text{maximize}} \quad \text{Corr}^2(U, V), \quad (3)$$

where $\mathcal{U} = \{\langle \langle \text{Log}_\mu y, \psi \rangle \rangle_\mu : \psi \in L^2(T\mu)\}$, $\mathcal{V} = \{X^\top \theta : \theta \in \mathbb{R}^p\}$, and $\bar{\mathcal{U}}$ and $\bar{\mathcal{V}}$ are the closures of \mathcal{U} and \mathcal{V} , respectively. This guarantees that an optimal canonical variable pair (U, V) does exist. However, they cannot necessarily be written in terms of the canonical vectors, i.e. it does not necessarily hold that $U \in \bar{\mathcal{U}}$ can be written as $U = \langle \langle \text{Log}_\mu y, \psi \rangle \rangle_\mu$ for some $\psi \in L^2(T\mu)$. To simplify the exposition, we defer the details of this formulation to Appendix A.

Given that the Riemannian-valued random process y and its associated representation $\text{Log}_\mu y \in L^2(T\mu)$ are infinite dimensional, we must resort to some form of dimension reduction. Specifically, we make the following assumption:

Assumption 2.1. *There exists a complete orthonormal system for $L^2(T\mu)$, $\{\phi_j\}_{j=1}^\infty$, and $d \in \mathbb{N}$ with $d \leq p$ such that $\text{Log}_\mu y = \sum_{j=1}^d Y_j \phi_j$, where the Y_j are random variables with $\text{Var}(Y_j) < \infty$.*

We refer to d as the rank of the functional data. Note that, to define the orthonormal system $\{\phi_j\}_{j=1}^\infty$, we could employ the principal components in Section 2.2, or alternatively, we could design its basis functions to capture specific features of interest.

Next, define $Y = (Y_1, \dots, Y_d)$ and let Σ_Y be the $d \times d$ covariance of Y . Let $\|\cdot\|_2$ denote the Euclidean 2-norm of a vector in \mathbb{R}^d . Without loss of generality, we suppose that X and Y are mean 0. When Assumption 2.1 is satisfied, the following theorem states that the infinite-dimensional canonical correlation problem in equation (3) is equivalent to solving a suitably formulated finite-dimensional regression problem.

Theorem 2.1. *Assume that condition 2.1 holds. Then, there are at most d nontrivial canonical variable pairs $\{(U_k, V_k)\}$, and each pair (U_k, V_k) can be written in terms of the associated canonical directions: $U_k = \langle \text{Log}_\mu y, \psi_k \rangle_\mu$ and $V_k = X^\top \theta_k$ for some $\psi_k \in L^2(T\mu)$ and $\theta_k \in \mathbb{R}^p$. Additionally, suppose $\Sigma_X \in \mathbb{R}^{p \times p}$ and $\Sigma_Y \in \mathbb{R}^{d \times d}$ are invertible. Let B be the solution to the multivariate least-squares problem*

$$\underset{B \in \mathbb{R}^{p \times d}}{\text{minimize}} \mathbb{E} \left[\|\Sigma_Y^{-1/2} Y - B^\top X\|_2^2 \right], \quad (4)$$

and let

$$B^\top \Sigma_X B = \tilde{H} D^2 \tilde{H}^\top \quad (5)$$

be an eigen-decomposition of $B^\top \Sigma_X B$. Define

$$T = B \tilde{H} D^{-1} \in \mathbb{R}^{p \times d}, \quad (6)$$

$$H = \Sigma_Y^{-1/2} \tilde{H} \in \mathbb{R}^{d \times d}. \quad (7)$$

Then, the k th column of H , η_k , characterizes the k th canonical function ψ_k through $\psi_k = \sum_{j=1}^d \eta_{kj} \phi_j$, and the k th column of T is the k th high-dimensional canonical vector θ_k . Moreover, the optimum values attained by the maximization problem in equation (3) are the diagonal entries of Γ^2 , which we denote by $\gamma_1^2, \dots, \gamma_d^2$.

The proof of Theorem 2.1 can be found in Appendix A.5. This suggests a novel methodology for deriving estimates of the canonical functions $\{\psi_k\}$ and the canonical vectors $\{\theta_k\}$. This entails defining a subspace, spanned by $\{\phi_j\}_{j=1, \dots, d}$, onto which the tangent space representations of the functional data are projected. Subsequently, the canonical functions and vectors can be characterized by the equations (4)-(7), using empirical estimates in place of the theoretical population values.

Crucially, as opposed to other methods in the literature (see, e.g., Chen et al. 2013; Gao et al. 2017), the proposed model circumvents the direct estimation of Σ_X^{-1} , i.e., the precision matrix of the variable X , which is a notoriously difficult problem in high dimensions as it can be estimated only under restrictive structural assumptions. Our strategy will yield interpretable results by enforcing sparsity directly on the canonical directions $\{\theta_k\}$ through an additional penalty term on the estimate of B . The complexity of the functional canonical direction is controlled by projecting the functional data on a finite-dimensional subspace. Such an approach leverages the smooth nature of the functional data (and its tangent space representation) — which is reflected in the eigenvalues of the covariance rapidly decaying to zero — suggesting that such a projection can serve as an efficient and interpretable approximation.

Remark 1. *In Assumption 2.1, the Riemannian-valued random function has been assumed to have a finite-dimensional representation. This is an important step in order to whiten the functional data and allow CCA to find patterns that are small in magnitude but nevertheless correlated with the high-dimensional data. However, Theorem 2.1 remains valid even when Assumption 2.1 is replaced with the weaker assumption that there exists a complete orthonormal system $\{\phi_i\}_{i=1}^\infty$ for $L^2(T\mu)$, and a set of indices $I \subset \{1, 2, \dots\}$, with finite cardinality $|I| = d$, such that*

$$\begin{aligned} \text{Corr}(X_k, \langle \text{Log}_\mu y, \phi_j \rangle_\mu) &= 0, & k = 1, \dots, p, \forall j \in I^c, \\ \text{Corr}(\langle \text{Log}_\mu y, \phi_i \rangle_\mu, \langle \text{Log}_\mu y, \phi_j \rangle_\mu) &= 0, & \forall i \in I, \forall j \in I^c, \end{aligned}$$

where I^c denotes the complement of I in $\{1, 2, \dots\}$. Intuitively, this implies that there is only a finite number of basis elements $\{\phi_i\}_{i \in I}$ that capture the correlation between X and $\text{Log}_\mu y$ through the scores $\{\langle \text{Log}_\mu y, \phi_i \rangle_\mu\}_{i \in I}$. For more details on this weaker assumption, see also Appendix A.3.

3 Estimation

Suppose we are given N observations

$$(y_i, x_i), \quad i = 1, \dots, N,$$

each being a realization of the pair (y, X) . We propose the following estimation procedure, outlined in four steps.

Step A: RFPCA

We first compute the sample version of the Fréchet mean, defined as

$$\hat{\mu}(t) = \arg \min_{x \in \mathcal{M}} \frac{1}{N} \sum_{i=1}^N d_{\mathcal{M}}^2(y_i(t), x).$$

We then estimate the tangent space representations of the functional data observations using $\text{Log}_{\hat{\mu}} y_i \in L^2(T\hat{\mu})$. Next, we define an orthonormal basis for the tangent space representations using the RFPCA framework proposed in Dai et al. (2017), Lin and Yao (2019), and Shao et al. (2022) to estimate a data-driven basis $\{\hat{\phi}_j\}_{j=1}^d$. Specifically, we estimate the tangent-space covariance operator \mathcal{C} using the sample covariance function $\hat{\mathcal{C}} = \frac{1}{N} \sum_{i=1}^N \text{Log}_{\hat{\mu}} y_i \otimes \text{Log}_{\hat{\mu}} y_i$. Each population loading function ϕ_j and associated eigenvalue ω_j can be estimated using the eigenfunction $\hat{\phi}_j$ and eigenvalue $\hat{\omega}_j$ of $\hat{\mathcal{C}}$. The empirical Principal Component expansion of $\{\text{Log}_{\hat{\mu}} y_i\}$ is then given by

$$\text{Log}_{\hat{\mu}} y_i = \sum_{j=1}^d \hat{Y}_{ij} \hat{\phi}_j,$$

where $\hat{Y}_{ij} = \langle \hat{\phi}_j, \text{Log}_{\hat{\mu}} y_i \rangle_{\hat{\mu}}$ are the PC scores. Here we assume that the rank of the Principal Component expansion d is such that $d < \min(p, N)$. For completeness, in Appendix E, we provide a detailed description of the RFPCA algorithm, including a computationally efficient explicit basis construction for the space of SPD matrices equipped with the affine invariant metric.

Step B: Regularized regression

Next, we use the scores \hat{Y}_{ij} to represent the manifold-valued functional data and estimate the canonical directions leveraging the characterization in Theorem 2.1. We let $\mathbb{X} \in \mathbb{R}^{N \times p}$ and $\hat{Y} \in \mathbb{R}^{N \times d}$ denote the data matrices $(x_{ij})_{ij}$ and $(\hat{Y}_{ij})_{ij}$, respectively, where our notation emphasizes that the entries of \hat{Y} are estimates.

Define $\hat{\Sigma}_Y = \frac{1}{N} \hat{Y}^T \hat{Y}$ and $\hat{\Sigma}_X = \frac{1}{N} \mathbb{X}^T \mathbb{X}$. We estimate the matrix B in equation (4) using \hat{B} , which is derived by solving the following group lasso problem:

$$\hat{B} = \arg \min_{B \in \mathbb{R}^{p \times d}} \frac{2}{N} \left\| \hat{Y} \hat{\Sigma}_Y^{-1/2} - \mathbb{X} B \right\|_F^2 + \lambda \|B\|_{\ell_1, \ell_2}, \quad (8)$$

where $\|\cdot\|_F$ denotes the Frobenius norm and $\|B\|_{\ell_1, \ell_2} = \sum_{i=1}^p \|b_i\|_2$ is a group lasso penalty. Here, b_i refers to the i th row of B .

Step C: Eigenanalysis

Given $\hat{\Sigma}_X^{1/2} \hat{B}$, we then compute its right singular vectors $\hat{H} \in \mathbb{R}^{d \times d}$ and singular values matrix $\hat{D} \in \mathbb{R}^{d \times d}$, that is,

$$\hat{B}^T \hat{\Sigma}_X \hat{B} = \hat{H} \hat{D}^2 \hat{H}^T.$$

Step D: Estimates computation

We define

$$\hat{T} = \hat{B}\hat{H}\hat{D}^{-1}, \quad (9)$$

$$\hat{H} = \hat{\Sigma}_Y^{-1/2}\hat{H}, \quad (10)$$

where $\hat{T} \in \mathbb{R}^{p \times d}$ and $\hat{H} \in \mathbb{R}^{d \times d}$ are estimates of T and H , respectively. Then, $\hat{T} = [\hat{\theta}_1, \dots, \hat{\theta}_d]$ is a matrix whose columns $\hat{\theta}_k$ are the estimates of θ_k , and $\hat{H} = [\hat{\eta}_1, \dots, \hat{\eta}_d]$ is a matrix whose columns $\hat{\eta}_k$ are the estimates of η_k . The estimated canonical functions are therefore given by $\hat{\psi}_k = \sum_{j=1}^d \hat{\eta}_{kj} \hat{\phi}_j$, for $k = 1, \dots, d$, resulting in the estimated canonical functions and vectors $(\hat{\psi}_k, \hat{\theta}_k)$.

Algorithm 1 Asymmetric Sparse-Functional CCA

Input: Pairs $(y_i, x_i)_{i=1, \dots, N}$ of manifold-valued functional data and high-dimensional data; rank of the manifold-valued functional data d .

1. Obtain $\hat{\phi}_j, \hat{\omega}_j$, for $j = 1, \dots, d$, and \hat{Y} applying Intrinsic RFPCA to $(y_i)_{i=1, \dots, N}$.
 2. Compute $\hat{\Sigma}_Y = \text{diag}(\hat{\omega}_j)$ and $\hat{\Sigma}_X = \frac{1}{N} \mathbb{X}^\top \mathbb{X}$.
 3. Compute \hat{B} solving the group lasso problem in equation (8) using the `glmnet` package (Friedman et al. 2010).
 4. Compute $\hat{H} = [\hat{\eta}_1, \dots, \hat{\eta}_d]$ and $\hat{T} = [\hat{\theta}_1, \dots, \hat{\theta}_d]$ in equations (9) and (10).
 5. Compute the estimated canonical functions $\hat{\psi}_k = \sum_{j=1}^d \hat{\eta}_{kj} \hat{\phi}_j$ for $k = 1, \dots, d$.
 6. Return $\{\hat{\theta}_k\}_{k=1}^d$, the estimated canonical vectors associated with X , and $\{\hat{\psi}_k\}_{k=1}^d$, the estimated canonical functions associated with y .
-

The sparsity-promoting regularization norm employed in equation (8) encourages entire rows of the matrix B to be set to zero. From the equation $\hat{T} = \hat{B}\hat{H}\hat{D}^{-1}$, it follows that the corresponding rows of \hat{T} will also be zero. This yields canonical vectors $\{\theta_k\}$ with a group sparsity structure, meaning they share identical sparsity patterns. The main steps of the estimation procedure are summarized in Algorithm 1.

3.1 Special instances

To demonstrate the versatility of our model, we present a few special cases. Even if some of these settings are simpler than the motivating neuroimaging application, the proposed method still provides an innovative approach to analyzing such data.

- In situations where $y_i \in \mathcal{M}$, meaning our imaging data are manifold-valued observations without a temporal dimension, Algorithm 1 can be adapted by using tangent-space PCA (Marron and Dryden 2021) rather than RFPCA, similar to the setting considered in Kim et al. (2014). This model is especially useful for studying static connectivity networks.
- When the imaging data take the form of classical functional data, that is $y_i(t) \in \mathcal{M} \subset \mathbb{R}$ for all $t \in \mathcal{T}$, one can apply Algorithm 1 by replacing RFPCA with classical FPCA (Ramsay and Silverman 2015; Yao et al. 2005). In addition, when $y_i(t) \in \mathcal{M} \subset \mathbb{R}^d$, multivariate FPCA can be employed (Happ and Greven 2018).

Central to the proposed methodology is a CCA model for pairs of observations (y_i, x_i) , where $y_i \in \mathbb{R}^d$, $x_i \in \mathbb{R}^p$, $d \ll N$, and the covariance of y_i is full-rank. In the imaging setting, we use a dimension reduction model to compute the low-dimensional component. However, this setting may

Algorithm 2 Asymmetric Sparse CCA

Input: Pairs $(y_i, x_i)_{i=1, \dots, N}$ of low- and high-dimensional data. Let $\mathbb{Y} = (y_{ij})_{ij}$ and $\mathbb{X} = (x_{ij})_{ij}$.

1. Compute $\hat{\Sigma}_Y = \frac{1}{N} \mathbb{Y}^\top \mathbb{Y}$ and $\hat{\Sigma}_X = \frac{1}{N} \mathbb{X}^\top \mathbb{X}$.
2. Compute \hat{B} solving the group lasso problem

$$\hat{B} = \arg \min_{B \in \mathbb{R}^{p \times d}} \frac{2}{N} \left\| \mathbb{Y} \hat{\Sigma}_Y^{-1/2} - \mathbb{X} B \right\|_F^2 + \lambda \|B\|_{\ell_1, \ell_2} \quad (11)$$

using the `glmnet` package (Friedman et al. 2010).

3. Compute $\hat{H} = [\hat{\eta}_1, \dots, \hat{\eta}_d]$ and $\hat{T} = [\hat{\theta}_1, \dots, \hat{\theta}_d]$ in equations (9) and (10).
 4. Return $\{\hat{\eta}_k\}_{k=1}^d$, the estimated canonical vectors associated with $\{y_i\}_{i=1}^N$, and $\{\hat{\theta}_k\}_{k=1}^d$, the estimated canonical functions associated with $\{x_i\}_{i=1}^N$.
-

also be of independent interest and plays a crucial role in the development of the theoretical results. Therefore, we outline the algorithm for this particular setting in Algorithm 2.

In this special case, our approach is closely related to the Eigenvector-CCA model proposed in Wang and Zhou (2021). Yet, notable differences exist between the two approaches. For example, we ensure that the estimated canonical vectors satisfy the correct orthogonality conditions $\hat{H}^\top \hat{\Sigma}_Y \hat{H} = I_d$ and $\hat{T}^\top \hat{\Sigma}_X \hat{T} = I_p$. Furthermore, our proposed model does not rely on the assumption that the data have been generated from a regression model.

4 Theory

Here we investigate the convergence properties of the proposed estimators. We first study the asymptotic properties of the asymmetric Sparse CCA model outlined in Algorithm 2, which sets the stage for studying the asymptotic convergence properties of the asymmetric Sparse-Functional CCA model outlined in Algorithm 1.

4.1 Estimation error rates for asymmetric Sparse CCA

In this section, we state error bounds for the asymmetric Sparse CCA model outlined in Algorithm 2. We assume the observations $y_i \in \mathbb{R}^d$ and $x_i \in \mathbb{R}^p$ are independent copies of the random variables Y and X , respectively. We denote with γ_k the k th canonical correlation attained in the population version of the problem and recall that $T = [\theta_1, \dots, \theta_d] \in \mathbb{R}^{p \times d}$. Moreover, we denote with $K = \max \{i \in \{1, \dots, d\} : \gamma_i > 0\}$ the number of nontrivial canonical vectors. To simplify the notation, we use the conventions $\gamma_{d+1}^2 = -\infty$ and $\gamma_0^2 = \infty$. We use $\text{cond}(A) = \|A\|_2 / \|A^{-1}\|_2$ to denote the condition number of an invertible matrix A , and $\|A\|_2$ denotes the largest singular value of A . The norm $\|A\|_{2, \infty}$ denotes the maximum Euclidean norm of the rows of A , and $\|A\|_{\ell_1, \ell_2} = \sum_{i=1}^p \|a_i\|_2$, where a_i is the i th row of A . The notation $a \lesssim b$ indicates inequality up to an absolute constant, i.e., there exists an absolute constant $C > 0$ such that $a \leq Cb$. Next, we introduce the main assumptions.

Assumption 4.1. *The random variables X and Y are strict sub-Gaussian random vectors with invertible covariance matrices Σ_X and Σ_Y , respectively. Strict sub-Gaussian random vectors are introduced in Definition B.2 of the appendix.*

Assumption 4.2. *It holds that $d \leq p$, $d \log(p) = o(N)$, $\text{cond}(\Sigma_Y)^2 d = o(N)$, and $\gamma_1 > \dots > \gamma_K$ are bounded from below and are distinct.*

Assumption 4.3. *The norms $\|\Sigma_X\|_{2, \infty}$, $\|T\|_{\ell_1, \ell_2}$ are bounded from above and are larger than 1, $\|\Sigma_X^{-1}\|_2, \|\Sigma_Y^{-1}\|_2 \geq 1$, and $\hat{\eta}^\top \hat{\Sigma}_Y^{-1/2} \Sigma_Y^{1/2} \eta \geq 0$ for $k = 1, \dots, K$.*

The sub-Gaussian condition in Assumption 4.1 ensures that X and Y do not have heavy tails, allowing us to use standard concentration results for the estimation of Σ_X and Σ_Y . Strict sub-Gaussianity (Kereta and Klock 2021) facilitates the proofs by allowing the sub-Gaussian norm of a random variable and its variance to be used interchangeably.

In Assumption 4.2, the condition that $d \log(p) = o(N)$ allows p to grow exponentially in N/d (i.e., $p \lesssim e^{N/d}$) while still retaining consistency of the estimator for the canonical vectors. The critical component of the condition $\text{cond}(\Sigma_Y)^2 d = o(N)$ is that $d = o(N)$, which ensures that Σ_Y can be estimated at a sufficiently fast rate by its sample estimator $\hat{\Sigma}_Y$. The presence of $\text{cond}(\Sigma_Y)^2$ allows us to show that $\|\hat{\Sigma}_Y\|_2 \lesssim \|\Sigma_Y\|_2$ and to ignore lower order terms of $\frac{d}{N}$, simplifying the theorem statement. We assume that the correlations $\gamma_1, \dots, \gamma_K$ are distinct in order to estimate each canonical vector separately instead of estimating entire subspaces.

Assumption 4.3 is not essential, and mainly serves to simplify the statement of the theorem. Since the canonical vectors are defined only up to a sign, we use condition $\hat{\eta}^\top \hat{\Sigma}_Y^{1/2} \Sigma_Y^{1/2} \eta \geq 0$ to account for the sign ambiguity of the CCA solutions, allowing us to compare the estimates of the canonical vectors with their population counterparts through the differences $\|\theta_k - \hat{\theta}_k\|_2$ and $\|\eta_k - \hat{\eta}_k\|_2$.

Theorem 4.1. *Suppose Assumptions 4.1-4.3 hold. Fix $\alpha \in (0, 1)$, and for some absolute constant $C > 0$, define the regularization parameter in Algorithm 2 as $\lambda = C \sqrt{\frac{d}{N} \log(p\alpha^{-1})}$. Then, with probability $1 - \alpha$, we have that, for $k = 1, \dots, K$,*

$$\|\theta_k - \hat{\theta}_k\|_2^2 \lesssim \left(\frac{d}{N} \log(p\alpha^{-1}) \right)^{1/2} \frac{\gamma_1^2 \|\Sigma_X\|_{2,\infty} \|T\|_{\ell_1, \ell_2}^2}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)^2} \frac{\|\Sigma_X^{-1}\|_2}{\gamma_k^2}, \quad (12)$$

$$\|\eta_k - \hat{\eta}_k\|_2^2 \lesssim \left(\frac{d}{N} \log(p\alpha^{-1}) \right)^{1/2} \frac{\gamma_1^2 \|\Sigma_X\|_{2,\infty} \|T\|_{\ell_1, \ell_2}^2}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)^2} \|\Sigma_Y^{-1}\|_2, \quad (13)$$

where θ_k and η_k denote the high- and low-dimensional population canonical vectors, respectively.

The proof follows directly from Theorem B.2 in the appendices. We refer to this bound as a “slow”-rate bound, as it makes fewer assumptions but results in slower convergence rates relative to the sample size N . Specifically, we make no sparsity assumptions on the high-dimensional canonical vectors. In Theorem B.3, we provide the “fast”-rate bound, where under more restrictive assumptions, the term $\left(\frac{d}{N} \log(p\alpha^{-1})\right)^{1/2}$ is replaced by $\frac{d}{N} \log(p\alpha^{-1})$, similar to what is observed in lasso regression problems (Hastie et al. 2015). The proof of Theorem 4.1 hinges on two key components: firstly, deterministic group lasso bounds for in-sample prediction error (Gaynanova 2020), and secondly, the rates at which $\|\Sigma_{XY} - \hat{\Sigma}_{XY}\|_{2,\infty}$ and $\|B^\top(\Sigma_X - \hat{\Sigma}_X)B\|_2$ converge to zero under the sub-Gaussian assumptions for X and Y . Here, $\hat{\Sigma}_X$ and $\hat{\Sigma}_{XY}$ represent the sample covariance matrices. As an intermediate step in the proof, we show Theorem B.1, which gives similar slow and fast rate bounds for the estimated canonical correlations $\hat{\gamma}_k$.

Under the stated assumptions the canonical vector estimates are consistent. Moreover, our rates of convergence depend on the dimension of the high-dimensional data, p , only through $\log(p)$. The bounds for the k th canonical directions depend on the nearest canonical correlation gaps, resembling those concerning the variance in the PCA literature.

We emphasize that our rates are dependent on Σ_X only through $\|\Sigma_X\|_{2,\infty}$, and not $\|\Sigma_X\|_2$. The norm $\|\Sigma_X\|_{2,\infty}$ can be much smaller than $\|\Sigma_X\|_2$, particularly when many of the X_j 's are correlated with one another. This property highlights the robustness of the proposed methodology in the high-dimensional setting, where highly correlated covariates are commonplace.

We are able to establish our error bounds for each canonical vector θ_k, η_k , independently, and these bounds depend on each other only through the norms of the canonical vectors $\|T\|_{\ell_1, \ell_2}^2$, and through the neighboring canonical correlation gaps. It is also worth noting that the error associated with θ depends on Σ_X^{-1} but not Σ_Y^{-1} . Similarly, the error associated with η depends on Σ_Y^{-1} but not Σ_X^{-1} . Hence, Y can be poorly behaved without impacting the estimation of θ , and vice-versa.

4.2 Estimation error rates for asymmetric Sparse-Functional CCA

In this section, we investigate the asymptotic properties of our proposed estimators $\hat{\psi}_k$ and $\hat{\theta}_k$, outlined in Algorithm 1, for the canonical functions ψ_k and canonical vectors θ_k . In this setting, the observations are pairs of Riemannian-valued functional data $y_i \in L^2(T\mu)$ and high-dimensional multivariate data $x_i \in \mathbb{R}^p$. Given the technical nature of many of the assumptions, we refer the reader to Assumptions C.1-C.4 in the appendices for a complete list.

As in the multivariate case, we denote with γ_k the k th canonical correlation attained in the population version of the problem, and we denote with $K = \max\{i \in \{1, \dots, d\} : \gamma_i > 0\}$ the number of nontrivial canonical vectors. We again use the conventions $\gamma_{d+1}^2 = -\infty$ and $\gamma_0^2 = \infty$. Recall that we denote by d the rank of the functional data and p the dimension of the multivariate data. We suppose that the canonical vectors $\{\theta_k\}$ are s -sparse with a consistent group structure. We let X_S denote the random vector where we omit covariates $\{X_j\}$ that do not contribute to the association structure with Y . For the high-dimensional terms to match the speed of convergence of the functional terms, we assume that $\Sigma_X^{1/2}$ satisfies the group restricted eigenvalue condition $\text{RE}(s, 3, d)$, introduced in Definition B.3, with parameter $\kappa = \kappa(s, d, \Sigma_X^{1/2})$, which yields ‘fast’-rate bounds.

Theorem 4.2. *For some absolute constant $C > 0$, define the regularization parameter in Algorithm 1 as $\lambda = C\sqrt{\frac{d}{N} \log(p)}$. Then, under Assumptions C.1-C.4, for $k = 1, \dots, K$, we have*

$$\|\psi_k - \Gamma_{\hat{\mu}, \mu} \hat{\psi}_k\|_{\mu}^2 = O_P \left(\frac{d^2 s \log(p)}{N} \frac{\|\psi_k\|_{\mu}^2 \kappa \|\Sigma_X\|_{2, \infty}}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2, \gamma_{k-1} - \gamma_k, \gamma_k - \gamma_{k+1})^2} \right), \quad (14)$$

$$\|\theta_k - \hat{\theta}_k\|_2^2 = O_P \left(\frac{ds \log(p)}{N} \frac{\|\Sigma_{X_S}^{-1}\|_2^{1/2} + \left(\frac{\gamma_1}{\gamma_k}\right)^2 \|\Sigma_X\|_{2, \infty} \|\Sigma_X^{-1}\|_2 \kappa^2}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2, \gamma_{k-1} - \gamma_k, \gamma_k - \gamma_{k+1})^2} \right), \quad (15)$$

where we have omitted the terms $\mathbb{E}[\|\text{Log}_{\mu} y\|_{\mu}^4]$ and $\text{Var}(\langle \phi_j, \text{Log}_{\mu} y \rangle_{\mu})$ for $j = 1, \dots, d$.

The theorem presented here is a special case of Theorems C.2 and C.3. As in Theorem 4.1, the rate of convergence depends on p only through the term $\log(p)$ and on Σ_X only through $\|\Sigma_X\|_{2, \infty}$. Our rate also depends on the dimensionality of the reduced representation of the functional data, d , linearly and quadratically in the estimation of θ_k and ψ_k , respectively. The quadratic term d^2 is most likely not tight but arises from our choice to estimate each ϕ_j via $\hat{\phi}_j$, individually, rather than estimating subspaces. As in Theorem 4.1, the convergence rates depend on the neighboring correlation gaps.

It follows from Theorem 4.2 that if terms other than d , s , p , and N are treated as constants, then, if $d^2 s \log(p) = o(N)$, we have that $\hat{\psi}_k$ and $\hat{\theta}_k$ are consistent estimators for ψ_k and θ_k , respectively. Thus, for the proposed methodology, p is allowed to grow exponentially with respect to $\frac{N}{d^2 s}$ (i.e., $p \lesssim e^{\frac{N}{d^2 s}}$) while consistency is retained.

5 Application to dynamic functional connectivity

5.1 Data and preprocessing

We analyze resting-state fMRI images from 1003 subjects in the Human Connectome Project dataset (Van Essen et al. 2012). Throughout the duration of these 15-minute fMRI scans, participants were at rest and not engaging in any specific activities. Details on the acquisition process can be found in Glasser et al. (2013) and Smith et al. (2013). The fMRI images have been pre-processed using the minimal pre-processing HCP pipeline (Glasser et al. 2013), including spatial artifact, distortion removal, and mapping onto a common reference template (Smith et al. 2013).

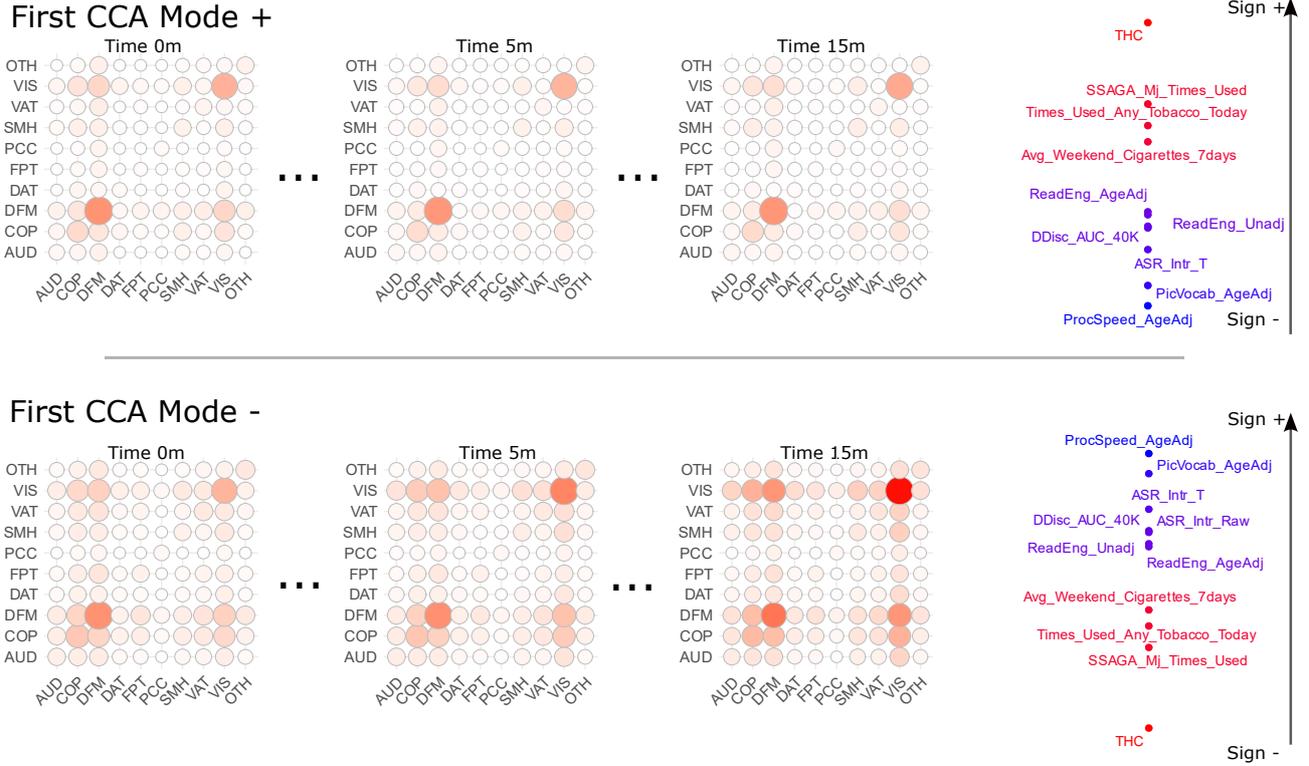


Figure 2: This figure illustrates the first mode of covariation between dynamic connectivity and behavioral measures. On the top panel, we show $(\text{Exp}_{\hat{\mu}}(-c\hat{\psi}_1), -c\hat{\theta}_1)$, which we refer to as ‘First CCA Mode +’, on the bottom panel we show $(\text{Exp}_{\hat{\mu}}(+c\hat{\psi}_1), +c\hat{\theta}_1)$, which we refer to as ‘First CCA Mode -’. These represent two extremities of the spectrum identified by the first mode of covariation. Within each panel, we show the canonical function of SPD covariances $\text{Exp}_{\hat{\mu}}(\pm\hat{\psi}_1)$ at three different times, and a subset of the selected entries of the canonical vector $\pm\hat{\theta}_1$. The depicted mode of covariation suggests that subjects with an increasing variance over time within the visual (VIS) and default mode (DFN) functional systems, as well as an increasing covariance between these systems, positively correlate with higher scores in ‘ProcSpeed_AgeAdj’ – assessing processing speed – and ‘PicVocab_AgeAdj’ – evaluating language/vocabulary comprehension and negatively correlate with using cannabis and opiates (variables THC, SSAGA_Mj-Use, and SSAGA_Times_Used.Opiates).

We define 360 spatially localized regions of interest (ROIs) using the multimodal parcellation proposed in Glasser et al. (2016). These 360 regions are further aggregated into 10 distinct functional systems following the definition in Power et al. (2011). These are the somatosensory/motor network (SMT), cingulo-opercular network (COP), auditory network (AUD), default mode network (DMN), visual network (VIS), frontoparietal network (FPT), salience network (SAL), ventral attention network (VAT), dorsal attention network (DAT), and a category for Other Regions (OTH), which includes areas that are not strictly classified within the aforementioned functional systems.

We partition the fMRI data into 20 time intervals of equal length. For each interval, we reduce the fMRI data to a ‘functional fingerprint’ representation that is a 10×10 SPD covariance that captures the temporal correlation between the fMRI signals of different functional systems within a specific time interval. These matrices are denoted as $y_i(t_j)$ where $i = 1, \dots, N = 1003$ represents the subject and $j = 1, \dots, 20$ denotes the time interval.

In addition, an extensive set of 150 subject traits of lifestyle, demographic, and psychometric measures are also provided for the same cohort of 1003 subjects. We denote these by x_i , with $i = 1, \dots, N = 1003$. To account for potential confounding factors, we regressed out of the 150 variables nine confounders identified in Smith et al. (2015), and the squares of the continuous ones, using

multivariable linear regression.

