

Nonlinear Multiview Analysis: Identifiability and Neural Network-assisted Implementation

Qi Lyu and Xiao Fu

Abstract—Multiview analysis aims at extracting shared latent components from data samples that are acquired in different domains, e.g., image, text, and audio. Classic multiview analysis, e.g., *canonical correlation analysis* (CCA), tackles this problem via matching the linearly transformed views in a certain latent domain. More recently, powerful nonlinear learning tools such as kernel methods and neural networks are utilized for enhancing the classic CCA. However, unlike linear CCA whose theoretical aspects are clearly understood, nonlinear CCA approaches are largely intuition-driven. In particular, it is unclear under what conditions the shared latent components across the views can be identified—while identifiability plays an essential role in many applications. In this work, we revisit nonlinear multiview analysis and address both the theoretical and computational aspects. Our work leverages a useful nonlinear model, namely, the post-nonlinear model, from the nonlinear mixture separation literature. Combining with multiview data, we take a nonlinear multiview mixture learning viewpoint, which is a natural extension of the classic generative models for linear CCA. From there, we derive a learning criterion. We show that minimizing this criterion leads to identification of the latent shared components up to certain ambiguities, under reasonable conditions. Our derivation and formulation also offer new insights and interpretations to existing deep neural network-based CCA formulations. On the computation side, we propose an effective algorithm with simple and scalable update rules. A series of simulations and real-data experiments corroborate our theoretical analysis.

Index Terms—Unsupervised learning, multiview analysis, neural networks, identifiability, post-nonlinear mixture model

I. INTRODUCTION

Multiview analysis has been an indispensable tool in statistical signal processing, machine learning, and data analytics. In the context of multiview learning, a view can be understood as measurements of data entities (e.g., a cat) in a certain domain (e.g., text, image, and audio). Most data entities naturally appear in different domains. Multiview analysis aims at extracting essential and common information from different views. Compared with single-view analysis tools like *principal component analysis* (PCA) [1], *independent component analysis* (ICA) [2], and *nonnegative matrix factorization* (NMF) [3], [4], multiview analysis tools such as *canonical correlation analysis* (CCA) [5] have an array of unique features. For

This work is supported in part by National Science Foundation under Projects NSF ECCS 1608961 and ECCS 1808159, and in part by the Army Research Office (ARO) under projects ARO W911NF-19-1-0247 and ARO W911NF-19-1-0407.

Q. Lyu and X. Fu are with the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR 97331, United States. email: (lyuqi, xiao.fu)@oregonstate.edu

example, CCA has been shown to be more robust to noise and view-specific strong interference [6], [7].

The classic CCA has been extensively studied in the literature, ever since its proposal in statistics in the 1930s [5], [8]. In a nutshell, the classic CCA seeks linear transformations for the views. The transformations are supposed to ‘project’ the views to a domain where the views share similar representations. Interestingly, the formulated optimization problem, although being nonconvex, can be recast into a generalized eigendecomposition problem and solved efficiently [9]. Since the 1990s, many attempts have been made towards scaling up classic CCA to handle big datasets; see [10]–[14]. Beyond the classic two-view CCA, a series of generalized CCA (GCCA) formulations [15], [16] for handling more views exist, whose scalable versions have also been considered; see, e.g., [17]–[19].

Linear transformation-based CCA/GCCA algorithms are elegant in computation. However, from a modeling viewpoint, restricting the transformations to be linear makes the ‘modeling power’ limited. For decades, a lot of effort has been invested to extending the CCA/GCCA ideas to the nonlinear regime—via incorporating a variety of nonlinear transformations. For example, kernel CCA has been popular since the 2000s [20], [21]. More recently, together with the success of deep learning, deep neural networks are also used to enhance CCA for unsupervised representation learning [22], [23]. Compared to kernel methods, deep neural networks are considered more flexible and more scalable.

Kernel CCA and deep CCA have demonstrated effectiveness in many real-world applications, e.g., image representation learning [22] and speech processing [24]. This is encouraging—it shows that incorporating nonlinearity in data analytics is indeed well-motivated. On the other hand, it is still largely unclear under what conditions these methods will work (fail) or how to improve existing schemes with theoretical supports. In fact, unlike linear CCA whose generative models, parameter identifiability properties, and computational aspects are fairly well understood, nonlinear CCA formulations are largely intuition-driven.

In this work, we revisit the nonlinear multiview analysis problem, and offer both theoretical understanding and theory-backed implementation. Our work is motivated by a recent work in linear CCA [7], where different views are modeled as mixtures of shared components and view-specific components (which are interference). This model is similar to the one considered in machine learning in the context of probabilistic CCA [6]. Under this model, interesting interpretation for the effectiveness of CCA can be obtained. Specifically, the work

in [7] shows that the classic linear CCA can extract the shared components' range space—even if the interference terms are much stronger than the shared components. The model is fairly simple and succinct, yet the insight is significant: it offers an interesting interpretation for the effectiveness of CCA for extracting shared information across views under possibly strong interference.

Building upon the intriguing perspectives in [7], we take a step further and consider the following problem: If the acquired views are mixtures of shared and view-specific components distorted by *unknown nonlinear functions*, is it still possible to extract the same shared information as in the linear case? This problem is well-motivated, since a large variety of acquired real data are subject to *unknown* nonlinear distortion effects. For example, in array processing, satellite and microwave communications, and biological signal processing, nonlinear distortions often happen; see the discussion in [25]–[27] and the references therein.

However, taking into consideration unknown nonlinear distortions makes the problem of interest much more challenging—both in theory and practice. Nonlinearity removal from mixture models was considered under some limited cases, which are mostly single-view analysis problems. Some notable ones are 1) *nonlinear independent component analysis* (nICA) [25], [28]–[31], where the components of interest are statistically independent random processes, and 2) *nonlinear mixture learning* (NML) [32], where the sought latent components reside in a confined manifold, i.e., the probability simplex. In both cases, the assumptions on the components of interest are leveraged on to come up with identification criteria. However, both statistical independence and the probability-simplex type structure are considered special assumptions, which may not hold in general. In addition, both nICA and NML assume “clean” mixtures without interference present. With multiple views available, can we circumvent these strong assumptions like independence or the probability simplex structure? More importantly, is fending against non-interesting strong interference components still possible under nonlinear settings?

Contributions Bearing these questions in mind, we address both the analytical and computational aspects of nonlinear multiview analysis. Our detailed contributions are as follows:

- **Model-Based Formulation and Analysis.** We propose a multiview nonlinear mixture model that is a natural extension of the mixture model-based linear multiview analysis [6], [7]. To be specific, we model each view as a nonlinear mixture of shared and view-specific interference components, where the nonlinear distortions are unknown continuous invertible functions imposed on each of the received components. This distortion model is reminiscent of the *post-nonlinear mixture* (PNL) model in the nonlinear blind source separation literature. Note that the PNL model is a special nonlinear model, yet is considered effective in modeling nonlinearity in many applications—especially applications with nonlinear distortions happening on the sensor end [25], [26], [30]. We propose an identification criterion for extracting the *shared* information across views. We show that solving the formulated

problem implicitly removes the unknown nonlinearity up to trivial ambiguities—making the nonlinear multiview analysis problem boil down to a linear CCA problem. This means that our formulation enjoys the same identifiability properties as those in the linear case [7], despite of working under a much more challenging scenario.

- **Neural Network-Based Algorithm Design.** Based on our formulation, we propose a neural network based implementation. The formulated optimization surrogate is delicately designed to realize the identification criterion in practice. In particular, possible trivial solutions revealed in the analysis are circumvented via a careful construction of the optimization objective and constraints. Based on the formulation, we propose a simple block coordinate descent (BCD) algorithm. The proposed implementation is compatible with existing popular neural network architectures (e.g., *convolutive neural networks* (CNN) and fully connected neural networks [33]) and is scalable for handling big data.

- **Extensive Experiments.** We test our method in a number of simulations under different scenarios to validate the identifiability theory and to showcase the effectiveness of the implementation. In addition, a couple of real datasets (i.e., a multiview brain imaging dataset and a multiview handwritten digit dataset) are employed to demonstrate the usefulness of the proposed approach for handling real-world problems.

Part of the work will be presented at IEEE SAM 2020 [34]. The journal version additionally contains derivations of the theoretical claims, detailed algorithm design and implementation, more comprehensive simulations, and more real experiments.

Notation We largely follow the established conventions in signal processing. To be specific, we use x , \mathbf{x} , \mathbf{X} to represent a scalar, vector, and matrix, respectively. $\|\cdot\|$ denotes the Euclidean norm, i.e., $\|\mathbf{x}\|_2$ and $\|\mathbf{X}\|_F$; $\|\cdot\|_0$ denotes the ℓ_0 pseudo-norm which counts the number of nonzero elements of its argument; \top and \dagger denote the transpose and Moore-Penrose pseudo-inverse operations, respectively; the shorthand notation f is used to denote a function $f(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$; f' and f'' denote the first-order and second-order derivatives of the function f , respectively; $\text{Tr}(\cdot)$ denotes the trace operator; $\mathbf{1}$ denotes an all-one vector with a proper length; \mathbf{I} denotes an identity matrix with a proper size; $f \circ g$ denotes the function composition operation; $\mathcal{X} \times \mathcal{Y}$ means the Cartesian product of the sets \mathcal{X} and \mathcal{Y} .

II. BACKGROUND

In this section, we briefly introduce some preliminaries pertaining to our work.

A. Mixture Models

Mixture models have proven very useful in signal processing and machine learning. The simplest *linear mixture model* (LMM) can be written as:

$$\mathbf{y}_\ell = \mathbf{A}\mathbf{s}_\ell, \quad \ell = 1, 2, \dots, N, \quad (1)$$

where $\mathbf{y}_\ell \in \mathbb{R}^M$ denotes the ℓ th observed signal, $\mathbf{A} \in \mathbb{R}^{M \times K}$ the mixing system, and $\mathbf{s}_\ell \in \mathbb{R}^K$ the K latent components (or,

the “sources”) measured at sample ℓ . Many applications are concerned with identifying \mathbf{A} and/or \mathbf{s}_ℓ ’s from the observed \mathbf{y}_ℓ ’s; see [27]. If both the mixing system and the sources are unknown, estimating \mathbf{A} and \mathbf{s}_ℓ simultaneously poses a very hard problem—which is known as the *blind source separation* (BSS) problem; see, e.g., [27], [35]–[39]. The BSS problem is ill-posed, since in general the model is not identifiable; i.e., even if there is no noise, one could have an infinite number of solutions that satisfy (1). Nonetheless, identifiability can be established via exploiting properties of \mathbf{s}_ℓ . For example, the seminal work of ICA [2] shows that the identifiability of \mathbf{s}_ℓ can be established (up to scaling and permutation ambiguities) leveraging on statistical independence between the elements of \mathbf{s}_ℓ . Later on, identifiability of \mathbf{A} and/or \mathbf{s}_ℓ were established through exploiting other properties of the latent components, e.g., convex geometry [35], [40], quasi-stationarity [41], spectral diversity [42], sparsity [39], and boundedness [38].

Beyond the LMM, *nonlinear* mixture models (NMMs) have also attracted much attention—since nonlinearity naturally happens in practice for numerous reasons. For example, starting from the 1990s, a line of work named nonlinear ICA [25], [28], [30], [31] considered the following model:

$$\mathbf{y}_\ell = \mathbf{g}(\mathbf{s}_\ell), \quad \ell = 1, 2, \dots, N, \quad (2)$$

where $\mathbf{g}(\cdot) : \mathbb{R}^K \rightarrow \mathbb{R}^M$ is a nonlinear mixing system. While it was shown that (2) is in general not identifiable, the so-called *post-nonlinear* (PNL) model can be identified, again under the premise that $s_{k,\ell}$ and $s_{j,\ell}$ are statistically independent random processes [25], [30], [43]. To be specific, the PNL model is defined as

$$\mathbf{y}_\ell = \mathbf{g}(\mathbf{A}\mathbf{s}_\ell), \quad \ell = 1, 2, \dots, N, \quad (3)$$

with $\mathbf{g}(\cdot) = [g_1(\cdot), \dots, g_M(\cdot)]^\top$ and $g_m(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$ being a univariate nonlinear function. This model, although less general relative to (2), is still very meaningful—which finds applications in cases where unknown nonlinear effects happen *individually* at multiple sensors/channels, e.g., in brain signal processing [44], [45] and hyperspectral imaging [46]. Under the PNL model, \mathbf{s}_ℓ was first shown to be identifiable via exploiting statistical independence [25], [30]. Recently, the work in [32] proved that even if the sources are dependent, some other properties, e.g., nonnegativity and a sum-to-one structure (i.e., $\mathbf{1}^\top \mathbf{s}_\ell = 1$), can be exploited to identify \mathbf{s}_ℓ under the PNL model.

B. Multiview Data and Mixture Learning

In practice, data representing the same entities are oftentimes acquired in different domains—leading to the so-called multiview data. Multiview data has been frequently connected to mixture models, since it is believed that multiple views of the same data sample have certain shared but latent components. In addition, it is commonly believed that using multiple views may have advantages over using a single view for learning problems, e.g., for combating noise and strong interference—since more information is available.

One way of modeling multiview data is using the following linear model [6]:

$$\mathbf{y}_\ell^{(q)} \approx \mathbf{A}^{(q)} \mathbf{s}_\ell, \quad q = 1, 2, \dots, Q, \quad (4)$$

where q is the index of the view and Q is the number of available views. In this work, we mainly consider $Q = 2$. But the techniques can be readily extended to cover the $Q \geq 3$ case. From the above, one can see that \mathbf{s}_ℓ can be used to model the latent shared components across different views, while $\mathbf{A}^{(q)}$ is the basis of the subspace where the q th view is observed. In other words, the different appearances of the views are caused by the differences of the observation subspaces, while the latent representations of the views are identical. It was shown in [6] that CCA is a maximum likelihood estimator (MLE) for $\text{range}(\mathbf{A}^{(q)})$ under the model $\mathbf{y}_\ell^{(q)} = \mathbf{A}^{(q)} \mathbf{s}_\ell + \mathbf{v}_\ell^{(q)}$, where $q = 1, 2$, if the noise $\mathbf{v}_\ell^{(q)}$ is Gaussian—even if the noise is *colored* and its covariance is *unknown*. This spells out the benefits of using a CCA-like learning criterion. If one concatenates the views (i.e., to form $\mathbf{y}_\ell = [(\mathbf{y}_\ell^{(1)})^\top, (\mathbf{y}_\ell^{(2)})^\top]^\top$ and applies single-view tools like PCA, it is MLE only when the noise follows the i.i.d. zero-mean *white* Gaussian distribution. Similar models have also been studied in the context of joint blind source separation (JBSS) and multiview biological signal analysis [47], [48]. Unlike the model in [6], in [47], the latent components in each view are not exactly the same, but closely correlated across views. In the recent work [7], the above model in (4) is further developed to incorporate view-specific components, where we have:

$$\mathbf{y}_\ell^{(q)} = \mathbf{A}^{(q)} [\mathbf{s}_\ell^\top, (\mathbf{c}_\ell^{(q)})^\top]^\top, \quad \ell = 1, \dots, N \quad (5)$$

where $\mathbf{A}^{(q)} \in \mathbb{R}^{M_q \times (K+R_q)}$, $\mathbf{s}_\ell \in \mathbb{R}^K$ collects the *shared components* and $\mathbf{c}_\ell^{(q)} \in \mathbb{R}^{R_q}$ are the view-specific components. This model is plausible, since it gives more flexibility for modeling the cross-view disparities. More importantly, it explains the enhanced robustness of multiview analysis against the single view ones, e.g., PCA, for extracting essential shared information (i.e., \mathbf{s}_ℓ up to certain ambiguities) about the data. More specifically, consider the following linear CCA formulation [8], [17]:

$$\begin{aligned} & \underset{\mathbf{B}^{(q)} \in \mathbb{R}^{K \times M_q}}{\text{maximize}} \quad \text{Tr} \left(\frac{1}{N} \sum_{\ell=1}^N \mathbf{B}^{(1)} \mathbf{y}_\ell^{(1)} \left(\mathbf{B}^{(2)} \mathbf{y}_\ell^{(2)} \right)^\top \right) \\ & \text{subject to} \quad \frac{1}{N} \sum_{\ell=1}^N \mathbf{B}^{(q)} \mathbf{y}_\ell^{(q)} (\mathbf{y}_\ell^{(q)})^\top (\mathbf{B}^{(q)})^\top = \mathbf{I}, \quad q = 1, 2 \end{aligned} \quad (6)$$

where $\mathbf{B}^{(q)} \in \mathbb{R}^{K \times M_q}$ ¹ and the constraints are employed to avoid degenerate solutions. CCA aims to find the most correlated linear projections of both views. Note that CCA was not developed under the particular multiview mixtures in (5). However, it was shown in [7] that, under the model

¹Note that in the literature, the size of $\mathbf{B}^{(q)}$ can be $R \times M_q$, where $R \in \{1, \dots, K\}$. For example, setting $R = 1$ means that one aims at extracting one shared component, which is similar to the convention in PCA. Nevertheless, for notational simplicity, we set $R = K$ in this paper. The proofs and results can be trivially extended to cover the general case $\mathbf{B}^{(q)} \in \mathbb{R}^{R \times M_q}$ where $1 \leq R \leq K$.

in Eq. (5), solving the above can identify $\hat{\mathbf{x}}_\ell = \Theta \mathbf{s}_\ell$, where $\Theta \in \mathbb{R}^{K \times K}$ is a nonsingular matrix—no matter how strong the view-specific components $\mathbf{c}_\ell^{(q)}$ are—the proof in [7] is restated in Appendix F using notations of this paper, for ease of exposition. On the contrary, PCA always extracts the energy-wise strongest components.