5.2 Analysis

We apply Algorithm 1 to the pairs $(y_i(\cdot), x_i)$. Specifically, we model the SPD-valued functional data $\{y_i(\cdot)\}$ using the affine-invariant Riemannian metric. The Fréchet mean $\hat{\mu}$ and tangent space representations $\{\text{Log}_{\hat{\mu}} y_i\}$ are computed. See Section 3.1 for details. Both the hyperparameters λ and d , the number of PCs used to reduce the dimension of the SPD-valued functional data, are chosen by cross-validation. Specifically, for every candidate d , the parameter λ is chosen to minimize the cross-validated prediction error of the regression model in equation (8), while d is chosen by examining the scree plot of the cross-validated canonical correlations. We chose the smallest d for which the cross-validated correlations appear to level off, that is, $d = 12$.

The outlined procedure results in a set of K estimated canonical directions $(\hat{\psi}_k, \hat{\theta}_k)_{k=1}^K$, where $\{\hat{\theta}_k\}$ are the canonical vectors associated with $\{x_i\}$, and $\{\hat{\psi}_k\}$ are the (tangent-space) representations of the canonical functions associated with $\{y_i\}$. After inspection of the cross-validated correlations and their associated variance, we decided to retain only the first pair of canonical directions.

5.3 Results and Discussion

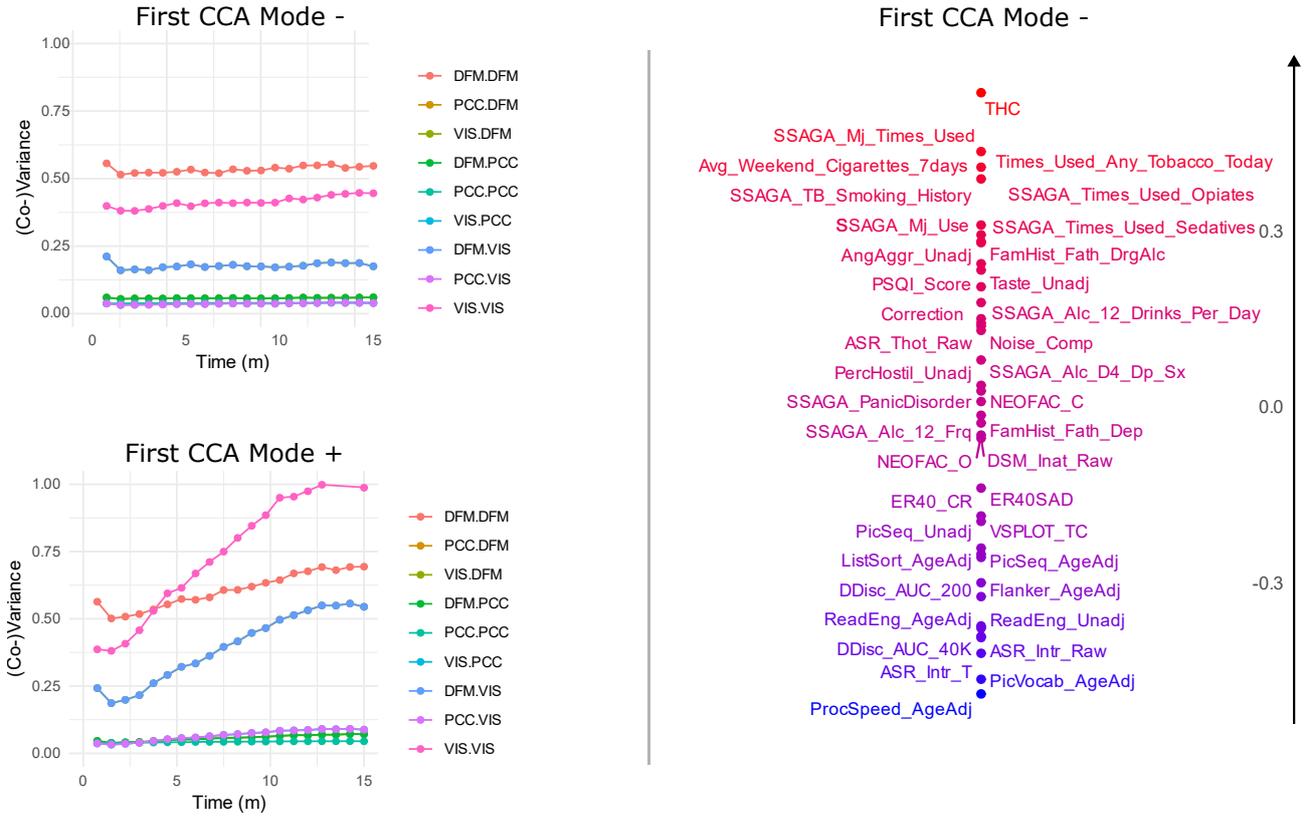


Figure 3: On the left panel, for both ‘First CCA Mode -’ and ‘First CCA Mode +’, we show the temporal dynamics of selected entries of the dynamic mode of connectivity shown in Figure‘2. Notably, some of these, e.g., the DFM-PCC covariance, remain stationary for both ‘First CCA Mode +’ and ‘First CCA Mode -’, while others, e.g., the DFM-VIS covariance, have markedly different patterns. On the right panel, we show a complete list of the 39 variables, of the canonical vector $\pm\hat{\theta}_1$, selected by the proposed model out of an initial set of 150, along with their relative importance.

In Figure 2, we display the first canonical direction $(\hat{\psi}_1, \hat{\theta}_1)$ by plotting

$$(\text{Exp}_{\hat{\mu}}(-c\hat{\psi}_1), -c\hat{\theta}_1), \quad (\text{Exp}_{\hat{\mu}}(+c\hat{\psi}_1), +c\hat{\theta}_1),$$

for a fixed positive constant c . In the figure, we refer to $(\text{Exp}_{\hat{\mu}}(-c\hat{\psi}_1), -c\hat{\theta}_1)$ as ‘First CCA Mode -’ and to $(\text{Exp}_{\hat{\mu}}(+c\hat{\psi}_1), +c\hat{\theta}_1)$ as ‘First CCA Mode +’. Intuitively, these represent the two extremities of the first mode of covariation between functional dynamic connectivity and lifestyle, demographic, and psychometric measures. The exponential map $\text{Exp}_{\hat{\mu}}(\cdot)$ allows us to map the canonical function back to the manifold of SPD-valued functions. The identified mode of covariation appears to link subjects with increasing variance over time within the visual and default mode functional systems, and increasing co-variance over time between these functional systems, to ‘positive’ lifestyle, demographic, and psychometric measures, such as better ‘ProcSpeed_AgeAdj’ score, which tests the ‘speed of processing’ and better ‘PicVocab_AgeAdj’ score, which tests the ability to match an audio recording of a word to the most closely related picture. On the other hand, a more ‘stationary’ connectivity pattern is associated with more ‘negative’ lifestyle traits, such as a positive test for THC (THC), whether the subject has ever used cannabis (SSAGA_Mj_Use), and the number of times the subject has used opiates (SSAGA_Times_Used_Opiates).

The multivariate component of the identified mode of covariation resembles the one found between static functional connectivity and lifestyle, demographic, and psychometric variables in Smith et al. (2015). However, as illustrated in Figure 3, our analysis reveals the non-stationary nature of this mode, with the latter portion of the scan emerging as the most informative in terms of functional connectivity. It is during this phase that the differences between the extremities of the mode of covariation become more evident.

It is plausible that a latent variable linked to both the identified dynamic connectivity and the behavioral components of the first mode of covariation is responsible for the observed correlation between them. This variable potentially reflects the subjects’ experience, such as growing impatience or distractions, during the 15-minute resting-state MRI session where they were instructed not to engage in specific tasks. Indeed, it appears that the ‘First CCA Mode -’ subjects (who are more likely to test positive for THC and have used opiates) maintain a consistent ‘wandering mind’, whereas the ‘First CCA +’ subjects (who are likely to have better pattern completion skills and language/vocabulary comprehension) show a behavioral drift. This results in a progressive activation of the visual cortex and default mode network, and their cooperation, which might reflect a growing unease and consequent search for external stimuli.

6 Simulations

We perform numerical experiments to investigate the finite sample performance of the proposed approach. First, we describe the data generation process. Then, we discuss the metrics utilized to evaluate the methods’ performance. Lastly, we introduce the alternative approaches for comparison with our method and comment on the results.

6.1 Data generation

Recall that $y : \mathcal{T} \rightarrow \mathcal{M}$ is a random Riemannian process, $X \in \mathbb{R}^p$ is a high dimensional random vector, and $\mu : \mathcal{T} \rightarrow \mathcal{M}$ is a fixed smooth curve on \mathcal{M} modeling the population mean of y . Here, we fix $p = 200$ and choose \mathcal{M} to be the manifold of $m \times m$ SPD matrices, with $m = 3$. We let the time domain of y be $\mathcal{T} = [-1, 1]$. In the following, we aim to generate realizations of (y, X) according to a model that ensures that the K population canonical vectors and canonical functions are prespecified vectors $\{\theta_k\}_{k=1}^K \subset \mathbb{R}^p$ and functions $\{\psi_k\}_{k=1}^K \subset L^2(T\mu)$, respectively, with $K = 2$. We apply our proposed method and alternative approaches to this data to estimate the canonical vectors and functions, and then compare these estimates with the prespecified population quantities.

The procedure to generate the data is as follows. Take a random vector $Y \in \mathbb{R}^d$, a set of vectors $\{\eta_k\}_k \subset \mathbb{R}^d$, with $d = 3$, and an orthonormal basis $\{\phi_j\} \subset L^2(T\mu)$. Moreover, define $\text{Log}_\mu y = \sum_{j=1}^d Y_j \phi_j$ and $y = \text{Exp}_\mu(\text{Log}_\mu y)$. It follows from Theorem 2.1 that if the multivariate data (Y, X) have population canonical vector (η_k, θ_k) then the functional/multivariate data (y, X) will have population

canonical pairs (ψ_k, θ_k) , with $\psi_k = \sum_{j=1}^d \eta_{kj} \phi_j$. Additionally, we impose a group-sparse structure on the canonical vectors $\{\theta_k\}$. To replicate a realistic setting, we add an extra mode of variation $W\phi_{d+1}$ to $\text{Log}_\mu y$, with W a random variable that is independent of X and Y , and with $\text{Var}(W) = 1/2$. This aims to contaminate the observations without affecting the canonical functions and vectors.

To generate the multivariate data (Y, X) given prespecified canonical pairs $\{(\eta_k, \theta_k)\}$, we use the model introduced in Chen et al. (2013). This involves setting $\Sigma_X \in \mathbb{R}^{p \times p}$, $\Sigma_Y \in \mathbb{R}^{d \times d}$, the canonical vector pairs $\{(\eta_k, \theta_k)\}$, and the canonical correlations $\gamma_1, \dots, \gamma_K$, for $k = 1, \dots, K$. Then, we define (Y, X) as

$$\begin{pmatrix} Y \\ X \end{pmatrix} \sim \mathcal{N}\left(0, \begin{pmatrix} \Sigma_Y & \Sigma_{YX} \\ \Sigma_{XY} & \Sigma_X \end{pmatrix}\right), \quad (16)$$

where \mathcal{N} denotes the multivariate normal distribution and $\Sigma_{YX} = \Sigma_Y \left(\sum_{k=1}^K \gamma_k \eta_k \theta_k^T\right) \Sigma_X$. It is easy to show that the population canonical vectors of (Y, X) are (η_k, θ_k) with correlations γ_k , for $k = 1, \dots, K$. The set of canonical vectors $\{\eta_k\}$ is defined by generating K orthogonal random vectors, which are then normalized to satisfy the constraint $\eta_k^T \Sigma_Y \eta_j = \delta_{kj}$. Similarly, the canonical vectors $\{\theta_k\}$ are randomly generated and constrained to satisfy the condition $\theta_k^T \Sigma_X \theta_j = \delta_{kj}$. The group sparsity assumption is enforced by ensuring that only $k_1 = 20$ elements of each canonical vector (the same elements across all vectors) are non-zero. Additionally, the variables X_j corresponding to the non-zero components of θ_j have marginal covariance matrix $\Sigma_{X_S} = \text{diag}(\underbrace{2, \dots, 2}_{10}, \underbrace{1, \dots, 1}_{10}) \in \mathbb{R}^{k_1 \times k_1}$. The covariance Σ_X is then

defined as

$$\Sigma_X = \begin{pmatrix} \Sigma_{X_S} & 0 \\ 0 & I_{p-k_1} \end{pmatrix}. \quad (17)$$

The covariance Σ_Y is set to be diagonal with diagonal values being 3, 2, 1. The true canonical correlations are chosen to be $\gamma_1 = .95$ and $\gamma_2 = .6$.

We let the mean curve μ at each $t \in \mathcal{T}$ be a SPD matrix $\mu(t) \in \mathbb{R}^{m \times m}$. We set $\mu(0)$ by randomly generating its eigenvectors and setting the associated eigenvalues equal to (1, 2, 3). The mean $\mu(t)$ at the other time-points $t \in \mathcal{T}$ is generated by applying a time-variant rotation to the eigenvectors of $\mu(0)$. We choose each principal component ϕ_j to take the form $\phi_j(t) = E(t)P(t)$, for $j = 1, \dots, d$, where E is chosen at random from a set of orthogonal basis vectors for $L^2(T\mu)$, and P is chosen at random from a basis of orthogonal polynomials on $[-1, 1]$. This ensures that $\{\phi_j\}$ are orthogonal to one another as elements of $L^2(T\mu)$.

In our experiments, we generate N i.i.d. pairs (Y_i, X_i) from the multivariate CCA model in equation (16), for different choices of N . Next, we generate y_i via $\text{Exp}_\mu\left(\sum_{j=1}^d Y_{ij} \phi_j\right)$ and evaluate it at $L = 50$ locations $t_l \in [-1, 1]$, yielding $y_i(t_l)$ for $i = 1, \dots, N$ and $l = 1, \dots, L$. The observations

$$(\{y_i(t_l)\}_l, X_i)_i, \quad i = 1, \dots, N,$$

are used to estimate the canonical vectors and functions and compare different approaches.

6.2 Metrics

We use the following metrics to compare the estimated accuracy of the models considered.

A. Normalized Euclidean error for the canonical vector

$$\left\| \theta_1 / \|\theta_1\|_2 - \hat{\theta}_1 / \|\hat{\theta}_1\|_2 \right\|_2$$

This is a natural metric for evaluating the estimation accuracy.

B. F1-score for the canonical vector

$$2 \cdot \frac{P \cdot R}{P + R},$$

where $P = \frac{TP}{TP+FP}$ is the precision, $R = \frac{TP}{TP+FN}$ is the recall, TP is the number of true positives, FP the number of false positives, FN the number of false negatives.

C. L^2 Parallel transport error for the canonical function

$$\|\Gamma_{\hat{\mu},\mu}\hat{\psi}_1 - \psi_1\|_\mu$$

This metric allows us to use the $L^2(T\mu)$ norm to compare the estimates to the true population analog, by parallel transporting $\hat{\psi}_1 \in L^2(T\hat{\mu})$ and defining $\Gamma_{\hat{\mu},\mu}\hat{\psi}_1 \in L^2(T\mu)$.

D. Tangent Correlation

Using a large test set $\{\tilde{y}_i, \tilde{x}_i\}$, generated from the same distribution as the training data, we compute the sample correlation as follows:

$$\text{Corr}\left(\left(\langle \text{Log}_\mu \tilde{y}_i, \Gamma_{\hat{\mu},\mu}\hat{\psi}_1 \rangle_\mu\right)_i, \left(\tilde{x}_i^\top \hat{\theta}_1\right)_i\right).$$

We refer to this metric as the ‘Tangent’ correlation as it respects the manifold structure of the data.

E. Euclidean Correlation

Using a large test set $\{\tilde{y}_i, \tilde{x}_i\}$, generated from the same distribution as the training data, we compute the sample correlation as follows:

$$\text{Corr}\left(\left(\text{vec}(\tilde{y}_i)^\top \text{vec}(\hat{\psi}_1)\right)_i, \left(\tilde{x}_i^\top \hat{\theta}_1\right)_i\right).$$

We refer to this metric as the ‘Euclidean’ correlation as it ignores the manifold structure of the data.

6.3 Approaches for comparison

We compare 4 different approaches, detailed below.

1. **Proposed approach.** We apply Algorithm 1 without modifications. We use cross-validation to choose the regularization parameter λ in the group-lasso regression step as implemented by the `glmnet` package (Friedman et al. 2010).
2. **Sparse PCA-based approach:** IRFPCA + sparse PCA + classical CCA. We use IRFPCA, as in our approach, to reduce the dimensionality of the functional data. We use sparse PCA, using the `elasticnet` R package (Zou et al. 2006), to reduce the dimensionality of the multivariate data. Then, we use the estimated PCA scores as input for classical multivariate CCA. We provide sparse PCA with the exact number of principal components that are correlated with the functional data, i.e., $k_1 = 20$, and restrict the number of non-zero principal loadings per principal component to be 2.
3. **Sparse CCA-based approach:** IRFPCA + sparse CCA. We again use IRFPCA to reduce the dimension of the functional data. Next, we use the Penalized Matrix Analysis (PMA) approach to sparse CCA proposed in Witten et al. (2009) to compute canonical pairs between the PC scores from IRFPCA and the high dimensional data. The PMA approach to sparse CCA assumes that the covariance matrices of the data are the identity matrices, giving it a slight disadvantage. We choose the amount of penalization for θ_1 using the suggested permutation-type approach (Witten et al. 2009), and choose the penalization parameter for η_1 to induce virtually no penalization.

4. **Multivariate FPCA-based approach:** Multivariate FPCA + Asymmetric sparse CCA in Algorithm 2. This approach is analogous to the one proposed, except that the IRFPCA step is replaced by multivariate FPCA (Happ and Greven 2018). Therefore, it disregards the SPD manifold structure of the data. Specifically, it transforms each SPD matrix into a vector extracting the lower triangular part of the matrix. Then it applies multivariate FPCA to the resulting vector-valued functions.

We have chosen these alternative approaches in order to dissect specific components of the CCA problem. Specifically, approach 2 isolates the effect of selecting important features and identifying correlated components in two separate stages, and approach 3 isolates the effect of not taking advantage of the group sparsity structure in the canonical vectors and making restrictive assumptions on the covariance of the high-dimensional data. Approach 4 isolates the effect of treating manifold data as if it were Euclidean. Note that Approach 4 is technically solving a different canonical correlation problem than approaches 1-3 as it aims to maximize the Euclidean correlation rather than the tangent space correlation. For this reason, the underlying population canonical vectors and functions differ from those in the proposed model. Therefore, we only use metric E when evaluating the performance of approach 4.

Moreover, depending on the choice of Σ_X , either the PMA sparse CCA approach or the sparse PCA approach is at a disadvantage. Assuming Σ_X to be the identity matrix meets the assumptions of PMA sparse CCA but renders dimension reduction through sparse PCA less effective. Conversely, choosing Σ_X to not be the identity matrix benefits sparse PCA at the expense of the PMA sparse CCA approach.

6.4 Results and Discussion

In our experiments, we set $p = 200$ and vary N . For each value of N , we run 15 trials. We provide the IRFPCA model with the true rank, indicating the number of functional principal components associated with the variable X , that is, $d = 3$.

In Figure 4, we present the performance of approaches 1-3 measured using all defined metrics A-E. As previously mentioned, Approach 4 is assessed using only metric E due to its differing underlying model. In the high-dimensional setting, where $N = 50$ and $p = 200$, all four approaches showed similar performance across all metrics. This setting likely identifies the detectability limits of CCA methods. However, when provided with more samples, our approach quickly outperforms the other approaches. Differences in performance were more notable in the estimation of the Euclidean error for the canonical vectors, F1-score, and out-of-sample correlation. This suggests that the most challenging aspect of the setting considered is estimating the canonical vectors as opposed to the canonical functions. This can be explained by the similar modeling strategies adopted for the functional data. Approach 4, while able to find correlated components in the data according to the Euclidean notion of correlation (E), it suffered from bias due to treating the functional data as Euclidean.

For approach 2, the differences in performance can be explained by its two-step strategy that involves first selecting the important features and reducing the dimension of the multivariate data, followed by identifying correlated components. Specifically, the sparse PCA step is based solely on the variance structure of X , and not on its correlation with the functional data. In our simulation, the variables of X correlated with Y have the same or smaller variance than those not correlated with Y . As a result, sparse PCA, which is unsupervised, struggles to tease them apart.

7 Discussion and Conclusion

In this paper, we introduce a novel statistical model for identifying shared variation patterns between manifold-valued functional data and high-dimensional data. The proposed asymmetric CCA approach is designed to control the complexity of the canonical directions associated with the functional data by using Riemannian FPCA. This facilitates the identification of a lower-dimensional, smooth subspace

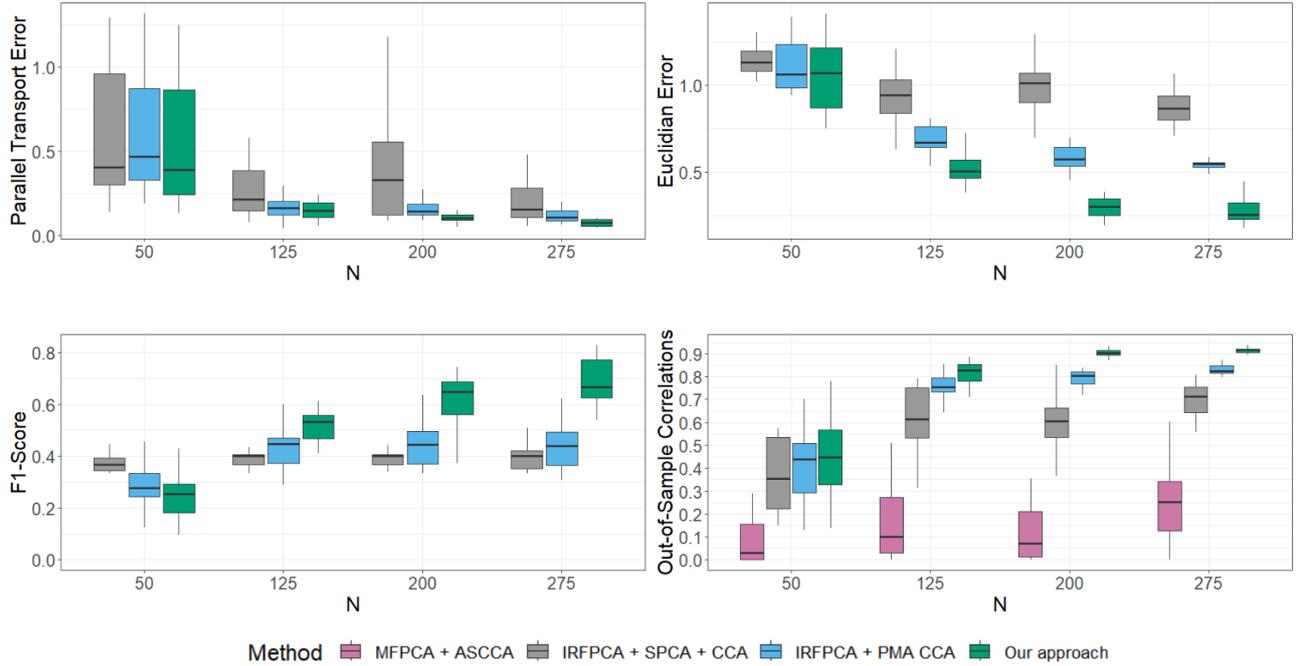


Figure 4: (Top left): Performance evaluation using metric A, which measures the normalized Euclidean error in the first high-dimensional canonical vector, on approaches 1-3. (Top right): Performance evaluation using metric C, which is the parallel transport error in the first canonical function, on approaches 1-3. (Bottom left): Performance evaluation using metric B, the F1-score of the first estimated high dimensional canonical vector compared to the associated population vector, on approaches 1-3. (Bottom right): Performance evaluation using out-of-sample correlations. We use out-of-sample tangent correlation (metric D) for approaches 1-3, and out-of-sample Euclidean correlation (metric E) for approach 4.

onto which these data can be projected. It controls the complexity of the high-dimensional canonical directions, which lack spatial structure, through a sparsity-promoting penalty that leads to the selection of the important variables. As opposed to other methods in the literature, this is achieved without requiring the estimation of the precision matrix of the high-dimensional data, which is in general prohibitive.

We apply asymmetric CCA to explore the association structure between resting-state dynamic functional connectivity, represented as time-indexed covariance matrices, and high-dimensional behavioral, lifestyle, and demographic features. Our analysis reveals a non-stationary pattern in functional connectivity, indicating that the usual assumption of temporal stationarity may not hold, even in resting-state studies. While this work focuses on an application in dynamic connectivity, the proposed method can be easily adapted to accommodate different Riemannian structures and to employ different data representation models, paving the way for several future extensions.

Appendices

The appendices are organized as follows. In Appendix A, we formalize the CCA problem for random elements of Hilbert spaces and prove Theorem 2.1. In Appendix B, we present intermediate results and associated proofs, and conclude with the proof of Theorem 4.1. In Appendix C, we prove Theorem 4.2, our main theoretical result on the asymptotic errors made in estimating the canonical vectors and functions. In Appendix D, we present several norm and matrix identities that are utilized throughout the appendices. Finally, in Appendix E, we provide additional details on the IRFPCA algorithm

(Lin and Yao 2019), which is used in the proposed Algorithm 1. We also present an explicit basis construction for the space of symmetric positive definite matrices equipped with the affine invariant metric, providing a computational speed-up compared to the Gram-Schmidt procedure proposed in Lin and Yao (2019).

A Canonical Correlation Analysis of Random Elements of Hilbert Spaces

In this section, we provide a more rigorous formalization of the CCA model. We mirror the development of Hsing and Eubank (2015) and Huang and Renaut (2015), but provide a less technical presentation by emphasizing the role of the canonical variables rather than the canonical vectors when formulating the general CCA problem. For an introduction to Hilbert space concepts and random elements taking values in Hilbert spaces, we refer to Hsing and Eubank (2015). For an introduction to classical CCA, we refer to Uurtio et al. (2018).

In Section A.1, we define the infinite-dimensional version of the CCA problem, and establish the existence of solutions in our asymmetric setting (Theorem A.1). In Section A.2, we state preliminary definitions and results for the subsequent sections. In Section A.3, we state the necessary assumption (Assumption A.1) to reduce the infinite-dimensional CCA problem to a finite-dimensional CCA problem (Theorem A.2). In Section A.4, we prove the results of the section and additionally prove Theorem 2.1.

A.1 Problem Statement

Let χ_1 and χ_2 be measurable functions from a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ to separable Hilbert spaces \mathbb{H}_1 and \mathbb{H}_2 , respectively (See Section 7.2. of Hsing and Eubank 2015). Here, \mathbb{H}_1 and \mathbb{H}_2 are arbitrary Hilbert spaces, but throughout the paper they correspond to $\mathbb{H}_1 = L^2(T\mu)$ and $\mathbb{H}_2 = \mathbb{R}^p$, and similarly χ_1 and χ_2 correspond to $\chi_1 = \text{Log}_\mu y$ and $\chi_2 = X$. Hilbert space inner products are denoted by $\langle \cdot, \cdot \rangle$, with associated norms $\|\cdot\|$. The specific choice of the norm or inner product will be clear from the context. We assume that $\mathbb{E}[\|\chi_i\|^2] < \infty$ so that the mean and covariance of χ_i are well-defined for $i = 1, 2$. The mean element of χ_i is defined as $h_i \equiv \mathbb{E}[\chi_i] \in \mathbb{H}_i$, and for simplicity, we assume $h_i = 0$ for $i = 1, 2$.

A seemingly natural way to formalize the canonical correlation problem for the infinite-dimensional case, which is analogous to the finite-dimensional case, is

$$\underset{f \in \mathbb{H}_1, g \in \mathbb{H}_2}{\text{maximize}} \text{Corr}^2(\langle \chi_1, f \rangle, \langle \chi_2, g \rangle) \quad (18)$$

where Corr is the usual correlation defined between two finite-variance real-valued random variables defined on $(\Omega, \mathcal{F}, \mathbb{P})$. If they exist, the solution (f, g) would be the first canonical vector pair. Equivalently, we can write this problem in terms of the canonical variables U, V as

$$\underset{U \in \mathcal{U}, V \in \mathcal{V}}{\text{maximize}} \text{Corr}^2(U, V) \quad (19)$$

where $\mathcal{U} = \{\langle \chi_1, f \rangle : f \in \mathbb{H}_1\}$, $\mathcal{V} = \{\langle \chi_2, g \rangle : g \in \mathbb{H}_2\}$. However, the maximum of this problem may not be attained by any $U \in \mathcal{U}$, $V \in \mathcal{V}$ (Cupidon et al. 2008).

It turns out we can amend this by simply taking the closures of \mathcal{U} and \mathcal{V} . Let $\mathbb{L}^2(\Omega, \mathcal{F}, \mathbb{P})$ denote the Hilbert space of square-integrable random variables on Ω with inner product $\langle U, V \rangle = \text{Cov}(U, V)$ for $U, V \in \mathbb{L}^2(\Omega, \mathcal{F}, \mathbb{P})$. Note that U being square-integrable means that $\text{Var}(U) < \infty$. If we replace \mathcal{U} , \mathcal{V} in the problem above with their closures as subsets of $\mathbb{L}^2(\Omega, \mathcal{F}, \mathbb{P})$, denoted $\bar{\mathcal{U}}$, $\bar{\mathcal{V}}$, (for further discussion of $\bar{\mathcal{U}}$, $\bar{\mathcal{V}}$, see the discussion following Example 7.6.5 of Hsing and Eubank (2015)) then it can be shown that the maximum *will* be attained for some $U \in \bar{\mathcal{U}}$, $V \in \bar{\mathcal{V}}$, provided that a certain linear operator is assumed to be compact. In our setting, this compactness condition holds because we use

$\mathbb{H}_2 = \mathbb{R}^p$, a finite-dimensional space. The result can be found in Theorem 10.1.2 of Hsing and Eubank (2015), which we restate in our context here.

The following result establishes the existence of solutions to the general CCA problem in our asymmetric setting where $\dim(\mathbb{H}_1) = \infty$ but $\dim(\mathbb{H}_2) = p < \infty$, and that there are at most p nontrivial solutions.

Theorem A.1. *If $\dim(\mathbb{H}_1) = \infty$ and $\dim(\mathbb{H}_2) = p < \infty$, then there exists $U_1 \in \overline{\mathcal{U}}$ and $V_1 \in \overline{\mathcal{V}}$ which are solutions to*

$$\sup_{U \in \overline{\mathcal{U}}, V \in \overline{\mathcal{V}}} \text{Corr}^2(U, V),$$

with $\text{Var}(U_1) = \text{Var}(V_1) = 1$. For $k = 2, \dots, p$, there exists $U_k \in \overline{\mathcal{U}}$ and $V_k \in \overline{\mathcal{V}}$, which are solutions to

$$\sup_{\substack{U \in \overline{\mathcal{U}}: \text{Cov}(U, U_i) = 0, i=1, \dots, k-1 \\ V \in \overline{\mathcal{V}}: \text{Cov}(V, V_i) = 0, i=1, \dots, k-1}} \text{Corr}^2(U, V),$$

with $\text{Var}(U_k) = \text{Var}(V_k) = 1$. Moreover, for all $U \in \overline{\mathcal{U}}$ and $V \in \overline{\mathcal{V}}$ which are uncorrelated with $U_1, \dots, U_p, V_1, \dots, V_p$, respectively, the pair (U, V) is a trivial solution, that is,

$$\sup_{\substack{U \in \overline{\mathcal{U}}: \text{Cov}(U, U_i) = 0, i=1, \dots, k \\ V \in \overline{\mathcal{V}}: \text{Cov}(V, V_i) = 0, i=1, \dots, k}} \text{Corr}^2(U, V) = 0.$$

The pairs $\{(U_k, V_k)\}_{k=1}^p$ are called the canonical variable pairs. We refer to the problem of finding $\{(U_k, V_k)\}_{k=1}^p$ as the population canonical correlation problem.

Remark 2. *Intuitively, we must cast the problem in terms of the canonical variables $U \in \overline{\mathcal{U}}$, $V \in \overline{\mathcal{V}}$ rather than canonical vectors $f \in \mathbb{H}_1$, $g \in \mathbb{H}_2$ because $\mathcal{U} = \{\langle \chi_1, f \rangle : f \in \mathbb{H}_1\}$ and $\mathcal{V} = \{\langle \chi_2, g \rangle : g \in \mathbb{H}_2\}$ are not large enough for the supremum of the CCA problem to be attained. To emphasize this point, given an optimal $U \in \overline{\mathcal{U}}$ of the form $U = \lim_{j \rightarrow \infty} \langle \chi_1, f_j \rangle$ for some sequence $(f_j)_j \in \mathbb{H}_1$, then U can not necessarily be written as an inner product $U = \langle f, \chi_1 \rangle$, because $\lim_{j \rightarrow \infty} f_j$ may not belong to \mathbb{H}_1 .*

In the next section, we introduce an assumption that allows us to make the infinite-dimensional CCA problem finite dimensional, and furthermore formulate the CCA problem in terms of the canonical vectors rather than the canonical variables. From now on, we assume $\dim(\mathbb{H}_2) = p < \infty$ as in Theorem A.1.

A.2 Background

We begin with preliminary definitions and properties of the random element χ_1 . In particular, we introduce the covariance operator \mathcal{K}_1 . The eigenvectors of \mathcal{K}_1 , also known as the principal components of χ_1 , are fundamental to our approach for two reasons: they provide a data-driven subspace for projecting χ_1 and their properties simplify our proofs.

Given that $\mathbb{E}[\|\chi_1\|^2] < \infty$ and χ_1 has mean 0, the covariance operator of χ_1 is well-defined as $\mathcal{K}_1 \equiv \mathbb{E}[\chi_1 \otimes \chi_1]$. Here, the tensor product $f \otimes g : \mathbb{H} \rightarrow \mathbb{H}$ between $f, g \in \mathbb{H}$ for a Hilbert space \mathbb{H} , is defined as $(f \otimes g)(h) = \langle f, h \rangle g$ for all $h \in \mathbb{H}$.