To understand the result in [7], it might be the best to view the CCA cost function in its equivalent form as follows [8]:

$$\underset{\mathbf{B}^{(q)} \in \mathbb{R}^{R \times M_q}}{\text{minimize}} \sum_{\ell=1}^N \left\| \mathbf{B}^{(1)} \mathbf{y}_\ell^{(1)} - \mathbf{B}^{(2)} \mathbf{y}_\ell^{(2)} \right\|_2^2 \quad (7)$$

$$\text{subject to } \frac{1}{N} \sum_{\ell=1}^N \mathbf{B}^{(q)} \mathbf{y}_\ell^{(q)} (\mathbf{y}_\ell^{(q)})^\top (\mathbf{B}^{(q)})^\top = \mathbf{I}, \quad q = 1, 2.$$

One can check that $\mathbf{B}^{(q)} = [\Theta, \mathbf{0}] (\mathbf{A}^{(q)})^\dagger \in \mathbb{R}^{K \times M_q}$ makes the cost in (7) zero under the model in (5). The work in [7] further shows that this type of solution is unique up to an arbitrary non-singular Θ . Since $\mathbf{B}^{(q)} \mathbf{y}_\ell^{(q)} = \Theta \mathbf{s}_\ell$, this means that the range space spanned by the shared components \mathbf{S}^\top (where $\mathbf{S} = [\mathbf{s}_1, \dots, \mathbf{s}_N]$) can be extracted. Note that the resulting projected views, i.e.,

$$\hat{\mathbf{x}}_\ell = \mathbf{B}^{(q)} \mathbf{y}_\ell^{(q)} = \Theta \mathbf{s}_\ell, \quad \ell = 1, \dots, N,$$

are again LMMs, and many techniques introduced in the previous section can be used to identify \mathbf{s}_ℓ —with the interference components $\mathbf{c}_\ell^{(q)}$ having been eliminated by CCA.

C. Nonlinear Multiview Learning

As in the single-view case, nonlinearity is natural to be considered in the multiview case. For example, nonlinear learning tools such as the kernel method and deep learning were combined with CCA, where interesting results were observed [21]–[23], [49]. For example, the idea of deep CCA is to employ deep neural networks, instead of the linear operators $\mathbf{B}^{(q)}$'s as in the classic CCA, to perform data transformation. The deep CCA formulation in [23] is based on maximizing the correlation between $\mathbf{f}^{(q)}(\mathbf{y}_\ell^{(q)})$'s across views, where $\mathbf{f}^{(q)}(\cdot)$ is a deep neural network-represented nonlinear transformation. Nevertheless, these works mostly focus on practical implementations, rather than theoretical aspects. It is unclear if the interesting identifiability results as in the classic multiview mixture model in [7] still hold when unknown nonlinearity is imposed onto the views. This is the starting point of our work.

III. PROPOSED APPROACH

A. A Nonlinear Multiview Model

In this section, we propose a nonlinear multiview analysis method that aims at learning shared information from the following model:

$$\mathbf{y}_\ell^{(q)} = \mathbf{g}^{(q)}(\mathbf{A}^{(q)} \mathbf{z}_\ell^{(q)}), \quad q = 1, 2, \quad (8)$$

where $\mathbf{z}_\ell^{(q)} = [\mathbf{s}_\ell^\top, (\mathbf{c}_\ell^{(q)})^\top]^\top$, the shared components are uncorrelated zero-mean stationary random processes, i.e.,

$$\mathbb{E}[\mathbf{s}_{i,\ell}] = 0, \quad \mathbb{E}[\mathbf{s}_\ell \mathbf{s}_\ell^\top] = \mathbf{\Sigma}_s : \text{diagonal},$$

the function $\mathbf{g}^{(q)}(\cdot) = [g_1^{(q)}(\cdot), \dots, g_{M_q}^{(q)}(\cdot)]^\top$ and $g_m^{(q)}(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$ represents the view specific nonlinear distortion at channel m . Note that each $g_m^{(q)}(\cdot)$ is assumed to be *invertible*—otherwise there is no hope to recover the underlying model. Under the nonlinear setting, we also assume that the latent variables are defined over continuous open sets \mathcal{S} and \mathcal{C}_q , i.e.,

$$\mathbf{s}_\ell \in \mathcal{S}, \quad \mathbf{c}_\ell^{(q)} \in \mathcal{C}_q, \quad q = 1, 2. \quad (9)$$

Note that this model is a natural extension of the multiview linear mixture models as in [6], [7] (especially the one in [7]) and the PNL model in single-view nonlinear mixture analysis [25], [30], [43].

As we have discussed, one major motivation behind the model in (8) is that in signal sensing and data acquisition, nonlinear distortions oftentimes happen at the sensor end. Hence, a PNL model is considered appropriate for such cases. Nonlinear effects are perhaps particularly severe for biosensors, since biological signals are hard to predict or calibrate. On the other hand, multiple views often exist in biology, e.g., electroencephalogram (EEG), magnetoencephalography (MEG), and functional magnetic resonance imaging (fMRI) measurements for the same stimuli [50]–[52]. Note that in practice, EEG, MEG, and fMRI images may contain closely correlated latent components other than sharing exactly the same latent components [48]. Nonetheless, the proposed model may serve as a reasonable approximation. The proposed model in (8) can be leveraged on to handle such biosensor-acquired multiview data.

We should mention that the model in (8) could also benefit more general data analytics applications such as image/text classification. The model in (8) is evolved from the classic linear mixture model that has been widely used in data science for pre-processing, in particular, dimensionality reduction via computational tools such as PCA, ICA, and CCA. Adding a layer of nonlinearity can help capture more dynamics that were not modeled by the classic tools, which likely leads to better performance. This is the basic intuition behind the deep autoencoder [53] and kernel/deep CCA [20], [21], [23]. Nonetheless, unlike deep autoencoder or deep CCA that are purely data-driven, we take a model-based route in this work to reveal the insights behind the effectiveness of nonlinear multiview learning—which may also shed some light on more principled design of learning criteria and algorithms.

B. Proposed Formulation

Our goal is to learn $\mathbf{x}_\ell = \Theta \mathbf{s}_\ell$, $\ell = 1, \dots, N$, with a certain nonsingular Θ under the nonlinear model (8). This goal is identical to that in [7] under the simpler linear case [cf. Eq. (5)]. Note that we do not impose strong structural assumptions on \mathbf{s}_ℓ , e.g., statistical independence between $s_{r,\ell}$ and $s_{j,\ell}$ as in [25], [30] or the stochasticity assumption (i.e., $\mathbf{1}^\top \mathbf{s}_\ell = 1, \mathbf{s}_\ell \geq \mathbf{0}$) as in [32]—existing nonlinear ICA and mixture learning techniques cannot be applied to the problem of interest.

Our idea is to seek a nonlinear mapping $f^{(q)}(\cdot) : \mathbb{R}^{M_q} \rightarrow \mathbb{R}^{M_q}$ and a linear operator $\mathbf{B}^{(q)}$ such that the following criterion is minimized:

$$\underset{\mathbf{B}^{(q)} \in \mathbb{R}^{K \times M_q}, \mathbf{f}^{(q)}}{\text{minimize}} \sum_{\ell=1}^N \left\| \mathbf{B}^{(1)} \mathbf{f}^{(1)} \left(\mathbf{y}_\ell^{(1)} \right) - \mathbf{B}^{(2)} \mathbf{f}^{(2)} \left(\mathbf{y}_\ell^{(2)} \right) \right\|_2^2.$$

subject to : $\mathbf{f}^{(q)} : \text{invertible}, q = 1, 2, \quad (10a)$

$$\frac{1}{N} \sum_{\ell=1}^N \left(\mathbf{B}^{(q)} \mathbf{f}^{(q)} \left(\mathbf{y}_\ell^{(q)} \right) \mathbf{f}^{(q)} \left(\mathbf{y}_\ell^{(q)} \right)^\top (\mathbf{B}^{(q)})^\top \right) = \mathbf{I}, \quad (10b)$$

The idea is very similar to that of the linear CCA reformulation as in (7), with nonlinearity taken into consideration. Ideally, we wish to obtain

$$\mathbf{B}^{(q)} = [\Theta, \mathbf{0}] (\mathbf{A}^{(q)})^\dagger \in \mathbb{R}^{K \times M_q}, \quad (11a)$$

$$\mathbf{f}^{(q)}(\cdot) = [f_1^{(q)}(\cdot), \dots, f_{M_q}^{(q)}(\cdot)]^\top : \mathbb{R}^{M_q} \rightarrow \mathbb{R}^{M_q}, \quad (11b)$$

$$f_m^{(q)}(\cdot) = (g_m^{(q)}(\cdot))^{-1}, m = 1, \dots, M_q, \quad (11c)$$

where $\Theta \in \mathbb{R}^{K \times K}$ is nonsingular. The above will extract the shared row subspace. Also note that the solution in (11) is a feasible solution to Problem (10), which always exists. The key question is: Is the solution in (11) the *only* solution? That is, does the formulation in (10) have identifiability for the shared subspace spanned by the rows of $\mathbf{S} = [s_1, \dots, s_N]$?

C. Nonlinearity Removal

To see how we approach the identifiability problem, we consider the following problem:

$$\text{find } \mathbf{B}^{(1)}, \mathbf{B}^{(2)}, \mathbf{f}^{(1)}, \mathbf{f}^{(2)} \quad (12a)$$

$$\text{s.t. } \mathbf{B}^{(1)} \mathbf{f}^{(1)} \left(\mathbf{y}_\ell^{(1)} \right) = \mathbf{B}^{(2)} \mathbf{f}^{(2)} \left(\mathbf{y}_\ell^{(2)} \right), \quad (12b)$$

$$\forall \mathbf{y}_\ell^{(q)} = \mathbf{g}^{(q)}(\mathbf{A}^{(q)} \mathbf{s}_\ell^{(q)}), \mathbf{s}_\ell \in \mathcal{S}, \mathbf{c}_\ell^{(q)} \in \mathcal{C}_q,$$

$$\mathbf{f}^{(q)} : \text{invertible}, q = 1, 2, \quad (12c)$$

$$\mathbb{E} \left[\mathbf{B}^{(q)} \mathbf{f}^{(q)} \left(\mathbf{y}_\ell^{(q)} \right) \mathbf{f}^{(q)} \left(\mathbf{y}_\ell^{(q)} \right)^\top (\mathbf{B}^{(q)})^\top \right] = \mathbf{I}. \quad (12d)$$

Note that the above is derived from (10) assuming that one has uncountably infinite $\mathbf{y}_\ell^{(q)}$'s such that all possible values of $\mathbf{y}_\ell^{(q)}$ are exhausted. We use an equality constraint in (12b) since when there is no noise, the optimal value in (10) should be zero under the model in Eq. (8)—attaining zero fitting cost is equivalent to satisfying an equality constraint. Like Problem (10), Problem (12) also at least admits one solution as in (11).

To see the reason why the formulation in (12) can remove nonlinearity and identify the desired signal subspace, let us start with a simple illustrative case where $\mathbf{A}^{(q)} \in \mathbb{R}^{2 \times 2}$ and there is only one shared component. This way, the generative model is simplified as follows:

$$\mathbf{y}_\ell^{(1)} = \left[g_1^{(1)} \left(A_{11}^{(1)} s_\ell + A_{12}^{(1)} c_\ell^{(1)} \right), g_2^{(1)} \left(A_{21}^{(1)} s_\ell + A_{22}^{(1)} c_\ell^{(1)} \right) \right]^\top$$

$$\mathbf{y}_\ell^{(2)} = \left[g_1^{(2)} \left(A_{11}^{(2)} s_\ell + A_{12}^{(2)} c_\ell^{(2)} \right), g_2^{(2)} \left(A_{21}^{(2)} s_\ell + A_{22}^{(2)} c_\ell^{(2)} \right) \right]^\top$$

For this simplified case, the solution $\mathbf{B}^{(q)}$ becomes a row vector $(\mathbf{b}^{(q)})^\top$.

To proceed, we will be using the following condition:

Definition 1 (Ubiquitously Unanchored) Consider a collection of d real-valued random components $\mathbf{v} = [v_1, \dots, v_d]^\top \in \mathbb{R}^d$, where v_i resides in a continuous and open set $\mathcal{V}_i \subseteq \mathbb{R}$ such that $\text{volume}(\mathcal{V}_i) > 0$. Denote \bar{v}_j as any fixed value from \mathcal{V}_j . Assume that for any $i \in \{1, \dots, d\}$ and any \bar{v}_j where $j \neq i$, the vectors

$$[\bar{v}_1, \dots, \bar{v}_{i-1}, v, \bar{v}_{i+1}, \dots, \bar{v}_d]^\top, \quad \forall v \in \mathcal{V}_i$$

are contained in the domain of \mathbf{v} . Then, the components in \mathbf{v} are called *ubiquitously unanchored*.

Remark 1 Intuitively speaking, the *ubiquitously unanchored* condition means that for any fixed $v_j = \bar{v}_j$ where $j \in \{1, \dots, d\}/\{i\}$, v_i is still allowed to take all the possible values in \mathcal{V}_i —thereby being “unanchored” by the other components. Of course, if v_1, \dots, v_d are statistically independent, then \mathbf{v} satisfies this condition. On the other hand, the ubiquitously unanchored condition is much milder than statistical independence. If \mathbf{v} is a continuous random vector, the ubiquitously unanchored condition boils down to that the joint probability density function (PDF) of \mathbf{v} takes positive values in the domain of $\mathcal{V} = \mathcal{V}_1 \times \dots \times \mathcal{V}_d$. This is much milder than requiring that the components of \mathbf{v} to be statistically independent. For example, consider a case where $d = 2$ and $\mathcal{V}_i = (-1, 1)$ for $i = 1, 2$. In this case, one can have $v_1 = v_2 + n$ where n is a random perturbation. Even if the perturbation admits a small variance and thus is weak (meaning that v_1 and v_2 are highly dependent), if the PDF of n admits positive values over $(-1, 1)$, the components in \mathbf{v} are still ubiquitously unanchored.

Theorem 1 (The 2×2 Case) Assume that the elements of $\mathbf{A}^{(q)} \in \mathbb{R}^{2 \times 2}$ are drawn from any jointly continuous distribution for $q = 1, 2$, and that $\mathbf{g}^{(q)}$ and the composition $\mathbf{f}^{(q)} \circ \mathbf{g}^{(q)}$ are both twice differentiable. Assume that every view admits one shared component ($s_\ell \in \mathbb{R}$) and one interference term ($c_\ell^{(q)} \in \mathbb{R}$), and that the components in $[s_\ell, c_\ell^{(1)}, c_\ell^{(2)}]^\top$ are ubiquitously unanchored. Suppose that $(\mathbf{b}^{(q)}, \mathbf{f}^{(q)})$ for $q = 1, 2$ are feasible solutions of Problem (12). Then, we have the following holds almost surely:

$$h_i^{(q)}(x) = f_i^{(q)} \circ g_i^{(q)}(x) = \alpha_i^{(q)} x + d_i^{(q)}, \quad \alpha_i^{(q)} \neq 0, \quad \forall i.$$

where $d_i^{(q)} \in \mathbb{R}$ and $\alpha_i^{(q)} \in \mathbb{R}$ for $i = 1, \dots, M_q$; i.e., the composition is an affine function with probability one.

Proof: The proof is relegated to Appendix A. \blacksquare

Theorem 1 shows that even if we have different view-specific components for multiview data, we can always identify the shared components under mild assumptions. This property is quite appealing since no strong assumptions like statistical independence, non-negativity or simplex structure are required for \mathbf{s}_ℓ —as is opposed to existing frameworks, e.g., those in [25], [28]–[30], [32]. In addition, there is no constraint

on the energy of each component—i.e., even if the energy of the shared component is significantly smaller compared to that of the view-specific component, the proposed criterion can still recover the shared subspace. This property is the same as that of the classic linear CCA [7], which we restated in Appendix F. Theorem 1 articulates the theoretical benefits of utilizing multiple views in nonlinear component analysis. This is consistent with previous study on linear multiview analysis [6], [7], [54].