In the lemma below, we collect the properties of χ_1 and \mathcal{K}_1 used in what follows. An orthonormal sequence of elements $\{e_j\}_{j=1}^\infty$ of a Hilbert space \mathbb{H} such that $\overline{\text{span}\{e_j\}} = \mathbb{H}$ is referred to as a complete orthonormal system (CONS) for \mathbb{H} .

Lemma A.1. *Let $\text{Im}(\mathcal{K}_1)$ denote the image of \mathcal{K}_1 and $\overline{\text{Im}(\mathcal{K}_1)} \subseteq \mathbb{H}_1$ denote its closure in \mathbb{H}_1 . Then, the following statements hold.*

1. *With probability 1, $\chi_1 \in \overline{\text{Im}(\mathcal{K}_1)}$, and for any $f \in \text{Im}(\mathcal{K}_1)^\perp$, $\langle f, \chi_1 \rangle = 0$*

2. \mathcal{K}_1 has the eigendecomposition $\mathcal{K}_1 = \sum_{j=1}^{\infty} \omega_j e_j \otimes e_j$, where $e_j \in \mathbb{H}_1$ for $j = 1, \dots, \infty$ and $\omega_1 \geq \omega_2 \geq \dots \geq 0$. $\{e_j\}_{j=1}^{\infty}$ forms a CONS of $\overline{\text{Im}(\mathcal{K}_1)}$, and $\omega_j \rightarrow 0$ as $j \rightarrow \infty$. We refer to $\{(e_j, \omega_j)\}_{j=1}^{\infty}$ as the eigensystem of \mathcal{K}_1 , with eigenvectors e_j and eigenvalues ω_j .
3. With probability 1, $\chi_1 = \sum_{j=1}^{\infty} \langle \chi_1, e_j \rangle e_j$. We refer to the $\{\langle \chi_1, e_j \rangle\}_{j=1}^{\infty}$ as the principal scores, and they are uncorrelated random variables with $\mathbb{E}[\langle \chi_1, e_j \rangle] = 0$ and $\text{Var}(\langle \chi_1, e_j \rangle) = \omega_j$.
4. $\{\langle \chi_1, e_j \rangle / \omega_j^{1/2}\}$ forms a CONS for \overline{U} and $\overline{U} = \{\sum_{j=1}^{\infty} a_j \langle \chi_1, e_j \rangle : \sum_{j=1}^{\infty} \omega_j a_j^2 < \infty\}$.

Remark 3. Item 1 elucidates the role of $\overline{\text{Im}(\mathcal{K}_1)}$ as the subspace of \mathbb{H}_1 where χ_1 resides. Item 2 shows that the eigenvectors $\{e_j\}$ are a CONS for $\overline{\text{Im}(\mathcal{K}_1)}$, which implies Item 3, that the principal scores $\{\langle \chi_1, e_j \rangle\}$ characterize χ_1 . Item 4 establishes that the set of potential canonical variables, \overline{U} , is equivalent to the set of linear combinations of the principal scores.

For an arbitrary CONS $\{e_j\}_{j=1}^{\infty}$ for $\overline{\text{Im}(\mathcal{K}_1)}$, the associated scores $\{\langle e_j, \chi_1 \rangle\}$ may not be orthogonal in \overline{U} , but as the following result shows, the associated scores still span \overline{U} . This is the property that allows us to not rely on the principal component basis and instead use an arbitrary CONS for \mathbb{H}_1 in Assumption A.1.

Lemma A.2. For any complete orthonormal system $\{e_j\}_{j=1}^{\infty}$ of $\overline{\text{Im}(\mathcal{K}_1)}$, $\overline{U} = \overline{\text{span}\{\langle e_j, \chi_1 \rangle\}}$. Thus, any element of $U \in \overline{U}$ can be written as $\sum_{j=1}^{\infty} a_j \langle \chi_1, e_j \rangle$ where the a_j are such that $\text{Var}(U) < \infty$.

A.3 Reduction to a finite dimensional problem

We can now introduce Assumption A.1 on the correlation structure between χ_1 and χ_2 .

Assumption A.1. There exists a complete orthonormal system $\{e_j\}_{j=1}^{\infty}$ for \mathbb{H}_1 and a set of indices $I \subset \{1, 2, \dots\}$, with finite cardinality $|I| = d$, such that

$$\text{Corr}(V, \langle \chi_1, e_j \rangle) = 0, \quad \forall V \in \overline{V}, j \in I^c, \quad (20)$$

$$\text{Corr}(\langle \chi_1, e_i \rangle, \langle \chi_1, e_j \rangle) = 0, \quad \forall i \in I, \forall j \in I^c, \quad (21)$$

where I^c denotes the complement of I in $\mathbb{N} = \{1, 2, \dots\}$.

Remark 4. The complete orthonormal system $\{e_j\}_j$ is not required to be the principal component basis. In the case when $\mathbb{H}_2 = \mathbb{R}^p$ and $\chi_2 = X$, equation (20) can be rewritten as

$$\text{Corr}(X_i, \langle \chi_1, e_j \rangle) = 0, \quad \forall i = 1, \dots, p, j \in I^c. \quad (22)$$

Intuitively, this assumption states that all elements ψ of \mathbb{H}_1 whose projections $\langle \chi_1, \psi \rangle$ are correlated with X belong to a d -dimensional subspace. This is the assumption stated in Remark 1 of the main manuscript.

Remark 5. Assumption A.1 is weaker than the assumption that χ_1 admits a finite-dimensional representation $\chi_1 = \sum_{j=1}^d \langle \chi_1, e_j \rangle e_j$, for a set of vectors $\{e_j\}_{j=1}^d \subset \mathbb{H}_1$. To see this, we first note that the elements $\{e_j\}_{j=1}^d \subset \mathbb{H}_1$ are orthonormal. Then, we complete $\{e_j\}_{j=1}^d \subset \mathbb{H}_1$ to form a CONS $\{e_j\}_{j=1}^{\infty}$ for \mathbb{H}_1 , and take $I = \{1, \dots, d\}$. Given that the elements $\{e_j\}$ are orthonormal, we have that $\langle \chi_1, e_j \rangle = 0$ for all $j \in I^c$, with probability 1. Hence, conditions (20) and (21) are satisfied.

Making this assumption enables us to reduce the infinite-dimensional CCA problem to a finite-dimensional CCA problem; moreover, it allows us to formulate the CCA problem in terms of the population quantities of interest, the canonical vectors, rather than the canonical variables (Theorem A.2). Theorem A.2 can be viewed as a generalization of Theorem 1 of Krzyśko and Waszak (2013). They assume that χ_1 has a finite-dimensional representation, whereas here we make the weaker Assumption A.1.

Theorem A.2. *Reorder the complete orthonormal system $\{e_j\}_{j=1}^\infty$ for \mathbb{H}_1 in Assumption A.1 so that $I = \{1, \dots, d\}$. Then, under Assumption A.1, the solution to the population canonical correlation problem in Theorem A.1 is found for a $U \in \overline{\mathcal{U}}_d = \{\sum_{i=1}^d a_i \langle \chi_1, e_i \rangle : a_i \in \mathbb{R}\}$. Moreover, when $\mathbb{H}_2 = \mathbb{R}^p$ and $\chi_2 = X$, the problem is equivalent to the following multivariate (finite-dimensional) canonical correlation problem, where Y is the d -dimensional random vector such that $Y_j = \langle \chi_1, e_j \rangle$, the j th score associated with e_j , for $j = 1, \dots, d$:*

$$(a_1, b_1) = \arg \max_{a \in \mathbb{R}^d, b \in \mathbb{R}^p, \text{Var}(a^\top Y) = \text{Var}(b^\top X) = 1} \text{Corr}^2(a^\top Y, b^\top X), \quad (23)$$

$$(a_k, b_k) = \arg \max_{\substack{a \in \mathbb{R}^d, b \in \mathbb{R}^p, \text{Var}(a^\top Y) = \text{Var}(b^\top X) = 1 \\ \text{Cov}(a^\top Y, a_i^\top Y) = 0, i=1, \dots, k \\ \text{Cov}(b^\top X, b_i^\top X) = 0, i=1, \dots, k}} \text{Corr}^2(a^\top Y, b^\top X), \quad k = 2, \dots, \min(p, d). \quad (24)$$

We call the pair $(\sum_{j=1}^d a_{kj} e_j, b_k)$ the k th canonical pair, since $a^\top Y = \langle \sum_{j=1}^d a_{kj} e_j, \chi_1 \rangle$, and $b^\top X = \langle b, \chi_2 \rangle$, where a_{kj} is the j th entry of a_k , for $k = 1, \dots, \min(p, d)$.

This result is central to the proof of Theorem 2.1.

A.4 Proofs

Proof of Lemma A.1:

The first item is part 3 of Theorem 7.2.5. of Hsing and Eubank (2015). The second and third items are Theorem 7.2.6. and Theorem 7.2.7. of Hsing and Eubank (2015), respectively. For the fourth item, orthonormality is clear from item 3 and the fact that they form a CONS for $\overline{\mathcal{U}}$ follows from equation (7.42) of Hsing and Eubank (2015). The last equality follows from the first part of item 4 and because we can calculate $\text{Var}(\sum_{j=1}^\infty a_j \langle \chi_1, e_j \rangle) = \sum_{j=1}^\infty \omega_j a_j^2$ by the continuity of the inner product. \square

Proof of Lemma A.2:

We begin by employing the fourth item of Lemma A.1, which states that $\{\langle \chi_1, e_j \rangle / \omega_j^{1/2}\}$ is a CONS for $\overline{\mathcal{U}}$, where the e_j are the eigenvectors of \mathcal{K}_1 , and ω_j are the corresponding eigenvalues. Therefore, given an arbitrary CONS for $\overline{\text{Im}(\mathcal{K}_1)}$, $\{f_j\}_{j=1, \dots, \infty}$, to complete the proof it suffices to show that $\text{span}\{\langle f_j, \chi_1 \rangle\} = \text{span}\{\langle e_j, \chi_1 \rangle\}$.

To show the \subseteq direction, it suffices to show that $\langle f_k, \chi_1 \rangle \in \overline{\text{span}\{\langle e_j, \chi_1 \rangle\}}$, for every k , by the definitions of closure and span of a set of vectors. Since the functions $\{e_j\}$ form a CONS for $\overline{\text{Im}(\mathcal{K}_1)}$, there exists a sequence $(a_j)_{j=1}^\infty$ of scalars such that $f_k = \sum_{j=1}^\infty a_j e_j$. Therefore, $\langle f_k, \chi_1 \rangle = \sum_{j=1}^\infty a_j \langle \chi_1, e_j \rangle$ by continuity of the inner product on \mathbb{H}_1 , and we have $\langle f_k, \chi_1 \rangle \in \overline{\text{span}\{\langle e_j, \chi_1 \rangle\}}$.

To show the \supseteq direction, we must show that $\langle e_k, \chi_1 \rangle \in \overline{\text{span}\{\langle f_j, \chi_1 \rangle\}}$ for every k , which follows by similar arguments. \square

Proof of Theorem A.2:

We prove the statement for the first canonical pair; the proof for the remaining canonical pairs follows from a similar argument. Let (U, V) be the first canonical pair of the population canonical correlation problem in Theorem A.1. We consider the CONS $\{e_j\}_{j=1}^\infty$ for \mathbb{H}_1 from Assumption A.1, and reorder its elements so that $I = \{1, \dots, d\}$. By Lemma A.2, we write $U \in \overline{\mathcal{U}}$ as $U = \sum_{j=1}^\infty a_j \langle \chi_1, e_j \rangle$. We will show that under Assumption A.1, we can find a $Q = \sum_{j=1}^\infty q_j \langle \chi_1, e_j \rangle \in \overline{\mathcal{U}}$, with $q_j = 0$ for $j > d$, that attains the same maximum value as U . Thus we will have $Q \in \overline{\mathcal{U}}_d \equiv \{\sum_{i=1}^d a_i \langle \chi_1, e_i \rangle : a_i \in \mathbb{R}\}$, completing the proof of the first statement of the Theorem. For random variables with variance 1, such as U and V , we have that $\text{Cov}(U, V) = \text{Corr}(U, V)$. We use these interchangeably throughout the proof.

If the optimum value of $\text{Corr}^2(U, V)$ is 0, then we can select any Q with $q_j = 0$ for $j > d$ and $\text{Var}(Q) = 1$. Therefore, from now on, we focus on the case where $\text{Corr}^2(U, V) \neq 0$. Let

$$U = \sum_{j=1}^d a_j \langle \chi_1, e_j \rangle + \sum_{j=d+1}^{\infty} a_j \langle \chi_1, e_j \rangle \equiv W + Z. \quad (25)$$

Then, from continuity of the inner product and Assumption A.1, it follows that $\text{Cov}(Z, V) = 0$ and $\text{Cov}(W, Z) = 0$, from conditions (20) and (21) respectively. Before constructing Q , we note that the variance of W must be less than or equal to 1. To see this, we use

$$1 = \text{Var}(U) = \text{Var}(W + Z) = \text{Var}(W) + 2\text{Cov}(W, Z) + \text{Var}(Z) = \text{Var}(W) + \text{Var}(Z), \quad (26)$$

since $\text{Cov}(W, Z) = 0$. Then, $\text{Var}(W) \leq 1$ since both $\text{Var}(W)$ and $\text{Var}(Z)$ are positive and sum to 1.

Now, we construct a canonical variable Q with the desired property. The optimal value of the CCA population problem in Theorem A.1 under assumption A.1 is

$$\text{Corr}^2(U, V) = \text{Cov}^2(W + Z, V) \quad (27)$$

$$= (\text{Cov}(W, V) + \text{Cov}(Z, V))^2 \quad (28)$$

$$= \text{Cov}^2(W, V) \quad (29)$$

$$= \text{Corr}^2(W, V) \text{Var}(W). \quad (30)$$

Having established $\text{Var}(W) \leq 1$, there are three cases, either $\text{Var}(W) = 0$, $0 < \text{Var}(W) < 1$, or $\text{Var}(W) = 1$. In the case $\text{Var}(W) = 1$, we take $Q = W$, and using equation (30), we see that the pair (W, V) attains the same maximum correlation as (U, V) . This completes the proof as W is of the desired form. Now, we will show that the other two cases, $\text{Var}(W) = 0$, $0 < \text{Var}(W) < 1$, are not possible. $\text{Var}(W) = 0$ cannot hold since, by equation (30), we would have $\text{Cov}(U, V) = 0$, which we have already ruled out. Assume towards a contradiction that $0 < \text{Var}(W) < 1$, let $c = \frac{1}{\text{Var}(W)^{1/2}} > 1$, and define $Q = cW$. Then, we have that $\text{Var}(Q) = c^2 \text{Var}(W) = 1$, and

$$\text{Corr}^2(U, V) = \text{Corr}^2(W, V) \text{Var}(W) < \text{Corr}^2(Q, V), \quad (31)$$

by equation (30), $\text{Corr}^2(W, V) = \text{Corr}^2(Q, V)$, and $\text{Var}(W) < 1$. However, this is a contradiction as it would imply that the pair (Q, V) attains a larger value of the objective than (U, V) . This completes the proof of the first statement.

Having established the existence of a solution of the stated form for Q , that we are able to reformulate the CCA problem in terms of the finite-dimensional vectors $\{a_k\}_k$ rather than $U \in \bar{U}$ follows from the definition of \bar{U}_d and the bilinearity of the inner product $\langle \cdot, \cdot \rangle$ on \mathbb{H}_1 . In the case that $\mathbb{H}_2 = \mathbb{R}^p$ and $\chi_2 = X$, that we are able to reformulate the problem in terms of the finite-dimensional vectors $\{b_k\}_k$ rather than $V \in \bar{V}$ is due to the following argument. We have $\mathcal{V} \equiv \{\langle \chi_2, g \rangle : g \in \mathbb{H}_2\} = \text{span}\{\langle \chi_2, e_j \rangle, j = 1, \dots, p\}$ (where the e_j here are the standard unit vectors for \mathbb{R}^p) is isomorphic to \mathbb{R}^p , which is complete. Thus, $\{\langle \chi_2, g \rangle : g \in \mathbb{H}_2\}$ is complete, so its completion in $\mathbb{L}^2(\Omega, \mathcal{F}, \mathbb{P})$ is itself, i.e. $\bar{\mathcal{V}} = \mathcal{V}$. Therefore, $\bar{V} = \text{span}\{\langle \chi_2, e_j \rangle, j = 1, \dots, p\} = \{g^\top X : g \in \mathbb{R}^p\}$, i.e. the set of linear combinations of X_1, \dots, X_p .

The number of nontrivial canonical variables has changed from p in Theorem A.1 to $\min(p, d)$. This is because, in a finite-dimensional CCA problem concerning random vectors of dimensions p and d , the smaller of the two dimensions is the upper limit for the number of nontrivial canonical variables (Uurtio et al. 2018). This completes the proof. \square

A.5 Proof of Theorem 2.1

Given that Assumption 2.1 is a special case of Assumption A.1, by applying Theorem A.2, we readily derive the first part of the theorem. This establishes that there are at most d nontrivial canonical

variable pairs (U_k, V_k) . Moreover, each pair (U_k, V_k) can be written in terms of the canonical directions $U_k = \langle \langle \text{Log}_\mu y, \psi_k \rangle \rangle_\mu$ and $V_k = X^\top \theta_k$, for some $\psi_k \in L^2(T\mu)$ and $\theta_k \in \mathbb{R}^p$. Additionally, $(\psi_k, \theta_k) = (\sum_{j=1}^d a_{kj} \phi_j, b_k)$, where the pairs (a_k, b_k) are defined in Theorem A.2 as the solution to a multivariate CCA problem, and the functions $\{\phi_j\}$ form the CONS for $L^2(T\mu)$, defined in Assumption 2.1.

It remains to be shown that the solutions (a_k, b_k) to the multivariate CCA problem can be characterized by the equations (4)-(7). We focus on the infinite-dimensional optimization problem, in equation (23), that defines the first canonical pair. This is equivalent to

$$\sup_{a_1^\top \Sigma_X a_1 = 1 = b_1^\top \Sigma_Y b_1} a_1^\top \Sigma_{XY} b_1.$$

Now using the assumption that Σ_X and Σ_Y are invertible, we make a change of variables $\tilde{a}_1 = \Sigma_X^{-1/2} a_1$, $\tilde{b}_1 = \Sigma_Y^{-1/2} b_1$ and obtain the equivalent problem

$$\sup_{\tilde{a}_1^\top \tilde{a}_1 = 1 = \tilde{b}_1^\top \tilde{b}_1} \tilde{a}_1^\top \Sigma_X^{-1/2} \Sigma_{XY} \Sigma_Y^{-1/2} \tilde{b}_1$$

Let $U\Gamma V^\top = \Sigma_X^{-1/2} \Sigma_{XY} \Sigma_Y^{-1/2}$ be a singular value decomposition of $\Sigma_X^{-1/2} \Sigma_{XY} \Sigma_Y^{-1/2}$, where $U \in \mathbb{R}^{p \times d}$, $\Gamma \in \mathbb{R}^{d \times d}$, $V \in \mathbb{R}^{d \times d}$, $U^\top U = I_d = V^\top V$, and where Γ is a diagonal matrix with the diagonal elements $\gamma_1, \dots, \gamma_d$, in descending order. Note that $p \geq d$. Then it follows from standard properties of the SVD that the first columns of U and V , denoted as u_1 and v_1 respectively, are the solutions to the above problem, i.e. $(\tilde{a}_1, \tilde{b}_1) = (u_1, v_1)$. Similarly, it can be shown that the k th columns of U and V , u_k and v_k respectively, are such that $(\tilde{a}_k, \tilde{b}_k) = (u_k, v_k)$, and that the optimal correlations are the singular values $\gamma_1, \dots, \gamma_d$. Undoing the change of variables, it can be seen that the solutions to the original problems in equation (23) are the pairs formed by the k th columns of the matrices $\Sigma_X^{-1/2} U$ and $\Sigma_Y^{-1/2} V$. The associated correlations are the diagonal entries of Γ^2 .

Now let B be the solution to the optimization problem

$$\underset{B \in \mathbb{R}^{p \times d}}{\text{minimize}} \quad E \left[\left\| \Sigma_Y^{-1/2} Y - B^\top X \right\|_F^2 \right]. \quad (32)$$

It is straightforward to show that $B = \Sigma_X^{-1} \Sigma_{XY} \Sigma_Y^{-1/2}$. Therefore, we have

$$\Sigma_X^{1/2} B = U\Gamma V^\top, \quad (33)$$

$$B^\top \Sigma_X B = V\Gamma^2 V^\top \quad (34)$$

and

$$BV\Gamma^{-1} = \Sigma_X^{-1/2} U. \quad (35)$$

Identifying \tilde{H} , H , and T in equations (4)-(7) with V , $\Sigma_Y^{-1/2} V$, and $\Sigma_X^{-1/2} U$, respectively, completes the proof.

B Asymmetric Sparse CCA: Proof of Theorem 4.1

B.1 Notation

For a vector $x \in \mathbb{R}^p$ with entries $\{x_j\}$ we define its infinity norm $\|x\|_\infty = \max_j(|x_j|)$, its Euclidean norm $\|x\|_2 = \sqrt{\sum_{j=1}^p x_j^2}$, and its ℓ_1 norm $\|x\|_1 = \sum_{j=1}^p |x_j|$. For a matrix $A \in \mathbb{R}^{p \times d}$ with singular values $\sigma_1, \dots, \sigma_d$, its operator norm is $\|A\|_2 = \max_i(|\sigma_i|)$. To denote the i th row of the matrix A , we use A_i , and for the entry in the i th row and j th column, we use a_{ij} . We define the matrix norms $\|A\|_F = \left(\sum_{i=1}^p \sum_{j=1}^d a_{ij}^2 \right)^{1/2}$, $\|A\|_{\ell_1, \ell_2} = \sum_{i=1}^p \|A_i\|_2$, and $\|A\|_{\max} = \max_{(i,j)} |a_{i,j}|$.

Given the normed spaces $(\mathbb{R}^d, \|\cdot\|_\alpha)$ and $(\mathbb{R}^p, \|\cdot\|_\beta)$, and a matrix $A \in \mathbb{R}^{p \times d}$, we define the matrix norm induced by $\|\cdot\|_\alpha$ and $\|\cdot\|_\beta$ as

$$\|A\|_{\alpha, \beta} = \sup_{\|x\|_\alpha=1} \|Ax\|_\beta. \quad (36)$$

For additional properties of the matrix norms used throughout the paper, we refer to Section D. We use the notation $x \lesssim y$ for $x, y \in \mathbb{R}$ to indicate that $x \leq Cy$, with C some positive absolute constant.

B.2 Sub-Gaussian random vectors

Now we briefly define sub-Gaussian random vectors and state basic properties that we use in the proofs. We refer the reader to Vershynin (2018) for a more comprehensive introduction to sub-Gaussian random variables and vectors.

A random variable X is sub-Gaussian if, for some constant $C > 0$, it satisfies

$$\mathbb{P}\{|X| \geq t\} \leq 2 \exp(-t^2/C) \quad \text{for all } t \geq 0. \quad (37)$$

The sub-Gaussian norm of X is defined as

$$\|X\|_{\psi_2} = \inf \{t > 0 : \mathbb{E} \exp(X^2/t^2) \leq 2\}. \quad (38)$$

A random vector $X \in \mathbb{R}^p$ is called sub-Gaussian if $\langle X, x \rangle$ is sub-Gaussian for all $x \in \mathbb{R}^p$. The sub-Gaussian norm of X is defined as

$$\|X\|_{\psi_2} = \sup_{x \in \mathbb{R}^p: \|x\|_2=1} \|\langle X, x \rangle\|_{\psi_2}. \quad (39)$$

From its definition, it is clear that $\|X_i\|_{\psi_2} \leq \|X\|_{\psi_2}$, where X_i is the i th element of X . To simplify our analysis, we will also assume that sub-Gaussian vectors X satisfy the variance-proxy condition defined below.

Definition B.1. *A sub-Gaussian random vector X satisfies the variance-proxy condition if there exists a constant K_X such that for any $x \in \mathbb{R}^p$, $\|\langle X, x \rangle\|_{\psi_2} \leq K_X \text{Var}(\langle X, x \rangle)^{1/2}$.*

Intuitively, this condition implies that the sub-Gaussian norms of the one-dimensional marginals of X can be used as proxies for their standard deviations. Note that the reverse inequality $\text{Var}(\langle X, x \rangle)^{1/2} \leq K \|\langle X, x \rangle\|_{\psi_2}$ for $K = \sqrt{2}$ is always satisfied when X has mean 0 (Proposition 2.5.2. (ii) of Vershynin (2018)). Moreover, for a Gaussian random vector X , this proxy assumption holds with $K_X = 1$. If X is a zero-mean sub-Gaussian random vector that satisfies the variance-proxy condition and has covariance matrix Σ_X , it follows from the definition above that $\|X\|_{\psi_2} \leq K_X \|\Sigma_X\|_2^{1/2}$. Additionally, it is straightforward to show that $\max_i (\|X_i\|_{\psi_2}) \leq K_X \|\Sigma_X\|_{2, \infty}^{1/2}$, where X_i is the i th entry of X . The proxy assumption allows us to compare sub-Gaussian norms of vectors to one another through their variances.

Throughout our proofs, we assume that the variance-proxy condition applies to the random vectors X , $B^\top X$, Y , $\Sigma_Y^{-1/2} Y$, and $\Sigma_Y^{-1} Y$. To simplify our assumptions, for the main theorems in this section, we conveniently assume that X and Y are strict sub-Gaussians, as defined in (Kereta and Klock 2021):

Definition B.2. *A sub-Gaussian random vector X is called strict sub-Gaussian if there exists a constant K_X such that for any matrix $U \in \mathbb{R}^{k \times p}$, the following inequality is satisfied:*

$$\|UX\|_{\psi_2} \leq K_X \|\Sigma_{UX}\|_2^{1/2}. \quad (40)$$

B.3 Proof of Theorem 4.1

Recall that the matrices $\mathbb{X} \in \mathbb{R}^{N \times p}$ and $\mathbb{Y} \in \mathbb{R}^{N \times d}$ consist of N samples of the random vectors $X \in \mathbb{R}^p$ and $Y \in \mathbb{R}^d$, respectively. We assume that Σ_X and Σ_Y are invertible, and without loss of generality, we assume that X and Y have mean 0.

To estimate Σ_X and Σ_Y , we use their respective sample covariance estimates $\hat{\Sigma}_Y = \mathbb{Y}^\top \mathbb{Y} / N$ and $\hat{\Sigma}_X = \mathbb{X}^\top \mathbb{X} / N$. Define $B = \Sigma_X^{-1} \Sigma_{XY} \Sigma_Y^{-1/2}$, let \hat{B} be the solution to the sample group lasso problem (8), and let λ be the associated penalization constant. In the setting of Theorem 2.1, if we define \tilde{H} by the eigendecomposition $B^\top \Sigma_X B \equiv \tilde{H} D^2 \tilde{H}^\top$, then by letting

$$T = B \tilde{H} D^{-1} \in \mathbb{R}^{p \times d}, \quad (41)$$

$$H = \Sigma_Y^{-1/2} \tilde{H} \in \mathbb{R}^{d \times d}, \quad (42)$$

it follows that the k th column of H , η_k , is the k th canonical vector associated with Y , and the k th column of T , θ_k , is the k th canonical vector associated with X . Moreover, the diagonal entries of D^2 are the squared population canonical correlations $\gamma_1 > \dots > \gamma_d$, which we assume are distinct. This allows us to focus on estimating individual canonical vectors rather than subspaces spanned by canonical vectors sharing identical correlations.

We denote the columns of \tilde{H} as $\tilde{\eta}_k$ and denote by $\{\hat{\theta}_k\}$ and $\{\hat{\eta}_k\}$ the estimates of the canonical vectors, and by $\{\hat{\gamma}_k\}$ the estimated canonical correlations, that is, the diagonal entries of \hat{D} . Note that by definition, the squared population correlations $\gamma_1^2 \dots \gamma_d^2$ are the eigenvalues of $B^\top \Sigma_X B$ and the estimated squared correlations $\hat{\gamma}_1^2, \dots, \hat{\gamma}_d^2$ are the eigenvalues of $\hat{B}^\top \hat{\Sigma}_X \hat{B}$. In the remainder of this section, we derive bounds on the estimation error for the canonical correlations, quantified by $|\gamma_k^2 - \hat{\gamma}_k^2|$, and the canonical vectors, quantified by $\|\eta_k - \hat{\eta}_k\|_2^2$ and $\|\theta_k - \hat{\theta}_k\|_2^2$.

B.3.1 Deterministic bounds

We begin by presenting our deterministic results. To establish fast-rate bounds, we use the Group restricted eigenvalue condition, analogously to the lasso regression problem (Hastie et al. 2015) and similar to Gaynanova (2020) in the context of penalized optimal scoring.

Definition B.3 (Group restricted eigenvalue condition). *A matrix $Q \in \mathbb{R}^{q \times p}$ satisfies the Group restricted eigenvalue condition $RE(s, c, d)$ with parameter κ if for all sets $S \subset \{1, \dots, p\}$ with $|S| \leq s$, we have that, for all $A \in \mathbb{R}^{p \times d}$ such that $\|A_{\bar{S}}\|_{\ell_1, \ell_2} \leq c \|A_S\|_{\ell_1, \ell_2}$,*

$$\|QA\|_F \geq \frac{\|A_S\|_F^2}{\kappa}. \quad (43)$$

Here, $|S|$ denotes the cardinality of S , and $\bar{S} = \{1, \dots, p\} \setminus S$.

The following lemma establishes a deterministic bound for the 2-norm of the difference between the linear operators $B^\top \Sigma_X B$ and $\hat{B}^\top \hat{\Sigma}_X \hat{B}$. In turn, this quantity will be used to bound the errors $|\gamma_k^2 - \hat{\gamma}_k^2|$, $\|\eta_k - \hat{\eta}_k\|_2^2$ and $\|\theta_k - \hat{\theta}_k\|_2^2$.

Lemma B.1. *The following inequality holds:*

$$\begin{aligned} \|\hat{B}^\top \hat{\Sigma}_X \hat{B} - B^\top \Sigma_X B\|_2 &\leq \frac{1}{\sqrt{N}} \|\mathbb{X}B\|_2 \frac{1}{\sqrt{N}} \|\mathbb{X}(\hat{B} - B)\|_F + \|B^\top (\Sigma_X - \hat{\Sigma}_X) B\|_2 + \gamma_1 \left\| (\hat{B} - B)^\top \Sigma_X^{1/2} \right\|_2 \\ &\quad + \frac{1}{N} \|\mathbb{X}(\hat{B} - B)\|_F^2 + \|\hat{B} - B\|_{\ell_1, \ell_2} \left\| (\Sigma_X - \hat{\Sigma}_X) B \right\|_{2, \infty}. \end{aligned}$$

In the equation above, the first-order terms appear on the first line while the second-order terms appear on the second line of the equation. In this section, wherever possible, we will keep the convention.

Let $E = \mathbb{Y} \hat{\Sigma}_Y^{-1/2} - \mathbb{X}B$. Next, we derive ‘slow’- and ‘fast’-rate deterministic bounds.

Lemma B.2. *If $\lambda \geq \frac{2}{N} \|\mathbb{X}^\top E\|_{2,\infty}$, then the following slow-rate bound holds:*

$$\begin{aligned} \|\hat{B}^T \hat{\Sigma}_X \hat{B} - B^\top \Sigma_X B\|_2 &\lesssim \frac{1}{\sqrt{N}} \|\mathbb{X}B\|_2 \sqrt{\lambda} \|B\|_{\ell_1, \ell_2}^{1/2} + \|B^\top (\Sigma_X - \hat{\Sigma}_X) B\|_2 + \gamma_1 \|B\|_{\ell_1, \ell_2}^{1/2} (\lambda \|B\|_{\ell_1, \ell_2} + \|\hat{\Sigma}_X - \Sigma_X\|_{\max})^{1/2} \\ &\quad + \lambda \|B\|_{\ell_1, \ell_2} + \|B\|_{\ell_1, \ell_2} \|(\Sigma_X - \hat{\Sigma}_X) B\|_{2,\infty}. \end{aligned}$$

If, additionally, B has at most s non-zero rows, and $\frac{1}{\sqrt{N}}\mathbb{X}$ satisfies the Group restricted eigenvalue condition $RE(s, 3, d)$ with parameter κ_X , then the following fast-rate bound holds:

$$\begin{aligned} \|\hat{B}^T \hat{\Sigma}_X \hat{B} - B^\top \Sigma_X B\|_2 &\lesssim \frac{1}{\sqrt{N}} \|\mathbb{X}B\|_2 \kappa_X^{1/2} s^{1/2} \lambda + \|B^\top (\Sigma_X - \hat{\Sigma}_X) B\|_2 + \gamma_1 \|\Sigma_X\|_2^{1/2} \kappa_X s^{1/2} \lambda \\ &\quad + \kappa_X s \lambda^2 + \|(\Sigma_X - \hat{\Sigma}_X) B\|_{2,\infty} \kappa_X s \lambda. \end{aligned}$$

Note that in the fast-rate bound $\sqrt{\lambda}$ and $\|B\|_{\ell_1, \ell_2}$ are replaced with λ and $\kappa_X s$, respectively. Next, we derive a bound for $\frac{1}{N} \|\mathbb{X}^\top E\|_{2,\infty}$.

Lemma B.3. *The following inequality holds:*

$$\begin{aligned} \frac{1}{N} \|\mathbb{X}^\top E\|_{2,\infty} &\leq \left\| (\hat{\Sigma}_{XY} - \Sigma_{XY}) \Sigma_Y^{-1/2} \right\|_{2,\infty} + \left\| \Sigma_{XY} (\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2}) \right\|_{2,\infty} + \|(\Sigma_X - \hat{\Sigma}_X) B\|_{2,\infty} \\ &\quad + \left\| (\hat{\Sigma}_{XY} - \Sigma_{XY}) (\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2}) \right\|_{2,\infty}. \end{aligned}$$

Denote the right-hand side of the equation in Lemma B.3 as λ_0 . Given that $\frac{2}{N} \|\mathbb{X}^\top E\|_{2,\infty} \leq 2\lambda_0$, choosing $\lambda \geq 2\lambda_0$ ensures that $\lambda \geq \frac{2}{N} \|\mathbb{X}^\top E\|_{2,\infty}$. Thus, we can replace the assumption $\lambda \geq \frac{2}{N} \|\mathbb{X}^\top E\|_{2,\infty}$ with the assumption $\lambda \geq 2\lambda_0$. Later, we will establish a high-probability bound for λ_0 .

Due to the fact that $\|(\Sigma_X - \hat{\Sigma}_X) B\|_{2,\infty} \leq \lambda_0$, we obtain the following simplification of Lemma B.2, where the fourth and fifth terms are combined.

Lemma B.4. *If $\lambda \geq 2\lambda_0$, then the following slow-rate bound holds:*

$$\begin{aligned} \|\hat{B}^T \hat{\Sigma}_X \hat{B} - B^\top \Sigma_X B\|_2 &\lesssim \frac{1}{\sqrt{N}} \|\mathbb{X}B\|_2 \sqrt{\lambda} \|B\|_{\ell_1, \ell_2}^{1/2} + \|B^\top (\Sigma_X - \hat{\Sigma}_X) B\|_2 + \gamma_1 \|B\|_{\ell_1, \ell_2}^{1/2} (\lambda \|B\|_{\ell_1, \ell_2} + \|\hat{\Sigma}_X - \Sigma_X\|_{\max})^{1/2} \\ &\quad + \lambda \|B\|_{\ell_1, \ell_2}. \end{aligned}$$

If, additionally, B has at most s non-zero rows, and $\frac{1}{\sqrt{N}}\mathbb{X}$ satisfies the Group restricted eigenvalue condition $RE(s, 3, d)$ with parameter κ_X , then the following fast-rate bound holds:

$$\begin{aligned} \|\hat{B}^T \hat{\Sigma}_X \hat{B} - B^\top \Sigma_X B\|_2 &\lesssim \frac{1}{\sqrt{N}} \|\mathbb{X}B\|_2 \kappa_X^{1/2} s^{1/2} \lambda + \|B^\top (\Sigma_X - \hat{\Sigma}_X) B\|_2 + \gamma_1 \kappa_X^{1/2} s^{1/2} (1 + \kappa_X s \|\hat{\Sigma}_X - \Sigma_X\|_{\max})^{1/2} \lambda \\ &\quad + \kappa_X s \lambda^2. \end{aligned}$$

Lemma B.4 shows that the rate of convergence will ultimately be determined by λ_0 , $\|B^\top (\Sigma_X - \hat{\Sigma}_X) B\|_2$, and $\|\hat{\Sigma}_X - \Sigma_X\|_{\max}$.

B.3.2 Probabilistic bounds

From now on, we assume that X and Y are sub-Gaussian random vectors and that the variance-proxy condition in Definition B.1 holds for the random vectors X , $B^\top X$, Y , $\Sigma_Y^{-1/2} Y$, and $\Sigma_Y^{-1} Y$. We will repeatedly use the union bound and omit for simplicity the absolute constants arising from its applications.

First, we present an intermediary result that will be used to derive a probabilistic upper bound for λ_0 .

Lemma B.5. Let $X \in \mathbb{R}^p$ and $Z \in \mathbb{R}^d$ be zero-mean random vectors with covariance matrices Σ_X and Σ_Z and cross-covariance matrix Σ_{XZ} . Assume the entries of X and Z are sub-Gaussian random variables with norms $\|X_i\|_{\psi_2} = g_i$ and $\|Z_j\|_{\psi_2} = h_j$, for $i = 1, \dots, p$ and $j = 1, \dots, d$. Let $g = \max(g_i)$ and $h = \max(h_j)$. Let $\mathbb{X} \in \mathbb{R}^{N \times p}$ and $\mathbb{Z} \in \mathbb{R}^{N \times d}$ be data matrices such that the pairs of rows $\{(\mathbb{X}_i^\top, \mathbb{Z}_i^\top)\}$ are independent samples from the joint distribution (X, Z) . If $d \leq p$ and $\log(p) = o(N)$, then for any fixed $\eta \in (0, 1)$, with probability at least $1 - \eta$,

$$\left\| \Sigma_{XZ} - \frac{1}{N} \mathbb{X}^\top \mathbb{Z} \right\|_{2, \infty} \lesssim gh \sqrt{\frac{d}{N} \log(p\eta^{-1})}. \quad (44)$$

Remark 6. In Lemma B.5, it is stated that for a fixed η , if $\lim(\log(p)/N) = 0$ as p and N go to infinity, then, eventually, the stated bound holds.

Next, we derive probabilistic upper bounds for λ_0 , bounding the terms in Lemma B.3.

Lemma B.6. If $d \leq p$ and $\log(p) = o(N)$, then for any fixed $\eta \in (0, 1)$, with probability $1 - \eta$,

$$\left\| (\hat{\Sigma}_{XY} - \Sigma_{XY}) \Sigma_Y^{-1/2} \right\|_{2, \infty} \lesssim \max_i(\|X_i\|_{\psi_2}) \sqrt{\frac{d}{N} \log(p\eta^{-1})} \quad (45)$$

and

$$\left\| (\Sigma_X - \hat{\Sigma}_X) B \right\|_{2, \infty} \lesssim \max_i(\|X_i\|_{\psi_2}) \gamma_1 \sqrt{\frac{d}{N} \log(p\eta^{-1})}. \quad (46)$$

Moreover, if $d = o(N)$, then

$$\left\| \Sigma_{XY} \left(\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2} \right) \right\|_{2, \infty} \lesssim \|\Sigma_X\|_{2, \infty}^{1/2} \gamma_1 \sqrt{\frac{d + \log(\eta^{-1})}{N}} \quad (47)$$

and

$$\left\| (\hat{\Sigma}_{XY} - \Sigma_{XY}) \left(\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2} \right) \right\|_{2, \infty} \lesssim \left\| (\hat{\Sigma}_{XY} - \Sigma_{XY}) \Sigma_Y^{-1/2} \right\|_{2, \infty}. \quad (48)$$

Remark 7. As noted in Section B.1, $\max_i(\|X_i\|_{\psi_2}) \lesssim \|\Sigma_X\|_{2, \infty}^{1/2}$.

Using Lemmas B.3 and B.6, and Remark 7, it is straightforward to derive the following result.

Lemma B.7. If $d \leq p$, $\log(p) = o(N)$, and $d = o(N)$, then for any fixed $\eta \in (0, 1)$, with probability $1 - \eta$,

$$\lambda_0 \lesssim \|\Sigma_X\|_{2, \infty}^{1/2} \sqrt{\frac{d}{N} \log(p\eta^{-1})}. \quad (49)$$

Next, we establish bounds on the other terms appearing in B.4.

Lemma B.8. If $\log(p) = o(N)$, then for any fixed $\eta \in (0, 1)$, with probability $1 - \eta$,

$$\left\| \Sigma_X - \hat{\Sigma}_X \right\|_{\max} \lesssim \max(\|X_i\|_{\psi_2}^2) \sqrt{\frac{\log(p\eta^{-1})}{N}}. \quad (50)$$

If $d = o(N)$, then for any fixed $\eta \in (0, 1)$, with probability $1 - \eta$,

$$\left\| B^\top (\Sigma_X - \hat{\Sigma}_X) B \right\|_2 \lesssim \gamma_1^2 \sqrt{\frac{d + \log(\eta^{-1})}{N}}, \quad (51)$$

and

$$\frac{1}{\sqrt{N}} \|\mathbb{X} B\|_2 \lesssim \gamma_1. \quad (52)$$

Before presenting our final bounds, we establish that the group-restricted eigenvalue condition holds for the design matrix $\frac{1}{\sqrt{N}}\mathbb{X}$, with high probability, assuming that the same condition holds for $\Sigma_X^{1/2}$.

Lemma B.9. *Suppose $\Sigma_X^{1/2}$ satisfies the group restricted eigenvalue condition $RE(s, 3, d)$ with parameter $\kappa = \kappa(s, d, \Sigma_X^{1/2})$. If $\max_{i=1, \dots, p} (\|X_i\|_{\psi_2}^4) \kappa^2 s^2 \log(p) = o(N)$ and $s^2 \log(p) = o(N)$, then for any fixed η , with probability $1 - \eta$, $\frac{1}{\sqrt{N}}\mathbb{X}$ satisfies the group restricted eigenvalue condition $RE(s, 3, d)$ with parameter κ_X , where*

$$0 < \kappa_X \leq 2\kappa. \quad (53)$$

Next, we state our probabilistic bound for $\|\hat{B}^T \hat{\Sigma}_X \hat{B} - B^T \Sigma_X B\|_2$. The proof of the slow-rate bound follows straightforwardly from Lemmas B.4, B.7 and B.8. The proof of the fast-rate bound follows similarly from Lemmas B.4, B.7 and B.8, with the addition of Lemma B.9.

Theorem B.1. *Assume X and Y are sub-Gaussian random vectors and that $X, B^T X, \Sigma_Y^{-1/2} Y$ satisfy the variance-proxy condition B.1. Moreover, assume that $d \leq p$, $\log(p) = o(N)$, and $d = o(N)$. Fix $\eta \in (0, 1)$, and for some absolute constant $C > 0$, let $\lambda = C \|\Sigma_X\|_{2, \infty}^{1/2} \sqrt{\frac{d}{N} \log(p\eta^{-1})}$. Then, with probability $1 - \eta$, the following slow-rate bound holds:*

$$\begin{aligned} & \|\hat{B}^T \hat{\Sigma}_X \hat{B} - B^T \Sigma_X B\|_2 \lesssim \\ & \left(\frac{d}{N} \log(p\eta^{-1})\right)^{1/4} \left[\gamma_1^2 + \|\Sigma_X\|_{2, \infty}^{1/2} \|B\|_{\ell_1, \ell_2} + \gamma_1 \|\Sigma_X\|_{2, \infty}^{1/4} \|B\|_{\ell_1, \ell_2}^{1/2} \left(1 + \|B\|_{\ell_1, \ell_2}^{1/2} + \|\Sigma_X\|_{2, \infty}^{1/4}\right) \right]. \end{aligned}$$

Under the additional assumption that $B \in \mathbb{R}^{p \times d}$ has at most s nonzero rows, $\Sigma_X^{1/2}$ satisfies the group restricted eigenvalue condition $RE(s, 3, d)$ with parameter $\kappa = \kappa(s, d, \Sigma_X^{1/2})$, $s^2 \log(p) = o(N)$, and $\|\Sigma_X\|_{2, \infty}^2 \kappa^2 s^2 \log(p) = o(N)$, then the following slow-rate bound holds:

$$\|\hat{B}^T \hat{\Sigma}_X \hat{B} - B^T \Sigma_X B\|_2 \lesssim \left(\frac{d}{N} \log(p\eta^{-1})\right)^{1/2} \left[\gamma_1 \left(\gamma_1 + \kappa^{1/2} s^{1/2} \|\Sigma_X\|_{2, \infty}^{1/2} \right) \right]. \quad (54)$$

Corollary B.1. *In the setting of Theorem B.1, under the additional assumption that $\|\Sigma_X\|_{2, \infty}, \|B\|_{\ell_1, \ell_2} \geq 1$, then the expression of the slow-rate bound simplifies as follows:*

$$\|\hat{B}^T \hat{\Sigma}_X \hat{B} - B^T \Sigma_X B\|_2 \lesssim \left(\frac{d}{N} \log(p\eta^{-1})\right)^{1/4} \left[\gamma_1 \|\Sigma_X\|_{2, \infty}^{1/2} \|B\|_{\ell_1, \ell_2} \right]. \quad (55)$$

Under the additional assumption that $\|\Sigma_X\|_{2, \infty}, \kappa \geq 1$, then the expression of the fast-rate bound simplifies as follows:

$$\|\hat{B}^T \hat{\Sigma}_X \hat{B} - B^T \Sigma_X B\|_2 \lesssim \left(\frac{d}{N} \log(p\eta^{-1})\right)^{1/2} \left[\gamma_1 \|\Sigma_X\|_{2, \infty}^{1/2} s^{1/2} \kappa^{1/2} \right]. \quad (56)$$

Remark 8. *The eigenvalues of the matrices $\hat{B}^T \hat{\Sigma}_X \hat{B}$ and $B^T \Sigma_X B$ are $\{\hat{\gamma}_k^2\}$ and $\{\gamma_k^2\}$ respectively. Then by Weyl's inequality (Bhatia (2013) Corollary III.2.6.), because we have bounded the operator norm of the difference between these two matrices, we immediately obtain bounds on the estimation error of γ_k^2 by $\hat{\gamma}_k^2$ in Algorithm 2.*

To establish bounds for $\|\theta_k - \hat{\theta}_k\|_2$ and $\|\eta_k - \hat{\eta}_k\|_2$, we first introduce a few supporting lemmas.

Lemma B.10. *Under the slow-rate bound assumptions stated in Corollary B.1, for any fixed $\eta \in (0, 1)$, with probability at least $1 - \eta$, we have*

$$\|B - \hat{B}\|_2 \lesssim \left(\frac{d}{N} \log(p\eta^{-1})\right)^{1/4} \left\| \Sigma_X^{-1/2} \right\|_2 \|B\|_{\ell_1, \ell_2} \|\Sigma_X\|_{2, \infty}^{1/2}. \quad (57)$$

Under the fast-rate bound assumptions stated in Corollary B.1, for any fixed $\eta \in (0, 1)$, with probability at least $1 - \eta$, we have

$$\|B - \hat{B}\|_2 \lesssim \left(\frac{d}{N} \log(p\eta^{-1}) \right)^{1/2} \kappa s^{1/2} \|\Sigma_X\|_{2,\infty}^{1/2}. \quad (58)$$

Lemma B.11. Suppose that $Y \in \mathbb{R}^d$ is a sub-Gaussian vector, $d = o(N)$ and that Y satisfies the variance-proxy condition. Then, for any fixed $\eta \in (0, 1)$, with probability $1 - \eta$, we have

$$\|\Sigma_Y - \hat{\Sigma}_Y\|_2 \lesssim \|\Sigma_Y\|_2 \sqrt{\frac{d \log(\eta^{-1})}{N}}. \quad (59)$$

Additionally, suppose that $\Sigma_Y^{-1}Y$ satisfies the variance-proxy condition, and $\|\Sigma_Y\|_2^2 d = o(N)$. Then for fixed $\eta \in (0, 1)$, with probability $1 - \eta$, we have

$$\|\Sigma_Y^{-1/2} - \hat{\Sigma}_Y^{-1/2}\|_2 \lesssim \|\Sigma_Y^{1/2}\|_2 \|\Sigma_Y^{-1/2}\|_2^2 \sqrt{\frac{d \log(\eta^{-1})}{N}}. \quad (60)$$

Studying the theoretical properties of CCA through the lens of regression, using the matrix B , has been convenient thus far. However, for our final results, we bound $B = \Sigma_X^{-1/2} \tilde{T} D \tilde{H} = T D \tilde{H}$ in terms of quantities that are more directly related to the CCA problem. Using identity 12 in Section D, along with the standard properties of the 2-norm, and noting that \tilde{T} , \tilde{H} are orthogonal matrices and D is a diagonal matrix with diagonal values no greater than 1, we observe that

$$\|B\|_{\ell_1, \ell_2} \leq \|T\|_{\ell_1, \ell_2} \leq \|\Sigma_X^{-1/2}\|_{\ell_1, \ell_2}, \quad (61)$$

$$\|B\|_2 \leq \|T\|_2 \leq \|\Sigma_X^{-1/2}\|_2. \quad (62)$$

Hence, in Corollary B.1, we can replace the assumption that $\|B\|_{\ell_1, \ell_2} \geq 1$ with the assumption that $\|T\|_{\ell_1, \ell_2} \geq 1$.

Next, we state our probabilistic bounds on the estimated canonical vectors. We denote with $K = \max\{i \in \{1, \dots, d\} : \gamma_i > 0\}$ the number of nontrivial canonical vectors. Moreover, to simplify the notation, we use the conventions $\gamma_{K+1}^2 = -\infty$ and $\gamma_0^2 = \infty$.

Theorem B.2. Under the slow-rate bound assumptions stated in Corollary B.1 and assuming that the canonical correlations $\gamma_1, \dots, \gamma_K$ are bounded from below, Y and $\Sigma_Y^{-1}Y$ satisfy the variance-proxy condition, and $\hat{\eta}^\top \hat{\Sigma}_Y^{1/2} \Sigma_Y^{1/2} \eta \geq 0$, for $k = 1, \dots, K$.

If $d \log(p) \|\Sigma_X\|_{2,\infty}^2 \|T\|_{\ell_1, \ell_2}^4 = o(N)$, then, for any fixed $\eta \in (0, 1)$, with probability $1 - \eta$,

$$\|\theta_k - \hat{\theta}_k\|_2 \lesssim \left(\frac{d}{N} \log(p\eta^{-1}) \right)^{1/4} \frac{\gamma_1 \|\Sigma_X\|_{2,\infty}^{1/2} \|T\|_{\ell_1, \ell_2} \|\Sigma_X^{-1/2}\|_2}{\gamma_k \min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)}, \quad k = 1, \dots, K. \quad (63)$$

If $\|\Sigma_Y\|_2^2 d = o(N)$ and $\|\Sigma_Y\|_2^2 \|\Sigma_Y^{-1}\|_2^2 d = o(N)$, then, for any fixed $\eta \in (0, 1)$, with probability $1 - \eta$,

$$\|\eta_k - \hat{\eta}_k\|_2 \lesssim \left(\frac{d}{N} \log(p\eta^{-1}) \right)^{1/4} \frac{\gamma_1 \|\Sigma_X\|_{2,\infty}^{1/2} \|T\|_{\ell_1, \ell_2} \|\Sigma_Y^{-1/2}\|_2}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)}, \quad k = 1, \dots, K. \quad (64)$$

Theorem B.3. Under the fast-rate bound assumptions stated in Corollary B.1 and assuming that the canonical correlations $\gamma_1, \dots, \gamma_K$ are bounded from below, Y and $\Sigma_Y^{-1}Y$ satisfy the variance-proxy condition, and $\hat{\eta}_k^\top \hat{\Sigma}_Y^{1/2} \Sigma_Y^{1/2} \eta_k \geq 0$, for $k = 1, \dots, K$.

If $d \log(p) \|\Sigma_X\|_{2,\infty} s \kappa = o(N)$, then for any fixed $\eta \in (0, 1)$, with probability $1 - \eta$,

$$\|\theta_k - \hat{\theta}_k\|_2 \lesssim \left(\frac{d}{N} \log(p\eta^{-1}) \right)^{1/2} \frac{\gamma_1}{\gamma_k \min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)} \|\Sigma_X\|_{2,\infty}^{1/2} s^{1/2} \kappa \|T\|_2, \quad (65)$$

with $k = 1, \dots, K$.

If $\|\Sigma_Y\|_2^2 d = o(N)$, then for any fixed $\eta \in (0, 1)$, with probability $1 - \eta$,

$$\|\eta_k - \hat{\eta}_k\|_2 \lesssim \left(\frac{d}{N} \log(p\eta^{-1}) \right)^{1/2} \left\| \Sigma_Y^{-1/2} \right\|_2 \max \left\{ \frac{\gamma_1 \|\Sigma_X\|_{2,\infty}^{1/2} s^{1/2} \kappa^{1/2}}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)}, \|\Sigma_Y\|_2^{1/2} \|\Sigma_Y^{-1}\|_2^{1/2} \right\} \quad (66)$$

with $k = 1, \dots, K$.

B.4 Proofs for the deterministic bounds in Section B.3.1 and for the probabilistic bounds in Section B.3.2

Proof of Lemma B.1:

The triangle inequality is used repeatedly without comment. By adding and subtracting $\hat{B}^\top \Sigma_X B$, we have

$$\left\| \hat{B}^\top \hat{\Sigma}_X \hat{B} - B^\top \Sigma_X B \right\|_2 \leq \left\| \hat{B}^\top (\hat{\Sigma}_X \hat{B} - \Sigma_X B) \right\|_2 + \left\| (\hat{B} - B)^\top \Sigma_X B \right\|_2. \quad (67)$$

Since

$$\hat{\Sigma}_X \hat{B} - \Sigma_X B = \hat{\Sigma}_X (\hat{B} - B) + (\hat{\Sigma}_X - \Sigma_X) B \quad (68)$$

by adding and subtracting $\hat{\Sigma}_X B$, we deduce that

$$\left\| \hat{B}^\top \hat{\Sigma}_X \hat{B} - B^\top \Sigma_X B \right\|_2 \leq \underbrace{\left\| \hat{B}^\top \hat{\Sigma}_X (\hat{B} - B) \right\|_2}_{\text{Term I}} + \underbrace{\left\| \hat{B}^\top (\hat{\Sigma}_X - \Sigma_X) B \right\|_2}_{\text{Term II}} + \underbrace{\left\| (\hat{B} - B)^\top \Sigma_X B \right\|_2}_{\text{Term III}}. \quad (69)$$

We bound each term individually. Recall that $\hat{\Sigma}_X = \frac{1}{N} \mathbb{X}^\top \mathbb{X}$.

Term I: We have

$$\left\| \hat{B}^\top \hat{\Sigma}_X (\hat{B} - B) \right\|_2 \leq \frac{1}{\sqrt{N}} \left\| \mathbb{X} \hat{B} \right\|_2 \left\| \mathbb{X} (\hat{B} - B) \right\|_F \frac{1}{\sqrt{N}} \quad (70)$$

using $\|AB\|_2 \leq \|A\|_2 \|B\|_2 \leq \|A\|_2 \|B\|_F$. Since

$$\left\| \mathbb{X} \hat{B} \right\|_2 \leq \left\| \mathbb{X} B \right\|_2 + \left\| \mathbb{X} (\hat{B} - B) \right\|_2, \quad (71)$$

we have

$$\left\| \hat{B}^\top \hat{\Sigma}_X (\hat{B} - B) \right\|_2 \leq \frac{1}{\sqrt{N}} \left\| \mathbb{X} B \right\|_2 \left\| \mathbb{X} (\hat{B} - B) \right\|_F \frac{1}{\sqrt{N}} + \frac{1}{N} \left\| \mathbb{X} (\hat{B} - B) \right\|_F^2. \quad (72)$$

Term II: We have

$$\left\| \hat{B}^\top (\hat{\Sigma}_X - \Sigma_X) B \right\|_2 \leq \left\| B^\top (\hat{\Sigma}_X - \Sigma_X) B \right\|_2 + \left\| (\hat{B} - B)^\top (\hat{\Sigma}_X - \Sigma_X) B \right\|_2 \quad (73)$$

$$\leq \left\| B^\top (\hat{\Sigma}_X - \Sigma_X) B \right\|_2 + \left\| \hat{B} - B \right\|_{\ell_1, \ell_2} \left\| (\hat{\Sigma}_X - \Sigma_X) B \right\|_{2, \infty} \quad (74)$$

using $\|A^\top B\|_2 \leq \|A\|_{\ell_1, \ell_2} \|B\|_{2, \infty}$.

Term 3: We have

$$\left\| (\hat{B} - B)^\top \Sigma_X B \right\|_2 \leq \left\| (\hat{B} - B)^\top \Sigma_X^{1/2} \right\|_2 \left\| \Sigma_X^{1/2} B \right\|_2 \leq \gamma_1 \left\| (\hat{B} - B)^\top \Sigma_X^{1/2} \right\|_2 \quad (75)$$

since $\Sigma_X^{1/2} B = \tilde{T} D \tilde{H}$ where \tilde{T} and \tilde{H} are orthogonal and D is diagonal.

Combining these results we obtain the statement of the lemma. \square

Proof of Lemma B.2:

For reference, Lemma B.1 gives the bound

$$\left\| \hat{B}^\top \hat{\Sigma}_X \hat{B} - B^\top \Sigma_X B \right\|_2 \lesssim \frac{1}{\sqrt{N}} \left\| \mathbb{X} B \right\|_2 \sqrt{\lambda} \left\| B \right\|_{\ell_1, \ell_2}^{1/2} + \left\| B^\top (\Sigma_X - \hat{\Sigma}_X) B \right\|_2 + \gamma_1 \left\| B \right\|_{\ell_1, \ell_2}^{1/2} \left(\lambda \left\| B \right\|_{\ell_1, \ell_2} + \left\| \hat{\Sigma}_X - \Sigma_X \right\|_{\max} \right)^{1/2} \quad (76)$$

$$+ \lambda \left\| B \right\|_{\ell_1, \ell_2} + \left\| B \right\|_{\ell_1, \ell_2} \left\| (\Sigma_X - \hat{\Sigma}_X) B \right\|_{2, \infty}. \quad (77)$$

Proof of slow-rate bound:

Assuming that $\lambda \geq \frac{2}{N} \|\mathbb{X}^\top E\|_{2,\infty}$, then by Theorem 1 and Corollary 1 of Gaynanova (2020) we have

$$\frac{1}{N} \|\mathbb{X}(\hat{B} - B)\|_F^2 \lesssim \lambda \|B\|_{\ell_1, \ell_2}, \quad (78)$$

$$\left\| (\hat{B} - B)^\top \Sigma_X^{1/2} \right\|_F^2 \lesssim \|B\|_{\ell_1, \ell_2} (\lambda + \|B\|_{\ell_1, \ell_2} \|\hat{\Sigma}_X - \Sigma_X\|_{\max}), \quad (79)$$

and

$$\|\hat{B}\|_{\ell_1, \ell_2} \lesssim \|B\|_{\ell_1, \ell_2}. \quad (80)$$

This last equation implies that $\|\hat{B} - B\|_{\ell_1, \ell_2} \lesssim \|B\|_{\ell_1, \ell_2}$ by the triangle inequality. Applying these bounds to the terms in Lemma B.1 establishes the slow-rate bound.

Proof of fast-rate bound:

Assuming that $\lambda \geq \frac{2}{N} \|\mathbb{X}^\top E\|_{2,\infty}$, that B has at most s nonzero rows, and assuming the group restricted eigenvalue condition on $\frac{1}{\sqrt{N}}\mathbb{X}$, then by Theorem 2 and Corollary 2 of Gaynanova (2020) we have

$$\frac{1}{N} \|\mathbb{X}(\hat{B} - B)\|_F^2 \lesssim \kappa_X s \lambda^2 \quad (81)$$

$$\left\| (\hat{B} - B)^\top \Sigma_X^{1/2} \right\|_F^2 \lesssim \kappa_X s (1 + \kappa_X s \|\hat{\Sigma}_X - \Sigma_X\|_{\max}) \lambda^2 \quad (82)$$

and

$$\|\hat{B} - B\|_{\ell_1, \ell_2} \lesssim \kappa_X s \lambda. \quad (83)$$

Applying these bounds to the terms in Lemma B.1 establishes the fast-rate bound. \square

Proof of Lemma B.3:

From the definition of E , $E = \mathbb{Y} \hat{\Sigma}_Y^{-1/2} - \mathbb{X}B$, we have

$$\frac{1}{N} \mathbb{X}^\top E = \frac{1}{N} \mathbb{X}^\top \mathbb{Y} \hat{\Sigma}_Y^{-1/2} - \frac{1}{N} \mathbb{X}^\top \mathbb{X}B \quad (84)$$

$$= \hat{\Sigma}_{XY} \hat{\Sigma}_Y^{-1/2} - \hat{\Sigma}_X B \quad (85)$$

$$= \left(\hat{\Sigma}_{XY} \hat{\Sigma}_Y^{-1/2} - \Sigma_{XY} \Sigma_Y^{-1/2} \right) + \left(\Sigma_{XY} \Sigma_Y^{-1/2} - \hat{\Sigma}_X B \right). \quad (86)$$

Now considering the first term in equation (86), by adding and subtracting $\Sigma_{XY} \hat{\Sigma}_Y^{-1/2}$, and subsequently adding and subtracting $\Sigma_Y^{-1/2}$ to $\hat{\Sigma}_Y^{-1/2}$, we have

$$\hat{\Sigma}_{XY} \hat{\Sigma}_Y^{-1/2} - \Sigma_{XY} \Sigma_Y^{-1/2} = (\hat{\Sigma}_{XY} - \Sigma_{XY}) \hat{\Sigma}_Y^{-1/2} + \Sigma_{XY} (\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2}) \quad (87)$$

$$= (\hat{\Sigma}_{XY} - \Sigma_{XY}) (\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2}) + (\hat{\Sigma}_{XY} - \Sigma_{XY}) \Sigma_Y^{-1/2} + \Sigma_{XY} (\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2}). \quad (88)$$

For the second term, we have

$$\Sigma_{XY} \Sigma_Y^{-1/2} - \hat{\Sigma}_X B = \Sigma_X \Sigma_X^{-1} \Sigma_{XY} \Sigma_Y^{-1/2} - \hat{\Sigma}_X B \quad (89)$$

$$= (\Sigma_X - \hat{\Sigma}_X) B. \quad (90)$$

Combining these equalities and using the triangle inequality completes the proof. \square

Proof of Lemma B.5:

The proof is adapted from Lemma 7 of Gaynanova (2020) but considers cross-covariance matrices rather than the covariance matrices. Let X_{ij} denote entry (i, j) of \mathbb{X} and Z_{ij} denote entry (i, j) of \mathbb{Z} . Then

$$\frac{1}{N} (\mathbb{X}^\top \mathbb{Z})_{kl} = \frac{1}{N} \sum_{i=1}^N X_{ik} Z_{il}, \quad k = 1, \dots, p, \quad l = 1, \dots, d. \quad (91)$$

Let Σ_{XZ} be the cross-covariance matrix of X and Z with (k, l) entry equal to $\sigma_{kl} = \mathbb{E}[X_{ik}Z_{il}]$. Then, $X_{ik}Z_{il} - \sigma_{kl}$ are each mean 0 subexponential random variables, since

$$\|X_{ik}Z_{il}\|_{\psi_1} \leq \|X_{ik}\|_{\psi_2} \|Z_{il}\|_{\psi_2} = g_k h_l \leq gh$$

by Lemma 2.7.7. of Vershynin (2018), and because

$$\|X_{ik}Z_{il} - \sigma_{kl}\|_{\psi_1} \lesssim \|X_{ik}Z_{il}\|_{\psi_1}$$

by Exercise 2.7.10. of Vershynin (2018). Thus, $(\Sigma_{XZ} - \frac{1}{N}\mathbb{X}^\top\mathbb{Z})_{kl}$ is a sum of independently and identically distributed (i.i.d.) subexponential random variables, since each for fixed k and l , the $X_{ik}Z_{il} - \sigma_{kl}$ are mean 0 subexponential i.i.d. random variables over $i = 1, \dots, N$.

By Corollary 2.8.3. of Vershynin (2018), (Bernstein's inequality), for each k and l , and for every $t > 0$,

$$P\left(\left|\left(\Sigma_{XZ} - \frac{1}{N}\mathbb{X}^\top\mathbb{Z}\right)_{kl}\right| \geq t\right) \leq 2 \exp\left[-cN \min\left(\frac{t^2}{K^2}, \frac{t}{K}\right)\right], \quad (92)$$

where $K = Cgh$, for absolute constants c and C . Applying a union bound, we have

$$P\left(\left\|\Sigma_{XZ} - \frac{1}{N}\mathbb{X}^\top\mathbb{Z}\right\|_{\max} \geq t\right) \leq 2dp \exp\left[-cN \min\left(\frac{t^2}{K^2}, \frac{t}{K}\right)\right]. \quad (93)$$

When $t \leq K$, we have

$$P\left(\left\|\Sigma_{XZ} - \frac{1}{N}\mathbb{X}^\top\mathbb{Z}\right\|_{\max} \geq t\right) \leq 2dp \exp\left[-cN \min\frac{t^2}{K^2}\right]. \quad (94)$$

since $t^2/K^2 \leq t/K$ if and only if $t \leq K$. Letting the right hand side of equation (94) be denoted as η , we solve for t in terms of η to obtain

$$t = \sqrt{\log(2dp\eta^{-1}) \frac{K^2}{cN}}, \quad (95)$$

and so that, for $\eta \in (0, 1)$, if $\sqrt{\log(2dp\eta^{-1}) \frac{K^2}{cN}} \leq Cgh$, then with probability at least $1 - \eta$ we have

$$\left\|\Sigma_{XZ} - \frac{1}{N}\mathbb{X}^\top\mathbb{Z}\right\|_{\max} \lesssim gh \sqrt{\log(2dp\eta^{-1}) \frac{1}{N}}. \quad (96)$$

$\log(2dp\eta^{-1}) \leq \log(2) + 2\log(p\eta^{-1})$ since $d \leq p$, and because we suppose $\log(p) = o(N)$, it follows that for fixed $\eta \in (0, 1)$, with probability $1 - \eta$,

$$\left\|\Sigma_{XZ} - \frac{1}{N}\mathbb{X}^\top\mathbb{Z}\right\|_{\max} \lesssim gh \sqrt{\log(p\eta^{-1}) \frac{1}{N}}. \quad (97)$$

Using $\|A\|_{2,\infty} \leq \sqrt{d}\|A\|_{\max}$ for any $A \in \mathbb{R}^{p \times d}$ completes the proof. \square

Proof of Lemma B.6:

To establish

$$\left\|\left(\hat{\Sigma}_{XY} - \Sigma_{XY}\right)\Sigma_Y^{-1/2}\right\|_{2,\infty} \lesssim \max_i(\|X_i\|_{\psi_2}) \sqrt{\frac{d}{N} \log(p\eta^{-1})}, \quad (98)$$

we can use Lemma B.5 where $X = X$ and $Z = \Sigma_Y^{-1/2}Y$, since $\Sigma_{XY}\Sigma_Y^{-1/2} = \Sigma_{X,\Sigma_Y^{-1/2}Y}$. Then, the sub-Gaussian norms are $g = \max_i(\|X_i\|_{\psi_2})$ and $h = \max_i\left(\left\|\left(\Sigma_Y^{-1/2}Y\right)_i\right\|_{\psi_2}\right)$, where $\left(\Sigma_Y^{-1/2}Y\right)_i$ is the i th entry of $\Sigma_Y^{-1/2}Y \in \mathbb{R}^d$. Using Definition B.1 for $\Sigma_Y^{-1/2}Y$, we have

$$h \leq K_{\Sigma_Y^{-1/2}Y} \left\|\Sigma_Y^{-1/2}\Sigma_Y\Sigma_Y^{-1/2}\right\|_2^{1/2} = K_{\Sigma_Y^{-1/2}Y}. \quad (99)$$

Treating $K_{\Sigma_Y^{-1/2}Y}$ as an absolute constant establishes equation (98).