In the proof of Theorem 1, we have implicitly used the fact that there at least exists a solution $\mathbf{b}^{(q)}$ to Problem (10) such that $\|\mathbf{b}^{(q)}\|_0 = 2$ (i.e., $\mathbf{b}^{(q)}$ is a “dense” vector). In fact, this holds for more general case where $\mathbf{B}^{(q)}$ is a $K \times M_q$ matrix. We have the following proposition:

Proposition 1 *Assume that $\mathbf{A}^{(q)} \in \mathbb{R}^{M_q \times (K+R_q)}$ is drawn from any joint absolutely continuous distribution, where $M_q \geq K + R_q$. Then, there exists a solution $\mathbf{B}^{(q)} \in \mathbb{R}^{K \times M_q}$ of Problem (12) that satisfies $\|\mathbf{B}^{(q)}\|_0 = KM_q$, with probability one.*

Proof: The proof is relegated to Appendix B. ■

With the insights gained from the 2×2 case and Proposition 1, we are ready to show the nonlinear distortion identifiability under the general problem setting, i.e., under $\mathbf{A}^{(q)} \in \mathbb{R}^{M_q \times (K+R_q)}$. We have the following theorem:

Theorem 2 (Nonlinearity Removal) *Consider the nonlinear model in (8). Assume that $M_q \geq K + R_q$, $q = 1, 2$, that the mixing matrices $\mathbf{A}^{(q)}$ for $q = 1, 2$ are drawn from any absolutely continuous distributions, and that the components in $[\mathbf{s}_\ell^\top, (\mathbf{c}_\ell^{(1)})^\top, (\mathbf{c}_\ell^{(2)})^\top]^\top$ are ubiquitously unanchored. Suppose that $(\mathbf{B}^{(q)}, \mathbf{f}^{(q)})$ for $q = 1, 2$ are solutions of (12) with $\|\mathbf{B}^{(q)}\|_0 = KM_q$. Then, the following holds almost surely:*

$$h_i^{(q)}(x) = f_i^{(q)} \circ g_i^{(q)}(x) = \alpha_i^{(q)}x + d_i^{(q)}, \quad \alpha_i^{(q)} \neq 0, \quad \forall i.$$

where $d_i^{(q)} \in \mathbb{R}$ and $\alpha_i^{(q)} \in \mathbb{R}$ for $i = 1, \dots, M_q$; i.e., the composition is an affine function with probability one.

Proof: The proof is relegated to Appendix C. ■

Note that the for the general case, we have an additional constraint $\|\mathbf{B}^{(q)}\|_0 = KM_q$, which arises because the general case is more challenging in terms of analysis. The condition $\|\mathbf{B}^{(q)}\|_0 = KM_q$ means that $\mathbf{B}^{(q)}$ contains no zero entries. Note that this condition was not needed in the special 2×2 case (cf. Theorem 1). This makes the conditions stated in Theorem 2 a bit more restrictive—having more latent components gives rise to a more challenging nonlinearity removal problem.

In Theorems 1–2, the shared components $\mathbf{s}_{k,\ell}$ and $\mathbf{s}_{j,\ell}$ can be heavily dependent—yet they cannot be completely dependent, since they have to satisfy the ubiquitously unanchored condition. However, if the dimensions of the interference terms $\mathbf{c}_\ell^{(q)}$ are large enough so that the views are sufficiently diverse, we show that \mathbf{s}_ℓ can even be completely dependent, without affecting nonlinearity removal:

Theorem 3 *Consider the nonlinear model in Eq. (8). Assume that $M_q \geq K + R_q$, $q = 1, 2$, and that the mixing*

matrices $\mathbf{A}^{(q)}$ for $q = 1, 2$ are drawn from any absolutely continuous distributions. Assume that the components in $[\mathbf{s}_{k,\ell}, (\mathbf{c}_\ell^{(1)})^\top, (\mathbf{c}_\ell^{(2)})^\top]^\top$ are ubiquitously unanchored for any k . Further assume that the dimensions of the components satisfy

$$\frac{R_q(R_q + 1)}{2} \geq M_q. \quad (13)$$

Suppose that $(\mathbf{B}^{(q)}, \mathbf{f}^{(q)})$ for $q = 1, 2$ are solutions of Eq. (12) with $\|\mathbf{B}^{(q)}\|_0 = KM_q$. Then, the composition $f_i^{(q)} \circ g_i^{(q)}(x)$ for all i, q are affine functions with probability one.

Proof: The proof is relegated to Appendix D. ■

Under Theorem 3, fully dependent $\mathbf{s}_{k,\ell}$ and $\mathbf{s}_{j,\ell}$ are allowed to exist; also see our experiments in Sec. V. In terms of shared component extraction, we have the following corollary:

Corollary 1 (Subspace Identifiability) *Under the generative model (8), assume that a feasible solution of Problem (12), denoted by $\mathbf{B}^{(q)}$ and $\mathbf{f}^{(q)}$ for $q = 1, 2$, can be found, where $\|\mathbf{B}^{(q)}\|_0 = KM_q$. Assume that $\mathbf{s}_\ell \in \mathbb{R}^K$ and $\mathbf{c}_\ell^{(q)} \in \mathbb{R}^{M_q}$ for $\ell = 1, \dots, N$ are drawn from certain jointly continuous distribution with $N \geq K + R_1 + R_2$. Also assume that the learned functions satisfy that the components of $\mathbf{B}^{(q)} \mathbf{f}^{(q)}(\mathbf{y}_\ell^{(q)})$ admit zero mean. Then, we have*

$$\mathbf{B}^{(q)} \mathbf{f}^{(q)}(\mathbf{y}_\ell^{(q)}) = \Theta \mathbf{s}_\ell, \quad q = 1, 2,$$

for a certain nonsingular $\Theta \in \mathbb{R}^{K \times K}$ almost surely.

Proof: The above can be shown by combining Theorem 2 and the identifiability theorem of linear CCA [7], which we restate in Appendix F in the supplementary materials. ■ We would like to remark that the assumption that $f_i^{(q)} \circ g_i^{(q)}$ has zero mean is natural, since in our model $\mathbf{s}_{i,\ell}$ has zero mean. In addition, from a CCA viewpoint, correlation is measured with centered data. However, neither of the formulations in (12) nor (10) can enforce this directly—from what we have proved in Theorems 1–2, the learned mapping $h_i^{(q)}$ is affine, not linear. That is, there is a constant term $d_i^{(q)}$ existing. Nonetheless, this term can be easily removed via adding a constraint in the formulation—more discussions will be seen in the next section. Another remark is that, linear CCA identifies $\Theta \mathbf{s}_\ell$ under the model in (5) with finite samples [7]. However, our nonlinearity removal result is established under the assumption that the equality constraint in (12) holds over a continuous domain, making the conditions in our case more stringent. Nevertheless, the proposed approach works well with large but finite N (i.e., number of samples), as will be seen in the experiments. Our result may also shed some light on the reason why nonlinear approaches are often in favor of “big data”.

D. Related Works

We should mention that there are a couple of works that are related to mixture models and nonlinear CCA. A classic work in [21] utilizes kernel CCA to solve the ICA problem under the linear mixture model. The method splits $\mathbf{y}_\ell = \mathbf{A}\mathbf{s}_\ell$ to $\mathbf{y}_\ell^{(1)} = \mathbf{A}^{(1)}\mathbf{s}_\ell$ and $\mathbf{y}_\ell^{(2)} = \mathbf{A}^{(2)}\mathbf{s}_\ell$, where $\mathbf{A} = [(\mathbf{A}^{(1)})^\top, (\mathbf{A}^{(2)})^\top]^\top$, and then applies kernel CCA to

match the transformed $\mathbf{y}_\ell^{(1)}$ and $\mathbf{y}_\ell^{(2)}$ —through which s_ℓ -identification can be achieved. The method uses nonlinear CCA to solve the problem of interest, while the model is not nonlinearly distorted. Another work in [43] considers the model in (3), i.e., $\mathbf{y}_\ell = \mathbf{g}(\mathbf{A}s_\ell)$. Algorithms were derived for removing nonlinearity, again, through splitting the received signals into two parts, i.e., $\mathbf{y}_\ell^{(1)}$ and $\mathbf{y}_\ell^{(2)}$, and maximizing the correlation between transformed $\mathbf{y}_\ell^{(1)}$ and $\mathbf{y}_\ell^{(2)}$. This is effectively multiview matching as we did. Nevertheless, the work does not consider the view-specific interference terms $\mathbf{c}_\ell^{(q)}$ —which is an important consideration in practice. The work also imposes strong assumptions, e.g., Gaussianity of \mathbf{x}_ℓ , where $\mathbf{x}_\ell = \mathbf{A}s_\ell$, for establishing identifiability of $\mathbf{g}(\cdot)$.

IV. PRACTICAL IMPLEMENTATION

The identifiability theorems indicate that the nonlinear functions $\mathbf{g}^{(q)}(\cdot)$ in the data generation process for $q = 1, 2$ may be removed up to affine transformations—if one can find a solution to Problem (12). Problem (12) addresses the population case that requires practical approximations under finite samples and workable parametrization for $\mathbf{f}^{(q)}$ and $\mathbf{g}^{(q)}$. In this section, we cast (12) into a numerical optimization-friendly form and propose an algorithm for tackling the reformulated problem.

A. Parametrization for Nonlinear Functions

To tackle Problem (12), we first parameterize $\mathbf{f}^{(q)}$ and $\mathbf{g}^{(q)}$ using two neural networks. Since we seek for continuous functions, neural networks are good candidates since they are the so-called ‘universal function approximators’. Note that one-hidden-layer networks can already represent all continuous functions over bounded domains to ϵ -accuracy with a finite number of neurons [55]. Nevertheless, we keep our parametrization flexible to incorporate multiple layers—which have proven effective in practice [33]. In addition, many established network configurations and architectures can be incorporated into our framework. This is particularly of interest since multiview data oftentimes are with quite diverse forms, e.g., image and text. Using some established neural network paradigms for different types of data (e.g., CNN for image data [33]) may enhance performance.

In our case, since we consider $f_m^{(q)}(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$ and $g_m^{(q)}(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$ as nonlinear functions, we parameterize them using neural networks. For example, we use the following to approximate $g_m^{(q)}(x)$ (and the same for $f_m^{(q)}(x)$):

$$g_m^{(q)}(x) \approx (\mathbf{w}_L^{(q)})^\top \times \sigma\left(\mathbf{W}_L^{(q)} \dots \sigma\left(\mathbf{W}_2^{(q)} \sigma\left(\mathbf{w}_1^{(q)} x + \gamma_1^{(q)}\right) + \gamma_2^{(q)}\right) \dots + \gamma_{L-1}^{(q)}\right),$$

where $\mathbf{w}_1^{(q)} \in \mathbb{R}^{N_1}$ and $\mathbf{w}_L^{(q)} \in \mathbb{R}^{N_L}$ are the network weights of the input and output layers, respectively, $\mathbf{W}_l^{(q)} \in \mathbb{R}^{N_l \times N_{l-1}}$ denotes the network weights from layer $l-1$ to layer l , $\gamma_l^{(q)} \in \mathbb{R}^{N_l}$ is the bias term of layer l , and $\sigma(x)$ is the so-called activation function. One typical activation function is the sigmoid function, i.e., $\sigma(x) = \frac{1}{1+e^{-x}}$. There exist many other choices, e.g., the tanh function, and the so-called

rectified linear unit ReLU function [33]. Note that different configurations of $\mathbf{W}_l^{(q)}$ and $\sigma(\cdot)$ lead to different types of neural networks; see [33].

B. Reformulation

With the neural network-based parametrization denoted as $\mathbf{f}_{\text{NN}}^{(q)}$ and $\mathbf{g}_{\text{NN}}^{(q)}$, one can use (10) as a working surrogate, with proper modifications. To be specific, we consider the following optimization problem:

$$\begin{aligned} & \underset{\mathbf{U}, \theta_{\text{NN}}, \mathbf{B}^{(q)}}{\text{minimize}} \sum_{q=1}^Q \sum_{\ell=1}^N \left\| \mathbf{u}_\ell - \mathbf{B}^{(q)} \mathbf{f}_{\text{NN}}^{(q)}\left(\mathbf{y}_\ell^{(q)}\right) \right\|_2^2 \\ & \quad + \lambda \sum_{q=1}^Q \sum_{\ell=1}^N \left\| \mathbf{y}_\ell^{(q)} - \mathbf{g}_{\text{NN}}^{(q)}\left(\mathbf{f}_{\text{NN}}^{(q)}\left(\mathbf{y}_\ell^{(q)}\right)\right) \right\|_2^2 \\ & \text{subject to } \frac{1}{N} \left[\sum_{\ell=1}^N \mathbf{u}_\ell \mathbf{u}_\ell^\top \right] = \mathbf{I}, \quad \frac{1}{N} \sum_{\ell=1}^N \mathbf{u}_\ell = \mathbf{0} \quad (14) \end{aligned}$$

where $Q = 2$ for the two-view case, and we introduce a slack variable $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_N] \in \mathbb{R}^{K \times N}$ that represents the extracted shared components (ideally we wish $\mathbf{u}_\ell = s_\ell$), $\mathbf{y}_\ell^{(q)}$ is the ℓ th input data for the q th view, $\mathbf{f}_{\text{NN}}^{(q)}(\cdot) = [f_{\text{NN},1}^{(q)}(\cdot), \dots, f_{\text{NN},M_q}^{(q)}(\cdot)]^\top$ is a neural network (NN)-parametrized element-wise non-linear mapping that we aim at learning for nonlinearity removal, and $\mathbf{g}_{\text{NN}}^{(q)}(\cdot) = [g_{\text{NN},1}^{(q)}(\cdot), \dots, g_{\text{NN},M_q}^{(q)}(\cdot)]^\top$ is another NN-parametrized nonlinear function for learning the generative function in (8). For conciseness, we use θ_{NN} to denote the network parameters in the NNs representing $\mathbf{f}_{\text{NN}}^{(q)}$ and $\mathbf{g}_{\text{NN}}^{(q)}$.

To explain the formulation, the first term in the cost function is by lifting the constraint (12b) to the cost function via introducing a slack variable \mathbf{U} . This is reminiscent of the fitting formulation in (10). Note that introducing \mathbf{U} is important, since it entails us to propose a lightweight algorithm. The second term is a regularization that promotes $\mathbf{f}_{\text{NN}}^{(q)}$ to be invertible [i.e., to reflect the invertibility constraints in (10) and (12)], at least for the available data samples $\mathbf{y}_\ell^{(q)}$ for $\ell = 1, \dots, N$. Note that if $\mathbf{f}_{\text{NN}}^{(q)}$ is invertible, then there exists $\mathbf{g}_{\text{NN}}^{(q)}$ such that the second term is zero—but the converse is not necessarily true. This is because the second term is only imposed on the available finite samples in practice, not all the possible $\mathbf{y}_\ell^{(q)}$ ’s. Nonetheless, using such a data reconstruction to prevent $\mathbf{f}_{\text{NN}}^{(q)}$ from being an irreversible function in general works, especially when the number of samples is large. This idea is known as the *autoencoder* [53], which is considered a nonlinear counterpart of PCA. Using autoencoder to promote invertibility is a simple heuristic, and can often strike a good balance between effectiveness and computational complexity. The term $(1/N)\mathbf{U}\mathbf{U}^\top = \mathbf{I}$ corresponds to (10b).

We should highlight the constraint $(1/N) \sum_{\ell=1}^N \mathbf{u}_\ell = \mathbf{0}$, i.e., the extracted matrix \mathbf{U} has a zero column mean, which is enforced for extracting zero-mean latent components (cf. our discussions following Corollary 1). In addition, adding this constraint is in fact vital for avoiding numerical

problems. To see this subtle point, recall that according to Theorems 1 – 2, we have

$$\mathbf{f}_{\text{NN}}^{(q)}(\mathbf{y}_\ell^{(q)}) = \mathbf{D}^{(q)} \mathbf{A}^{(q)} \mathbf{s}_\ell^{(q)} + \mathbf{d}^{(q)}, \quad (15)$$

where $\mathbf{D}^{(q)} = \text{Diag}(\alpha_1^{(q)}, \dots, \alpha_M^{(q)})$. If we directly match $\mathbf{f}_{\text{NN}}^{(1)}(\mathbf{y}_\ell^{(1)})$ with $\mathbf{f}_{\text{NN}}^{(2)}(\mathbf{y}_\ell^{(2)})$, i.e., enforcing $\mathbf{B}^{(1)} \mathbf{f}_{\text{NN}}^{(1)}(\mathbf{y}_\ell^{(1)}) = \mathbf{B}^{(2)} \mathbf{f}_{\text{NN}}^{(2)}(\mathbf{y}_\ell^{(2)}) = \mathbf{u}_\ell$, then, one trivial solution is to simply making $\mathbf{D} \approx \mathbf{0}$ and $\mathbf{d}^{(1)} = \mathbf{d}^{(2)}$, with $\mathbf{B}^{(1)} = \mathbf{B}^{(2)}$ —i.e., the constant $\mathbf{d}^{(q)}$ can easily dominate. Hence, we enforce \mathbf{u}_ℓ to be zero mean, which will automatically take out $\mathbf{d}^{(q)}$.

Remark 2 We should mention that we do not constrain $\|\mathbf{B}^{(q)}\|_0 = M_q K$ in our working formulation since it seems that $\mathbf{B}^{(q)}$ with zero entries rarely happens if $\mathbf{B}^{(q)}$ is randomly initialized (it never happened in our extensive experiments). Hence, incorporating such a hard constraint just for ‘safety’ in theory may not be worthy—a lot more complex algorithms may be required for handling this constraint.

C. Proposed Algorithm

We propose a *block coordinate descent* (BCD)-based algorithm to handle Problem (14).

1) *The $(\theta_{\text{NN}}, \mathbf{B})$ -Subproblem:* We first consider the problem of updating the neural networks when fixing \mathbf{U} . This is an unconstrained optimization problems and can be handled by gradient descent. Denote the loss function as $\mathcal{L}(\theta, \mathbf{U}) = \sum_{\ell=1}^N \mathcal{L}_\ell(\theta, \mathbf{U})$, where θ collects $\theta_{\text{NN}}^{(q)}$ and $\mathbf{B}^{(q)}$ for $q = 1, 2$, and $\mathcal{L}_\ell(\theta, \mathbf{U}) = \sum_{q=1}^Q \|\mathbf{u}_\ell - \mathbf{B}^{(q)} \mathbf{f}_{\text{NN}}^{(q)}(\mathbf{y}_\ell^{(q)})\|_2^2 + \lambda \sum_{q=1}^Q \|\mathbf{y}_\ell^{(q)} - \mathbf{g}_{\text{NN}}^{(q)}(\mathbf{f}_{\text{NN}}^{(q)}(\mathbf{y}_\ell^{(q)}))\|_2^2$. At iteration t , the update rule is simply

$$\theta_{t+1} \leftarrow \theta_t - \gamma_t \nabla_\theta \mathcal{L}(\theta_t, \mathbf{U}_t). \quad (16)$$

where γ_t is the step size chosen for the t th update.

Note that computing $\nabla_\theta \mathcal{L}(\theta_t, \mathbf{U}_t)$ is normally not easy, since computing the gradient of neural networks is a resource-consuming process—the gradient normally requires *backpropagation* (BP) based algorithms to compute in a sample-by-sample manner if multiple layers are involved. Hence, instead of using the full gradient to update θ , one can also use stochastic gradient, i.e.,

$$\theta_{t+1} \leftarrow \theta_t - \gamma_t \sum_{\ell \in \mathcal{I}_t} \nabla_\theta \mathcal{L}_\ell(\theta_t, \mathbf{U}_t), \quad (17)$$

where \mathcal{I}_t is an index set randomly sampled at iteration t .

2) *\mathbf{U} -Subproblem:* To update \mathbf{U} , we solve the following subproblem:

$$\begin{aligned} \mathbf{U}_{t+1} &\leftarrow \arg \min_{\mathbf{U}} \sum_{q=1}^Q \sum_{\ell=1}^N \|\mathbf{u}_\ell - \mathbf{B}_{t+1}^{(q)} \mathbf{f}_{t+1}^{(q)}(\mathbf{y}_\ell^{(q)})\|_2^2 \\ &\text{subject to } (1/N) \mathbf{U} \mathbf{U}^\top = \mathbf{I}, \quad \frac{\mathbf{U} \mathbf{1}}{N} = \mathbf{0}, \end{aligned} \quad (18)$$

where $\mathbf{f}_{t+1}^{(q)}$ is $\mathbf{f}_{\text{NN}}^{(q)}$ parametrized by the corresponding network parameters from θ_{t+1} . This problem is seemingly difficult since it has two constraints, in which one is nonconvex. However, this problem turns out to have a semi-algebraic

solution. One can see that the \mathbf{U} -subproblem can be re-expressed as

$$\begin{aligned} &\underset{\mathbf{U}}{\text{minimize}} \left\| \mathbf{U} - \frac{1}{2} \sum_{q=1}^Q \mathbf{B}_{t+1}^{(q)} \mathbf{F}_{t+1}^{(q)} \right\|_F^2 \\ &\text{subject to } (1/N) \mathbf{U} \mathbf{U}^\top = \mathbf{I}, \quad \frac{\mathbf{U} \mathbf{1}}{N} = \mathbf{0}, \end{aligned} \quad (19)$$

where $\mathbf{F}_{t+1}^{(q)} = [\mathbf{f}_{t+1}^{(q)}(\mathbf{y}_1^{(q)}), \dots, \mathbf{f}_{t+1}^{(q)}(\mathbf{y}_N^{(q)})]$. The above can be shown by expanding (18):

$$\begin{aligned} \sum_{q=1}^Q \sum_{\ell=1}^N \|\mathbf{u}_\ell - \mathbf{B}_{t+1}^{(q)} \mathbf{f}_{t+1}^{(q)}(\mathbf{y}_\ell^{(q)})\|_2^2 &= \sum_{q=1}^Q \|\mathbf{U} - \mathbf{B}_{t+1}^{(q)} \mathbf{F}_{t+1}^{(q)}\|_F^2 \\ &= \sum_{q=1}^Q \left(\|\mathbf{U}\|_F^2 + \|\mathbf{B}_{t+1}^{(q)} \mathbf{F}_{t+1}^{(q)}\|_F^2 \right) - 2 \text{Tr}(\mathbf{U}^\top \mathbf{B}_{t+1}^{(q)} \mathbf{F}_{t+1}^{(q)}). \end{aligned}$$

Since the first and the last terms are constants, we have the following optimization problem:

$$\underset{\substack{(1/N) \mathbf{U} \mathbf{U}^\top = \mathbf{I} \\ \frac{\mathbf{U} \mathbf{1}}{N} = \mathbf{0}}}{\text{minimize}} \sum_{q=1}^Q -\text{Tr}(\mathbf{U}^\top \mathbf{B}_t^{(q)} \mathbf{F}_t^{(q)}),$$

which is equivalent to that in (19). Problem (19) is seemingly hard, since it has two constraints—and one constraint is nonconvex. Nonetheless, we show the above can be solved in closed form. To this end, we show the following Lemma:

Lemma 1 Consider the following optimization problem

$$\underset{\mathbf{U}}{\text{minimize}} \|\mathbf{U} - \mathbf{Z}\|_F^2 \quad (20a)$$

$$\text{subject to } (1/N) \mathbf{U} \mathbf{U}^\top = \mathbf{I}, \quad \frac{\mathbf{U} \mathbf{1}}{N} = \mathbf{0}. \quad (20b)$$

An optimal solution is $\mathbf{U}_* = \sqrt{N} \mathbf{P} \mathbf{Q}^\top$, where \mathbf{P} and \mathbf{Q} are left and right singular vectors of $\mathbf{Z} \mathbf{W}$ respectively, with $\mathbf{W} = \mathbf{I} - \frac{1}{N} \mathbf{1} \mathbf{1}^\top$.

Proof: The proof is relegated to Appendix E. ■

The overall algorithm is summarized in Algorithm 1. One can see that the algorithm does not have computationally heavy updates. The most resource-consuming step is the SVD used for updating \mathbf{U} . However, this step only takes $\mathcal{O}(NK^2)$ flops, which is still linear in the number of samples. Note that K is the number of shared components sought, which is often small in practice. A side note is that we also observe that one can update θ multiple times until switching to the next block—which often improves the speed of convergence.

Remark 3 We should mention that although our technical part was developed under $q = 1, 2$, the theorems and algorithm naturally hold when one has $Q \geq 3$ views. Another remark is that many off-the-shelf tricks for speeding up training neural networks, e.g., adaptive step size and momentum [56], can also be employed to determine $\gamma_{t,s}$, which can normally accelerate convergence.

Algorithm 1: Nonlinear Multiview Component Analysis (NMCA)

Data: Data $\mathbf{Y}^{(q)} \in \mathbb{R}^{M_q \times N}$, estimated dimension of the latent space and network structures.

Result: \mathbf{U} and $\boldsymbol{\theta}$

- 1 Initialize $\mathbf{U}_0 \leftarrow \sqrt{N} \mathbf{P} \mathbf{Q}^\top$ where $\mathbf{P} \mathbf{D} \mathbf{Q}^\top = (\sum_{q=1}^Q \mathbf{F}_{t+1}^{(q)}) \mathbf{W}$;
- 2 $t \leftarrow 1$;
- 3 **while** stopping criterion is not reached **do**
- 4 $\boldsymbol{\theta}_{s,0} \leftarrow \boldsymbol{\theta}_t$;
- 5 $s \leftarrow 1$;
- 6 **while** stopping criterion is not reached **do**
- 7 $\boldsymbol{\theta}_{t,s+1} \leftarrow \boldsymbol{\theta}_{t,s} - \gamma_{t,s} \sum_{\ell \in \mathcal{I}_{t,s}} \nabla_{\boldsymbol{\theta}} \mathbf{L}_\ell(\boldsymbol{\theta}_{t,s}, \mathbf{U}_t)$;
- 8 $s \leftarrow s + 1$;
- 9 **end**
- 10 $\boldsymbol{\theta}_{t+1} \leftarrow \boldsymbol{\theta}_{t,s}$;
- 11 $\mathbf{U}_{t+1} \leftarrow \sqrt{N} \mathbf{P} \mathbf{Q}^\top$ where $\mathbf{P} \mathbf{D} \mathbf{Q}^\top = (\sum_{q=1}^Q \mathbf{F}_{t+1}^{(q)}) \mathbf{W}$;
- 12 $t \leftarrow t + 1$;
- 13 **end**

Remark 4 In the proposed approach, we use one neural network to parameterize an $f_{\text{NN},m}^{(q)}$ or $g_{\text{NN},m}^{(q)}$. Hence, M_q individual networks are used to approximate $\mathbf{f}_{\text{NN}}^{(q)}$. This follows our PNL-based nonlinear signal model. In practice, if the number of channels (M_q) is large, such a parametrization could be costly and adds considerable memory and computational burdens to the implementation. One workaround is to use a single fully connected, multiple-input-multiple-output (MIMO) neural network to approximate $\mathbf{f}_{\text{NN}}^{(q)} : \mathbb{R}^{M_q} \rightarrow \mathbb{R}^{M_q}$ (and the same applies to $g_{\text{NN}}^{(q)}$). This effectively means that all the individual mappings $f_{\text{NN},i}^{(q)}$ for $i = 1, \dots, M_q$ share neurons in the parameterization (as is opposed to each $f_{\text{NN},i}^{(q)}$ using an individual network). Using a single fully connected network can reduce the computational burden substantially. It also sometimes outputs better results, perhaps because the associated optimization problem is better solved. Using a single MIMO network is a valid approximation for our model in (12) since the MIMO neural network can also approximate any MIMO continuous function in principle; see [57, Proposition 1], which is a simple extension of the universal approximation theory for multiple-input-single-output (MISO) functions [55].

D. Connection to Deep CCA Approaches

An interesting observation is that, although started from very different perspectives, the proposed formulation in (14) and the line of work, namely, deep CCA [22], [23], [58] end up with similar formulations. In particular, the deep canonically correlated autoencoder (DCCAE) [22] formulation is as follows:

$$\begin{aligned} & \underset{\boldsymbol{\theta}_{\text{NN}}, \mathbf{B}^{(q)}}{\text{minimize}} \sum_{\ell=1}^N \left\| \mathbf{B}^{(1)} \mathbf{f}_{\text{NN}}^{(1)} \left(\mathbf{y}_\ell^{(1)} \right) - \mathbf{B}^{(2)} \mathbf{f}_{\text{NN}}^{(2)} \left(\mathbf{y}_\ell^{(2)} \right) \right\|_2^2 \\ & \quad + \lambda \sum_{q=1}^2 \sum_{\ell=1}^N \left\| \mathbf{y}_\ell^{(q)} - \mathbf{g}_{\text{NN}}^{(q)} \left(\mathbf{B}^{(q)} \mathbf{f}_{\text{NN}}^{(q)} \left(\mathbf{y}_\ell^{(q)} \right) \right) \right\|_2^2 \\ & \text{subject to } \frac{1}{N} \sum_{\ell=1}^N \left(\mathbf{B}^{(q)} \mathbf{f}_{\text{NN}}^{(q)} \left(\mathbf{y}_\ell^{(q)} \right) \mathbf{f}_{\text{NN}}^{(q)} \left(\mathbf{y}_\ell^{(q)} \right)^\top (\mathbf{B}^{(q)})^\top \right) \\ & \quad = \mathbf{I}, \quad q = 1, 2. \end{aligned} \quad (21)$$

If one changes the second term in (14) from using $\mathbf{g}_{\text{NN}}^{(q)}(\mathbf{f}_{\text{NN}}^{(q)}(\mathbf{y}_\ell^{(q)}))$ to $\mathbf{g}_{\text{NN}}^{(q)}(\mathbf{B}^{(q)} \mathbf{f}_{\text{NN}}^{(q)}(\mathbf{y}_\ell^{(q)}))$ for reconstruction, and removing the zero-mean constraint in (14), the two formulations are conceptually equivalent. The seemingly subtle differences, however, are quite vital—which make DCCAE inappropriate for the model in (8). First, the reconstruction term in DCCAE’s cost function (i.e., the second term) is based on reconstructing $\mathbf{y}^{(q)}$ from $\mathbf{B}^{(q)} \mathbf{f}_{\text{NN}}^{(q)}$. This could be problematic under the model of interest in (8), since the $\mathbf{c}_\ell^{(q)}$ part can be nullified by $\mathbf{B}^{(q)} \mathbf{f}_{\text{NN}}^{(q)}$ —which means that such reconstruction can never be attained, making the training objective awkward. In addition, perhaps more importantly, the learned signals $\mathbf{B}^{(q)} \mathbf{f}_{\text{NN}}^{(q)}$ in DCCAE are not enforced to have zero mean. As we have seen in our analysis, nonlinear removal can only be done up to an affine transformation. Without the zero-mean constraint, trivial or numerically pathological solutions could easily happen [cf. the discussion in Eq. (15)]. We will also show this effect in the simulation.

Similarly, if one uses $Q \geq 3$ and removes the second term and the zero-mean constraints in our formulation, the formulation becomes the same as the generalized GCCA criterion in [58]. In particular, the GCCA work in [58] also employs a slack variable \mathbf{U} to represent the shared latent components as in our work. However, as we have seen in the analysis, both the reconstruction (for invertibility) and the zero-mean (for taking away constants) parts are quite critical for avoiding trivial solutions.

V. NUMERICAL RESULTS

In this section, we apply the proposed algorithm to a number of synthetic-data and real-data experiments to showcase the effectiveness of our method and to validate our analysis.

A. Setup

Baselines We use a number of baselines throughout this section:

- **Principal Component Analysis (PCA)** [1]: PCA is widely used for feature extraction and dimensionality reduction. The first principal component has the largest variance.
- **Canonical Correlation Analysis (CCA)** [5]: CCA finds linear transformations of the two views that exhibit maximum correlation with each other.
- **Kernel CCA (KCCA)** [49]: KCCA matches the two views in a high dimensional feature space via certain nonlinear kernel functions. The baseline from [49] is a scalable version of KCCA.
- **Deep CCA (DCCA)** [23]: DCCA employs deep neural networks to represent the nonlinear transformations of the two views.
- **Deep Canonically Correlated AutoEncoders (DCCAE)** [22]: DCCAE utilizes autoencoder on top of DCCA, to avoid trivial solutions.

Performance Metric: Per Corollary 1, the nonlinear multiview analysis recovers the range space of \mathbf{S}^\top , where $\mathbf{S} = [s_1, \dots, s_N]$. Therefore, we employ the following *subspace distance* measure $\text{dist}(\mathcal{S}, \mathcal{U}) = \|\mathbf{P}_s^\perp \mathbf{Q}_u^\top\|_2$ as the performance metric, where $\mathcal{S} = \text{range}(\mathbf{S}^\top)$ and $\mathcal{U} = \text{range}(\mathbf{U}^\top)$, \mathbf{P}_s^\perp

is defined as $\mathbf{P}_s^\perp = \mathbf{I} - \mathbf{S}^\top(\mathbf{S}\mathbf{S}^\top)^{-1}\mathbf{S}$, and \mathbf{Q}_u is the orthogonal basis of \mathbf{U} . Note that $\|\mathbf{X}\|_2$ here denotes the matrix 2-norm, i.e., the largest singular value of \mathbf{X} . The metric is bounded within $[0, 1]$. If $\mathbf{U} \approx \Theta\mathbf{S}$, then $\|\mathbf{P}_s^\perp\mathbf{Q}_u^\top\|_2$ should be close to zero.