Establishing

$$\|(\Sigma_X - \hat{\Sigma}_X) B\|_{2,\infty} \lesssim \max_i(\|X_i\|_{\psi_2}) \gamma_1 \sqrt{\frac{d}{N} \log(p\eta^{-1})} \quad (100)$$

follows an identical argument except that we let $Z = B^\top X$, so

$$h \leq K_{B^\top X} \|B^\top \Sigma_X B\|_2^{1/2}$$

in the final step. The identity $\|B^\top \Sigma_X B\|_2 = \gamma_1^2$ establishes equation (100).

To deduce that

$$\|\Sigma_{XY} (\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2})\|_{2,\infty} \lesssim \|\Sigma_X\|_{2,\infty}^{1/2} \gamma_1 \sqrt{\frac{d + \log(\eta^{-1})}{N}}, \quad (101)$$

we begin by using $\|AB\|_{2,\infty} \leq \|A\|_{2,\infty} \|B\|_2$ and $\Sigma_Y^{-1/2} \Sigma_Y^{1/2} = I_d$ to obtain

$$\|\Sigma_{XY} (\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2})\|_{2,\infty} \leq \underbrace{\|\Sigma_{XY} \Sigma_Y^{-1/2}\|_{2,\infty}}_{\text{Term I}} \underbrace{\|\Sigma_Y^{1/2} (\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2})\|_2}_{\text{Term II}}. \quad (102)$$

Considering Term I, we have

$$\|\Sigma_{XY} \Sigma_Y^{-1/2}\|_{2,\infty} = \|\Sigma_X^{1/2} \Sigma_X^{-1/2} \Sigma_{XY} \Sigma_Y^{-1/2}\|_{2,\infty} \quad (103)$$

$$\leq \|\Sigma_X^{1/2}\|_{2,\infty} \|\Sigma_X^{-1/2} B\|_2 \quad (104)$$

$$= \|\Sigma_X^{1/2}\|_{2,\infty} \gamma_1. \quad (105)$$

In the below, we are able to bound Term II without incurring unnecessary factors of $\|\Sigma_Y^{-1/2}\|_2$ using the results of Kereta and Klock (2021) which pertain to precision matrix estimation along subspaces. The main idea is to bound $\|\Sigma_Y^{1/2} (\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2})\|_2$ in terms of $\|\Sigma_Y^{1/2} (\hat{\Sigma}_Y^{-1} - \Sigma_Y^{-1}) \Sigma_Y^{1/2}\|_2$, to which the results of Kereta and Klock (2021) can be applied. That $\|\Sigma_Y^{1/2} (\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2})\|_2$ is not simply equal to $\|\Sigma_Y^{1/2} (\hat{\Sigma}_Y^{-1} - \Sigma_Y^{-1}) \Sigma_Y^{1/2}\|_2$ is due to Σ_Y and $\hat{\Sigma}_Y$ not necessarily commuting with one another. We begin with

$$\|\Sigma_Y^{1/2} (\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2})\|_2 = \|\Sigma_Y^{1/2} \hat{\Sigma}_Y^{-1/2} - I\|_2. \quad (106)$$

Using identity 15 in Section D and that both $\Sigma_Y^{1/2}$ and $\hat{\Sigma}_Y^{1/2}$ are positive definite along with their inverses,

$$\text{Term II} = \left\| \left(\Sigma_Y^{1/2} \hat{\Sigma}_Y^{-1} \Sigma_Y^{1/2} \right)^{1/2} - I \right\|_2 = \left\| \left(\Sigma_Y^{1/2} \hat{\Sigma}_Y^{-1} \Sigma_Y^{1/2} \right)^{1/2} - \left(\Sigma_Y^{1/2} \Sigma_Y^{-1} \Sigma_Y^{1/2} \right)^{1/2} \right\|_2. \quad (107)$$

Using $\|A^{1/2} - B^{1/2}\|_2 \leq \frac{1}{2} \max(\|A^{-1}\|_2, \|B^{-1}\|_2)^{1/2} \|A - B\|_2$ for positive definite matrices A and B , we deduce that

$$\text{Term II} \leq \frac{1}{2} \max \left(\left\| \left(\Sigma_Y^{1/2} \hat{\Sigma}_Y^{-1} \Sigma_Y^{1/2} \right)^{-1} \right\|_2, \|I_d^{-1}\|_2 \right)^{1/2} \left\| \Sigma_Y^{1/2} \hat{\Sigma}_Y^{-1} \Sigma_Y^{1/2} - \Sigma_Y^{1/2} \Sigma_Y^{-1} \Sigma_Y^{1/2} \right\|_2 \quad (108)$$

$$= \frac{1}{2} \underbrace{\max(\|\Sigma_Y^{-1} \hat{\Sigma}_Y\|_2, 1)}_{\text{Term II.I}} \underbrace{\left\| \Sigma_Y^{1/2} (\hat{\Sigma}_Y^{-1} - \Sigma_Y^{-1}) \Sigma_Y^{1/2} \right\|_2}_{\text{Term II.I}}. \quad (109)$$

We bound $\|\Sigma_Y^{-1}\hat{\Sigma}_Y\|_2$ in Term II.I with

$$\|\Sigma_Y^{-1}\hat{\Sigma}_Y\|_2 \leq \|\Sigma_Y^{-1}(\hat{\Sigma}_Y - \Sigma_Y)\|_2 + \|\Sigma_Y^{-1}\Sigma_Y\|_2 \quad (110)$$

$$= \|\Sigma_Y^{-1/2}(\hat{\Sigma}_Y - \Sigma_Y)\Sigma_Y^{-1/2}\|_2 + 1. \quad (111)$$

We apply a result concerning covariance estimation along subspaces, Lemma 2 of Kereta and Klock (2021), to deduce that for fixed $\eta \in (0, 1)$ with probability $1 - \eta$,

$$\|\Sigma_Y^{-1/2}(\hat{\Sigma}_Y - \Sigma_Y)\Sigma_Y^{-1/2}\|_2 \lesssim \|\Sigma_Y^{-1/2}Y\|_{\psi_2}^2 \max\left(\sqrt{\frac{2d + \log(\eta^{-1})}{N}}, \frac{2d + \log(\eta^{-1})}{N}\right). \quad (112)$$

From Definition B.1 for $\Sigma_Y^{-1/2}Y$ combined with the assumption that $d = o(N)$, we have that in the limit this term is bounded by 1. Therefore, for fixed $\eta \in (0, 1)$, with probability $1 - \eta$, $\|\Sigma_Y^{-1}\hat{\Sigma}_Y\|_2 \lesssim 1$. From this, Term II.I $\lesssim 1$ as well.

Having bounded $\|\Sigma_Y^{1/2}(\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2})\|_2$ in terms of Term II.II = $\|\Sigma_Y^{1/2}(\hat{\Sigma}_Y^{-1} - \Sigma_Y^{-1})\Sigma_Y^{1/2}\|_2$, we will bound the latter term. We use Theorem 10 of Kereta and Klock (2021) directly, implying that for fixed $\eta \in (0, 1)$, if $N \gtrsim (d + \log(\eta^{-1}))\|\Sigma_Y^{-1/2}Y\|_{\psi_2}^4$, then with probability $1 - \eta$,

$$\|\Sigma_Y^{1/2}(\hat{\Sigma}_Y^{-1} - \Sigma_Y^{-1})\Sigma_Y^{1/2}\|_2 \lesssim \|\Sigma_Y^{-1/2}Y\|_{\psi_2}^2 \sqrt{\frac{\text{rank}(\Sigma_Y) + \log(\eta^{-1})}{N}}. \quad (113)$$

By Definition B.1 for $\Sigma_Y^{-1/2}Y$, $\|\Sigma_Y^{-1/2}Y\|_{\psi_2} \lesssim 1$, so that the assumption $d = o(N)$ ensures that for fixed η , $N \gtrsim (d + \log(\eta^{-1}))\|\Sigma_Y^{-1/2}Y\|_{\psi_2}^4$ eventually.

With our bounds for both Term II.I and Term II.II, we deduce that for fixed $\eta \in (0, 1)$, with probability $1 - \eta$,

$$\text{Term II} = \|\Sigma_Y^{1/2}(\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2})\|_2 \lesssim \sqrt{\frac{d + \log(\eta^{-1})}{N}}. \quad (114)$$

Now having bounded both Term I and Term II, we finally establish that for fixed $\eta \in (0, 1)$, if $d = o(N)$, with probability $1 - \eta$,

$$\|\Sigma_{XY}(\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2})\|_{2,\infty} \lesssim \|\Sigma_X\|_{2,\infty}^{1/2} \gamma_1 \sqrt{\frac{d + \log(\eta^{-1})}{N}}, \quad (115)$$

completing the proof of (101).

To show

$$\|(\hat{\Sigma}_{XY} - \Sigma_{XY})(\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2})\|_{2,\infty} \lesssim \|(\hat{\Sigma}_{XY} - \Sigma_{XY})\Sigma_Y^{-1/2}\|_{2,\infty}, \quad (116)$$

we begin with

$$\|(\hat{\Sigma}_{XY} - \Sigma_{XY})(\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2})\|_{2,\infty} \leq \|(\hat{\Sigma}_{XY} - \Sigma_{XY})\Sigma_Y^{-1/2}\|_{2,\infty} \|\Sigma_Y^{1/2}(\hat{\Sigma}_Y^{-1/2} - \Sigma_Y^{-1/2})\|_2, \quad (117)$$

using $\Sigma_Y^{-1/2}\Sigma_Y^{1/2} = I_d$ and $\|AB\|_{2,\infty} \leq \|A\|_{2,\infty}\|B\|_2$. From equation (114) we obtain that with probability $1 - \eta$, the second factor in (117) is bounded by an absolute constant as $d = o(N)$, completing the proof. \square

Proof of Lemma B.8:

That for fixed $\eta \in (0, 1)$, if $\log(p) = o(N)$, then with probability $1 - \eta$,

$$\|\Sigma_X - \hat{\Sigma}_X\|_{\max} \lesssim \max(\|X_i\|_{\psi_2}^2) \sqrt{\frac{\log(p\eta^{-1})}{N}} \quad (118)$$

follows from Lemma 7 of Gaynanova (2020).

That for fixed $\eta \in (0, 1)$, if $d = o(N)$, then with probability $1 - \eta$,

$$\|B^\top(\Sigma_X - \hat{\Sigma}_X)B\|_2 \lesssim \gamma_1^2 \sqrt{\frac{\text{rank}(B^\top \Sigma_X) + \log(\eta^{-1})}{N}} \quad (119)$$

follows from Lemma 2 of Kereta and Klock (2021) in addition to Definition B.1 applied to $B^\top X$, which implies that $\|B^\top X\|_{\psi_2} \leq K_{B^\top X} \|B^\top \Sigma_X B\|_2^{1/2}$. Using $\text{rank}(AB) \leq \min(\text{rank}(A), \text{rank}(B))$, we have

$$\text{rank}(B^\top \Sigma_X) = \text{rank}(\Sigma_X B) = \text{rank}(\Sigma_{XY} \Sigma_Y^{-1/2}) \leq d, \quad (120)$$

which establishes the desired result.

To show that for fixed $\eta \in (0, 1)$, if $d = o(N)$, then with probability $1 - \eta$,

$$\frac{1}{\sqrt{N}} \|\mathbb{X}B\|_2 \lesssim \gamma_1, \quad (121)$$

we begin with

$$\frac{1}{\sqrt{N}} \|\mathbb{X}B\|_2 = \frac{1}{\sqrt{N}} \|B^\top \mathbb{X}^\top \mathbb{X}B\|_2^{1/2} = \|B^\top \hat{\Sigma}_X B\|_2^{1/2}, \quad (122)$$

which holds since $\|A\|_2 = \|A^\top A\|_2^{1/2}$. Adding and subtracting $B^\top \Sigma_X B$ and using the triangle inequality, we obtain that for fixed $\eta \in (0, 1)$, with probability $1 - \eta$,

$$\|B^\top \hat{\Sigma}_X B\|_2 \leq \|B^\top \Sigma_X B\|_2 + \|B^\top (\hat{\Sigma}_X - \Sigma_X) B\|_2 \lesssim \gamma_1^2. \quad (123)$$

In the last inequality, we have used $d = o(N)$ and equation (119) to deduce that $\|B^\top (\hat{\Sigma}_X - \Sigma_X) B\|_2$ becomes smaller than γ_1^2 eventually. This completes the proof. \square

Proof of Lemma B.9:

By Lemma 6 of Gaynanova (2020), it suffices to show that under the condition $s^2 \log(p) = o(N)$, that for fixed η , with probability $1 - \eta$, we have $s \|\Sigma_X - \hat{\Sigma}_X\|_{\max} \leq (32\kappa)^{-1}$. By the first item of Lemma B.8, we then have that for fixed η with probability $1 - \eta$,

$$\kappa s \|\Sigma_X - \hat{\Sigma}_X\|_{\max} \lesssim \kappa s \max_i (\|X_i\|_{\psi_2}^2) \sqrt{\frac{\log(p\eta^{-1})}{N}}. \quad (124)$$

It therefore suffices that $\max_i (\|X_i\|_{\psi_2}^4) \kappa^2 s^2 \log(p) = o(N)$, since then with probability $1 - \eta$, $\kappa s \|\Sigma_X - \hat{\Sigma}_X\|_{\max}$ is arbitrarily small. But this is the assumed condition, so the proof is complete. \square

Proof of Lemma B.10:

Proof of slow-rate bound:

Under the slow-rate assumptions of Corollary B.1 and by Corollary 1 in Gaynanova (2020), for fixed η , with probability $1 - \eta$, we have

$$\|B - \hat{B}\|_2^2 \lesssim \|\Sigma_X^{-1}\|_2 \|B\|_{\ell_1, \ell_2} (\lambda + \|B\|_{\ell_1, \ell_2} \|\hat{\Sigma}_X - \Sigma_X\|_{\max}). \quad (125)$$

Then

$$\lambda \lesssim \|\Sigma_X\|_{2, \infty}^{1/2} \left(\frac{d}{N} \log(p\eta^{-1}) \right)^{1/2}, \quad (126)$$

and by Lemma B.8,

$$\|\Sigma_X - \hat{\Sigma}_X\|_{\max} \lesssim \|\Sigma_X\|_{2, \infty} \left(\frac{d}{N} \log(p\eta^{-1}) \right)^{1/2}. \quad (127)$$

The last two equations bound $\|B - \hat{B}\|_2$, and since $\|\Sigma_X\|_{2,\infty}, \|B\|_{\ell_1, \ell_2} \geq 1$, the proof of the slow-rate bound is complete.

Proof of fast-rate bound:

Under the slow-rate assumptions of Corollary B.1 and by Theorem 2 in Gaynanova (2020), for fixed η , with probability $1 - \eta$, we have

$$\|B - \hat{B}\|_2 \lesssim \kappa s^{1/2} \lambda. \quad (128)$$

Again using $\lambda \lesssim \|\Sigma_X\|_{2,\infty}^{1/2} \left(\frac{d}{N} \log(p\eta^{-1})\right)^{1/2}$, the result is shown. \square

Proof of Lemma B.11:

The proof of the first statement follows from Lemma 2 of Kereta and Klock (2021) and Definition B.1 applied to Y .

To show the bound on $\|\Sigma_Y^{-1/2} - \hat{\Sigma}_Y^{-1/2}\|_2$, we begin by using identity 13 in Section D applied to Σ_Y^{-1} and $\hat{\Sigma}_Y^{-1}$. We have

$$\|\Sigma_Y^{-1/2} - \hat{\Sigma}_Y^{-1/2}\|_2 \leq \frac{1}{2} \max(\|\Sigma_Y\|_2, \|\hat{\Sigma}_Y\|_2)^{1/2} \|\Sigma_Y^{-1} - \hat{\Sigma}_Y^{-1}\|_2. \quad (129)$$

To bound $\|\Sigma_Y^{-1} - \hat{\Sigma}_Y^{-1}\|_2$, we use Corollary 11 of Kereta and Klock (2021): if

$$(\text{rank}(\Sigma_Y) + \log(\eta^{-1})) \left\| \Sigma_Y^{1/2} Y \right\|_{\psi_2}^4 \lesssim N$$

eventually, then with probability $1 - \eta$,

$$\|\Sigma_Y^{-1} - \hat{\Sigma}_Y^{-1}\|_2 \lesssim \left\| \Sigma_Y^{-1} Y \right\|_{\psi_2}^2 \sqrt{\frac{\text{rank}(\Sigma_Y) + \log(\eta^{-1})}{N}}. \quad (130)$$

The variance proxy condition on $\Sigma_Y^{1/2} Y$ (Definition B.1) implies that $\left\| \Sigma_Y^{-1/2} Y \right\|_{\psi_2}$ is bounded by an absolute constant. The condition $d = o(N)$ then ensures $(\text{rank}(\Sigma_Y) + \log(\eta^{-1})) \left\| \Sigma_Y^{1/2} Y \right\|_{\psi_2}^4 \lesssim N$ eventually. The variance proxy condition on $\Sigma_Y^{-1} Y$ implies $\left\| \Sigma_Y^{-1} Y \right\|_{\psi_2} \lesssim \left\| \Sigma_Y^{-1/2} \right\|_2$.

To bound $\|\hat{\Sigma}_Y\|_2$, we use the triangle inequality to obtain $\|\hat{\Sigma}_Y\|_2 \leq \|\hat{\Sigma}_Y - \Sigma_Y\|_2 + \|\Sigma_Y\|_2$. Then, using the first statement of the lemma and with the additional assumption that $\|\Sigma_Y\|_2^2 d = o(N)$, we have with probability $1 - \eta$,

$$\|\hat{\Sigma}_Y\|_2 \lesssim \|\Sigma_Y\|_2. \quad (131)$$

Combining these results together establishes the statement of the lemma and the proof is complete. \square

Proof of Theorems B.2 and B.3:

Proof of bounds for θ :

By definition, we have that

$$\theta_k = B \tilde{\eta}_k \gamma_k^{-1} \quad (132)$$

$$\hat{\theta}_k = \hat{B} \hat{\eta}_k \hat{\gamma}_k^{-1}. \quad (133)$$

We bound $\|\theta_k - \hat{\theta}_k\|_2$ by bounding all three of $\|B - \hat{B}\|_2$, $\|\tilde{\eta}_k - \hat{\eta}_k\|_2$, and $|\gamma_k^{-1} - \hat{\gamma}_k^{-1}|$. For ease of notation, we denote $\tilde{\eta}_k$ by v . We begin with

$$\|\theta_k - \hat{\theta}_k\|_2 \leq \|Bv\gamma_k^{-1} - Bv\hat{\gamma}_k^{-1}\|_2 + \|Bv\hat{\gamma}_k^{-1} - \hat{B}\hat{v}\hat{\gamma}_k^{-1}\|_2 \quad (134)$$

$$\leq |\gamma_k^{-1} - \hat{\gamma}_k^{-1}| \|Bv\|_2 + |\hat{\gamma}_k^{-1}| \|Bv - \hat{B}\hat{v}\|_2. \quad (135)$$

Examining $\|Bv - \hat{B}\hat{v}\|_2$ on the right-hand side of (135):

$$\|Bv - \hat{B}\hat{v}\|_2 \leq \|Bv - B\hat{v}\|_2 + \|B\hat{v} - \hat{B}\hat{v}\|_2 \quad (136)$$

$$\leq \|B\|_2 \|v - \hat{v}\|_2 + \|B - \hat{B}\|_2 \|\hat{v}\|_2. \quad (137)$$

Using (135) and since $\|\hat{v}\|_2 = 1$, we have

$$\|\theta_k - \hat{\theta}_k\|_2 \leq |\gamma_k^{-1} - \hat{\gamma}_k^{-1}| \|Bv\|_2 + |\hat{\gamma}_k^{-1}| (\|B\|_2 \|v - \hat{v}\|_2 + \|B - \hat{B}\|_2). \quad (138)$$

where we note that we now have bounded $\|\theta_k - \hat{\theta}_k\|_2$ in terms of $\|B - \hat{B}\|_2$, $\|\tilde{\eta}_k - \hat{\eta}_k\|_2$, and $|\gamma_k^{-1} - \hat{\gamma}_k^{-1}|$. To bound $|\gamma_k^{-1} - \hat{\gamma}_k^{-1}|$, we can use identity 16 from Section D, giving us that

$$|\gamma_k^{-1} - \hat{\gamma}_k^{-1}| \leq \min(\gamma_k, \hat{\gamma}_k)^{-3} |\gamma_k^2 - \hat{\gamma}_k^2|. \quad (139)$$

Bounding the two factors in equation (139) amounts to establishing that γ_k is close to $\hat{\gamma}_k$. For this, we apply Weyl's inequality (Bhatia (2013) Corollary III.2.6.) to obtain

$$|\gamma_k^2 - \hat{\gamma}_k^2| \leq \|B^\top \Sigma_X B - \hat{B}^\top \hat{\Sigma}_X \hat{B}\|_2. \quad (140)$$

since γ_k^2 is the k th eigenvalue of $B^\top \Sigma_X B$, and $\hat{\gamma}_k^2$ is the k th eigenvalue of $\hat{B}^\top \hat{\Sigma}_X \hat{B}$. We then deduce that

$$\gamma_k^2 - \|\hat{B}^\top \hat{\Sigma}_X \hat{B}\|_2 \leq \hat{\gamma}_k^2. \quad (141)$$

From equation (141) we establish that $\frac{1}{2}\gamma_k^2 \leq \hat{\gamma}_k^2$ by using

$$\|B^\top \Sigma_X B - \hat{B}^\top \hat{\Sigma}_X \hat{B}\|_2 \lesssim \frac{1}{2}\gamma_k^2, \quad (142)$$

which holds asymptotically in both the fast and slow rate cases under our assumption that the γ_k are bounded from below and from Corollary B.1. From $\frac{1}{2}\gamma_k^2 \leq \hat{\gamma}_k^2$ we also deduce

$$\min(\gamma_k, \hat{\gamma}_k)^{-2} \lesssim \frac{1}{\gamma_k^2}, \quad (143)$$

and additionally that

$$\frac{1}{\hat{\gamma}_k} \lesssim \frac{1}{\gamma_k}. \quad (144)$$

We now use our results thus far regarding γ_k and $\hat{\gamma}_k$ to obtain a simplified equation (138):

$$\|\theta_k - \hat{\theta}_k\|_2 \lesssim \frac{1}{\gamma_k^3} \|B^\top \Sigma_X B - \hat{B}^\top \hat{\Sigma}_X \hat{B}\|_2 \|Bv\|_2 + \frac{1}{\gamma_k} (\|B\|_2 \|v - \hat{v}\|_2 + \|B - \hat{B}\|_2) \quad (145)$$

To bound $\|v - \hat{v}\|_2$, we can use the Davis-Kahan theorem (Corollary 3 of Yu et al. (2015)). Assuming that $\tilde{\eta}_k^\top \hat{\eta}_k \geq 0$, then

$$\|\tilde{\eta}_k - \hat{\eta}_k\|_2 \leq \frac{2^{3/2} \|B^\top \Sigma_X B - \hat{B}^\top \hat{\Sigma}_X \hat{B}\|_2}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)}, \quad (146)$$

because $\tilde{\eta}_k$ is the k th eigenvector of $B^\top \Sigma_X B$, and $\hat{\eta}_k$ is the k th eigenvector of $\hat{B}^\top \hat{\Sigma}_X \hat{B}$. From this and equation (145), we obtain

$$\|\theta_k - \hat{\theta}_k\|_2 \lesssim \frac{1}{\gamma_k^2} \|B^\top \Sigma_X B - \hat{B}^\top \hat{\Sigma}_X \hat{B}\|_2 \|\theta_k\|_2 + \frac{1}{\gamma_k} \left(\|B\|_2 \frac{\|B^\top \Sigma_X B - \hat{B}^\top \hat{\Sigma}_X \hat{B}\|_2}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)} + \|B - \hat{B}\|_2 \right), \quad (147)$$

where we have also used the definition of θ_k , $\theta_k = Bv\gamma_k^{-1}$. Rearranging this expression, we have

$$\|\theta_k - \hat{\theta}_k\|_2 \lesssim \left(\frac{\|\theta_k\|_2}{\gamma_k^2} + \frac{\|B\|_2}{\gamma_k \min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)} \right) \|B^\top \Sigma_X B - \hat{B}^\top \hat{\Sigma}_X \hat{B}\|_2 + \frac{\|B - \hat{B}\|_2}{\gamma_k}. \quad (148)$$

Now we use our bounds for $\|B^\top \Sigma_X B - \hat{B}^\top \hat{\Sigma}_X \hat{B}\|_2$ and $\|B - \hat{B}\|_2$ depending on if we are in the slow or fast rate case. In the slow rate case, we also use

$$\|\theta_k\|_2 = \left\| \Sigma_X^{-1/2} \tilde{\theta}_k \right\|_2 \leq \left\| \Sigma_X^{-1/2} \right\|_2 \|\tilde{\theta}_k\|_2 = \left\| \Sigma_X^{-1/2} \right\|_2, \quad (149)$$

where we have used the standard result of classical CCA that $\theta_k = \Sigma_X^{-1/2} \tilde{\theta}_k$ for some unit vector $\tilde{\theta}_k$ Uurtio et al. 2018.

The proofs for both the fast and slow rate then follow directly from Lemma B.10, Theorem B.1, and rearranging of terms. This completes the proof of the bounds on $\|\theta_k - \hat{\theta}_k\|_2$.

Proof of bounds for η :

By definition, we have

$$\eta_k = \Sigma_Y^{-1/2} \tilde{\eta}_k, \quad (150)$$

$$\hat{\eta}_k = \hat{\Sigma}_Y^{-1/2} \hat{\tilde{\eta}}_k. \quad (151)$$

We bound $\|\eta_k - \hat{\eta}_k\|_2$ with the triangle inequality:

$$\|\eta_k - \hat{\eta}_k\|_2 \leq \left\| \Sigma_Y^{-1/2} \tilde{\eta}_k - \Sigma_Y^{-1/2} \hat{\tilde{\eta}}_k \right\|_2 + \left\| \Sigma_Y^{-1/2} \hat{\tilde{\eta}}_k - \hat{\Sigma}_Y^{-1/2} \hat{\tilde{\eta}}_k \right\|_2 \quad (152)$$

$$\leq \left\| \Sigma_Y^{-1/2} \right\|_2 \|\tilde{\eta}_k - \hat{\tilde{\eta}}_k\|_2 + \left\| \Sigma_Y^{-1/2} - \hat{\Sigma}_Y^{-1/2} \right\|_2 \|\hat{\tilde{\eta}}_k\|_2. \quad (153)$$

To simplify this expression, we can use $\|\hat{\tilde{\eta}}_k\|_2 = 1$, the Davis-Kahan Theorem for $\|\tilde{\eta}_k - \hat{\tilde{\eta}}_k\|_2$ as in the proof for the θ bounds, and the second statement of Lemma B.11. For clarity, we state the assumptions required for these results: $d = o(N)$, $\|\Sigma_Y\|_2^2 d = o(N)$, the variance proxy condition (Definition B.1) for Y and $\Sigma_Y^{-1} Y$, and $\tilde{\eta}_k^\top \hat{\tilde{\eta}}_k \geq 0$ for $k = 1, \dots, K$. Then, for $\eta \in (0, 1)$, we have

$$\|\eta_k - \hat{\eta}_k\|_2 \lesssim \left\| \Sigma_Y^{-1/2} \right\|_2 \left[\frac{\|\hat{B}^\top \hat{\Sigma}_X \hat{B} - B^\top \Sigma_X B\|_2}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)} + \|\Sigma_Y\|_2^{1/2} \left\| \Sigma_Y^{-1/2} \right\|_2 \sqrt{\frac{d \log(\eta^{-1})}{N}} \right]. \quad (154)$$

Now we apply Corollary B.1 to equation (154) under the fast and slow rate assumptions. In the slow rate case, we have

$$\|\eta_k - \hat{\eta}_k\|_2 \lesssim \left\| \Sigma_Y^{-1/2} \right\|_2 \left[\frac{\gamma_1 \|\Sigma_X\|_{2,\infty}^{1/2} \|B\|_{\ell_1, \ell_2}}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)} \left(\frac{d}{N} \log(p\eta^{-1}) \right)^{1/4} + \|\Sigma_Y\|_2^{1/2} \left\| \Sigma_Y^{-1/2} \right\|_2 \left(\frac{d \log(\eta^{-1})}{N} \right)^{1/2} \right]. \quad (155)$$

Factoring out $\left(\frac{d}{N}\right)^{1/4}$ we obtain

$$\|\eta_k - \hat{\eta}_k\|_2 \lesssim \left\| \Sigma_Y^{-1/2} \right\|_2 \left(\frac{d}{N} \right)^{1/4} \left[\frac{\gamma_1 \|\Sigma_X\|_{2,\infty}^{1/2} \|B\|_{\ell_1, \ell_2}}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)} \log(p\eta^{-1})^{1/4} + \|\Sigma_Y\|_2^{1/2} \left\| \Sigma_Y^{-1/2} \right\|_2 \log(\eta^{-1})^{1/2} \left(\frac{d}{N} \right)^{1/4} \right]. \quad (156)$$

In the bracketed expression we are able to combine the first and second terms, since we assume γ_1 is bounded from below, using the additional assumption that $\|\Sigma_Y\|_2^2 \|\Sigma_Y^{-1}\|_2^2 d = o(N)$, and because the other terms in the first term are greater than or equal to 1. Then, with probability $1 - \eta$,

$$\|\eta_k - \hat{\eta}_k\|_2 \lesssim \frac{\gamma_1 \|\Sigma_X\|_{2,\infty}^{1/2} \|B\|_{\ell_1, \ell_2} \left\| \Sigma_Y^{-1/2} \right\|_2}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)} \left(\frac{d \log(p\eta^{-1})}{N} \right)^{1/4}, \quad (157)$$

completing the proof of the slow-rate bound for η .

In the fast-rate case, under the fast-rate bound assumptions of Corollary B.1 and applying Corollary B.1 to equation (154), we establish

$$\|\eta_k - \hat{\eta}_k\|_2 \lesssim \|\Sigma_Y^{-1/2}\|_2 \max\left(\frac{\gamma_1 \|\Sigma_X\|_{2,\infty}^{1/2} s^{1/2} \kappa^{1/2}}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2)}, \|\Sigma_Y\|_2^{1/2} \|\Sigma_Y^{-1}\|_2^{1/2}\right) \left(\frac{d \log(p\eta^{-1})}{N}\right)^{1/2}, \quad (158)$$

completing the proof of the fast-rate bound for η . \square

C Asymmetric Sparse-Functional CCA: Proof of Theorem 4.2

In this section, we prove Theorem 4.2. Recall that we assume condition 2.1 holds, that is, the functional data admit a finite d -dimensional expansion $\text{Log}_\mu y = \sum_{j=1}^d Y_j \phi_j$, where ϕ_j are the PC functions of $\text{Log}_\mu y$ and $Y_j = \langle \text{Log}_\mu y, \phi_j \rangle_\mu$. Let $Y = (Y_1, \dots, Y_n)^\top$. Then, the canonical pairs $\{(\psi_k, \theta_k)\}$ can be computed by solving the multivariate CCA problem

$$\underset{\text{Var}(\eta_1^\top Y) = \text{Var}(\theta_1^\top X) = 1}{\text{maximize}} \quad \text{Corr}^2(\eta_1^\top Y, \theta_1^\top X), \quad (159)$$

with subsequent canonical pairs defined analogously. They are given by the pair $(\psi_k = \sum_{j=1}^d \phi_j \eta_{kj}, \theta_k)$. Recall that $\mathcal{C} \equiv \mathbb{E}[\text{Log}_\mu y \otimes \text{Log}_\mu y]$ admits the expansion $\mathcal{C} = \sum_{j=1}^\infty \omega_j \phi_j \otimes \phi_j$, where $\{\phi_j\}$ are the eigenfunctions of \mathcal{C} with associated eigenvalues $\{\omega_j\}$. We let $\gamma_1^2 \dots \gamma_d^2$ denote the squared canonical correlations attained by the pairs $(\psi_1, \theta_1), \dots, (\psi_d, \theta_d)$.

We denote by $\hat{\psi}_k$ and $\hat{\theta}_k$ the canonical vectors estimated using our proposed Algorithm 1. In practice, we are given a sample of N independent pairs

$$(y_i, x_i), \quad i = 1, \dots, N, \quad (160)$$

where each pair (y_i, x_i) is an independent observation of the pair (y, X) . Here, the functions $\{y_i\}$ are assumed to be fully observed on \mathcal{T} . We denote $\tau \equiv |\mathcal{T}|$, the length of the time interval of the functional data. We store the observations $\{x_i\}$ in a matrix $\mathbb{X} \in \mathbb{R}^{N \times p}$.

In Algorithm 1, we estimate μ using the sample Fréchet mean, denoted as $\hat{\mu}$, and estimate the eigenfunctions $\{\phi_j\}$ using $\hat{\phi}_j$, which are the eigenfunctions of the sample covariance function $\hat{\mathcal{C}} \equiv \frac{1}{N} \sum_{i=1}^N \text{Log}_{\hat{\mu}} y_i \otimes \text{Log}_{\hat{\mu}} y_i$.

Hence the functional data can be represented using the vector $Z \in \mathbb{R}^d$, where its j th element is $Z_j \equiv \langle \text{Log}_{\hat{\mu}} y_1, \hat{\phi}_j \rangle_{\hat{\mu}}$. We note that in the definition of Z , both $\hat{\mu}$ and $\hat{\phi}_j$ depend on y_1 since their estimation depends on the full sample y_1, \dots, y_N . We also note that the distribution of Z depends on the sample size N . In practice, we solve the following problem for the first pair of canonical variables:

$$\underset{\text{Var}(a_1^\top Z) = \text{Var}(b_1^\top X) = 1}{\text{maximize}} \quad \text{Corr}^2(a_1^\top Z, b_1^\top X). \quad (161)$$

The subsequent canonical pairs can be defined analogously. We denote the solutions to these problems as $(a_1, b_1), \dots, (a_d, b_d)$. We let $\tilde{\gamma}_1^2 \dots \tilde{\gamma}_d^2$ denote the squared canonical correlations attained by the pairs $(a_1, b_1), \dots, (a_d, b_d)$. We expect that, under appropriate assumptions, a_k and b_k will closely approximate η_k and θ_k , respectively, provided that $\hat{\mu}$ and $\{\hat{\phi}_j\}$ closely approximates μ and $\{\phi_j\}$, respectively.