Algorithm 1 is implemented as follows. The θ update is implemented with multiple updates in each iteration (i.e., 100 times for each fixed \mathbf{U}_t). The batch size $|\mathcal{I}_t|$ is 1,000. The Adam algorithm [56] is employed for updating θ . The initial step size is set to be 10^{-3} in all the experiments. The algorithm is stopped after 10,000 epochs in total. Our algorithm is implemented using PyTorch.

B. Synthetic-Data Experiments

In the first experiment, we construct the views as follows. The shared components $\mathbf{S} = [s_1, \dots, s_N] \in \mathbb{R}^{2 \times N}$ are sampled from a parabola [i.e., $s_{2,\ell} = (s_{1,\ell})^2$, where $s_{1,\ell} \in (-1, 1)$]. In this case, the shared components are completely dependent. Besides, the sample mean is subtracted. The shared components $s_\ell \in \mathbb{R}^2$ for $\ell = 1, \dots, N$ are shown in Fig. 1. View-specific components $\mathbf{c}_\ell^{(q)} \in \mathbb{R}^3$ for $q = 1, 2$ are set to be i.i.d. Gaussian components with means -0.5 and 0.8 , respectively, and variances 1.0^2 and 1.5^2 , respectively. Hence, we have $M_1 = 5$ and $M_2 = 5$. Note that under this setting, the conditions in Theorem 3 are satisfied. The sample size for each view is $N = 1,000$. The elements of the mixing matrices $\mathbf{A}^{(1)} \in \mathbb{R}^{5 \times 5}$, $\mathbf{A}^{(2)} \in \mathbb{R}^{5 \times 5}$ follow zero-mean unit-variance i.i.d. Gaussian distribution. The nonlinear functions employed are as follows: $g_1^{(1)}(x) = 3\text{sigmoid}(x) + 0.1x$, $g_2^{(1)}(x) = 5\text{sigmoid}(x) + 0.2x$, $g_3^{(1)}(x) = 0.2\exp(x)$, $g_4^{(1)}(x) = -4\text{sigmoid}(x) - 0.3x$, $g_5^{(1)}(x) = -3\text{sigmoid}(x) - 0.2x$; $g_1^{(2)}(x) = 5\tanh(x) + 0.2x$, $g_2^{(2)}(x) = 2\tanh(x) + 0.1x$, $g_3^{(2)}(x) = 0.1x^3 + x$, $g_4^{(2)}(x) = -5\tanh(x) - 0.4x$, $g_5^{(2)}(x) = -6\tanh(x) - 0.3x$. Fig. 2 visualizes the observed views (via t-SNE [59]). One can see the views are severely distorted. In the experiment, our method uses 5 individual neural networks in order to learn the inverse of the 5 nonlinear mappings from $g_1^{(q)}$ to $g_5^{(q)}$. Every network has one-hidden-layer with 256 neurons. We adopt the ReLU activation function in the networks.

The learned $\hat{\mathbf{x}}_\ell = \mathbf{B}^{(q)}\mathbf{f}_{\text{NN}}^{(q)}(\mathbf{y}_\ell^{(q)})$ by the proposed approach and the outputs of the baselines are visualized in Fig. 3. One can see that the proposed approach yields a scaled and permuted version of the parabola in Fig. 1—which is as expected. The baselines are not as promising: DCCA and DCCAE also work reasonably—but the scatter plots are much noisier. The linear mapping based methods PCA and CCA clearly fail. One particular observation is that if we do not impose the zero-mean constraint in (14), the result is much worse compared to the one with this constraint. This supports our claim that trivial or noninteresting solutions may easily happen if the zero-mean constraint is not imposed.

Fig. 4 shows the learned $\hat{f}_m^{(q)} \circ g_m^{(q)}$ for $q = 1, 2$ with 3 randomly selected dimensions. One can see that in this simulation, all the composition functions are visually affine.

Table I shows the averaged dist from 10 random trials under the same settings as those for in Figs. 3-4. One can see

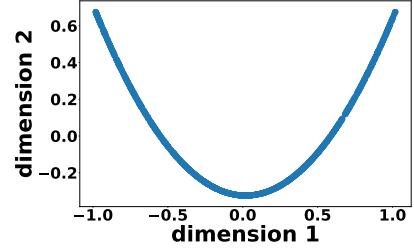


Fig. 1. The shared 2-dimensional components.

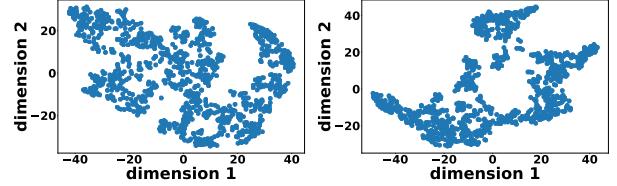


Fig. 2. Visualization of view 1 (left) and view 2 (right).

TABLE I
THE dist VALUE [MEAN (STANDARD DEVIATION)] OF THE APPROACHES UNDER TEST.

method	Proposed	Proposed w/o zero-mean	DCCA
dist	0.035(8.95e ⁻³)	0.998(7.35e ⁻⁴)	0.523(1.51e ⁻¹)
method	DCCAE	KCCA	CCA
dist	0.530(1.63e ⁻¹)	0.998(4.57e ⁻³)	0.993(8.56e ⁻³)
method	PCA		
dist	0.999(1.52e ⁻³)		

TABLE II
THE dist VALUE [MEAN (STANDARD DEVIATION)] OF THE APPROACHES AGAINST NUMBER OF SAMPLES.

# samples	500	1000	2000	5000	10000
Proposed	0.101(1.65e ⁻²)	0.035(8.95e ⁻³)	0.030(3.61e ⁻³)	0.017(2.00e ⁻³)	0.009(1.60e ⁻³)
DCCA	0.799(3.31e ⁻²)	0.523(1.51e ⁻¹)	0.316(3.85e ⁻²)	0.165(9.82e ⁻³)	0.129(1.02e ⁻²)
DCCAE	0.817(3.31e ⁻²)	0.530(1.63e ⁻¹)	0.310(3.39e ⁻²)	0.183(1.21e ⁻²)	0.118(1.11e ⁻²)
KCCA	0.998(1.70e ⁻³)	0.998(4.57e ⁻³)	0.999(5.01e ⁻³)	0.989(5.10e ⁻³)	0.986(1.50e ⁻²)
CCA	0.987(5.32e ⁻³)	0.993(8.56e ⁻³)	0.988(3.67e ⁻³)	0.998(7.23e ⁻³)	0.978(2.36e ⁻²)
PCA	0.987(4.12e ⁻²)	0.999(1.52e ⁻³)	0.993(3.37e ⁻²)	0.998(7.33e ⁻³)	0.992(1.27e ⁻²)

that the proposed approach admits a dist value that is almost zero, validating our subspace identifiability claim. It is much lower than those of the baselines, perhaps because our method exploits model information. Again, subtracting the mean in our formulation [cf. the constraint $(1/N) \sum_{\ell=1}^N \mathbf{u}_\ell = \mathbf{0}$] is indeed crucial, since the result without this constraint is much worse.

Table II shows the performance of the algorithms when the number of samples N changes. Again, we have $s_\ell \in \mathbb{R}^2$ and each view has three interference components in this simulation. Note that the nonlinearity removal theorems are derived under the population case, i.e., the functional equations in (12b) hold for all $s_\ell \in \mathcal{S}$ and $\mathbf{c}_\ell^{(q)} \in \mathcal{C}_q$. This is not attainable in practice, where we only have a finite number of samples. Hence, it is of interest to observe how the performance scales with N . From Table II one can see that the dist value of the proposed approach clearly decreases when the number of samples increases from 500 to 10,000. On the other hand, using 500 samples, the proposed approach already exhibits good performance.

To observe the impact of the view-specific interference, we define the *Shared Component to Interference Ratio* (SCIR)

$$\text{SCIR} = 10 \log_{10} \left(\frac{\|\mathbf{S}\|_F^2 / K}{\frac{1}{Q} \sum_{q=1}^Q \|\mathbf{C}^{(q)}\|_F^2 / R_q} \right) \text{ dB.}$$

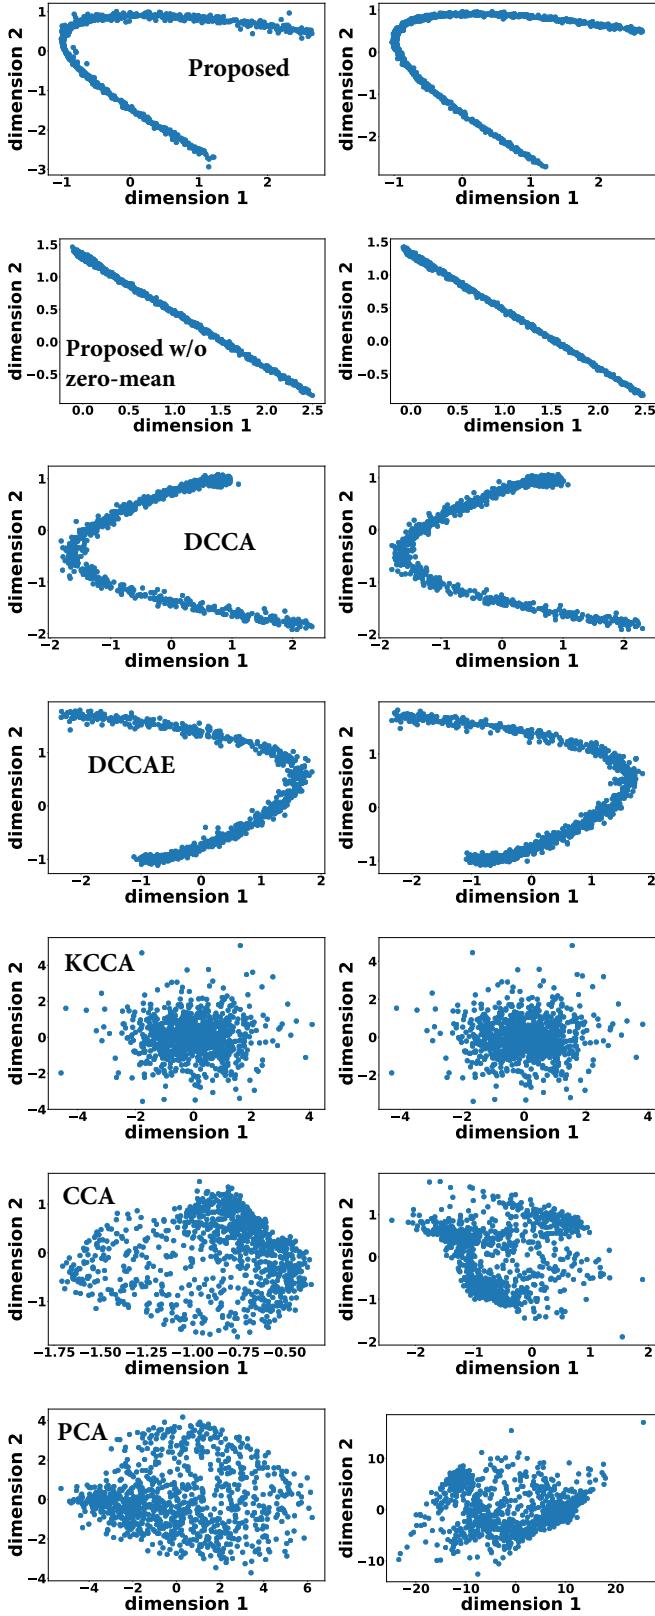


Fig. 3. Example of the recovered shared component. Top to bottom: Proposed, Proposed w/o zero-mean constraint, DCCA, DCCAE, KCCA, CCA and PCA.

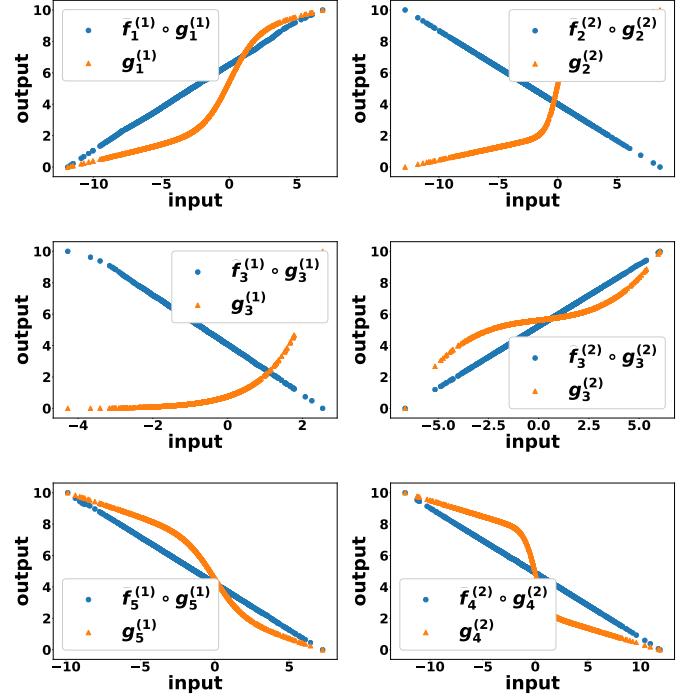


Fig. 4. The nonlinear distortion and the learned composite functions of two views. Left: view 1; Right: view 2. From top to bottom: randomly selected 3 dimensions.

TABLE III
THE dist VALUE [MEAN (STANDARD DEVIATION)] OF THE APPROACHES SCIR.

SCIR	-10 dB	-20 dB
Proposed	0.035($8.95e^{-3}$)	0.064($3.40e^{-2}$)
DCCA	0.523($1.51e^{-1}$)	0.999($2.16e^{-3}$)
DCCAE	0.530($1.63e^{-1}$)	0.998($1.71e^{-3}$)
KCCA	0.998($4.57e^{-3}$)	0.999($2.32e^{-3}$)
CCA	0.993($8.56e^{-3}$)	0.998($1.09e^{-3}$)
PCA	0.999($1.52e^{-3}$)	0.997($2.13e^{-3}$)

TABLE IV
THE dist VALUE [MEAN (STANDARD DEVIATION)] UNDER DIFFERENT STRUCTURES OF NEURAL NETWORKS.

# params of each dim	structure	dist
≈ 512	[1,256,1]	0.035($8.95e^{-3}$)
	[1,16,16,1]	0.058($9.32e^{-3}$)
≈ 2048	[1,1024,1]	0.039($1.05e^{-2}$)
	[1,32,32,32,1]	0.048($7.68e^{-3}$)
≈ 8192	[1,4096,1]	0.053($2.13e^{-2}$)
	[1,64,64,64,1]	0.058($2.25e^{-2}$)

Table III shows the result under different SCIRs. Other settings are the same as those in the previous simulation. One can see that even if the ratio is -20 dB, the performance of the proposed approach is still very good—the average subspace distance metric is 0.064, whereas all the other baselines fail to work. This observation is consistent with our analysis, as well as the interference-robust property of linear CCA [7].

Table IV shows the dist value against different network structures. Observing this is of interest, since in machine learning tuning the structure of networks to fit a task can sometimes be tedious. The proposed framework is unsupervised, which means that there may not even be training samples for us to tune the network structure. Fortunately, the observation here is that the result is not heavily affected by the choice of neural

TABLE V
DIFFERENT DIMENSIONS OF SOURCE COMPONENTS

source dim	1	2	3	4
dist	0.032(1.00e ⁻³)	0.016(2.56e ⁻³)	0.015(2.62e ⁻³)	0.015(3.34e ⁻³)

network structures—although using more parameters and more layers could slightly improve performance.

In Fig. 5, we test to what extent can the nonlinearity be solved by the proposed model. Specifically, we set the nonlinear function $g_i^{(q)}(x) = \alpha x^3 + x$, then we try different values of α , i.e., $\alpha = 0.01, 0.05, 0.1, 0.5$, and 1 . In this simulation, the shared components are zero-mean i.i.d. Gaussian variables and $K = 3$, and the view-specific components for both views also follow this distribution with $R_1 = R_2 = 2$. One can see that when α increases, the performance of the proposed approach indeed deteriorates, given the same computational resources, as expected. Nonetheless, even when $\alpha = 0.5$, one can see that the nonlinearity imposed on the views is largely removed.

In Table V, we present the results of the algorithms under different K , i.e., the dimension of the shared components. We fix the observation dimension to be 5 and the K ranges from 1 to 4. One can see that different K settings essentially do not affect the results—which is consistent with CCA identifiability analyses in [7] and this work.

C. Real Data Experiments

In this section, we test the proposed algorithm on two real datasets, namely, an EEG/MEG dataset and a multiview handwritten digits dataset.