We assume that the canonical vectors $\{\theta_k\}$ are group s -sparse, and that the associated vectors $\{b_k\}$ are also group s -sparse. Let the support $S \subseteq \{1, \dots, p\}$ represent the indices of non-zero elements of θ_k or b_k , with cardinality $|S| \leq 2s$. This sparsity condition allows us to simplify equations (159) and (161) by replacing X with $X_S \in \mathbb{R}^{|S|}$, the random vector consisting of only the entries $\{X_j : j \in S\}$. Moreover, θ_k and b_k can be replaced with $\theta_{k,S}$ and $b_{k,S}$, respectively.

We begin by deriving an error bound for the estimation of ψ_k using $\hat{\psi}_k$. Since the population quantity ψ_k belongs to $L^2(T\mu)$ and our estimate $\hat{\psi}_k$ belongs to $L^2(T\hat{\mu})$, we use the parallel transport operator $\Gamma_{\mu, \hat{\mu}}$ to define estimation error, as proposed in Lin and Yao (2019). For ease of notation, we denote $\Gamma_{f,g} U - V$ as $U \delta_\Gamma V \in L^2(Tg)$, for any vector fields $U \in L^2(Tf)$ and $V \in L^2(Tg)$.

C.1 Bounding the canonical function error

Next, we derive a bound for $\hat{\psi}_k \delta_\Gamma \psi_k \in L^2(T\mu)$. Recall that, by definition, $\psi_k = \sum_{j=1}^d \phi_j \eta_{kj}$ and $\hat{\psi}_k = \sum_{j=1}^d \hat{\phi}_j \hat{\eta}_{kj}$.

Lemma C.1. *The following inequality holds:*

$$\|\hat{\psi}_k \delta_\Gamma \psi_k\|_\mu^2 \lesssim \underbrace{\|\hat{\eta}_k - a_k\|_2^2}_{\text{Term I}} + \underbrace{\|a_k - \eta_k\|_2^2}_{\text{Term II}} + \underbrace{\left(\|\eta_k\|_\infty \sum_{j=1}^d \|\hat{\phi}_j \delta_\Gamma \phi_j\|_\mu \right)^2}_{\text{Term III}}. \quad (162)$$

Remark 9. *Term II, that is $\|a_k - \eta_k\|_2^2$, captures differences between the population CCA problem described in equations (159) and that in (161). The non-random nature of this term complicates the analysis, as it requires deriving bounds for the expectation rather than establishing probability bounds.*

The first term, $\|\hat{\eta}_k - a_k\|_2$, will be bounded using our multivariate CCA arguments. The second and third terms will be bounded in the following section.

C.1.1 Bounding Terms II and III of Lemma C.1

Assuming that Σ_Y and Σ_Z are invertible, and from the definitions of a_k and η_k as the solutions to the problems in equations (161) and (159) respectively, we have that

$$a_k = \Sigma_Z^{-1/2} \tilde{a}_k, \quad (163)$$

$$\eta_k = \Sigma_Y^{-1/2} \tilde{\eta}_k, \quad (164)$$

where \tilde{a}_k is the k th eigenvector (unit vector) of $A^\top A$ and $\tilde{\eta}$ is the k th eigenvector (unit vector) of $C^\top C$, where we have

$$A = \Sigma_{X_S}^{-1/2} \Sigma_{X_S Z} \Sigma_Z^{-1/2}, \quad (165)$$

$$C = \Sigma_{X_S}^{-1/2} \Sigma_{X_S Y} \Sigma_Y^{-1/2}. \quad (166)$$

Applying inequality 17 in Section D, we have that

$$\|a_k - \eta_k\|_2^2 \lesssim \left\| \Sigma_Z^{-1/2} - \Sigma_Y^{-1/2} \right\|_2^2 + \|\Sigma_Y^{-1}\|_2 \|\tilde{a}_k - \tilde{\eta}_k\|_2^2. \quad (167)$$

Noting that $A^\top A$ has the same eigenvectors as $|A|$, where $|A| \equiv (A^\top A)^{1/2}$, we can apply the Davis-Kahan theorem (Corollary 3 of Yu et al. (2015)) and obtain

$$\|\tilde{a}_k - \tilde{\eta}_k\|_2 \lesssim \frac{\||A| - |C|\|_2}{\min(\gamma_{k-1} - \gamma_k, \gamma_k - \gamma_{k+1})}, \quad k = 1, \dots, d, \quad (168)$$

where we assume $\tilde{a}_k^\top \tilde{\eta}_k \geq 0$ and we use the conventions $\gamma_{d+1} = -\infty$, and $\gamma_0 = \infty$. Next, note that

$$\||A| - |C|\|_2 \leq \||A| - |C|\|_F \leq \sqrt{2} \|A - C\|_F, \quad (169)$$

where the second inequality is identity 19 from Section D. We can then bound $\|A - C\|_F^2$ in terms of $\left\| \Sigma_Z^{-1/2} - \Sigma_Y^{-1/2} \right\|_2^2$ and $\mathbb{E}[\|Y - Z\|_2^2]$ using the following lemma.

Lemma C.2. *Under the group s -sparsity assumptions on b_k and θ_k , we have that*

$$\|A - C\|_F^2 \leq 2s \left[\mathbb{E}[\|Z\|_2^2] \left\| \Sigma_Z^{-1/2} - \Sigma_Y^{-1/2} \right\|_2^2 + \|\Sigma_Y^{-1}\|_2 \mathbb{E}[\|Y - Z\|_2^2] \right]. \quad (170)$$

Next, we bound $\|\Sigma_Z^{-1/2} - \Sigma_Y^{-1/2}\|_2^2$ in terms of $\mathbb{E}[\|Z - Y\|_2^2]$ using the following lemma.

Lemma C.3. *It can be shown that*

$$\|\Sigma_Z^{-1/2} - \Sigma_Y^{-1/2}\|_2^2 \lesssim \mathbb{E}[\|Z - Y\|_2^2] \max(\|\Sigma_Y^{-1}\|_2, \|\Sigma_Z^{-1}\|_2)^3 \max(\mathbb{E}[\|Z\|_2^2], \mathbb{E}[\|Y\|_2^2]). \quad (171)$$

Additionally, we have that

$$\|\Sigma_Y - \Sigma_Z\|_2^2 \lesssim \max(\mathbb{E}[\|Z\|_2^2], \mathbb{E}[\|Y\|_2^2]) \mathbb{E}[\|Z - Y\|_2^2]. \quad (172)$$

Combining the results of this section we obtain the following bound on $\|a_k - \eta_k\|_2^2$ in terms of $\mathbb{E}[\|Z - Y\|_2^2]$.

Lemma C.4. *Under the group s -sparsity assumptions on b_k and θ_k , we have*

$$\|a_k - \eta_k\|_2^2 \lesssim \mathbb{E}[\|Z - Y\|_2^2] \times \left[\left(1 + \frac{s \|\Sigma_Y^{-1}\|_2 \mathbb{E}[\|Z\|_2^2]}{\min(\gamma_{k-1} - \gamma_k, \gamma_k - \gamma_{k+1})^2} \right) \max(\|\Sigma_Y^{-1}\|_2, \|\Sigma_Z^{-1}\|_2)^3 \max(\mathbb{E}[\|Z\|_2^2], \mathbb{E}[\|Y\|_2^2]) + \frac{s \|\Sigma_Y^{-1}\|_2^2}{\min(\gamma_{k-1} - \gamma_k, \gamma_k - \gamma_{k+1})^2} \right]$$

Additionally, if we assume that $\|\Sigma_Z^{-1}\|_2 \lesssim \|\Sigma_Y^{-1}\|_2$, $\mathbb{E}[\|Z\|_2^2] \lesssim \mathbb{E}[\|Y\|_2^2]$, and $\|\Sigma_Y^{-1}\|_2, \mathbb{E}[\|Y\|_2^2] \geq 1$, then the statement simplifies as follows:

$$\|a_k - \eta_k\|_2^2 \lesssim \mathbb{E}[\|Z - Y\|_2^2] \frac{s \|\Sigma_Y^{-1}\|_2^4 \mathbb{E}[\|Y\|_2^2]^2}{\min(\gamma_{k-1} - \gamma_k, \gamma_k - \gamma_{k+1})^2}. \quad (173)$$

Remark 10. *In subsequent discussions, we will detail the conditions necessary for the additional assumptions stated here to hold.*

Having established a bound for $\|a_k - \eta_k\|_2^2$ in terms of $\mathbb{E}[\|Z - Y\|_2^2]$, we now turn our attention to bounding $\mathbb{E}[\|Z - Y\|_2^2]$ using 'known' quantities. In the process, we will derive a bound for $\mathbb{E}[\|\hat{\phi}_j \delta_\Gamma \phi_j\|_\mu]$, which will enable us to derive a probabilistic bound. This, in turn, will be used to bound Term III in Lemma C.1.

To establish a bound for $\mathbb{E}[\|Z - Y\|_2^2]$, we first begin by introducing a lemma to bound $\|Z - Y\|_2^2$.

Lemma C.5. *It can be shown that*

$$\|Z - Y\|_2^2 \leq 2d \|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^2 + 2 \sum_{j=1}^d \|\text{Log}_\mu y_i\|_\mu^2 \|\hat{\phi}_j \delta_\Gamma \phi_j\|_\mu^2. \quad (174)$$

Next, we aim to establish a bound for $\|\hat{\phi}_j \delta_\Gamma \phi_j\|_\mu^2$. Consider an operator $\hat{\mathcal{C}}$ on $L^2(T\hat{\mu})$. We use parallel transport to define $\Phi \hat{\mathcal{C}}$ as the operator on $L^2(T\mu)$ such that $\Phi \hat{\mathcal{C}}(V) = \Gamma_{\hat{\mu}, \mu}(\hat{\mathcal{C}}(\Gamma_{\mu, \hat{\mu}} V)) \in L^2(T\mu)$, for every $V \in L^2(T\mu)$. We also define the operator $\hat{\mathcal{C}}_\mu \equiv \frac{1}{N} \sum_{i=1}^N \text{Log}_\mu y_i \otimes \text{Log}_\mu y_i$ on $L^2(T\mu)$. Moreover, we use $\|\cdot\|_{\text{op}}$ to denote the operator norm on $L^2(T\mu)$.

Lemma C.6. *For any $j \geq 1$, we have that*

$$\|\hat{\phi}_j \delta_\Gamma \phi_j\|_\mu^2 \lesssim \frac{\|\mathcal{C} - \hat{\mathcal{C}}\|_{\text{op}}^2 + \|\hat{\mathcal{C}}_\mu - \Phi \hat{\mathcal{C}}\|_{\text{op}}^2}{\min(\omega_{j-1} - \omega_j, \omega_j - \omega_{j+1})^2}. \quad (175)$$

We introduce the following lemma to bound the expectation of the terms in the previous lemma.

Lemma C.7. *If $\mathbb{E} \left[\|\text{Log}_\mu y_1\|_\mu^4 \right] < \infty$, then*

$$\mathbb{E} \left[\|\mathcal{C} - \hat{\mathcal{C}}\|_{\text{op}}^2 \right] \leq \frac{1}{N} \mathbb{E} \left[\|\text{Log}_\mu y_1\|_\mu^4 \right]. \quad (176)$$

Additionally, under the assumption that $\mathbb{E} \left[\|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^4 \right]^{1/2} \lesssim \mathbb{E} \left[\|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^2 \right]$, we have

$$\mathbb{E} \left[\|\hat{\mathcal{C}}_\mu - \Phi \hat{\mathcal{C}}\|_{\text{op}}^2 \right] \lesssim \left(\mathbb{E} \left[\|\text{Log}_\mu y_i\|_\mu^4 \right]^{1/2} + \mathbb{E} \left[\|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^2 \right] \right) \mathbb{E} \left[\|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^2 \right]. \quad (177)$$

The lemma above shows that, in order to bound $\mathbb{E} \left[\|\hat{\phi}_j \delta_\Gamma \phi_j\|_\mu^2 \right]$, it is first necessary to bound $\mathbb{E} \left[\|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^2 \right]$. To do this, we need the following more technical results. But first, we state some preliminary definitions. Let $T\mathcal{M}$ denote the tangent bundle of \mathcal{M} , and let ∇ denote the Riemannian connection on \mathcal{M} . The next result, which is a mean value theorem for the parallel transport operation, is used in the proof of Lemma C.9.

Lemma C.8. *For a smooth vector field $U : \mathcal{M} \rightarrow T\mathcal{M}$, $x, y \in \mathcal{M}$, with minimizing geodesic $\gamma(t)$ between x and y (so that $\gamma(0) = x$ and $\gamma(d(x, y)) = y$), we have that*

$$\|\mathcal{P}_{y,x} U(y) - U(x)\|_{x \leq d(y,x)} \sup_{c \in [0, d(x,y)]} \|\nabla_{\gamma'(c)} U(\gamma(c))\|_{\gamma(c)}. \quad (178)$$

Remark 11. *When the vector field U has bounded Hessian H , defined below, we can use this result to bound the parallel transport error by the geodesic distance $d(x, y)$ between the base points of the vector field.*

For $x \in \mathcal{M}$ and $t \in \mathcal{T}$, define $f_t(x) \equiv \frac{1}{2} d^2(x, y_1(t))$. Let H_t be the Riemannian Hessian of f_t , i.e. $H_t(x) : T_x \mathcal{M} \rightarrow T_x \mathcal{M}$ such that for all $x \in \mathcal{M}$, $t \in \mathcal{T}$, and $v \in T_x \mathcal{M}$, $H_t(x)(v) = \nabla_v \text{grad } f_t(x)$. For a mapping $A(x)$, which for each x is an operator $A(x) : T_x \mathcal{M} \rightarrow T_x \mathcal{M}$, we define the operator norm at x of A : $\|A\|_{\text{op}, x} \equiv \sup_{v \in T_x \mathcal{M}, \|v\|_x=1} \|A(x)(v)\|_x$. Recall that $\tau = |\mathcal{T}|$ is the length of the time interval of the functional data.

Lemma C.9. *Assume that*

1. \mathcal{M} is a complete, simply-connected Riemannian manifold with nonpositive sectional curvature.
2. $\sup_{t \in \mathcal{T}} \mathbb{E} \left[d(y_1(t), y_2(t))^3 \right] < \infty$.
3. $\sup_{t \in \mathcal{T}, x \in \mathcal{M}} \|H_t(x)\|_{\text{op}, x}^2 \lesssim 1$ with probability 1.

Then,

$$\mathbb{E} \left[\|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^2 \right] \lesssim \frac{\tau}{N}. \quad (179)$$

The next lemma combines the results of Lemmas C.6, C.7, and C.9.

Lemma C.10. *Under the assumptions of Lemma C.9, and additionally assuming that $\mathbb{E} \left[\|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^4 \right]^{1/2} \lesssim \mathbb{E} \left[\|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^2 \right]$, and $\mathbb{E} \left[\|\text{Log}_\mu y_1\|_\mu^4 \right] \geq 1$, we have*

$$\mathbb{E} \left[\|\hat{\phi}_j \delta_\Gamma \phi_j\|_\mu^2 \right] \lesssim \frac{1}{N} \frac{\tau \mathbb{E} \left[\|\text{Log}_\mu y_1\|_\mu^4 \right]}{\min(\omega_{j-1} - \omega_j, \omega_j - \omega_{j+1})^2}. \quad (180)$$

From Lemma C.10 and Markov's inequality, we obtain the following inequality, which we can use to bound Term III in Lemma C.1.

Corollary C.1. *Under the assumptions of Lemma C.10, we have that*

$$\|\hat{\phi}_j \delta_\Gamma \phi_j\|_\mu^2 = O_P \left(\frac{1}{N} \frac{\tau \mathbb{E} \left[\|\text{Log}_\mu y_1\|_\mu^4 \right]}{\min(\omega_{j-1} - \omega_j, \omega_j - \omega_{j+1})^2} \right). \quad (181)$$

Now we can combine the results of Lemmas C.5, C.9 and C.10 to obtain a bound on $\mathbb{E}[\|Z - Y\|^2]$.

Lemma C.11. *Under the assumptions of Lemma C.10 and additionally assuming that*

$\mathbb{E} \left[\|\hat{\phi}_j \delta_\Gamma \phi_j\|_\mu^4 \right]^{1/2} \lesssim \mathbb{E} \left[\|\hat{\phi}_j \delta_\Gamma \phi_j\|_\mu^2 \right]$, *we have that*

$$\mathbb{E}[\|Z - Y\|_2^2] \lesssim \frac{\tau d}{N} \mathbb{E} \left[\|\text{Log}_\mu y_1\|_\mu^4 \right]^{3/2} \max_{j=1, \dots, d} \left(\frac{1}{\omega_j - \omega_{j+1}} \right)^2. \quad (182)$$

The proof follows from applying the Cauchy-Schwarz inequality and collecting similar terms.

Now we can state the conditions under which the additional assumptions of Lemma C.4 will hold. From now on, we keep tracking the terms τ and $\mathbb{E} \left[\|\text{Log}_\mu y_1\|_\mu^4 \right]$ in the error bounds, but we assume they are constant.

Lemma C.12. *We have that $\|\Sigma_Y^{-1}\|_2 = \omega_d^{-1}$, $\|\Sigma_Y\|_2 = \omega_1$, and $\mathbb{E}[\|Y\|_2^2] = \sum_{j=1}^d \omega_j$. Additionally, if $\mathbb{E}[\|Y\|_2^2], \|\Sigma_Y^{-1}\|_2 \geq 1$, and*

$$\frac{d \sum_{j=1}^d \omega_j}{N \omega_d^2} \max_{j=1, \dots, d} \left(\frac{1}{\omega_j - \omega_{j+1}} \right)^2 = o(1), \quad (183)$$

then $\mathbb{E}[\|Z - Y\|_2^2] = o(1)$, $\|\Sigma_Z^{-1}\|_2 \lesssim \|\Sigma_Y^{-1}\|_2$, $\|\Sigma_Z\|_2 \lesssim \|\Sigma_Y\|_2$ and $\mathbb{E}[\|Z\|_2^2] \lesssim \mathbb{E}[\|Y\|_2^2]$.

Now we can combine Lemmas C.4, C.11, and C.12 to obtain a final bound on $\|a_k - \eta_k\|_2^2$.

Theorem C.1. *Under the assumptions of Lemmas C.11 and C.12, we have that*

$$\|a_k - \eta_k\|_2^2 \lesssim \frac{\tau s d}{N} \max_{j=1, \dots, d} \left(\frac{1}{\omega_j - \omega_{j+1}} \right)^2 \left(\frac{\sum_{j=1}^d \omega_j}{\omega_d^2} \right)^2 \frac{\mathbb{E} \left[\|\text{Log}_\mu y_1\|_\mu^4 \right]^{3/2}}{\min(\gamma_{k-1} - \gamma_k, \gamma_k - \gamma_{k+1})^2}. \quad (184)$$

Before we combine the bounds for the three components detailed in Lemma C.1, we show that $\tilde{\gamma}_k$ and γ_k are asymptotically equivalent, which allows us to simplify the expression of our final bounds. Recall that the correlations $\{\tilde{\gamma}_k\}$ are defined using equation (161) as the canonical correlations between X and Z , while the correlations $\{\gamma_k\}$ are the canonical correlations between X and Y . To this end, we first state the following bound, which follows directly from Lemmas C.2 and C.3. This will also be used later to derive a bound for $\|\theta_k - \hat{\theta}_k\|_2$.

Lemma C.13. *Under the assumptions of Lemma C.12, we have that*

$$\|A - C\|_F^2 \lesssim s \|\Sigma_Y^{-1}\|_2^3 \mathbb{E}[\|Y\|_2^2]^2 \mathbb{E}[\|Z - Y\|_2^2]. \quad (185)$$

We can prove the next lemma by using Lemmas C.11, C.12, and C.13.

Lemma C.14. *Under the assumptions of Lemmas C.11 and C.12, and further assuming that τ and $\mathbb{E} \left[\|\text{Log}_\mu y_1\|_\mu^4 \right]$ are absolute constants, that the canonical correlations $\{\gamma_k\}$ and $\{\tilde{\gamma}_k\}$ are bounded from below and that*

$$\frac{d s \left(\sum_{j=1}^d \omega_j \right)^2}{N \omega_d^3} \max_{j=1, \dots, d} \left(\frac{1}{\omega_j - \omega_{j+1}} \right)^2 = o(1), \quad (186)$$

then γ_k and $\tilde{\gamma}_k$ are asymptotically equivalent, i.e., $\gamma_k \lesssim \tilde{\gamma}_k$ and $\tilde{\gamma}_k \lesssim \gamma_k$. Additionally, $\gamma_k^2 \lesssim \tilde{\gamma}_k^2$ and $\tilde{\gamma}_k^2 \lesssim \gamma_k^2$.

Now we are now in a position to establish bounds for all three terms in Lemma C.1, using Theorem B.3 (applied to X and Z and using the bound for η), Theorem C.1, and Corollary C.1 for Term I, II, and III, respectively. This will yield our final bound for the canonical functions $\|\hat{\psi}_k \delta_\Gamma \psi_k\|_\mu^2$, which is presented in Section C.3.

C.2 Bounding high-dimensional canonical vector error

In this section, we derive a bound for the estimation error of the high-dimensional canonical vectors, $\|\theta_k - \hat{\theta}_k\|_2$. Having already derived a bound for the canonical functions, the proof is straightforward. We start with an application of the triangle inequality. For $k = 1, \dots, d$,

$$\|\theta_k - \hat{\theta}_k\|_2^2 \lesssim \|\theta_k - b_k\|_2^2 + \|b_k - \hat{\theta}_k\|_2^2, \quad (187)$$

where b_k is the high dimensional canonical vector given by the solution to problem (161). Assuming s -group sparsity for θ_k and b_k , we represent the associated vectors with non-zero entries as $\theta_{k,S}$ and $b_{k,S}$. Recall that these are at most $2s$ -dimensional. By definition, we then have $\|\theta_k - b_k\|_2 = \|\theta_{k,S} - b_{k,S}\|_2$, hence

$$\|\theta_k - \hat{\theta}_k\|_2^2 \lesssim \|\theta_{k,S} - b_{k,S}\|_2^2 + \|b_k - \hat{\theta}_k\|_2^2. \quad (188)$$

To bound the second term, we can use Theorem B.3, applied to the random vectors X and Z . To bound the second term, we can make the following argument which is similar to the one made to bound $\|\eta_k - a_k\|_2$. Let

$$A = \Sigma_{X_S}^{-1/2} \Sigma_{X_S Z} \Sigma_Z^{-1/2}, \quad (189)$$

$$C = \Sigma_{X_S}^{-1/2} \Sigma_{X_S Y} \Sigma_Y^{-1/2}. \quad (190)$$

Then, from classical CCA (Uurtio et al. 2018), we have that $\theta_{k,S} = \Sigma_{X_S}^{-1/2} \tilde{\theta}_{k,S}$, where $\tilde{\theta}_{k,S}$ is the k th eigenvector of CC^\top , and $b_k = \Sigma_{X_S}^{-1/2} \tilde{b}_{k,S}$, where $\tilde{b}_{k,S}$ is the k th eigenvector of AA^\top . Therefore,

$$\|\theta_{k,S} - b_{k,S}\|_2 = \left\| \Sigma_{X_S}^{-1/2} \tilde{\theta}_{k,S} - \Sigma_{X_S}^{-1/2} \tilde{b}_{k,S} \right\|_2 \quad (191)$$

$$= \left\| \Sigma_{X_S}^{-1/2} (\tilde{\theta}_{k,S} - \tilde{b}_{k,S}) \right\|_2 \quad (192)$$

$$\leq \left\| \Sigma_{X_S}^{-1/2} \right\|_2 \|\tilde{\theta}_{k,S} - \tilde{b}_{k,S}\|_2. \quad (193)$$

Since $A^\top A$ has the same eigenvectors as $|A| = (AA^\top)^{1/2}$, we can then use the Davis-Kahan theorem (Yu et al. 2015) to derive the following bound. If $\tilde{\theta}_{k,S}^\top \tilde{b}_{k,S} \geq 0$, then

$$\|\tilde{\theta}_{k,S} - \tilde{b}_{k,S}\|_2 \lesssim \frac{\||A^\top| - |C^\top|\|_2}{\min(\gamma_{k-1} - \gamma_k, \gamma_k - \gamma_{k+1})}. \quad (194)$$

The term $\||A^\top| - |C^\top|\|_2$ can be bounded as follows:

$$\||A^\top| - |C^\top|\|_2 \leq \||A^\top| - |C^\top|\|_F \leq \sqrt{2} \|A^\top - C^\top\|_F = \sqrt{2} \|A - C\|_F, \quad (195)$$

where the second inequality is identity 19 from Section D. Having established a bound for $\|\tilde{\theta}_{k,S} - \tilde{b}_{k,S}\|_2$ in terms of $\|A - C\|_F$, we can apply similar arguments to those used to bound $\|a_k - \eta_k\|_2$ in order to derive a bound for $\|\tilde{\theta}_{k,S} - \tilde{b}_{k,S}\|_2$ in terms of $\mathbb{E}[\|Z - Y\|_2^2]$. Combining the results of this section, using Lemmas C.11 and C.13, we establish the following result.

Lemma C.15. *Under the assumptions of Lemmas C.11 and C.12, and if $\tilde{\theta}_{k,S}^\top \tilde{b}_{k,S} \geq 0$, we have that*

$$\|\theta_{k,S} - b_{k,S}\|_2^2 \lesssim \frac{\left\| \Sigma_{X_S}^{-1/2} \right\|_2^2 s \left\| \Sigma_Y^{-1} \right\|_2^3 \mathbb{E}[\|Y\|_2^2]^2}{\min(\gamma_{k-1} - \gamma_k, \gamma_k - \gamma_{k+1})^2} \frac{\tau d}{N} \mathbb{E} \left[\left\| \text{Log}_\mu y_1 \right\|_\mu^4 \right]^{3/2} \max_{j=1, \dots, d} \left(\frac{1}{\omega_j - \omega_{j+1}} \right)^2. \quad (196)$$

Next, by applying Theorem B.3, Lemma C.15, and equation (188), we can establish the final bound for the high-dimensional canonical vector error $\|\theta_k - \hat{\theta}_k\|_2^2$, which is presented in Section C.3.

C.3 Final rates

In this section, we presented our final results. Recall that we denote as $K = \max\{i \in \{1, \dots, d\} : \gamma_i > 0\}$ the number of nontrivial canonical vectors and we use the conventions $\gamma_{d+1}^2 = -\infty$, $\tilde{\gamma}_{d+1}^2 = -\infty$, $\gamma_0^2 = +\infty$ and $\tilde{\gamma}_0^2 = +\infty$.

For clarity, we first provide a comprehensive list of our assumptions. For the definitions of the quantities that appear below, please see the beginning of Section C.

Assumption C.1 (Manifold Properties).

1. The manifold \mathcal{M} is a complete simply-connected Riemannian manifold with nonpositive sectional curvature.
2. The curvature $\sup_{t \in \mathcal{T}, x \in \mathcal{M}} \|H_t(x)\|_{\text{op}, x}^2$ is bounded with probability 1.

Assumption C.2 (Distributional Assumptions).

1. The random vectors X and Z are strict subgaussian random vectors (Definition B.2).
2. The covariance matrices Σ_X , Σ_Y , and Σ_Z are invertible.
3. The group s -sparsity assumption holds for $\{b_k\}$ and $\{\theta_k\}$.
4. The matrix $\Sigma_X^{1/2}$ satisfies the group restricted eigenvalue condition $RE(s, 3, d)$ (Definition B.3) with parameter $\kappa = \kappa(s, d, \Sigma_X^{1/2})$.
5. The functional data are such that $\sup_{t \in \mathcal{T}} \mathbb{E} [d(y_1(t), y_2(t))^3] < \infty$.

Assumption C.3 (Rate Assumptions).

1. $d \leq p$;
2. $\text{cond}(\Sigma_Y)^2 d = o(N)$;
3. $\kappa^2 s^2 d \log(p) = o(N)$;
4. $ds \frac{(\sum_{j=1}^d \omega_j)^2}{\omega_d^3} \max_{j=1, \dots, d} \left(\frac{1}{\omega_j - \omega_{j+1}} \right)^2 = o(N)$;
5. The correlations $\gamma_1, \dots, \gamma_K$ are bounded from below and distinct from one another, as well as $\tilde{\gamma}_1, \dots, \tilde{\gamma}_K$.
6. $\mathbb{E} \left[\|\hat{\phi}_j \delta_\Gamma \phi_j\|_\mu^4 \right]^{1/2} \lesssim \mathbb{E} \left[\|\hat{\phi}_j \delta_\Gamma \phi_j\|_\mu^2 \right]$;
7. $\mathbb{E} \left[\|\text{Log}_\mu y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^4 \right]^{1/2} \lesssim \mathbb{E} \left[\|\text{Log}_\mu y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^2 \right]$.

Assumption C.4 (Minor Assumptions).

1. The quantities $\|\Sigma_X\|_{2, \infty}$, $\|T\|_{\ell_1, \ell_2}$ are bounded from above and are ≥ 1 .
2. The variables τ and $\mathbb{E} \left[\|\text{Log}_\mu y_1\|_\mu^4 \right]$ are constants.
3. The following quantities are larger than 1: $\kappa, \omega_1, \omega_d^{-1}, \|\eta\|_\infty, \|\Sigma_X^{-1}\|_2, \|\Sigma_Z^{-1}\|_2$.
4. $a_k^\top \Sigma_Z^{1/2} \Sigma_Y^{-1/2} \eta_k \geq 0$ and $\hat{\eta}_k^\top \hat{\Sigma}_Z^{-1/2} \Sigma_Z^{1/2} a_k \geq 0$ for $k = 1, \dots, K$.

Next, we present our main results.

Theorem C.2 (Canonical Function Error Bound). *Under Assumptions C.1-C.4, we have*

$$\|\hat{\psi}_k \delta_{\Gamma} \psi_k\|_{\mu}^2 = O_P \left(\frac{d^2 s \log(p)}{N} \tau \kappa \|\Sigma_X\|_{2, \infty} \|\eta\|_{\infty}^2 \mathbb{E} \left[\|\text{Log}_{\mu} y\|_{\mu}^4 \right]^{3/2} \right) \quad (197)$$

$$\cdot \left(\frac{\sum_{j=1}^d \omega_j}{\omega_d^2} \right)^2 \max_{j=1, \dots, d} \left(\frac{1}{\omega_j - \omega_{j+1}} \right)^2 \frac{1}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2, \gamma_{k-1} - \gamma_k, \gamma_k - \gamma_{k+1})^2}, \quad (198)$$

with $k = 1, \dots, K$.

Theorem C.3 (Canonical Vector Error Bound). *Under Assumptions C.1-C.4, and additionally assuming that $\theta_{k,S}^{\top} \Sigma_{X_S} b_{k,S} \geq 0$, we have*

$$\|\theta_k - \hat{\theta}_k\|_2^2 = O_P \left(\frac{d \tau s \log(p)}{N} \right) \frac{1}{\min(\gamma_{k-1}^2 - \gamma_k^2, \gamma_k^2 - \gamma_{k+1}^2, \gamma_{k-1} - \gamma_k, \gamma_k - \gamma_{k+1})^2} \quad (199)$$

$$\cdot \left[\left\| \Sigma_{X_S}^{-1/2} \right\|_2 \frac{\left(\sum_{j=1}^d \omega_j \right)^2}{\omega_d^3} \max_{j=1, \dots, d} \left(\frac{1}{\omega_j - \omega_{j+1}} \right)^2 \mathbb{E} \left[\|\text{Log}_{\mu} y\|_{\mu}^4 \right]^{3/2} + \left(\frac{\gamma_1}{\gamma_k} \right)^2 \|\Sigma_X\|_{2, \infty} \kappa^2 \left\| \Sigma_X^{-1} \right\|_2 \right], \quad (200)$$

with $k = 1, \dots, K$.

Here we make a few remarks on the assumptions of Theorems C.2 and C.3. Some of the more technical assumptions arise from avoiding overly simplifying assumptions. For instance, with the exception of the curvature-related quantity $H_t(x)$, we do not assume that the random variables/functions are bounded. Furthermore, we avoid assuming Gaussianity of the random variables of interest and instead assume these are sub-Gaussians. As in Lin and Yao (2019), we do not assume that the Fréchet mean μ is known and instead estimate it using its sample version $\hat{\mu}$; this choice introduces significant complexity, making it necessary to use the parallel transport operator to compare estimates and estimands, which are defined in different tangent spaces. These challenges are further compounded by the dimensionality reduction step that takes place before CCA, which requires that we derive bounds in *expectation* rather than in probability. Specifically, this requires showing that $\mathbb{E} \left[d(\hat{\mu}(t), \mu(t))^2 \right] \lesssim \frac{1}{N}$, i.e., equation (251), using the results of Schötz (2019) as opposed to those of Lin and Yao (2019). See also Remark 9.

Remarks on Assumption C.1:

Assumption C.1 is required to bound the term $\mathbb{E} \left[d(\hat{\mu}(t), \mu(t))^2 \right]$. See also Lemma C.9. Here, $H_t(x)$ is the Riemannian Hessian of the random function $d^2(x, y(t))$, and is related to curvature on the manifold (Pennec 2017). Assuming this quantity is bounded allows us to bound the parallel transport distance in terms of the geodesic distance, as shown in Lemma C.8.

Remarks on Assumption C.2:

Items 1-4 of Assumption C.2 are used to facilitate the application of our multivariate CCA results in Section 4.1 to the sparse-functional setting considered here. Specifically, Items 3-4 are used in particular to get fast-rate bounds, which match the root- n estimation rate of the functional quantities. We note that in item 3, we do not require that θ_k and b_k have the same sparsity structure, but only that they are both s -sparse. Item 4 is a generalization of the standard restricted-eigenvalue condition in Lasso theory (Hastie et al. 2015) and is equivalent to the one proposed in Gaynanova (2020). Item 5 is a weak assumption about the boundedness of the variance of $y(t)$ on the manifold and along with Assumption C.1 is necessary to show that $\mathbb{E} \left[d(\hat{\mu}(t), \mu(t))^2 \right]$ is root- n consistent.

Remarks on Assumption C.3:

Item 1 of Assumption C.3 formalizes our asymmetrical treatment of the data, where we assume that the rank of the functional data is smaller than the dimension of the high-dimensional data. Items 2 - 5 are mainly used to simplify the theorem statements by allowing us to bound norms of estimated quantities using the corresponding population quantities. In particular, item 2 is used in

conjunction with Lemma C.12 to show that $\|\hat{\Sigma}_Z\|_2 \lesssim \|\Sigma_Y\|_2$. Item 3 allows us to only make a group restricted eigenvalue assumption on $\Sigma_X^{1/2}$ rather than the data matrix $\frac{1}{\sqrt{N}}\mathbb{X}$ (see Lemma B.9). Item 4 states that the variances ω_j of the principal scores Y_j should not shrink too quickly as N and d grow. Note that if d is assumed constant, the condition reduces to $s = o(N)$. This is used to establish the simplifying bounds given in Lemmas C.12 and C.14. Item 5 is used to replace $\hat{\gamma}_k$ and $\tilde{\gamma}_k$ with the population canonical correlations γ_k (see Lemma C.14 and the discussion in the proof of Theorem B.2). Items 6 and 7 are technical conditions. These mainly arise from the complexity of the setting considered here. The assumption $\mathbb{E}\left[\|\hat{\phi}_j\delta_\Gamma\phi_j\|_\mu^4\right]^{1/2} \lesssim \mathbb{E}\left[\|\hat{\phi}_j\delta_\Gamma\phi_j\|_\mu^2\right]$ could be replaced with a boundness assumption on $\|\text{Log}_\mu y_i\|_\mu$, or alternatively, by adopting a sample splitting strategy to estimate $\hat{\mu}$, the principal components $\{\hat{\phi}_j\}$, and to carry out CCA. This can be seen from the second term of equation (174), where the absence of one of these assumptions requires that we use the Cauchy-Schwarz inequality, introducing fourth moments. Note that if we were to assume that $\hat{\mu} = \mu$, then $\mathbb{E}\left[\|\text{Log}_{\hat{\mu}} y_i\delta_\Gamma \text{Log}_\mu y_i\|_\mu^4\right]^{1/2} = 0$, immediately satisfying the condition in item 7.

Remarks on Assumption C.4:

Items 1-3 are not critical and only serve the purpose of simplifying the theorem statements. Item 4 is introduced to account for the sign ambiguity of the CCA solutions.

C.4 Proofs for results in Section C

Proof of Lemma C.1:

We define $\tilde{\psi}_k = \sum_{j=1}^d \hat{\phi}_j a_{kj}$. For ease of notation, we drop the k in writing $\hat{\psi}_k, \psi_k$ and $\tilde{\psi}_k, \eta_k$, etc. We have

$$\|\hat{\psi}\delta_\Gamma\psi\|_\mu^2 = \|\Gamma_{\hat{\mu},\mu}\hat{\psi} - \psi\|_\mu^2 \leq 2\|\Gamma_{\hat{\mu},\mu}\hat{\psi} - \Gamma_{\hat{\mu},\mu}\tilde{\psi}\|_\mu^2 + 2\|\Gamma_{\hat{\mu},\mu}\tilde{\psi} - \psi\|_\mu^2. \quad (201)$$

The first term in equation (201) is

$$2\|\Gamma_{\hat{\mu},\mu}(\hat{\psi} - \tilde{\psi})\|_\mu^2 = \|\Gamma_{\mu,\hat{\mu}}(\Gamma_{\hat{\mu},\mu}(\hat{\psi} - \tilde{\psi}))\|_{\hat{\mu}}^2 = 2\|\hat{\psi} - \tilde{\psi}\|_{\hat{\mu}}^2.$$

Define $\bar{\psi}$ as $\sum_{j=1}^d \hat{\phi}_j \eta_j$. Then the second term in equation (201) is

$$\leq 4\|\Gamma_{\hat{\mu},\mu}\tilde{\psi} - \Gamma_{\hat{\mu},\mu}\bar{\psi}\|_\mu^2 + 4\|\Gamma_{\hat{\mu},\mu}\bar{\psi} - \psi\|_\mu^2.$$

Therefore,

$$\|\hat{\psi}\delta_\Gamma\psi\|_\mu^2 \lesssim \|\hat{\psi} - \tilde{\psi}\|_{\hat{\mu}}^2 + \|\tilde{\psi} - \bar{\psi}\|_{\hat{\mu}}^2 + \|\bar{\psi}\delta_\Gamma\psi\|_\mu^2. \quad (202)$$

The first term in equation (202) is

$$\|\hat{\psi} - \tilde{\psi}\|_{\hat{\mu}}^2 = \left\| \sum_{j=1}^d \hat{\phi}_j \hat{\eta}_j - \sum_{j=1}^d a_j \hat{\phi}_j \right\|_{\hat{\mu}}^2 = \left\| \sum_{j=1}^d \hat{\phi}_j (\hat{\eta}_j - a_j) \right\|_{\hat{\mu}}^2 = \|\hat{\eta} - a\|_2^2, \quad (203)$$

where in the third equality we have used that the $\hat{\phi}_j$ are orthonormal in $L^2(T\hat{\mu})$. Similarly, the second term is

$$\|\tilde{\psi} - \bar{\psi}\|_{\hat{\mu}}^2 = \|a - \eta\|_2^2. \quad (204)$$

By the triangle inequality, the third term is

$$\|\bar{\psi}\delta_\Gamma\psi\|_\mu^2 \leq \left(\|\eta\|_\infty \sum_{j=1}^d \|\hat{\phi}_j\delta_\Gamma\phi_j\|_\mu \right)^2, \quad (205)$$

completing the proof. \square

Proof of Lemma C.2:

For ease of notation, we write X_S as X throughout the proof of the lemma. From the definition of A and C , we have

$$\|A - C\|_F = \left\| \Sigma_X^{-1/2} \Sigma_{XZ} \Sigma_Z^{-1/2} - \Sigma_X^{-1/2} \Sigma_{XY} \Sigma_Y^{-1/2} \right\|_F = \left\| \Sigma_X^{-1/2} \left(\mathbb{E}[XZ^\top] \Sigma_Z^{-1/2} - \mathbb{E}[XY^\top] \Sigma_Y^{-1/2} \right) \right\|_F. \quad (206)$$

From the linearity of expectation, this is

$$= \left\| \mathbb{E} \left[\Sigma_X^{-1/2} X \left(Z^\top \Sigma_Z^{-1/2} - Y^\top \Sigma_Y^{-1/2} \right) \right] \right\|_F = \left\| \mathbb{E} \left[\Sigma_X^{-1/2} X \left(\Sigma_Z^{-1/2} Z - \Sigma_Y^{-1/2} Y \right)^\top \right] \right\|_F. \quad (207)$$

By Theorem 2.6.7 of Hsing and Eubank (2015), using the Frobenious norm of an outer product, and the Cauchy-Schwarz inequality, we have

$$\|A - C\|_F \leq \mathbb{E} \left[\left\| \Sigma_X^{-1/2} X \left(\Sigma_Z^{-1/2} Z - \Sigma_Y^{-1/2} Y \right)^\top \right\|_F \right] \quad (208)$$

$$= \mathbb{E} \left[\left\| \Sigma_X^{-1/2} X \right\|_2 \left\| \Sigma_Z^{-1/2} Z - \Sigma_Y^{-1/2} Y \right\|_2 \right] \quad (209)$$

$$\leq \mathbb{E} \left[\left\| \Sigma_X^{-1/2} X \right\|_2^2 \right]^{1/2} \mathbb{E} \left[\left\| \Sigma_Z^{-1/2} Z - \Sigma_Y^{-1/2} Y \right\|_2^2 \right]^{1/2}. \quad (210)$$

We have that $\mathbb{E} \left[\left\| \Sigma_X^{-1/2} X \right\|_2^2 \right] = \mathbb{E} \left[\text{tr} \left(\Sigma_X^{-1} X X^\top \right) \right] = \text{tr} \left(I_s \right) = s$. Using this along with identity 17 from Section D to upper bound $\mathbb{E} \left[\left\| \Sigma_Z^{-1/2} Z - \Sigma_Y^{-1/2} Y \right\|_2^2 \right]^{1/2}$ completes the proof.

Proof of Lemma C.3:

The first statement follows from the second statement, by identity 14 from Section D and Lemma C.3. To show the second statement, we begin with

$$\|\Sigma_Y - \Sigma_Z\|_2 = \left\| \mathbb{E}[YY^\top] - \mathbb{E}[ZZ^\top] \right\|_2 = \left\| \mathbb{E}[YY^\top - ZZ^\top] \right\|_2, \quad (211)$$

and using Theorem 2.6.7 of Hsing and Eubank (2015), this is

$$\leq \mathbb{E} \left[\|YY^\top - ZZ^\top\|_2 \right] = \mathbb{E} \left[\|YY^\top - ZY^\top + ZY^\top - ZZ^\top\|_2 \right] = \mathbb{E} \left[\|Z(Z - Y)^\top + (Z - Y)Y^\top\|_2 \right]. \quad (212)$$

From this, the triangle inequality, and using the two-norm of an outer product, we have

$$\|\Sigma_Y - \Sigma_Z\|_2 \leq \mathbb{E} \left[\|Z(Z - Y)^\top\|_2 \right] + \mathbb{E} \left[\|(Z - Y)Y^\top\|_2 \right] = \mathbb{E} \left[\|Z\|_2 \|Z - Y\|_2 \right] + \mathbb{E} \left[\|Z - Y\|_2 \|Y\|_2 \right] \quad (213)$$

By the Cauchy-Schwarz inequality, the right hand side is

$$= \mathbb{E} \left[(\|Z\|_2 + \|Y\|_2) \|Z - Y\|_2 \right] \leq \mathbb{E} \left[(\|Z\|_2 + \|Y\|_2)^2 \right]^{1/2} \mathbb{E} \left[\|Z - Y\|_2^2 \right]^{1/2} \quad (214)$$

The second statement of the lemma follows, and the proof is complete. \square

Proof of Lemma C.5:

Define $W \in \mathbb{R}^d$ as the random vector with $W_j \equiv \langle \text{Log}_{\hat{\mu}} y_1, \Gamma_{\hat{\mu}, \mu} \hat{\phi}_j \rangle_{\hat{\mu}}$ for $j = 1, \dots, d$. Then

$$\|Z - Y\|_2^2 \leq 2 \|Z - W\|_2^2 + 2 \|W - Y\|_2^2. \quad (215)$$

We have

$$Z_j - W_j = \langle \text{Log}_{\hat{\mu}} y_1, \hat{\phi}_j \rangle_{\hat{\mu}} - \langle \text{Log}_{\mu} y_1, \Gamma_{\hat{\mu}, \mu} \hat{\phi}_j \rangle_{\mu} \quad (216)$$

$$= \langle \text{Log}_{\hat{\mu}} y_1, \hat{\phi}_j \rangle_{\hat{\mu}} - \langle \Gamma_{\mu, \hat{\mu}} \text{Log}_{\mu} y_1, \hat{\phi}_j \rangle_{\hat{\mu}} \quad (217)$$

$$= \langle \text{Log}_{\hat{\mu}} y_1 - \Gamma_{\mu, \hat{\mu}} \text{Log}_{\mu} y_1, \hat{\phi}_j \rangle_{\hat{\mu}}, \quad (218)$$

and therefore, by the Cauchy-Schwarz inequality, and because the $\hat{\phi}_j$ are orthonormal along $\hat{\mu}$,

$$|Z_j - W_j| \leq \|\text{Log}_{\hat{\mu}} y_1 - \Gamma_{\mu, \hat{\mu}} \text{Log}_{\mu} y_1\|_{\hat{\mu}} \|\hat{\phi}_j\|_{\hat{\mu}} \quad (219)$$

$$= \|\text{Log}_{\mu} y_1 \delta_{\Gamma} \text{Log}_{\hat{\mu}} y_1\|_{\hat{\mu}} \quad (220)$$

$$= \|\text{Log}_{\hat{\mu}} y_1 \delta_{\Gamma} \text{Log}_{\mu} y_1\|_{\mu}, \quad (221)$$

where in the last equality we have used that $\|U \delta_{\Gamma} V\|_{\mu} = \|V \delta_{\Gamma} U\|_{\hat{\mu}}$ for $U \in L^2(T\hat{\mu})$ and $V \in L^2(T\mu)$. We also have

$$W_j - Y_j = \langle \text{Log}_{\mu} y_1, \Gamma_{\hat{\mu}, \mu} \hat{\phi}_j \rangle_{\mu} - \langle \text{Log}_{\mu} y_1, \phi_j \rangle_{\mu} \quad (222)$$

$$= \langle \text{Log}_{\mu} y_1, \hat{\phi}_j \delta_{\Gamma} \phi_j \rangle_{\mu}, \quad (223)$$

so that again by the Cauchy-Schwarz inequality,

$$|W_j - Y_j| \leq \|\text{Log}_{\mu} y_1\|_{\mu} \|\hat{\phi}_j \delta_{\Gamma} \phi_j\|_{\mu}. \quad (224)$$

Thus,

$$\|Z - Y\|_2^2 \leq 2 \sum_{j=1}^d |Z_j - W_j|^2 + |W_j - Y_j|^2 \quad (225)$$

$$\leq 2 \sum_{j=1}^d \|\text{Log}_{\hat{\mu}} y_1 \delta_{\Gamma} \text{Log}_{\mu} y_1\|_{\mu}^2 + \|\text{Log}_{\mu} y_1\|_{\mu}^2 \|\hat{\phi}_j \delta_{\Gamma} \phi_j\|_{\mu}^2, \quad (226)$$

from which the statement of the lemma follows, and the proof is complete. \square

Proof of Lemma C.6:

From Lin and Yao (2019) (page 3551), $\Gamma_{\hat{\mu}, \mu} \hat{\phi}_j$ are the eigenvectors of $\Phi \hat{\mathcal{C}}$. By definition, the ϕ_j are the eigenvectors of \mathcal{C} . Thus, we can use the Davis-Kahan Theorem for Hilbert spaces (Jirak and Wahl 2020) to obtain that

$$\Gamma_{\hat{\mu}, \mu} \hat{\phi}_j - \phi_j \leq 2\sqrt{2} \max((\omega_{j-1} - \omega_j)^{-1}, (\omega_j - \omega_{j+1})^{-1}) \|\mathcal{C} - \Phi \hat{\mathcal{C}}\|_{\text{op}}, \quad (227)$$

where $\omega_0 \equiv \infty$ and $\|\mathcal{C}\|_{\text{op}} \equiv \sup_{U \in L^2(T\mu), \|U\|_{\mu}=1} \|\mathcal{C}U\|_{\mu}$. Adding and subtracting $\hat{\mathcal{C}}_{\mu}$ in $\|\mathcal{C} - \Phi \hat{\mathcal{C}}\|_{\text{op}}$ and using identity 18 from Section D, the proof is complete. \square

Proof of Lemma C.7:

The first statement is equivalent to Lemma 5.2 of Cardot et al. (1999). To show the second statement, we begin with the definition of $\hat{\mathcal{C}}$ and Proposition 2 item 5 of Lin and Yao (2019), from which we deduce that

$$\Phi \hat{\mathcal{C}} = \frac{1}{N} \sum_{i=1}^N (\Gamma_{\hat{\mu}, \mu} \text{Log}_{\hat{\mu}} y_i) \otimes (\Gamma_{\hat{\mu}, \mu} \text{Log}_{\hat{\mu}} y_i). \quad (228)$$

This implies

$$\hat{\mathcal{C}}_{\mu} - \Phi \hat{\mathcal{C}} = \frac{1}{N} \sum_{i=1}^N a_i \otimes a_i - b_i \otimes b_i, \quad (229)$$

where we denote $a_i \equiv \text{Log}_{\mu} y_i$ and $b_i \equiv \Gamma_{\hat{\mu}, \mu} \text{Log}_{\hat{\mu}} y_i$. It is straightforward to show that $a \otimes a - b \otimes b = a \otimes (a - b) + (a - b) \otimes (b - a) + (a - b) \otimes a$, where $a, b \in L^2(T\mu)$. From Theorem 3.4.7. of Hsing and Eubank (2015), $\|a \otimes b\|_{\text{op}} = \|a\|_{\mu} \|b\|_{\mu}$, and therefore,

$$\|a \otimes a - b \otimes b\|_{\text{op}} \leq 2 \|a\|_{\mu} \|a - b\|_{\mu} + \|a - b\|_{\mu}^2. \quad (230)$$

Then,

$$\|\hat{\mathcal{C}}_\mu - \Phi\hat{\mathcal{C}}_{\text{op}}\| \leq \frac{1}{N} \sum_{i=1}^N 2 \|\text{Log}_\mu y_i\|_\mu \|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu + \|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^2, \quad (231)$$

so that

$$\|\hat{\mathcal{C}}_\mu - \Phi\hat{\mathcal{C}}_{\text{op}}\|^2 \lesssim \left(\frac{1}{N} \sum_{i=1}^N 2 \|\text{Log}_\mu y_i\|_\mu \|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu \right)^2 + \left(\frac{1}{N} \sum_{i=1}^N \|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu^2 \right)^2. \quad (232)$$

It is straightforward to show that, for any i.i.d. random variables W_i with finite variance we have

$$\mathbb{E} \left[\left(\frac{1}{N} \sum_{i=1}^N W_i \right)^2 \right] \leq \mathbb{E} [W_1^2], \quad (233)$$

where the W_i are not required to have mean 0. Taking the expectation in equation (232), applying this last result on each term, using the Cauchy-Schwarz inequality, and subsequently using the assumption stated in the Lemma concerning $\|\text{Log}_{\hat{\mu}} y_i \delta_\Gamma \text{Log}_\mu y_i\|_\mu$, the proof is complete. \square

Proof of Lemma C.8:

We start with a mean value theorem result for smooth functions f from $[0, t_1]$ into a normed vector space V , where $t_1 \in \mathbb{R}$. By Theorem 1.1.1. of Hörmander (2015), we have

$$\|f(t_1) - f(0)\| = |b - a| \sup_{c \in [0, t_1]} f'(c). \quad (234)$$

We set

$$f(t) = \Gamma_{\gamma(t), \gamma(0)} U(\gamma(t)) - U(\gamma(0)), \quad (235)$$

where U is a smooth vector field on \mathcal{M} , $U : \mathcal{M} \rightarrow T\mathcal{M}$, and $\gamma(t)$ is the minimizing geodesic between two points $x, y \in \mathcal{M}$ ($\gamma(0) = x, \gamma((x, y)) = y$). Letting $t_1 = d(x, y)$, then $\gamma : [0, t_1] \rightarrow \mathcal{M}$, and $f : [0, t_1] \rightarrow T_x\mathcal{M}$. Letting $\|W\|_x$ denote the norm of $W \in T_x\mathcal{M}$, and using $f(t)$ in equation (234), we have

$$\|\Gamma_{y,x} U(y) - U(x)\| \leq d(x, y) \sup_{c \in [0, d(x, y)]} \|f'(c)\|. \quad (236)$$

We can determine $f'(c)$:

$$f'(c) = \lim_{t \rightarrow 0^+} \frac{f(c+t) - f(c)}{t} \quad (237)$$

$$= \lim_{t \rightarrow 0^+} \frac{\Gamma_{\gamma(c+t), x} U(\gamma(c+t)) - U(x) - \Gamma_{\gamma(c), x} U(\gamma(c)) + U(x)}{t} \quad (238)$$

$$= \lim_{t \rightarrow 0^+} \frac{\Gamma_{\gamma(c+t), x} U(\gamma(c+t)) - \Gamma_{\gamma(c), x} U(\gamma(c))}{t}. \quad (239)$$

Using that $\Gamma_{z,x} U(z) = \Gamma_{y,x} (\Gamma_{z,y} U(z))$, where $x = x, y = \gamma(c)$, and $z = \gamma(c+t)$, we have

$$f'(c) = \lim_{t \rightarrow 0^+} \frac{\Gamma_{\gamma(c), x} [\Gamma_{\gamma(c+t), \gamma(c)} U(\gamma(c+t)) - U(\gamma(c))]}{t} \quad (240)$$

$$= \Gamma_{\gamma(c), x} \left[\lim_{t \rightarrow 0^+} \frac{\Gamma_{\gamma(c+t), \gamma(c)} U(\gamma(c+t)) - U(\gamma(c))}{t} \right] \quad (241)$$

$$= \Gamma_{\gamma(c), x} [\nabla_{\gamma'(c)} U]. \quad (242)$$

Since for any smooth vector field W , $\|\Gamma_{y,x}(W)\|_x = \|W\|_y$, the proof is complete. \square

Proof of Lemma C.9:

We apply Lemma C.8 to the vector $V_t(p) \equiv \text{Log}_p y_1(t)$ for fixed t . We have that (equation (25) of Kendall and Le (2011))

$$V_t(p) = \text{grad} \left(-\frac{1}{2} d(p, y_1(t))^2 \right) = \text{grad}(f_t)(p). \quad (243)$$

Then, by the definition of the Riemannian Hessian H_t of f_t , for a smooth curve γ on \mathcal{M} at time c ,

$$\|\nabla_{\gamma'(c)} V_t(\gamma(c))\|_{\gamma(c)} = \|\nabla_{\gamma'(c)} \text{grad}(f_t)(p)\|_{\gamma(c)} = \|H_t(\gamma(c))(\gamma'(c))\|_{\gamma(c)}. \quad (244)$$

Using Lemma C.8, we choose $x = \mu(t)$ and $y = \hat{\mu}(t)$, and we denote $\gamma_t(s)$ as the minimizing geodesic between $\mu(t)$ and $\hat{\mu}(t)$:

$$\|\Gamma_{\hat{\mu}(t), \mu(t)} \text{Log}_{\mu(\hat{t})} y_1(t) - \text{Log}_{\mu(t)} y_1(t)\| \leq d(\hat{\mu}(t), \mu(t)) \sup_{c \in [0, d(\hat{\mu}(t), \mu(t))]} \|H_t(\gamma_t(c))(\gamma_t'(c))\|_{\gamma_t(c)} \quad (245)$$

$$\lesssim d(\hat{\mu}(t), \mu(t)) \quad (246)$$

for all $t \in \mathcal{T}$ with probability one, where in the last inequality we have used Assumption 3 in the Lemma statement along with the fact that, since $\gamma_t(s)$ is a minimizing geodesic, $\|\gamma_t'(s)\|_{\gamma_t(s)} = 1$ for all s . Then,

$$\mathbb{E} \left[\|\text{Log}_{\hat{\mu}} y_1 \delta_\Gamma \text{Log}_\mu y_1\|_\mu^2 \right] = \mathbb{E} \left[\int_{\mathcal{T}} \|\Gamma_{\hat{\mu}(t), \mu(t)} \text{Log}_{\mu(\hat{t})} y_1(t) - \text{Log}_{\mu(t)} y_1(t)\|_\mu^2 dt \right] \quad (247)$$

$$\lesssim \mathbb{E} \left[\int_{\mathcal{T}} d(\hat{\mu}(t), \mu(t))^2 dt \right]. \quad (248)$$

By Tonelli's theorem,

$$\mathbb{E} \left[\int_{\mathcal{T}} d(\hat{\mu}(t), \mu(t))^2 dt \right] = \int_{\mathcal{T}} \mathbb{E} [d(\hat{\mu}(t), \mu(t))^2] dt. \quad (249)$$

Next, we will apply Corollary 4 of Schötz (2019) to bound $\mathbb{E} [d(\hat{\mu}(t), \mu(t))^2]$. To do so, we need the following definitions related to the metric entropy of geodesic balls in \mathcal{M} . Let $B_\delta(a) \equiv \{x \in \mathcal{M} : d(x, a) \leq \delta\}$ be the ball of radius r centered at a on \mathcal{M} , and $N(B, r) \equiv \min(k \in \mathbb{N} | \exists q_1, \dots, q_k \in \mathcal{M} : B \subseteq \cup_{j=1}^k B_r(q_j))$ be the covering number of a set B using radius r . Then, the entropy assumption in the statement of Corollary 4 of Schötz (2019), that there exists $0 < \beta < 1$ such that, for all $\delta, r > 0$, $\log(N(B_\delta(\mu(t)), r))^{1/2} \lesssim (\frac{\delta}{r})^\beta$, is satisfied (Ahidar-Coutrix et al. (2020) Example 2.3).

Now additionally using Assumption 1 made in the statement of the Lemma, we can apply Corollary 4 of Schötz (2019) with $\varepsilon = 1$ to obtain

$$\mathbb{E} [d(\hat{\mu}(t), \mu(t))^2] \lesssim \mathbb{E} [d(y_1(t), y_2(t))^3]^{2/3} \frac{1}{N} \quad (250)$$

for all $t \in \mathcal{T}$.

Using Assumption 2 made in the statement of the Lemma, we have

$$\mathbb{E} [d(\hat{\mu}(t), \mu(t))^2] \lesssim \frac{1}{N}, \quad (251)$$

which we can combine with equation (248) to establish

$$\mathbb{E} \left[\|\text{Log}_{\hat{\mu}} y_1 \delta_\Gamma \text{Log}_\mu y_1\|_\mu^2 \right] \lesssim \int_{\mathcal{T}} \frac{1}{N} dt = \frac{\tau}{N}, \quad (252)$$

completing the proof. \square

Proof of Lemma C.12:

Note that $\mathbb{E} [\|Y\|_2^2] = \mathbb{E} [Y^\top Y] = \mathbb{E} [\text{tr}(Y Y^\top)] = \text{tr}(\mathbb{E}[Y Y^\top]) = \text{tr}(\Sigma_Y)$. Then, the first two statements

of the Lemma follow from observing that Σ_Y is diagonal with entries ω_j , the eigenvalues of the covariance operator of $\text{Log}_\mu y$, by the results of Lemma A.1.

To show $\mathbb{E}[\|Z\|_2^2] \lesssim \mathbb{E}[\|Y\|_2^2]$, we begin with the triangle inequality:

$$\mathbb{E}[\|Z\|_2^2] \lesssim \mathbb{E}[\|Z - Y\|_2^2] + \mathbb{E}[\|Y\|_2^2]. \quad (253)$$

Using Lemma C.11 and $\mathbb{E}[\|Y\|_2^2] = \sum_{j=1}^d \omega_j$, we observe for $\mathbb{E}[\|Z\|_2^2] \lesssim \mathbb{E}[\|Y\|_2^2]$ to hold, it is necessary that $\frac{d}{N} \max_{j=1, \dots, d} \left(\frac{1}{\omega_j - \omega_{j+1}} \right)^2 \left(\sum_{j=1}^d \omega_j \right)^{-1} \lesssim 1$. This follows from assumption (183) provided in the Lemma, because under the other assumptions provided in the Lemma, both $(\sum_{j=1}^d \omega_j) / \omega_d^2$ and $\sum_{j=1}^d \omega_j$ are greater than or equal to 1. From this, we also deduce that $\mathbb{E}[\|Z - Y\|_2^2] = o(1)$.

To show that $\|\Sigma_Z^{-1}\|_2 \lesssim \|\Sigma_Y^{-1}\|_2$, a more involved argument is required. We again begin with the triangle inequality:

$$\|\Sigma_Z^{-1}\|_2^2 \lesssim \|\Sigma_Z^{-1} - \Sigma_Y^{-1}\|_2^2 + \|\Sigma_Y^{-1}\|_2^2 \quad (254)$$

$$\leq \|\Sigma_Z^{-1}\|_2^2 \|\Sigma_Z - \Sigma_Y\|_2^2 \|\Sigma_Y^{-1}\|_2^2 + \|\Sigma_Y^{-1}\|_2^2. \quad (255)$$

This implies

$$\|\Sigma_Z^{-1}\|_2^2 - \|\Sigma_Y^{-1}\|_2^2 \|\Sigma_Z - \Sigma_Y\|_2^2 \|\Sigma_Y^{-1}\|_2^2 \lesssim \|\Sigma_Y^{-1}\|_2^2, \quad (256)$$

so that

$$\|\Sigma_Z^{-1}\|_2^2 \left(1 - \|\Sigma_Z - \Sigma_Y\|_2^2 \|\Sigma_Y^{-1}\|_2^2 \right) \leq \|\Sigma_Y^{-1}\|_2^2, \quad (257)$$

and therefore, if $\|\Sigma_Z - \Sigma_Y\|_2^2 \|\Sigma_Y^{-1}\|_2^2 = o(1)$ held, then we would have $\|\Sigma_Z^{-1}\|_2 \lesssim \|\Sigma_Y^{-1}\|_2$ and the proof would be complete. Thus, to show $\|\Sigma_Z^{-1}\|_2 \lesssim \|\Sigma_Y^{-1}\|_2$, it suffices to show that $\frac{d}{N} \frac{\sum_{j=1}^d \omega_j}{\omega_d^2} \max_{j=1, \dots, d} \left(\frac{1}{\omega_j - \omega_{j+1}} \right)^2 = o(1)$ implies $\|\Sigma_Z - \Sigma_Y\|_2^2 \|\Sigma_Y^{-1}\|_2^2 = o(1)$. From the second item of Lemma C.3 and from $\mathbb{E}[\|Z\|_2^2] \lesssim \mathbb{E}[\|Y\|_2^2]$ which has already been shown, we have

$$\|\Sigma_Z - \Sigma_Y\|_2^2 \leq \max(\mathbb{E}[\|Z\|_2^2], \mathbb{E}[\|Y\|_2^2]) \mathbb{E}[\|Z - Y\|_2^2] \quad (258)$$

$$\lesssim \mathbb{E}[\|Z - Y\|_2^2] \sum_{j=1}^d \omega_j. \quad (259)$$

Thus, for $\|\Sigma_Z - \Sigma_Y\|_2^2 \|\Sigma_Y^{-1}\|_2^2 = o(1)$ to hold, it suffices to show that $\frac{1}{\omega_d^2} \mathbb{E}[\|Z - Y\|_2^2] \sum_{j=1}^d \omega_j = o(1)$. From Lemma C.11, we have made exactly the assumption so that $\frac{1}{\omega_d^2} \mathbb{E}[\|Z - Y\|_2^2] \sum_{j=1}^d \omega_j = o(1)$ holds, completing the proof that $\|\Sigma_Z^{-1}\|_2 \lesssim \|\Sigma_Y^{-1}\|_2$.

To show the final statement of the Lemma, that $\|\Sigma_Z\|_2 \lesssim \|\Sigma_Y\|_2$, we can make an argument similar to the one used to show $\|\Sigma_Z^{-1}\|_2 \lesssim \|\Sigma_Y^{-1}\|_2$. From this, it suffices to show that $\frac{d}{N} \frac{\sum_{j=1}^d \omega_j}{\omega_1^2} \max_{j=1, \dots, d} \left(\frac{1}{\omega_j - \omega_{j+1}} \right)^2 = o(1)$. Since $\omega_d \leq \omega_1$, $\frac{d}{N} \frac{\sum_{j=1}^d \omega_j}{\omega_1^2} \max_{j=1, \dots, d} \left(\frac{1}{\omega_j - \omega_{j+1}} \right)^2 = o(1)$ holds under our assumptions and the proof is complete. \square

Proof of Lemma C.14:

To bound $|\gamma_k - \tilde{\gamma}_k|$, we can use Weyl's inequality (Bhatia (2013) Corollary III.2.6.), because γ_k and $\tilde{\gamma}_k$ are the k eigenvalues of the matrices $|C|$ and $|A|$ respectively (Uurtio et al. 2018):

$$|\gamma_k - \tilde{\gamma}_k| \leq \||A| - |C|\|_2 \leq \||A| - |C|\|_F \lesssim \|A - C\|_F, \quad (260)$$

where in the second inequality we have used identity 19 from Section D. Combining Lemmas C.11, C.12, and C.13, we observe that the stated assumption in the Lemma implies $|\gamma_k - \tilde{\gamma}_k| \lesssim \min\{\gamma_k, \tilde{\gamma}_k\}$, establishing the stated conclusion using the assumption that γ_k and $\tilde{\gamma}_k$ are bounded from below. That γ_k^2 and $\tilde{\gamma}_k^2$ are asymptotically equivalent follows from the same argument, in addition to the function $f(x) = x^2$ being Lipschitz continuous on the interval $[0, 1]$.

D Additional identities and inequalities

In the proofs, we use several identities and inequalities involving matrices. For definitions of the various matrix operations used below we refer to Section B.1. In the following, A and B denote matrices for which the specified matrix multiplications are valid.

1. For $x \in \mathbb{R}^p$, $\|x\|_\infty \leq \|x\|_2 \leq \|x\|_1$.
2. $\|x\|_2 \leq \sqrt{p} \|x\|_\infty$, $\|x\|_1 \leq \sqrt{p} \|x\|_2$, and $\|x\|_1 \leq p \|x\|_\infty$.
3. $\|A\|_{\max} \leq \|A\|_{2,\infty} \leq \|A\|_2 \leq \|A\|_F \leq \|A\|_{\ell_1,\ell_2}$
4. For induced norms, $\|AB\|_{\beta,\alpha} \leq \|A\|_{\gamma,\alpha} \|B\|_{\beta,\gamma}$ (Trefethen and Bau (2022) equation (3.14)). In particular, $\|AB\|_2 = \|AB\|_{2,2} \leq \|A\|_{\infty,2} \|B\|_{2,\infty}$.
5. For any matrix A , $\|A\|_2 = \|A^\top A\|_2^{1/2}$.
6. In addition to its definition as the largest singular value, $\|A\|_2$ is the norm induced by $\|\cdot\|_2$ and $\|\cdot\|_2$.
7. In addition to its definition as the norm induced by the $\|\cdot\|_2$ and $\|\cdot\|_\infty$ norms, $\|A\|_{2,\infty} = \max_i (\|A_i\|_2)$, where A_i is the i th row of A . (Cape et al. (2019) Proposition 6.1.)
8. $\|AB\|_{2,\infty} \leq \|A\|_{2,\infty} \|B\|_2$. (Cape et al. (2019) Proposition 6.5.)
9. For $B \in \mathbb{R}^{p \times d}$, $\|B\|_{2,\infty} \leq \sqrt{d} \|B\|_{\max}$.
10. For the induced norm $\|\cdot\|_{\infty,2}$, $\|B^\top\|_{\infty,2} \leq \|B\|_{\ell_1,\ell_2}$.
11. If $A \in \mathbb{R}^{d \times p}$, and $B \in \mathbb{R}^{p \times d}$, then $\|AB\|_2 \leq \|A^\top\|_{2,\infty} \|B\|_{\ell_1,\ell_2}$. Additionally, $\|AB\|_2 \leq \|A^\top\|_{\ell_1,\ell_2} \|B\|_{2,\infty}$.
12. $\|AB\|_{\ell_1,\ell_2} \leq \|A\|_{\ell_1,\ell_2} \|B\|_2$.
13. If $A, B \in \mathbb{R}^{d \times d}$ are positive definite, then $\|A^{1/2} - B^{1/2}\|_2 \leq \frac{1}{2} \max(\|A^{-1}\|_2, \|B^{-1}\|_2)^{1/2} \|A - B\|_2$.
14. If $A, B \in \mathbb{R}^{d \times d}$ are positive definite, then $\|A^{-1/2} - B^{-1/2}\|_2 \leq \frac{1}{2} \max(\|A^{-1}\|_2, \|B^{-1}\|_2)^{3/2} \|A - B\|_2$.
15. For any positive definite matrices A, B of the same size, the product AB is diagonalizable with positive eigenvalues. Additionally, AB has the same eigenvalues as $(AB^2A)^{1/2}$. In particular, $\|AB - I\|_2 = \|(AB^2A)^{1/2} - I\|_2$. Note that AB may not be symmetric, and therefore is not necessarily positive definite.
16. For positive real numbers x and y , $|x^{-1} - y^{-1}| \leq \min(x, y)^{-3} |x^2 - y^2|$.
17. For $A, B \in \mathbb{R}^{d \times d}$, $a, b \in \mathbb{R}^d$, $\|Aa - Bb\|_2^2 \lesssim \|A - B\|_2^2 \|a\|_2^2 + \|B\|_2^2 \|a - b\|_2^2$
18. For any norm $\|\cdot\|$, $\|a + b\|^2 \leq 3(\|a\|^2 + \|b\|^2)$.
19. For matrices A, B of the same dimensions, $\||A| - |B|\|_F \leq \sqrt{2} \|A - B\|_F$, where $|A| \equiv (A^\top A)^{1/2}$. (Bhatia (2013) equation VII.39)

Proof. We prove the non-standard or non-straightforward results.