1) *EEG/MEG Dataset*: This dataset is called Multi-modal Face Dataset². It contains EEG and MEG measurements data of the same subjects when they are shown with real faces and scrambled faces. The task is to train a classifier using the EEG and MEG measurements to predict if the human subject sees a real face or a scrambled face.

The feature sizes of an EEG sample and an MEG sample are $M_1 = 130$ and $M_2 = 151$, respectively. There are 172 events during the experiment, among which 86 are real faces and the rest are scrambled faces. For each event, 161 snapshots are recorded for both EEG and MEG, respectively, resulting in $172 \times 161 = 27,692$ samples for each view. One could use the raw data to train classifiers, but the raw EEG/MEG signals are noisy, and are considered mixtures of components [44], [45]. Also, nonlinear distortions have likely happened in the sensing process. This makes the dataset suitable for evaluating the proposed approach.

We train two linear classifiers, namely, the linear support vector machines (SVM) and the logistic regression classifier, based on the learned low-dimensional representations u_ℓ for $\ell = 1, \dots, N$. The whole dataset is split into a training set, a validation set and a testing set. The validation set is employed to stop the learning procedures for algorithms. It is used for training the linear classifiers as well.

Since the dimension of the input samples are large, we adopt fully connected networks to parameterize the nonlinear

TABLE VI
CLASSIFICATION ERROR RATE (%) ON EEG/MEG DATA.

train/validation/test		Raw	PCA	CCA	KCCA	DCCA	DCCAE	Proposed
90%/5%/5%	SVM	45.79	47.07	47.59	43.65	40.55	41.40	28.91
	Logistic	46.58	46.48	46.84	43.46	38.61	38.58	28.44
80%/5%/15%	SVM	43.70	41.47	40.79	39.99	39.43	40.06	32.20
	Logistic	41.73	42.29	41.30	40.67	40.47	41.01	31.19

mappings, as we discussed in Remark 4. We use a three-layer network for $f_{NN}^{(q)}$, each having 256 neurons. We use a single-layer network for $g_{NN}^{(q)}$, which also has 256 neurons in the hidden layer. We set $\lambda = 10^{-6}$. Dropout [60] is also used in optimization. In this experiment, we use $K = 5$ for our approach. This parameter is selected over $\{2, 5, 10, 20, 50\}$ by tuning on the validation set. For PCA, the number of components is automatically selected so that 90% energy is retrained. For DCCA, the network structure is the same as ours for $f_{NN}^{(q)}$. DCCAE shares the same network structure of ours. All the results are averaged over 5 different random splits.

The results can be found in Table VI. One can see that using nonlinear dimensionality reduction techniques indeed improves upon the performance of classic techniques, which proves the necessity for taking nonlinear distortion into consideration. In addition, the proposed approach is the most promising one among the baselines. This is perhaps because the biological signal sensing problem fits our model well.

2) *Multiview Digit Dataset*: This dataset from [61] consists of features of handwritten digits (0 to 9) extracted from a collection of Dutch utility maps. There are 200 patterns per class (2,000 patterns in total) that have been digitized in binary images. These digits are represented in different views.

We employ three views for this experiment. The three views consist of 76 Fourier coefficients of the character shapes, 64 Karhunen-Loëve coefficients, and 47 Zernike moments, respectively. The training set contains 1,200 samples, and the validation and testing sets both have 400 samples. To validate the performance of the learned representation, we use spectral clustering [62] as the downstream task tool. After the models are learned, the testing set is fed to it and then clustering is performed on the output of the model.

All the settings are identical to those in the previous experiment. The latent dimension is selected over $\{2, 5, 10, 20, 50\}$ by tuning on the validation set. Similarly, λ is chosen over $\{1e^{-3}, 1e^{-4}, 1e^{-5}, 1e^{-6}\}$. The results are averaged over 5 different random splits.

The clustering accuracy results on the testing set are shown in Table. VII. In the first three columns, we test the algorithms using two of the three views. In the last column, we test our algorithm (since our approach can naturally handle more than two views) and GCCA on all of the three views. Here, we employ the SUMCOR GCCA formulation and the algorithm in [63]. Since KCCA, DCCA and DCCAE can only deal with two views, they do not have results in the last column. The observation is a bit different to what we have seen for the EEG/MEG experiment. The classic methods sometimes outperform the nonlinear baselines. This may be due to the fact that purely data-driven models may require more data to train. The proposed approach again shows promising results. Note that although our approach has a similar implementation

²<https://www.fil.ion.ucl.ac.uk/spm/data/mffaces/>

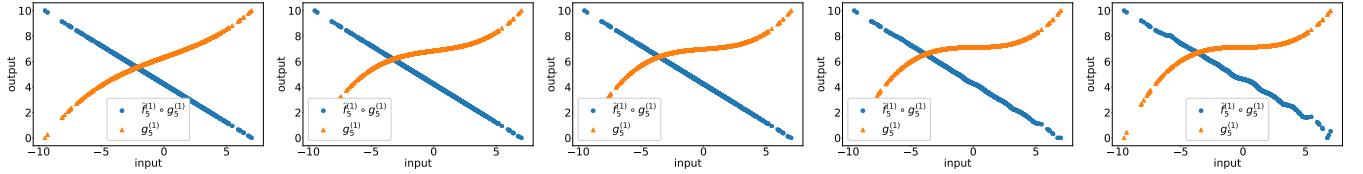


Fig. 5. The nonlinear distortion and the learned composite functions of one example dimension. From left to right: α is 0.01, 0.05, 0.1, 0.5, 1 and the corresponding recovered subspace metric is 0.011, 0.014, 0.019, 0.158, 0.345

TABLE VII
SPECTRAL CLUSTERING ACC (%) OF DIFFERENT ALGORITHMS.

	View1	View2	View1	View3	View2	View3	View1	View2	View3
Raw	51.55	64.15	51.55	50.05	64.15	50.05	51.55	64.15	50.05
PCA	52.10	65.75	52.10	51.05	65.75	51.05	52.10	65.75	51.05
(G)CCA	69.45	65.55	72.70	68.05	68.60	64.15	62.85	71.50	66.20
KCCA	71.80	71.65	73.55	72.80	63.65	59.45	-	-	-
DCCA	72.90	72.10	69.70	72.00	60.55	62.15	-	-	-
DCCAE	69.80	71.30	67.35	68.80	62.15	59.20	-	-	-
Proposed	73.40	72.90	71.15	71.05	70.35	66.15	74.20	77.05	73.60

as DCCA and DCCAE, the subtle differences exist because we utilize the information of the underlying signal model. Hence, our method can be understood as a mixed model and data-driven approach. It appears that using model information could help in practice, especially when the amount of data is not very large.

VI. CONCLUSION

In this work, we revisited the multiview analysis problem under a nonlinear setting. We showed that even under a class of unknown nonlinear distortions, the appealing identifiability guarantees for the classic linear multiview analysis still hold—if a proper identification criterion is used. We also proposed a neural network-based implementation for the proposed criterion, as well as a BCD-based optimization algorithm. A suite of numerical results were presented to corroborate our analysis. The proposed work also offers theoretical explanations to popular deep learning based canonical correlation analysis schemes, e.g., DCCAE and deep GCCA, from a model-based and identifiability-driven point of view.

Several future directions are in order: First, it is of great interest to study a model where different views do not share exactly the same latent components, while having similar ones. Such a model may be closer to real-world data. Second, offering sample complexity for nonlinear multiview learning is also very meaningful. Third, it is well-motivated to study the shared component identifiability under more general nonlinear distortion, since PNL is a special class of nonlinear model. In some case, e.g., image classification problems as in our last experiment, the PNL model may not be the most accurate. Using a more general nonlinear model may lead to better performance.

APPENDIX A PROOF OF THEOREM 1

Let us denote $\mathbf{h}^{(q)} = \mathbf{f}^{(q)} \circ \mathbf{g}^{(q)} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ for $q = 1, 2$ as the composition of the learned $\mathbf{f}^{(q)}$ and $\mathbf{g}^{(q)}$ in the generative model. Hence, from Eq. (10), we have

$$\begin{aligned} & b_1^{(1)} h_1^{(1)} \left(A_{11}^{(1)} s + A_{12}^{(1)} c^{(1)} \right) + b_2^{(1)} h_2^{(1)} \left(A_{21}^{(1)} s + A_{22}^{(1)} c^{(1)} \right) \\ & = b_1^{(2)} h_1^{(2)} \left(A_{11}^{(2)} s + A_{12}^{(2)} c^{(2)} \right) + b_2^{(2)} h_2^{(2)} \left(A_{21}^{(2)} s + A_{22}^{(2)} c^{(2)} \right) \end{aligned}$$

where we have omitted the subscript “ ℓ ” for notational simplicity. Note that due to (12d), the composition $h_i^{(q)} \neq 0$ —and thus such trivial solutions cannot happen.

Denote $\phi(c^{(1)}, c^{(2)}, s)$ as

$$\begin{aligned} \phi(c^{(1)}, c^{(2)}, s) &= \sum_{i=1}^2 b_i^{(1)} h_i^{(1)} \left(A_{i1}^{(1)} s + A_{i2}^{(1)} c^{(1)} \right) \\ &\quad - \sum_{j=1}^2 b_j^{(2)} h_j^{(2)} \left(A_{j1}^{(2)} s + A_{j2}^{(2)} c^{(2)} \right) = 0. \end{aligned} \quad (22)$$

To proceed, taking derivative w.r.t. $c^{(1)}$, we have $\frac{\partial \phi(c^{(1)}, c^{(2)}, s)}{\partial c^{(1)}} = 0$, which means

$$b_1^{(1)} A_{12}^{(1)} (h_1^{(1)})' + b_2^{(1)} A_{22}^{(1)} (h_2^{(1)})' = 0.$$

Note that we have used the fact that $\partial c^{(2)}/\partial c^{(1)} = 0$ and $\partial s/\partial c^{(1)} = 0$. To see this, first notice that the equality $\phi(c^{(1)}, c^{(2)}, s) = 0$ holds for all possible values of $[s, c^{(1)}, c^{(2)}]^\top$. By the ubiquitously unanchored condition on $[s, c^{(1)}, c^{(2)}]^\top$, it implies that for any $\bar{c}^{(2)}, \bar{s}$, all possible $c^{(1)}$ ’s satisfy the above equality—i.e., the *changing* of $c^{(1)}$ will not affect any $\bar{c}^{(2)}, \bar{s}$ satisfying the equality. Hence, even if $c^{(2)}$ and s are functions of $c^{(1)}$ (since they could be dependent), we still always have $\partial c^{(2)}/\partial c^{(1)} = 0$ and $\partial s/\partial c^{(1)} = 0$ (see more details in the supplementary material).

By further taking the second-order derivatives and invoking the ubiquitously unanchored assumption again, we have $\frac{\partial^2 \phi(c^{(1)}, c^{(2)}, s)}{\partial (c^{(1)})^2} = 0$ and $\frac{\partial^2 \phi(c^{(1)}, c^{(2)}, s)}{\partial c^{(1)} \partial s} = 0$, implying

$$\begin{aligned} & b_1^{(1)} (A_{12}^{(1)})^2 (h_1^{(1)})'' + b_2^{(1)} (A_{22}^{(1)})^2 (h_2^{(1)})'' = 0, \\ & b_1^{(1)} A_{11}^{(1)} A_{12}^{(1)} (h_1^{(1)})'' + b_2^{(1)} A_{21}^{(1)} A_{22}^{(1)} (h_2^{(1)})'' = 0. \end{aligned}$$

From the above, one can establish the following system of linear equations:

$$\underbrace{\begin{bmatrix} b_1^{(1)} A_{11}^{(1)} A_{12}^{(1)} & b_2^{(1)} A_{21}^{(1)} A_{22}^{(1)} \\ b_1^{(1)} (A_{12}^{(1)})^2 & b_2^{(1)} (A_{22}^{(1)})^2 \end{bmatrix}}_{\mathbf{H}} \begin{bmatrix} (h_1^{(1)})'' \\ (h_2^{(1)})'' \end{bmatrix} = \mathbf{0}$$

In order to prove that $h_i^{(1)}$ ’s are affine functions, we need to show that $(h_i^{(1)})'' = 0$. In fact, it is sufficient to show that the left matrix \mathbf{H} has full column rank. Since we can re-express \mathbf{H} in the following form

$$\mathbf{H} = \begin{bmatrix} A_{11}^{(1)} & A_{12}^{(1)} \\ A_{21}^{(1)} & A_{22}^{(1)} \end{bmatrix}^\top \begin{bmatrix} b_1^{(1)} A_{12}^{(1)} & 0 \\ 0 & b_2^{(1)} A_{22}^{(1)} \end{bmatrix}, \quad (23)$$

it is readily seen that $\text{rank}(\mathbf{H}) = 2$. Indeed, for the left matrix, it is full rank almost surely since $\mathbf{A}^{(1)}$ is drawn from an absolutely continuous distribution. By the same reason,

neither $A_{12}^{(1)}$ nor $A_{22}^{(1)}$ is zero with probability one. Hence, if $\|\mathbf{b}^{(1)}\|_0 = 2$, \mathbf{H} has to be full rank.

To show $\|\mathbf{b}^{(1)}\|_0 = 2$, suppose that, on the contrary, $b_1^{(1)} = 0$. One can see that on the left hand side of (12b) we have $\text{LHS} = b_2^{(1)}h_2^{(1)}\left(A_{21}^{(1)}s + A_{22}^{(1)}c^{(1)}\right)$ while the right-hand side is a function of s and $c^{(2)}$. Note that $h_2^{(1)} = f_2^{(1)} \circ g_2^{(1)}$ is invertible. Hence, LHS has to be a function of s and $c^{(1)}$. Otherwise, $b_2^{(1)}h_2^{(1)}$ cannot eliminate the contribution from $c^{(1)}$, which is a contradiction to the fact that (12b) holds. Therefore, the product of these two matrices on the right hand side of Eq. (23) is still full rank which leads to the fact that $(\mathbf{h}^{(1)})'' = \mathbf{0}$. By symmetry, one can readily show that $(h_i^{(2)})'' = 0$.

As a final step, we also need to show that $\alpha_i^{(q)} \neq 0$. This is trivial since (12c) holds. That is, $h_i^{(q)}(x) = f_i^{(q)} \circ g_i^{(q)}(x)$ is a composition of two invertible function—thereby being invertible itself. This excludes the possibility of $\alpha_i^{(q)}$ being zero—which would have made $h_i^{(q)}(\cdot)$ non-invertible.

APPENDIX B PROOF OF PROPOSITION 1

To show this proposition, it suffices to construct a solution such that $\mathbf{B}^{(q)}$ has no zero elements. Let $\mathbf{f}^{(q)} \circ \mathbf{g}^{(q)} = \mathbf{I}$. Then, one solution of $\mathbf{B}^{(q)}$ is $\mathbf{B}^{(q)} = [\Theta, \mathbf{0}](\mathbf{A}^{(q)})^\dagger$. The above solution is nothing but selecting out some columns of $(\mathbf{A}^{(q)})^\dagger$. Recall that we have assumed that $\mathbf{A}^{(q)}$ is drawn from a certain joint absolutely continuous distribution, which means that $(\mathbf{A}^{(q)})^\dagger$ also follows a certain absolutely continuous distribution [64]—thereby the elements of $\mathbf{B}^{(q)}$. Hence, we have $\text{Prob}(\mathbf{B}^{(q)}(i, j) = 0) = 0$.

APPENDIX C PROOF OF THEOREM 2

We show the case that $M_1 = K + R_1$ and $M_2 = K + R_2$, i.e., the mixing matrices $\mathbf{A}^{(1)}$ and $\mathbf{A}^{(2)}$ are square nonsingular matrices. The case where $M_q > K + R_q$ can then be shown via a straightforward approach: The M_q observations output by $\mathbf{g}^{(q)}(\cdot)$ can be divided to (possibly overlapped) groups, each of which has $K + R_q$ channels.

It suffices to show that the nonlinearity can be removed by considering any row of $(\mathbf{B}^{(q)})^\top$, denoted as $(\mathbf{b}^{(q)})^\top$. In this case, we consider only extracting one shared component. However, since we further restricted that the extracted components are uncorrelated [cf. the constraint in Eq. (12d)], the K extracted components by $\mathbf{B}^{(q)}$ will not have repeated ones.

From Eq. (10), we have the following equation

$$\begin{aligned} & \sum_{m=1}^{M_1} b_m^{(1)}h_m^{(1)} \left(\sum_{k=1}^K A_{m,k}^{(1)}s_k + \sum_{r=1}^{R_1} A_{m,K+r}^{(1)}c_r^{(1)} \right) \\ &= \sum_{j=1}^{M_2} b_j^{(2)}h_j^{(2)} \left(\sum_{l=1}^K A_{j,l}^{(2)}s_l + \sum_{s=1}^{R_2} A_{j,K+s}^{(2)}c_s^{(2)} \right). \end{aligned}$$

Denote $\phi(\mathbf{c}^{(1)}, \mathbf{c}^{(2)}, \mathbf{s})$ as

$$\begin{aligned} \phi(\mathbf{c}^{(1)}, \mathbf{c}^{(2)}, \mathbf{s}) = & \sum_{m=1}^{M_1} b_m^{(1)}h_m^{(1)} \left(\sum_{k=1}^K A_{m,k}^{(1)}s_k + \sum_{k=1}^{R_1} A_{m,K+k}^{(1)}c_k^{(1)} \right) \\ & - \sum_{j=1}^{M_2} b_j^{(2)}h_j^{(2)} \left(\sum_{l=1}^K A_{j,l}^{(2)}s_l + \sum_{l=1}^{R_2} A_{j,K+l}^{(2)}c_l^{(2)} \right) = 0. \end{aligned}$$

Let us first consider $\mathbf{h}^{(1)}$. Taking derivatives on both sides of the above, we have $\frac{\partial^2 \phi(\mathbf{c}^{(1)}, \mathbf{c}^{(2)}, \mathbf{s})}{\partial c_1^{(1)} \partial s_k} = 0$, $k = 1, \dots, K$ —which leads to K linear equations w.r.t. $(h_m^{(1)})''$ for $m = 1, \dots, M_1$. In addition $\frac{\partial^2 \phi(\mathbf{c}^{(1)}, \mathbf{c}^{(2)}, \mathbf{s})}{\partial c_1^{(1)} \partial c_r^{(1)}} = 0$, $r = 1, \dots, R_1$, yields another R_1 linear equations. In total, we have collected $K + R_1$ linear equations as follows:

$$\sum_{m=1}^{M_1} b_m^{(1)} A_{m,k}^{(1)} A_{m,K+1}^{(1)} (h_m^{(1)})'' = 0, \quad \text{for } k = 1, \dots, K + R_1.$$

In the above derivation, we have used the assumption that $[\mathbf{s}_\ell^\top, (\mathbf{c}_\ell^{(1)})^\top, (\mathbf{c}_\ell^{(2)})^\top]^\top$ is ubiquitously unanchored, and thus all the “cross-derivatives” are zeros.

It can be rewritten as the following linear system:

$$(\mathbf{A}^{(1)})^\top \begin{bmatrix} b_1^{(1)} A_{1,K+1}^{(1)} \\ \vdots \\ b_{M_1}^{(1)} A_{M_1,K+1}^{(1)} \end{bmatrix} \begin{bmatrix} (h_1^{(1)})'' \\ \vdots \\ (h_{M_1}^{(1)})'' \end{bmatrix} = \mathbf{0},$$

which holds for all $\mathbf{s}_\ell \in \mathcal{S}$ and $\mathbf{c}_\ell^{(q)} \in \mathcal{C}_q$, where $\mathbf{A}^{(1)} \in \mathbb{R}^{M_1 \times M_1}$ is the mixing matrix, which is full rank since it is drawn from an absolutely continuous distribution.

Now what remains is to show that the diagonal matrix in the middle is full rank which is equivalently that all the diagonal elements are non-zero. Since $\mathbf{A}^{(1)}$ is assumed to be continuously random variable, the probability that any element of it being zero is zero.

Therefore, we only need to show that $b_1^{(1)}, \dots, b_{M_1}^{(1)}$ are not zeros. By the assumption that $\|\mathbf{B}^{(q)}\|_0 = KM_q$, none of $\mathbf{b}^{(1)}$ will be zero. As a result, the matrix in the middle is full rank. Consequently, $(\mathbf{h}^{(1)})''$ has to be a all-zero vector. In addition, Proposition 1 guarantees that there exists such a solution $\mathbf{b}^{(1)}$. By applying the same proof to the second view, one can show that $\mathbf{h}^{(2)}$ is also an affine mapping.

APPENDIX D PROOF OF THEOREM 3

For the same $\phi(\mathbf{c}^{(1)}, \mathbf{c}^{(2)}, \mathbf{s})$ defined in previous theorems, by taking second order derivative $\frac{\partial^2 \phi(\mathbf{c}^{(1)}, \mathbf{c}^{(2)}, \mathbf{s})}{\partial c_i^{(1)} \partial c_j^{(1)}}$, we have $\frac{R_1(R_1+1)}{2}$ linear equations:

$$\sum_{m=1}^{M_1} b_m^{(1)} A_{m,K+i}^{(1)} A_{m,K+j}^{(1)} (h_m^{(1)})'' = 0, \quad (24)$$

for $i, j = 1, \dots, R_1$ and $i \leq j$. Note that since s_k and s_j could be fully dependent, we do not take cross-derivatives involving s_k or s_j . The reason is that $\frac{\partial s_k}{\partial s_j} \neq 0$ could happen if s_k and s_j

are completely dependent (or, if $[s_k, s_j]^\top$ is not ubiquitously unanchored). Eq. (24) can be re-expressed as:

$$\underbrace{\begin{bmatrix} \left(\mathbf{a}_{K+1}^{(1)} \circledast \mathbf{a}_{K+1}^{(1)}\right)^\top \\ \vdots \\ \left(\mathbf{a}_{K+R_1}^{(1)} \circledast \mathbf{a}_{K+R_1}^{(1)}\right)^\top \\ \left(\mathbf{a}_{K+1}^{(1)} \circledast \mathbf{a}_{K+2}^{(1)}\right)^\top \\ \vdots \\ \left(\mathbf{a}_{K+R_1-1}^{(1)} \circledast \mathbf{a}_{K+R_1}^{(1)}\right)^\top \end{bmatrix}}_{\mathbf{A}} \begin{bmatrix} b_1^{(1)} & & & \\ & \ddots & & \\ & & b_{M_1}^{(1)} & \end{bmatrix} \begin{bmatrix} (h_1^{(1)})'' \\ \vdots \\ (h_{M_1}^{(1)})'' \end{bmatrix} = \mathbf{0},$$

where \circledast denotes Hadamard product. Let $\mathbf{a}_j^{(1)}$ denote the i th column of $\mathbf{A}^{(1)}$. We need to show that \mathbf{A} has full column rank—i.e., there are M_1 rows that are linearly independent.

Since \mathbf{A} is a random matrix, it suffices to show that one of the $M_1 \times M_1$ submatrices admits nonzero determinant somewhere—which implies that the determinant is nonzero everywhere, almost surely [65]. To this end, consider a special case where $\mathbf{A}^{(1)}$ is a Vandermonde matrix, i.e., $\mathbf{a}_i^{(1)} = [1, x_i, x_i^2, \dots, x_i^{M_1-1}]^\top$, and $x_i \neq x_j$. Let us consider any R columns of the R_1 columns from $\mathbf{a}_{K+1}^{(1)}, \dots, \mathbf{a}_{K+R_1}^{(1)}$. Note that we require $R \leq R_1$ to satisfy $M_1 = \lfloor \frac{R(R+1)}{2} \rfloor$ (for notational simplicity, we assume $M_1 = \frac{R(R+1)}{2}$ in the sequel). Hence, the corresponding rows in the matrix of interest exhibit the following form:

$$\begin{bmatrix} 1 & x_1^2 & \dots & x_1^{2(M_1-1)} \\ \dots & \dots & & \dots \\ 1 & x_R^2 & \dots & x_R^{2(M_1-1)} \\ 1 & x_1 x_2 & \dots & (x_1 x_2)^{M_1-1} \\ \dots & \dots & & \dots \\ 1 & x_{R-1} x_R & \dots & (x_{R-1} x_R)^{M_1-1} \end{bmatrix}$$

Note that one can always construct such a sequence—e.g., $x_1 = 1, x_2 = 1.1, x_3 = 1.11, \dots$ —to make sure that the matrix above is full rank. This means that the determinant of the above matrix is nonzero somewhere. Hence, the matrix \mathbf{A} has full column rank almost surely [65]. This also means that $(h_i^{(1)})'' = 0$ for all $i = 1, \dots, M_1$. The same proof applies to $h_i^{(2)}$ for $i = 1, \dots, M_2$.

APPENDIX E PROOF OF LEMMA 1

First, notice that the rows of \mathbf{U} live in the null space of $\mathbf{1}$. Hence, we can re-express \mathbf{U} by $\mathbf{U} = \mathbf{V}\mathbf{W}$, where $\mathbf{W} = \mathbf{I} - \frac{1}{N}\mathbf{1}\mathbf{1}^\top$. Hence, Problem (20) can be written as the following equivalent problem:

$$\begin{aligned} & \underset{\mathbf{V}}{\text{minimize}} \|\mathbf{V}\mathbf{W} - \mathbf{Z}\|_F^2 \\ & \text{subject to } (1/N)\mathbf{V}\mathbf{W}\mathbf{W}^\top\mathbf{V}^\top = \mathbf{I}. \end{aligned}$$

By expanding the above, we have:

$$\begin{aligned} & \underset{\mathbf{V}}{\text{minimize}} \|\mathbf{V}\mathbf{W}\|_F^2 - 2\text{Tr}(\mathbf{V}\mathbf{W}\mathbf{Z}^\top) + \|\mathbf{Z}\|_F^2 \\ & \text{subject to } (1/N)\mathbf{V}\mathbf{W}\mathbf{W}^\top\mathbf{V}^\top = \mathbf{I}. \end{aligned}$$

The last term $\|\mathbf{Z}\|_F^2$ is a constant, it can be replaced with $\|\mathbf{Z}\mathbf{W}\|_F^2$. The middle term also can be written as $-2\text{Tr}(\mathbf{V}\mathbf{W}\mathbf{W}^\top\mathbf{Z}^\top)$ since \mathbf{W} is an orthogonal complement projector with $\mathbf{W} = \mathbf{W}^\top$ and $\mathbf{W} = \mathbf{W}^2$. Thus, we have the reformulated problem:

$$\begin{aligned} & \underset{\mathbf{V}}{\text{minimize}} \|\mathbf{V}\mathbf{W}\|_F^2 - 2\text{Tr}(\mathbf{V}\mathbf{W}\mathbf{W}^\top\mathbf{Z}^\top) + \|\mathbf{Z}\mathbf{W}\|_F^2 \\ & \text{subject to } (1/N)\mathbf{V}\mathbf{W}\mathbf{W}^\top\mathbf{V}^\top = \mathbf{I}. \end{aligned}$$

Plugging $\mathbf{U} = \mathbf{V}\mathbf{W}$ back, we have:

$$\begin{aligned} & \underset{\mathbf{U}}{\text{minimize}} \|\mathbf{U} - \mathbf{Z}\mathbf{W}\|_F^2 \\ & \text{subject to } (1/N)\mathbf{U}\mathbf{U}^\top = \mathbf{I}. \end{aligned}$$

An optimal solution can be obtained via invoking the *Procrustes projection*; i.e., a solution is simply taking SVD of $\mathbf{Z}\mathbf{W} = \mathbf{P}\mathbf{D}\mathbf{Q}^\top$ and let $\mathbf{U} = \sqrt{N}\mathbf{P}\mathbf{Q}^\top$.

REFERENCES

- [1] S. Wold, K. Esbensen, and P. Geladi, “Principal component analysis,” *Chemometrics and Intelligent Laboratory Systems*, vol. 2, no. 1-3, pp. 37–52, 1987.
- [2] P. Comon, “Independent component analysis, a new concept?” *Signal Processing*, vol. 36, no. 3, pp. 287 – 314, 1994.
- [3] P. Paatero and U. Tapper, “Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values,” *Environmetrics*, vol. 5, no. 2, pp. 111–126, 1994.
- [4] X. Fu, K. Huang, N. D. Sidiropoulos, and W.-K. Ma, “Nonnegative matrix factorization for signal and data analytics: Identifiability, algorithms, and applications,” *IEEE Signal Process. Mag.*, vol. 36, no. 2, pp. 59–80, March 2019.
- [5] H. Hotelling, “Relations between two sets of variates,” *Biometrika*, vol. 28, no. 3/4, pp. 321–377, 1936.
- [6] F. R. Bach and M. I. Jordan, “A probabilistic interpretation of canonical correlation analysis,” in *Proc. IEEE SPAWC 2019*, 2019, pp. 1097–1105.
- [7] M. S. Ibrahim and N. D. Sidiropoulos, “Cell-edge interferometry: Reliable detection of unknown cell-edge users via canonical correlation analysis,” in *Proc. IEEE SPAWC 2019*, 2019, pp. 1097–1105.
- [8] D. R. Hardoon, S. Szedmak, and J. Shawe-Taylor, “Canonical correlation analysis: An overview with application to learning methods,” *Neural Computation*, vol. 16, no. 12, pp. 2639–2664, 2004.
- [9] T. W. Anderson, *An introduction to multivariate statistical analysis*, no. 519.9 A53, 1962.
- [10] W. Wang, J. Wang, D. Garber, and N. Srebro, “Efficient globally convergent stochastic optimization for canonical correlation analysis,” in *Proceedings of NIPS 2016*, 2016, pp. 766–774.
- [11] Z. Ma, Y. Lu, and D. Foster, “Finding linear structure in large datasets with scalable canonical correlation analysis,” in *Proceedings of ICML 2015*, 2015, pp. 169–178.
- [12] R. Ge, C. Jin, S. M. Kakade, P. Netrapalli, and A. Sidford, “Efficient algorithms for large-scale generalized eigenvector computation and canonical correlation analysis,” in *Proceedings of ICML 2016*, 2016, pp. 2741–2750.
- [13] Y. Lu and D. P. Foster, “Large scale canonical correlation analysis with iterative least squares,” in *Advances in Neural Information Processing Systems*, 2014, pp. 91–99.
- [14] G. H. Golub and H. Zha, “The canonical correlations of matrix pairs and their numerical computation,” in *Linear Algebra For Signal Processing*. Springer, 1995, pp. 27–49.
- [15] J. R. Kettenring, “Canonical analysis of several sets of variables,” *Biometrika*, vol. 58, no. 3, pp. 433–451, 1971.
- [16] J. D. Carroll, “Generalization of canonical correlation analysis to three or more sets of variables,” in *Proceedings of the 76th annual convention of the American Psychological Association*, vol. 3, 1968, pp. 227–228.
- [17] X. Fu, K. Huang, M. Hong, N. D. Sidiropoulos, and A. M.-C. So, “Scalable and flexible multiview MAX-VAR canonical correlation analysis,” *IEEE Trans. Signal Process.*, vol. 65, no. 16, pp. 4150–4165, 2017.
- [18] X. Fu, K. Huang, E. Papalexakis, H. A. Song, P. P. Talukdar, C. Faloutsos, N. Sidiropoulos, and T. Mitchell, “Efficient and distributed generalized canonical correlation analysis for big multiview data,” *IEEE Trans. Knowl. Data Eng.*, vol. 31, no. 12, pp. 2304–2318, 2018.

[19] C. I. Kanatsoulis, X. Fu, N. D. Sidiropoulos, and M. Hong, "Structured SUMCOR multiview canonical correlation analysis for large-scale data," *IEEE Trans. Signal Process.*, vol. 67, no. 2, pp. 306–319, 2018.

[20] K. Fukumizu, F. R. Bach, and A. Gretton, "Statistical consistency of kernel canonical correlation analysis," *Journal of Machine Learning Research*, vol. 8, no. Feb, pp. 361–383, 2007.

[21] F. R. Bach and M. I. Jordan, "Kernel independent component analysis," *Journal of Machine Learning Research*, vol. 3, no. Jul, pp. 1–48, 2002.

[22] W. Wang, R. Arora, K. Livescu, and J. Bilmes, "On deep multi-view representation learning," in *Proceedings of ICML 2015*, 2015, pp. 1083–1092.

[23] G. Andrew, R. Arora, J. Bilmes, and K. Livescu, "Deep canonical correlation analysis," in *Proceedings of ICML 2013*, vol. 28, no. 3, 17–19 Jun 2013, pp. 1247–1255.

[24] W. Wang, R. Arora, K. Livescu, and J. A. Bilmes, "Unsupervised learning of acoustic features via deep canonical correlation analysis," in *Proceedings of IEEE ICASSP 2015*, 2015, pp. 4590–4594.

[25] A. Taleb and C. Jutten, "Source separation in post-nonlinear mixtures," *IEEE Trans. Signal Process.*, vol. 47, no. 10, pp. 2807–2820, 1999.

[26] C. Jutten and J. Karhunen, "Advances in blind source separation (BSS) and independent component analysis (ICA) for nonlinear mixtures," *International Journal of Neural Systems*, vol. 14, no. 05, pp. 267–292, 2004.

[27] P. Comon and C. Jutten, *Handbook of Blind Source Separation*. Elsevier, 2010.

[28] A. Hyvärinen and P. Pajunen, "Nonlinear independent component analysis: Existence and uniqueness results," *Neural Networks*, vol. 12, no. 3, pp. 429–439, 1999.