Proof of 10:

To show $\|B^\top\|_{\infty,2} \leq \|B\|_{\ell_1,\ell_2}$, we use the definition:

$$\|B^\top\|_{\infty,2} \equiv \sup_{\|x\|_\infty=1} \|B^\top x\|_2. \quad (261)$$

Without loss of generality, let B be an element of $\mathbb{R}^{p \times d}$. $\|B^\top x\|_2 = \left\| \sum_{i=1}^p b_i x_i \right\|_2$ where b_i is the i th row of B , and x_i is the i th entry of $x \in \mathbb{R}^p$. For $x \in \mathbb{R}^p$ with $\|x\|_\infty = 1$, we have

$$\left\| \sum_{i=1}^p b_i x_i \right\|_2 \leq \sum_{i=1}^p |x_i| \|b_i\|_2 \leq \sum_{i=1}^p \|b_i\|_2 = \|B\|_{\ell_1,\ell_2}, \quad (262)$$

using the triangle inequality and because $\|x\|_\infty = 1$. This completes the proof.

Proof of 11:

To show $\|AB\|_2 \leq \|A^\top\|_{2,\infty} \|B\|_{\ell_1,\ell_2}$, we begin by using item 4 to obtain

$$\|AB\|_2 = \|B^\top A^\top\|_2 \leq \|B^\top\|_{\infty,2} \|A^\top\|_{2,\infty}. \quad (263)$$

Using item 10 we have

$$\|B^\top\|_{\infty,2} \|A^\top\|_{2,\infty} \leq \|B\|_{\ell_1,\ell_2} \|A^\top\|_{2,\infty}, \quad (264)$$

completing the proof of the first statement. To show the second statement, we proceed similarly but apply item 4 to $\|AB\|_2$ instead of $\|B^\top A^\top\|_2$.

Proof of 12:

To show $\|AB\|_{\ell_1,\ell_2} \leq \|A\|_{\ell_1,\ell_2} \|B\|_2$, we begin with the definition. Without loss of generality, $A \in \mathbb{R}^{p \times q}$ and $B \in \mathbb{R}^{p \times r}$. We have

$$\|AB\|_{\ell_1,\ell_2} = \sum_{i=1}^p \|(AB)_i\|_2 = \sum_{i=1}^p \|A_i^\top B\|_2 \leq \sum_{i=1}^p \|a_i\|_2 \|B\|_2 = \|A\|_{\ell_1,\ell_2} \|B\|_2, \quad (265)$$

where $(C)_i$ denotes the i th row of a matrix C and we have used item 6.

Proof of 13:

The statement follows from equation X.46 of Bhatia (2013) by choosing $r = 1/2$ and since the 2-norm is unitarily invariant.

Proof of 14:

We begin with the equality $A^{-1} - B^{-1} = A^{-1} (B - A) B^{-1}$, which holds for any square invertible matrices A and B of the same size. This implies

$$\|A^{-1} - B^{-1}\|_2 \leq \|A^{-1}\|_2 \|A - B\|_2 \|B^{-1}\|_2. \quad (266)$$

We apply this to the matrices $A^{1/2}$ and $B^{1/2}$ to obtain

$$\|A^{-1/2} - B^{-1/2}\|_2 \leq \|A^{-1/2}\|_2 \|A^{1/2} - B^{1/2}\|_2 \|B^{-1/2}\|_2. \quad (267)$$

Using item 13 on $\|A^{1/2} - B^{1/2}\|_2$ in the last inequality, we deduce that

$$\|A^{-1/2} - B^{-1/2}\|_2 \leq \frac{1}{2} \max(\|A^{-1}\|_2, \|B^{-1}\|_2)^{1/2} \|A - B\|_2 \|A^{-1/2}\|_2 \|B^{-1/2}\|_2 \quad (268)$$

$$= \frac{1}{2} \|A^{-1/2}\|_2 \|B^{-1/2}\|_2 \max(\|A^{-1/2}\|_2, \|B^{-1/2}\|_2) \|A - B\|_2 \quad (269)$$

One of $\|A^{-1/2}\|_2$ and $\|B^{-1/2}\|_2$ is larger, and in either case, the statement to be proven holds.

Proof of 15:

The first claim, that AB is diagonalizable with positive eigenvalues, follows from Proposition 6.1

of Serre (2010). To show that AB has the same eigenvalues as $(AB^2A)^{1/2}$, let $C = AB$, so that $(AB^2A)^{1/2} = (CC^\top)^{1/2}$. Letting $U\Sigma V^\top = C$ be a singular value decomposition of C , we have

$$(CC^\top)^{1/2} = (U\Sigma V^\top V\Sigma U^\top)^{1/2} = (U\Sigma^2 U^\top)^{1/2} = U\Sigma U^\top, \quad (270)$$

where in the last line we have used that C has positive singular values from the first claim. The last claim follows from the previous claim, since for diagonalizable matrix C where $C = WDW^{-1}$ with D diagonal, we have

$$C - I = WDW^{-1} - WIW^{-1} = W(D - I)W^{-1}, \quad (271)$$

so that $C - I$ has eigenvalues equal to those of C minus 1.

Proof of 16:

Letting $f(x) = x^{-1/2}$, the mean value theorem implies

$$|f(x) - f(y)| \leq \min(x, y)^{-3/2} |x - y|. \quad (272)$$

Plugging in x^2 for x and y^2 for y , we obtain the result.

Proof of 17:

The statement follows from adding and subtracting Ba and the inequality $\|a + b\|_2^2 \leq 2\|a\|_2^2 + 2\|b\|_2^2$.

Proof of 18:

From the triangle inequality, we have

$$\|a + b\|_2^2 \leq (\|a\|_2 + \|b\|_2)^2 = \|a\|_2^2 + 2\|a\|_2\|b\|_2 + \|b\|_2^2 \quad (273)$$

One of $\|a\|_2$ or $\|b\|_2$ is larger than the other, and in either case, we deduce that $\|a\|_2^2 + 2\|a\|_2\|b\|_2 + \|b\|_2^2 \leq 3(\|a\|_2^2 + \|b\|_2^2)$, completing the proof. \square

E Intrinsic RFPCA algorithm

For completeness, in this section, we outline the main steps of the Intrinsic RFPCA algorithm. Recall that M is the dimension of \mathcal{M} , d is the number of principal components to compute, N is the number of observations, and L is the number of time-steps.

1. Estimate $\mu : \mathcal{T} \rightarrow \mathcal{M}$, the functional Fréchet mean of y , by computing the Fréchet mean of $\{y_i(t_l)\}_i$ separately for each point in $\{t_l\}_{l=1, \dots, L}$ where the functional data are observed.
2. Compute the linear representations $\text{Log}_{\hat{\mu}} y_i \in L^2(T\hat{\mu})$, for $i = 1, \dots, N$. In practice, this can be done by computing $\text{Log}_{\hat{\mu}(t_l)} y_i(t_l) \in T_{\hat{\mu}(t_l)}\mathcal{M}$ for every t_l by using the Log map on \mathcal{M} .
3. Let $E(x)$ be an orthonormal frame for the tangent space centered at $x \in \mathcal{M}$ (see Section E.1 for an example). An orthonormal frame is a collection of tangent bases $E(x) = \{E_1(x), \dots, E_M(x)\}$ for $T_x\mathcal{M}$, with $x \in \mathcal{M}$, which varies smoothly with x , i.e. for each $k = 1, \dots, M$, $E_k(x)$ is a smooth map from \mathcal{M} to $T\mathcal{M}$. For any fixed $x \in \mathcal{M}$, the functions $\{E_k(x)\}$ are orthonormal with respect to the inner product on $T_x\mathcal{M}$, that is, $\langle \cdot, \cdot \rangle_x$.

Compute the functional ‘coefficients’ $\hat{Z}_i : \mathcal{T} \rightarrow \mathbb{R}^M$ of the expansion of $\text{Log}_{\hat{\mu}} y_i : \mathcal{T} \rightarrow T\mathcal{M}$ relative to E , for each $i = 1, \dots, N$. In practice, $\hat{Z}_i(t_l) \in \mathbb{R}^M$ is computed separately for every $l = 1, \dots, L$. The k th entry of $\hat{Z}_i(t_l) \in \mathbb{R}^M$ is computed as $\langle \text{Log}_{\hat{\mu}(t_l)} y_i(t_l), E_k(\hat{\mu}(t_l)) \rangle_{\hat{\mu}(t_l)}$ for $k = 1, \dots, M$. The resulting $\{\hat{Z}_i\}$ are estimates of the realizations of a real vector-valued random process $Z : \mathcal{T} \rightarrow \mathbb{R}^M$, with k th component $Z_k : \mathcal{T} \rightarrow \mathbb{R}$ given by $Z_k(t) = \langle \text{Log}_{\mu(t)} y(t), E_k(\mu(t)) \rangle_{\mu(t)}$.

The process Z is a real vector-valued process with the same principal scores as the process $\text{Log}_{\mu} y$, and the j th principal component of Z , $\pi_j : \mathcal{T} \rightarrow \mathbb{R}^M$, is related to the j th principal component of $\text{Log}_{\mu} y$, $\phi_j \in L^2(T\mu)$, via $\phi_j = \sum_{k=1}^M \pi_{jk} E_k$. Here, π_{jk} denotes the k th entry of π_j for $k = 1, \dots, M$.

4. Apply Multivariate Functional Principal Component Analysis (MFPCA) (Happ and Greven 2018) to the functions $\{\hat{Z}_i\}$ to estimate d principal component functions $\hat{\pi}_j$ of the functional coefficients.

Estimate the principal component functions $\{\phi_j\}$ of the Log representations of the functional data as $\hat{\phi}_j = \sum_{k=1}^M \hat{\pi}_{jk} E_k(\hat{\mu})$, for $j = 1, \dots, d$. Then, estimate the associated scores as $\hat{Y}_{ij} = \frac{1}{L} \sum_{l=1}^L Z_i(t_l)^\top \hat{\pi}_j(t_l)$, for $j = 1, \dots, d$ and $i = 1, \dots, N$. The score estimates do not depend on the choice of the orthonormal frame E (Proposition 5, item 2 of Lin and Yao (2019)). Compute variance estimates $\hat{\lambda}_j = \text{Var}(\hat{Y}_{ij})$.

5. Return the estimated scores $\{\hat{Y}_{ij}\}$, variances $\{\hat{\lambda}_j\}$ and principal component functions $\{\hat{\phi}_j\}$.

E.1 An orthonormal frame for a manifold of SPD matrices

Here we provide an explicit construction of an orthonormal frame in the setting where \mathcal{M} is the manifold of $\mathbb{R}^{m \times m}$ symmetric positive matrices equipped with the affine invariant metric, as in our application setting. Recall that the maps Log and Exp maps are defined as $\text{Log}_F(G) = F^{1/2} \log(F^{-1/2} G F^{-1/2}) F^{1/2}$ and $\text{Exp}_F(W) = F^{1/2} \exp(F^{-1/2} W F^{-1/2}) F^{1/2}$. Moreover, the inner product at $F \in \mathcal{M}$ between $W, Z \in T_F \mathcal{M}$ is defined as $\langle W, Z \rangle_{\mathcal{M}} = \text{tr}(F^{-1} W F^{-1} Z)$. In this setting, we can use the result from Section 3.3.3.3. of Pennec et al. (2019), which provides an explicit construction of an orthonormal frame $E(F)$ for the tangent bundle, evaluated at an arbitrary $F \in \mathcal{M}$. This can be defined as:

$$E_{ij}(F) = \begin{cases} (F^{1/2} e_i) (F^{1/2} e_i)^\top & (1 \leq i = j \leq m) \\ \frac{1}{\sqrt{2}} \left((F^{1/2} e_i) (F^{1/2} e_j)^\top + (F^{1/2} e_j) (F^{1/2} e_i)^\top \right) & (1 \leq i < j \leq m), \end{cases} \quad (274)$$

where e_i denotes the i th standard unit vector in \mathbb{R}^m , and $F \in \mathbb{R}^{m \times m}$. For a fixed $F \in \mathcal{M}$, there are $M = m(m+1)/2$ unique $E_{ij}(F)$, where M is the dimension of \mathcal{M} . We also note that the iterative algorithm proposed in Cheng et al. (2016) can be used to estimate the Fréchet mean μ is detailed in equations (13) and (14) of Cheng et al. (2016).

References

- Ahidar-Coutrix, Adil, Le Gouic, Thibaut, and Paris, Quentin (2020). “Convergence rates for empirical barycenters in metric spaces: curvature, convexity and extendable geodesics”. In: *Probability theory and related fields* 177.1, pp. 323–368.
- Bhatia, Rajendra (2013). *Matrix analysis*. Vol. 169. Springer Science & Business Media.
- Bhattacharjee, Satarupa and Müller, Hans-Georg (2021). “Single Index Fréchet Regression”. In.
- Bhattacharya, Rabi and Patrangenaru, Vic (2003). “Large Sample Theory of Intrinsic and Extrinsic Sample Means on Manifolds. I”. In: *Annals of Statistics* 31.1, pp. 1–29. ISSN: 00905364. DOI: 10.1214/aos/1046294456.
- Cai, T. Tony, Ren, Zhao, and Zhou, Harrison H. (2016). “Estimating Structured High-Dimensional Covariance and Precision Matrices: Optimal Rates and Adaptive Estimation”. In: *Electronic Journal of Statistics* 10.1, pp. 1–59. ISSN: 1935-7524, 1935-7524. DOI: 10.1214/15-EJS1081.
- Cape, Joshua, Tang, Minh, and Priebe, Carey E (2019). “The two-to-infinity norm and singular subspace geometry with applications to high-dimensional statistics”. In.
- Cardot, Hervé, Ferraty, Frédéric, and Sarda, Pascal (1999). “Functional Linear Model”. In: *Statistics and Probability Letters* 45.1, pp. 11–22. ISSN: 01677152. DOI: 10.1016/S0167-7152(99)00036-X.
- Carmichael, Iain (2020). “Learning Sparsity and Block Diagonal Structure in Multi-View Mixture Models”. In: pp. 1–39.
- Chen, Mengjie, Gao, Chao, Ren, Zhao, and Zhou, Harrison H. (2013). “Sparse CCA via Precision Adjusted Iterative Thresholding”. In.

- Cheng, Guang, Ho, Jeffrey, Salehian, Hesamoddin, and Vemuri, Baba C (2016). “Recursive computation of the Fréchet mean on non-positively curved Riemannian manifolds with applications”. In: *Riemannian Computing in Computer Vision*. Springer, pp. 21–43.
- Cho, Min Ho, Kurtek, Sebastian, and Bharath, Karthik (2021). “Tangent Functional Canonical Correlation Analysis for Densities and Shapes, with Applications to Multimodal Imaging Data”. In: ISSN: 0047259X. DOI: 10.1016/j.jmva.2021.104870.
- Cupidon, J., Eubank, R., Gilliam, D., and Ruymgaart, F. (2008). “Some properties of canonical correlations and variates in infinite dimensions”. en. In: *Journal of Multivariate Analysis* 99.6, pp. 1083–1104. ISSN: 0047259X. DOI: 10.1016/j.jmva.2007.07.007.
- Dai, Xiongtao and Müller, Hans Georg (2018). “Principal Component Analysis for Functional Data on Riemannian Manifolds and Spheres”. In: *Annals of Statistics* 46.6B, pp. 3334–3361. ISSN: 00905364. DOI: 10.1214/17-AOS1660.
- Dai, Xiongtao, Müller, Hans Georg, and Yao, Fang (2017). “Optimal Bayes Classifiers for Functional Data and Density Ratios”. In: *Biometrika* 104.3, pp. 545–560. ISSN: 14643510. DOI: 10.1093/biomet/asx024.
- Dubey, Paromita and Müller, Hans Georg (2020). “Functional Models for Time-Varying Random Objects”. In: *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 82.2, pp. 275–327. ISSN: 14679868. DOI: 10.1111/rssb.12337.
- (2021). “Modeling Time-Varying Random Objects and Dynamic Networks”. In: *Journal of the American Statistical Association*. ISSN: 1537274X. DOI: 10.1080/01621459.2021.1917416.
- Feng, Qing, Jiang, Meilei, Hannig, Jan, and Marron, J. S. (2018). “Angle-Based Joint and Individual Variation Explained”. In: *Journal of Multivariate Analysis* 166, pp. 241–265. ISSN: 10957243. DOI: 10.1016/j.jmva.2018.03.008.
- Fletcher, P Thomas and Joshi, Sarang (2007). “Riemannian geometry for the statistical analysis of diffusion tensor data”. In: *Signal Processing* 87.2, pp. 250–262.
- Friedman, Jerome, Tibshirani, Robert, and Hastie, Trevor (2010). “Regularization Paths for Generalized Linear Models via Coordinate Descent”. In: *Journal of Statistical Software* 33.1, pp. 1–22. DOI: 10.18637/jss.v033.i01.
- Gao, Chao, Ma, Zongming, and Zhou, Harrison H. (2017). “Sparse CCA: Adaptive Estimation and Computational Barriers”. In: *Annals of Statistics* 45.5, pp. 2074–2101. ISSN: 00905364. DOI: 10.1214/16-AOS1519.
- Gaynanova, Irina (2020). “Prediction and Estimation Consistency of Sparse Multi-Class Penalized Optimal Scoring”. In: *Bernoulli* 26.1, pp. 286–322. ISSN: 13507265. DOI: 10.3150/19-BEJ1126.
- Gaynanova, Irina, Booth, James G., and Wells, Martin T. (2016). “Simultaneous Sparse Estimation of Canonical Vectors in the $p \geq N$ Setting”. In: *Journal of the American Statistical Association* 111.514, pp. 696–706. ISSN: 1537274X. DOI: 10.1080/01621459.2015.1034318.
- Glasser, Matthew F. et al. (2016). “A Multi-Modal Parcellation of Human Cerebral Cortex”. In: *Nature* 536.7615, pp. 171–178. ISSN: 14764687. DOI: 10.1038/nature18933.
- Glasser, Matthew F. et al. (2013). “The Minimal Preprocessing Pipelines for the Human Connectome Project”. In: *NeuroImage* 80, pp. 105–124. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.04.127.
- Happ, Clara and Greven, Sonja (2018). “Multivariate Functional Principal Component Analysis for Data Observed on Different (Dimensional) Domains”. In: *Journal of the American Statistical Association* 113.522, pp. 649–659. ISSN: 1537274X. DOI: 10.1080/01621459.2016.1273115.
- Hastie, Trevor, Tibshirani, Robert, and Wainwright, Martin (2015). *Statistical learning with sparsity: the lasso and generalizations*. CRC press.
- He, Guozhong, Müller, Hans-Georg, Wang, Jane-Ling, and Yang, Wenjing (2010). “Functional Linear Regression via Canonical Analysis”. In: *Bernoulli* 16.3. ISSN: 1350-7265. DOI: 10.3150/09-BEJ228.
- Hörmander, Lars (2015). *The analysis of linear partial differential operators I: Distribution theory and Fourier analysis*. Springer.

- Hotelling, H. (1936). “Relations Between Two Sets of Variates”. In: *Biometrika* 28.3-4, pp. 321–377. ISSN: 0006-3444. DOI: 10.1093/biomet/28.3-4.321.
- Hsing, Tailen and Eubank, Randall (2013). *Theoretical Foundations of Functional Data Analysis, with an Introduction to Linear Operators*. Chichester, UK: John Wiley & Sons, Ltd. ISBN: 978-1-118-76254-7. DOI: 10.1002/9781118762547.
- (2015). *Theoretical foundations of functional data analysis, with an introduction to linear operators*. Vol. 997. John Wiley & Sons.
- Huang, Qing and Renaut, Rosemary (2015). “Functional Partial Canonical Correlation”. In: *Bernoulli* 21.2. ISSN: 1350-7265. DOI: 10.3150/14-BEJ597.
- Hutchison, R. Matthew et al. (2013). “Dynamic Functional Connectivity: Promise, Issues, and Interpretations”. In: *NeuroImage* 80, pp. 360–378. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.05.079.
- Jirak, Moritz and Wahl, Martin (2020). “Perturbation bounds for eigenspaces under a relative gap condition”. In: *Proceedings of the American Mathematical Society* 148.2, pp. 479–494.
- Kendall, Wilfrid S and Le, Huiling (2011). “Limit theorems for empirical Fréchet means of independent and non-identically distributed manifold-valued random variables”. In.
- Kereta, Željko and Klock, Timo (2021). “Estimating covariance and precision matrices along subspaces”. In.
- Kessler, Daniel and Levina, Elizaveta (2023). *Computational Inference for Directions in Canonical Correlation Analysis*.
- Kim, Hyunwoo J., Adluru, Nagesh, Bendlin, Barbara B., Johnson, Sterling C., Vemuri, Baba C., and Singh, Vikas (2014). “Canonical Correlation Analysis on Riemannian Manifolds and Its Applications”. In: *Computer Vision – ECCV 2014*. Ed. by David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars. Vol. 8690. Cham: Springer International Publishing, pp. 251–267. ISBN: 978-3-319-10604-5 978-3-319-10605-2. DOI: 10.1007/978-3-319-10605-2_17.
- Kokoszka, Piotr and Reimherr, Matthew (2017). *Introduction to Functional Data Analysis*. New York, New York, USA: Chapman and Hall/CRC. ISBN: 978-1-315-11741-6. DOI: 10.1201/9781315117416.
- Krzyśko, Mirosław and Waszak, Łukasz (2013). “Canonical correlation analysis for functional data”. In: *Biometrical Letters* 50.2, pp. 95–105.
- Lee, John M (2018). *Introduction to Riemannian manifolds*. Vol. 2. Springer.
- Lee, John M. (2012). *Smooth Manifolds*. New York, NY: Springer New York. ISBN: 978-1-4419-9982-5. DOI: 10.1007/978-1-4419-9982-5_1.
- Liégeois, Raphaël, Li, Jingwei, Kong, Ru, Orban, Csaba, Van De Ville, Dimitri, Ge, Tian, Sabuncu, Mert R., and Yeo, B. T. Thomas (2019). “Resting Brain Dynamics at Different Timescales Capture Distinct Aspects of Human Behavior”. In: *Nature Communications* 10.1. ISSN: 20411723. DOI: 10.1038/s41467-019-10317-7.
- Lin, Dongdong, Zhang, Jigang, Li, Jingyao, Calhoun, Vince D, Deng, Hong-Wen, and Wang, Yu-Ping (2013). “Group Sparse Canonical Correlation Analysis for Genomic Data Integration”. In: *BMC Bioinformatics* 14.1, p. 245. ISSN: 1471-2105. DOI: 10.1186/1471-2105-14-245.
- Lin, Zhenhua and Yao, Fang (2019). “Intrinsic Riemannian Functional Data Analysis”. In: *The Annals of Statistics* 47.6. ISSN: 0090-5364. DOI: 10.1214/18-AOS1787.
- Liu, Zhangdaihong, Whitaker, Kirstie J., Smith, Stephen M., and Nichols, Thomas E. (2022). “Improved Interpretability of Brain-Behavior CCA With Domain-Driven Dimension Reduction”. In: *Frontiers in Neuroscience* 16, p. 851827. ISSN: 1662-453X. DOI: 10.3389/fnins.2022.851827.
- Liu, Zhiyuan, Schulz, Jörn, Taheri, Mohsen, Styner, Martin, Damon, James, Pizer, Stephen, and Marron, J. S. (2021). *Non-Euclidean Analysis of Joint Variations in Multi-Object Shapes*.
- Lock, Eric F., Hoadley, Katherine A., Marron, J. S., and Nobel, Andrew B. (2013). “Joint and Individual Variation Explained (JIVE) for Integrated Analysis of Multiple Data Types”. In: *Annals of Applied Statistics* 7.1, pp. 523–542. ISSN: 19326157. DOI: 10.1214/12-A0AS597.
- Marron, J.S. and Dryden, Ian L. (2021). *Object Oriented Data Analysis*. 1st ed. Boca Raton: Chapman and Hall/CRC. ISBN: 978-1-351-18967-5. DOI: 10.1201/9781351189675.

- McKeague, Ian W and Zhang, Xin (2022). “Significance Testing for Canonical Correlation Analysis in High Dimensions”. In: *Biometrika* 109.4, pp. 1067–1083. ISSN: 0006-3444, 1464-3510. DOI: 10.1093/biomet/asab059.
- Murden, Raphiel J., Zhang, Zhengwu, Guo, Ying, and Risk, Benjamin B. (2022). “Interpretive JIVE: Connections with CCA and an Application to Brain Connectivity”. In: *Frontiers in Neuroscience* 16. ISSN: 1662-453X.
- Pennec, Xavier (2017). “Hessian of the Riemannian squared distance”. In: *Preprint*. <https://www.sop.inria.fr/members/Xavier.Pennec/AOS-DiffRiemannianLog.pdf>.
- Pennec, Xavier, Sommer, Stefan, and Fletcher, Tom (2019). *Riemannian geometric statistics in medical image analysis*. Academic Press.
- Petersen, Alexander and Müller, Hans Georg (2019). “Fréchet Regression for Random Objects with Euclidean Predictors”. In: *Annals of Statistics* 47.2, pp. 691–719. ISSN: 00905364. DOI: 10.1214/17-AOS1624.
- Petersen, Alexander, Zhang, Chao, and Kokoszka, Piotr (2022). “Modeling Probability Density Functions as Data Objects”. In: *Econometrics and Statistics* 21, pp. 159–178. ISSN: 24523062. DOI: 10.1016/j.ecosta.2021.04.004.
- Power, Jonathon D. et al. (2011). “Functional Network Organization of the Human Brain”. In: *Neuron* 72.4, pp. 665–678. ISSN: 08966273. DOI: 10.1016/j.neuron.2011.09.006.
- Ramsay, James O. and Silverman, Bernard W. (2015). *Functional Data Analysis*. New York: Springer-Verlag. ISBN: 978-0-08-097087-5. DOI: 10.1016/B978-0-08-097086-8.42046-5.
- Schötz, Christof (2019). “Convergence rates for the generalized Fréchet mean via the quadruple inequality”. In.
- Serre, Denis (2010). *Matrices: Theory and Applications*. Springer.
- Shao, Lingxuan, Lin, Zhenhua, and Yao, Fang (2022). “Intrinsic Riemannian Functional Data Analysis for Sparse Longitudinal Observations”. In: *The Annals of Statistics* 50.3. ISSN: 0090-5364. DOI: 10.1214/22-AOS2172.
- Shin, Hyejin and Lee, Seokho (2015). “Canonical Correlation Analysis for Irregularly and Sparsely Observed Functional Data”. In: *Journal of Multivariate Analysis* 134, pp. 1–18. ISSN: 0047259X. DOI: 10.1016/j.jmva.2014.10.001.
- Shu, Hai, Wang, Xiao, and Zhu, Hongtu (2020). “D-CCA: A Decomposition-Based Canonical Correlation Analysis for High-Dimensional Datasets”. In: *Journal of the American Statistical Association* 115.529, pp. 292–306. ISSN: 1537274X. DOI: 10.1080/01621459.2018.1543599.
- Smith, Stephen M., Nichols, Thomas E., Vidaurre, Diego, Winkler, Anderson M., Behrens, Timothy E.J., Glasser, Matthew F., Ugurbil, Kamil, Barch, Deanna M., Van Essen, David C., and Miller, Karla L. (2015). “A Positive-Negative Mode of Population Covariation Links Brain Connectivity, Demographics and Behavior”. In: *Nature Neuroscience* 18.11, pp. 1565–1567. ISSN: 15461726. DOI: 10.1038/nn.4125.
- Smith, Stephen M. et al. (2013). “Resting-State fMRI in the Human Connectome Project”. In: *NeuroImage* 80, pp. 144–168. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.05.039.
- Stöcker, Almond and Greven, Sonja (2021). “Functional Additive Regression on Shape and Form Manifolds of Planar Curves”. In.
- Trefethen, Lloyd N and Bau, David (2022). *Numerical linear algebra*. Vol. 181. Siam.
- Uurtio, Viivi, Monteiro, João M., Kandola, Jaz, Shawe-Taylor, John, Fernandez-Reyes, Delmiro, and Rousu, Juho (2018). “A Tutorial on Canonical Correlation Methods”. In: *ACM Computing Surveys* 50.6, pp. 1–33. ISSN: 0360-0300, 1557-7341. DOI: 10.1145/3136624.
- Van Essen, D. C. et al. (2012). “The Human Connectome Project: A Data Acquisition Perspective”. In: *NeuroImage* 62.4, pp. 2222–2231. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2012.02.018.
- Vershynin, Roman (2018). *High-dimensional probability: An introduction with applications in data science*. Vol. 47. Cambridge university press.

- Wang, Wenjia and Zhou, Yi Hui (2021). “Eigenvector-Based Sparse Canonical Correlation Analysis: Fast Computation for Estimation of Multiple Canonical Vectors”. In: *Journal of Multivariate Analysis* 185. ISSN: 10957243. DOI: 10.1016/j.jmva.2021.104781.
- Witten, Daniela M., Tibshirani, Robert, and Hastie, Trevor (2009). “A Penalized Matrix Decomposition, with Applications to Sparse Principal Components and Canonical Correlation Analysis”. In: *Biostatistics* 10.3, pp. 515–534. ISSN: 14654644. DOI: 10.1093/biostatistics/kxp008.
- Xia, Cedric Huchuan et al. (2018). “Linked Dimensions of Psychopathology and Connectivity in Functional Brain Networks”. In: *Nature Communications* 9.1. ISSN: 20411723. DOI: 10.1038/s41467-018-05317-y.
- Yang, Xinghao, Liu, Weifeng, Liu, Wei, and Tao, Dacheng (2021). “A Survey on Canonical Correlation Analysis”. In: *IEEE Transactions on Knowledge and Data Engineering* 33.6, pp. 2349–2368. ISSN: 1041-4347, 1558-2191, 2326-3865. DOI: 10.1109/TKDE.2019.2958342.
- Yang, Yanrong and Pan, Guangming (2015). “Independence Test for High Dimensional Data Based on Regularized Canonical Correlation Coefficients”. In: *The Annals of Statistics* 43.2, pp. 467–500. ISSN: 0090-5364.
- Yao, Fang, Müller, Hans Georg, and Wang, Jane Ling (2005). “Functional Data Analysis for Sparse Longitudinal Data”. In: *Journal of the American Statistical Association* 100.470, pp. 577–590. ISSN: 01621459. DOI: 10.1198/016214504000001745.
- Yoon, Grace, Carroll, Raymond J., and Gaynanova, Irina (2020). “Sparse Semiparametric Canonical Correlation Analysis for Data of Mixed Types”. In: *Biometrika* 107.3, pp. 609–625. ISSN: 14643510. DOI: 10.1093/biomet/asaa007.
- Yu, Y., Wang, T., and Samworth, R. J. (2015). “A Useful Variant of the Davis-Kahan Theorem for Statisticians”. In: *Biometrika* 102.2, pp. 315–323. ISSN: 14643510. DOI: 10.1093/biomet/asv008.
- Yuan, Dongbang and Gaynanova, Irina (2021). *Double-Matched Matrix Decomposition for Multi-View Data*.
- Zhang, Jingfei, Sun, Will Wei, and Li, Lexin (2020). “Mixed-Effect Time-Varying Network Model and Application in Brain Connectivity Analysis”. In: *Journal of the American Statistical Association* 115.532, pp. 2022–2036. DOI: 10.1080/01621459.2019.1677242.
- Zhao, Yi, Caffo, Brian S., and Luo, Xi (2020). *Principal Regression for High Dimensional Covariance Matrices*.
- Zhao, Yi, Wang, Bingkai, Mostofsky, Stewart H, Caffo, Brian S, and Luo, Xi (2021). “Covariate Assisted Principal Regression for Covariance Matrix Outcomes”. In: *Biostatistics* 22.3, pp. 629–645. ISSN: 1465-4644, 1468-4357. DOI: 10.1093/biostatistics/kxz057.
- Zhou, Yidong and Müller, Hans-Georg (2021). “Dynamic Network Regression”. In.
- Zhu, Hongtu, Li, Tengfei, and Zhao, Bingxin (2023). “Statistical Learning Methods for Neuroimaging Data Analysis with Applications”. In: *Annual Review of Biomedical Data Science* 6.1, pp. 73–104. ISSN: 2574-3414, 2574-3414. DOI: 10.1146/annurev-biodatasci-020722-100353.
- Zhuang, Xiaowei, Yang, Zhengshi, and Cordes, Dietmar (2020). “A Technical Review of Canonical Correlation Analysis for Neuroscience Applications”. In: *Human Brain Mapping* 41.13, pp. 3807–3833. ISSN: 1065-9471. DOI: 10.1002/hbm.25090.
- Zou, Hui, Hastie, Trevor, and Tibshirani, Robert (2006). “Sparse Principal Component Analysis”. In: *Journal of Computational and Graphical Statistics* 15.2, pp. 265–286. ISSN: 10618600. DOI: 10.1198/106186006X113430.