[29] A. Hyvärinen and H. Morioka, "Unsupervised feature extraction by time-contrastive learning and nonlinear ICA," in *Advances in Neural Information Processing Systems*, 2016, pp. 3765–3773.

[30] S. Achard and C. Jutten, "Identifiability of post-nonlinear mixtures," *IEEE Signal Process. Lett.*, vol. 12, no. 5, pp. 423–426, 2005.

[31] E. Oja, "The nonlinear PCA learning rule in independent component analysis," *Neurocomputing*, vol. 17, no. 1, pp. 25–45, 1997.

[32] B. Yang, X. Fu, N. D. Sidiropoulos, and K. Huang, "Learning nonlinear mixtures: Identifiability and algorithm," *arXiv preprint arXiv:1901.01568*, 2019.

[33] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, p. 436, 2015.

[34] Q. Lyu and X. Fu, "Nonlinear multiview analysis: Identifiability and neural network-based implementation," in *IEEE SAM 2020*, to appear, 2020.

[35] X. Fu, K. Huang, B. Yang, W. Ma, and N. D. Sidiropoulos, "Robust volume minimization-based matrix factorization for remote sensing and document clustering," *IEEE Trans. Signal Process.*, vol. 64, no. 23, pp. 6254–6268, Dec 2016.

[36] X. Fu, W.-K. Ma, K. Huang, and N. D. Sidiropoulos, "Blind separation of quasi-stationary sources: Exploiting convex geometry in covariance domain," *IEEE Trans. Signal Process.*, vol. 63, no. 9, pp. 2306–2320, May 2015.

[37] L. De Lathauwer, J. Castaing, and J.-F. Cardoso, "Fourth-order cumulant-based blind identification of underdetermined mixtures," *IEEE Trans. Signal Process.*, vol. 55, no. 6, pp. 2965–2973, June 2007.

[38] S. Cruces, "Bounded component analysis of linear mixtures: A criterion of minimum convex perimeter," *IEEE Trans. Signal Process.*, vol. 58, no. 4, pp. 2141–2154, April 2010.

[39] M. Zibulevsky and B. A. Pearlmutter, "Blind source separation by sparse decomposition in a signal dictionary," *Neural Computation*, vol. 13, no. 4, pp. 863–882, 2001.

[40] X. Fu, K. Huang, and N. D. Sidiropoulos, "On identifiability of nonnegative matrix factorization," *IEEE Signal Process. Lett.*, vol. 25, no. 3, pp. 328–332, 2018.

[41] L. De Lathauwer and J. Castaing, "Blind identification of underdetermined mixtures by simultaneous matrix diagonalization," *IEEE Trans. Signal Process.*, vol. 56, no. 3, pp. 1096–1105, Mar. 2008.

[42] A. Belouchrani, K. A. Meraim, J.-F. Cardoso, and E. Moulines, "A blind source separation technique using second order statistics," *IEEE Trans. Signal Process.*, vol. 45, no. 2, pp. 434–444, Feb. 1997.

[43] A. Ziehe, M. Kawanabe, S. Harmeling, and K. R. Müller, "Blind separation of post-nonlinear mixtures using linearizing transformations and temporal decorrelation," *Journal of Machine Learning Research*, vol. 4, no. Dec, pp. 1319–1338, 2003.

[44] A. Ziehe, K. R. Müller, G. Nolte, B. M. Mackert, and G. Curio, "Artifact reduction in magnetoneurography based on time-delayed second-order correlations," *IEEE Trans. Biomed. Eng.*, vol. 47, no. 1, pp. 75–87, Jan 2000.

[45] F. Oveisi, S. Oveisi, A. Efranian, and I. Patras, "Nonlinear Independent Component Analysis for EEG-Based Brain-Computer Interface Systems," *Independent Component Analysis for Audio and Biosignal Applications*, Edited by Ganesh R. Naik, p. 165, 2012.

[46] N. Dobigeon, J.-Y. Tourneret, C. Richard, J. C. M. Bermudez, S. McLaughlin, and A. O. Hero, "Nonlinear unmixing of hyperspectral images: Models and algorithms," *IEEE Signal Process. Mag.*, vol. 31, no. 1, pp. 82–94, 2014.

[47] Y.-O. Li, T. Adali, W. Wang, and V. D. Calhoun, "Joint blind source separation by multiset canonical correlation analysis," *IEEE Trans. Signal Process.*, vol. 57, no. 10, pp. 3918–3929, 2009.

[48] N. M. Correa, Y.-O. Li, T. Adali, and V. D. Calhoun, "Canonical correlation analysis for feature-based fusion of biomedical imaging modalities and its application to detection of associative networks in schizophrenia," *IEEE J. Sel. Topics Signal Process.*, vol. 2, no. 6, pp. 998–1007, 2008.

[49] D. Lopez-Paz, S. Sra, A. Smola, Z. Ghahramani, and B. Schölkopf, "Randomized nonlinear component analysis," in *Proceedings of ICML 2014*, 2014, pp. 1359–1367.

[50] I. Rustandi, M. Just, and T. Mitchell, "Integrating multiple-study multiple-subject fMRI datasets using canonical correlation analysis," in *Proceedings of the MICCAI 2009 Workshop: Statistical modeling and detection issues in intra-and inter-subject functional MRI data analysis*, 2009.

[51] W. D. Penny, K. J. Friston, J. T. Ashburner, S. J. Kiebel, and T. E. Nichols, *Statistical parametric mapping: the analysis of functional brain images*. Elsevier, 2011.

[52] X. Fu, K. Huang, O. Stretcu, H. A. Song, E. Papalexakis, P. Talukdar, T. Mitchell, N. Sidiropoulos, C. Faloutsos, and B. Poczos, "Brainzoom: High resolution reconstruction from multi-modal brain signals," in *Proceedings SIAM SDM 2017*, 2017, pp. 216–227.

[53] G. E. Hinton and R. S. Zemel, "Autoencoders, minimum description length and Helmholtz free energy," in *Advances in Neural Information Processing Systems*, 1994, pp. 3–10.

[54] J. Vía, M. Anderson, X.-L. Li, and T. Adali, "Joint blind source separation from second-order statistics: Necessary and sufficient identifiability conditions," in *Proceedings of IEEE ICASSP 2011*, 2011, pp. 2520–2523.

[55] G. Cybenko, "Approximation by superpositions of a sigmoidal function," *Mathematics of Control, Signals and Systems*, vol. 2, no. 4, pp. 303–314, 1989.

[56] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[57] A. S. Zamzam, X. Fu, and N. D. Sidiropoulos, "Data-driven learning-based optimization for distribution system state estimation," *IEEE Trans. Power Syst.*, vol. 34, no. 6, pp. 4796–4805, 2019.

[58] A. Benton, H. Khayrallah, B. Gujral, D. Reisinger, S. Zhang, and R. Arora, "Deep generalized canonical correlation analysis," in *RepL4NLP at ACL*, 2017.

[59] L. van. der. Maaten and G. Hinton, "Visualizing data using t-SNE," *Journal of Machine Learning Research*, vol. 9, no. Nov, pp. 2579–2605, 2008.

[60] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.

[61] M. van Breukelen, R. P. Duin, D. M. Tax, and J. Den Hartog, "Handwritten digit recognition by combined classifiers," *Kybernetika*, vol. 34, no. 4, pp. 381–386, 1998.

[62] A. Y. Ng, M. I. Jordan, and Y. Weiss, "On spectral clustering: Analysis and an algorithm," in *Advances in Neural Information Processing Systems*, 2002, pp. 849–856.

[63] X. Fu, K. Huang, E. E. Papalexakis, H.-A. Song, P. P. Talukdar, N. D. Sidiropoulos, C. Faloutsos, and T. Mitchell, "Efficient and distributed algorithms for large-scale generalized canonical correlations analysis," in *Proc. IEEE ICDM 2016*, 2016, pp. 871–876.

[64] I. Olkin, "The density of the inverse and pseudo-inverse of a random matrix," *Statistics & Probability Letters*, vol. 38, no. 2, pp. 131–135, 1998.

[65] R. Caron, and T. Traynor, "The zero set of a polynomial," *WSMR Report (2005)*: 05-02.

Supplementary Materials for ‘‘Nonlinear Multiview Component Analysis: Identifiability and Neural Network-Assisted Implementation’’

Qi Lyu and Xiao Fu

APPENDIX F PROOF OF LINEAR CCA IDENTIFIABILITY

Since the theorem in [7] is an important component to our proof, we restate the theorem and its proof in [7] using the notations in this paper, for ease of exposition.

Theorem 4 (Linear CCA Identifiability) *Assume that the model in (5) holds. Denote $\mathbf{S} = [\mathbf{s}_1, \dots, \mathbf{s}_N]$ and $\mathbf{C}^{(q)} = [\mathbf{c}_1^{(q)}, \dots, \mathbf{c}_N^{(q)}]$ for $q = 1, 2$. Also assume that matrix $\mathbf{T} = [\mathbf{S}^\top, (\mathbf{C}^{(1)})^\top, (\mathbf{C}^{(2)})^\top] \in \mathbb{R}^{N \times (K+R_1+R_2)}$ has full column rank, and that $\mathbf{A}^{(q)} \in \mathbb{R}^{M_q \times (K+R_q)}$ has full column rank. Then the range space spanned by the shared components, i.e., $\text{range}(\mathbf{S}^\top)$, can be identified via solving the linear CCA problem in (7).*

Proof: According to our generative model, in the linear case,

$$\mathbf{y}_\ell^{(q)} = \mathbf{A}^{(q)}[\mathbf{s}_\ell^\top, (\mathbf{c}_\ell^{(q)})^\top]^\top, \quad \ell = 1, \dots, N, \quad (25)$$

Note that equation above is exactly the same as model Eq. (5) in [7].

First consider the vector case where $s_\ell \in \mathbb{R}$. In this case, $\mathbf{B}^{(q)}$ becomes a vector $\mathbf{b}^{(q)}$. By solving the problem in (7), we have the following when the cost reaches zero:

$$(\mathbf{Y}^{(1)})^\top \mathbf{b}^{(1)} = (\mathbf{Y}^{(2)})^\top \mathbf{b}^{(2)}.$$

Without loss of generality, we can set $\mathbf{b}^{(q)} = \mathbf{A}^{(q)}((\mathbf{A}^{(q)})^\top \mathbf{A}^{(q)})^{-1} \mathbf{r}^{(q)}$ where $\mathbf{r}^{(q)}$ is a vector that is in \mathbb{R}^{1+R_q} . This is because that the $\mathbf{b}^{(q)}$ vector can always be decomposed into the orthogonal projection on $\mathbf{A}^{(q)}$ and on its complement projection as well. The complement one will always be canceled since there is a $(\mathbf{A}^{(q)})^\top$ term in $(\mathbf{Y}^{(q)})^\top$. Consequently, we have

$$[\mathbf{s}^\top, (\mathbf{C}^{(1)})^\top] \mathbf{r}^{(1)} = [\mathbf{s}^\top, (\mathbf{C}^{(2)})^\top] \mathbf{r}^{(2)},$$

which can be rewritten as:

$$\begin{aligned} & [\mathbf{s}^\top, (\mathbf{C}^{(1)})^\top, (\mathbf{C}^{(2)})^\top] \begin{bmatrix} \mathbf{r}^{(1)} \\ \mathbf{0}_{R_2} \end{bmatrix} \\ &= [\mathbf{s}^\top, (\mathbf{C}^{(1)})^\top, (\mathbf{C}^{(2)})^\top] \begin{bmatrix} \mathbf{r}^{(2)}(1) \\ \mathbf{0}_{R_1} \\ \mathbf{r}^{(2)}(2 : \text{end}) \end{bmatrix} \end{aligned}$$

where $\mathbf{0}_{R_1}$ is all-zero vector with size R_1 and $\mathbf{0}_{R_2}$ is that of size R_2 . Rearranging terms, we have:

$$\mathbf{T}\mathbf{r} = \mathbf{0}$$

where $\mathbf{T} = [\mathbf{s}^\top, (\mathbf{C}^{(1)})^\top, (\mathbf{C}^{(2)})^\top] \in \mathbb{R}^{N \times (1+R_1+R_2)}$ and $\mathbf{r} = [(\mathbf{r}^{(1)}(1) - \mathbf{r}^{(2)}(1))^\top, \mathbf{r}^{(1)}(2 : \text{end})^\top, -\mathbf{r}^{(2)}(2 : \text{end})^\top]^\top$ where $\mathbf{r}^{(1)}(2 : \text{end})$ containing all except the first entries in $\mathbf{r}^{(1)}$.

It is obvious that if \mathbf{T} has full column rank, then \mathbf{r} has to be an all-zero vector which means that $\mathbf{r}^{(1)} = c\mathbf{e}_1$ and

$\mathbf{r}^{(2)} = ce'_1$ where \mathbf{e}_1 and \mathbf{e}'_1 are the first columns of identity matrices of different sizes, respectively. As a result, the shared component is identified as $\pm s/\|\mathbf{s}\|_2$.

For the $K > 1$ case, let $\mathbf{B}^{(q)} = \mathbf{A}^{(q)}((\mathbf{A}^{(q)})^\top \mathbf{A}^{(q)})^{-1} \mathbf{R}^{(q)}$. Similarly, define

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}^{(1)}(1 : K, :) - \mathbf{R}^{(2)}(1 : K, :) \\ \mathbf{R}^{(1)}(K+1 : \text{end}, :) \\ -\mathbf{R}^{(2)}(K+1 : \text{end}, :) \end{bmatrix} \in \mathbb{R}^{(K+R_1+R_2) \times K}$$

where $\mathbf{R}^{(1)}(1 : K, :)$ denotes rows 1 to K of $\mathbf{R}^{(1)}$. We have the following holds:

$$\mathbf{T}\mathbf{R} = \mathbf{0}.$$

If \mathbf{T} has full column rank, then the solution is unique with $\mathbf{R}^{(1)}(1 : K, :) = \mathbf{R}^{(2)}(1 : K, :) = \Theta$, $\mathbf{R}^{(1)}(K+1 : \text{end}, :) = \mathbf{0}$ and $\mathbf{R}^{(2)}(K+1 : \text{end}, :) = \mathbf{0}$. This means that the solution is

$$\mathbf{B}^{(q)}\mathbf{Y}^{(q)} = \Theta\mathbf{S}, \quad q = 1, 2,$$

which completes the proof. \blacksquare

One remark here is that if $s_\ell, \mathbf{c}_\ell^{(q)}$ are drawn from any jointly continuous distribution and the condition $N \geq R_1 + R_2 + K$ holds, then the matrix \mathbf{R} has full column rank almost surely.

APPENDIX G THE UBIQUITOUSLY UNANCHORED CONDITION AND CROSS-DERIVATIVES

The key consequence of having v_i ’s to be ubiquitously unanchored is that the cross-derivatives such as $\partial v_i / \partial v_j$ are all zeros under this condition.

Recall that the derivative of a continuous function $\phi(x)$ is defined as follows:

$$\varphi'(x) = \lim_{\Delta x \rightarrow 0} \frac{\varphi(x + \Delta x) - \varphi(x)}{\Delta x}.$$

Let us treat v_j as a function of v_i (since the two could be dependent). Consider the partial derivative as follows:

$$\frac{\partial v_j}{\partial v_i} = \lim_{\Delta v_i \rightarrow 0} \frac{v_j(v_i + \Delta v_i) - v_j(v_i)}{\Delta v_i}$$

where the notation $v_j(v_i)$ means that v_j is a function of v_i .

Suppose that the components of \mathbf{v} are ubiquitously unanchored. Hence, for any v_i , any fixed \bar{v}_j can appear together with it in the domain of \mathbf{v} . Let us denote $\mathcal{V}(v_i, \bar{v}_j) \subseteq \mathcal{V}$ as the set of \mathbf{v} such that $[\mathbf{v}]_i = v_i$ and $[\mathbf{v}]_j = \bar{v}_j$. Over this subset, we have

$$v_j(v_i + \Delta v_i) = \bar{v}_j, \quad v_j(v_i) = \bar{v}_j, \quad \forall \mathbf{v} \in \mathcal{V}(v_i, \bar{v}_j)$$

for all possible Δv_i . Hence, the above leads to

$$\begin{aligned} \frac{\partial v_j}{\partial v_i} &= \lim_{\Delta v_i \rightarrow 0} \frac{v_j(v_i + \Delta v_i) - v_j(v_i)}{\Delta v_i} \\ &= \lim_{\Delta v_i \rightarrow 0} \frac{\bar{v}_j - \bar{v}_j}{\Delta v_i} \\ &= 0, \quad \forall \mathbf{v} \in \mathcal{V}(v_i, \bar{v}_j), \end{aligned}$$

for all $j \neq i$. Repeating the above for all possible v_i and \bar{v}_j , one can reach the conclusion that $\frac{\partial v_j}{\partial v_i} = 0$.