



Multiview Variational Graph Autoencoders for Canonical Correlation Analysis

Yacouba Kaloga, Pierre Borgnat, Sundeep Prabhakar Chepuri, Patrice Abry, Amaury Habrard

► To cite this version:

Yacouba Kaloga, Pierre Borgnat, Sundeep Prabhakar Chepuri, Patrice Abry, Amaury Habrard. Multiview Variational Graph Autoencoders for Canonical Correlation Analysis. ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Jun 2021, Toronto, Canada. pp.5320-5324, 10.1109/ICASSP39728.2021.9414466 . hal-03435998

HAL Id: hal-03435998

<https://hal.science/hal-03435998>

Submitted on 19 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

MULTIVIEW VARIATIONAL GRAPH AUTOENCODERS FOR CANONICAL CORRELATION ANALYSIS

Yacouba Kaloga[§], Pierre Borgnat[§], Sundeep Prabhakar Chepuri^{*}, Patrice Abry[§] and Amaury Habrard[†]

[§]Univ Lyon, Ens de Lyon, Univ. Claude Bernard, CNRS, Laboratoire de Physique, Lyon, France

^{*}Department of Electrical and Communication Engineering, Indian Institute of Science, Bangalore, India

[†] University of Lyon, UJM-Saint-Etienne, CNRS, Laboratoire Hubert Curien, UMR 5516, France

ABSTRACT

We present a novel multiview canonical correlation analysis model based on a variational approach. This is the first nonlinear model that takes into account the available graph-based geometric constraints while being scalable for processing large scale datasets with multiple views. It is based on an autoencoder architecture with graph convolutional neural network layers. We experiment with our approach on classification, clustering, and recommendation tasks on real datasets. The algorithm is competitive with state-of-the-art multiview representation learning techniques.

Index Terms— Canonical correlation analysis, dimensionality reduction, multiview, graph neural networks, variational inference

1. INTRODUCTION

Interconnected societies generate large amounts of structured data, that frequently stem from observing a common set of objects (or sources) through different modalities. Such multiview datasets are encountered in many different fields like computational biology [2], acoustics [3], surveillance [4], or social networks [5], to list a few. Although there exist many tools to analyze and study multiview datasets [6], it is more than ever necessary to develop methods for analyzing large-scale and structured multiview datasets.

Canonical Correlation Analysis (CCA) [7] can be used for multiview representation learning, by seeking latent low-dimensional representations that are common to all the different views. This common representation that encodes information from different datasets can be leveraged to improve the performance of machine learning tasks, e.g., clustering [8]. *Algebraic* approaches to CCA obtain this latent low-dimensional manifold by maximizing correlations between the projections on the different views onto it. Being nonparametric, these approaches are powerful but do not scale well to

large datasets. Instead, *probabilistic* approaches to CCA scale more easily but, being based on models, are less versatile to adjust to model mismatch or to the data structure. The present work attempts to reconcile scalability and versatility in CCA.

It has then been extended to the multiview setting, e.g., see [9], and to account for nonlinear dependencies (beyond correlations), e.g., see Kernel CCA [10], Deep CCA [11], or Autoencoder CCA [12]. Despite significant improvements in performance, many of these approaches suffer from scalability issues [13, 14], mainly due to the prohibitive costs of the underlying eigendecomposition. Alternatively, it has been shown that CCA can be cast equivalently into a Bayesian inference problem [15]. As recent advances in Variational autoencoders [16] made Bayesian inference scalable, probabilistic CCA approaches gained popularity because of their potential (e.g., to deal with missing views) and scalability. Concomitantly, it was shown in [17, 18] that incorporating the available graph-induced knowledge about the common source into multiview CCA improves performance of various machine learning tasks. We refer to this graph-aware multiview CCA method from [18] as GMCCA. However, GMCCA suffers from the involved eigendecomposition costs. However, there is no CCA method that both can use some prior graph-based structure in the latent space, and is scalable. Such a method is proposed here.

The present work develops a scalable multiview variational graph autoencoder for CCA (MVGCCA). Section 2 recalls preliminaries in Multiview CCA. Section 3 describes the proposed approach and its key contributions. In particular, the way in which the graph structure is enforced in the common latent space preserving the scalability. Section 4 describes the datasets that are used for numerical experiments in Section 5.

2. MULTIVIEW-CCA BACKGROUND

Setting. We consider M -view datasets where the instance space X has M view, spaces in \mathbb{R}^{d_m} , $m = 1, \dots, M$. We denote n the sample size, and $X_m \in \mathbb{R}^{d_m \times n}$ the m -view data matrix. An instance i of view m is denoted as $x_m^i \in \mathbb{R}^{d_m}$.

Supported by the IFCAM project MA/IFCAM/19/56, SERB SRG/2019/000619 grant, the ACADEMICS Grant of IDEXLYON, Univ. Lyon, PIA ANR-16-IDEX-0005, and the CBP IT test platform (ENS de Lyon, France) for ML facilities, operating the SIDUS solution [1].

2.1. Algebraic approaches: linear CCA and GMCCA

Given 2-views, seek the best projectors $U_1 \in \mathbb{R}^{d \times d_1}$, $U_2 \in \mathbb{R}^{d \times d_2}$ with $d \ll \min(d_1, d_2)$ such that the correlation between $U_1 X_1$ and $U_2 X_2$ is maximized. This can be formulated as the following optimization problem:

$$\min_{U_1, U_2} \|U_1 X_1 - U_2 X_2\|_2^2 \text{ s.t. } U_m (X_m X_m^T) U_m^T = I_{d_m}. \quad (1)$$

The solution is obtained via an eigendecomposition (e.g., see [11, 18]). Extending this to multiview data is difficult as maximizing the pairwise correlation between the M views is NP-hard [19]. A relaxation is to seek a common low-dimensional representation $S \in \mathbb{R}^{d \times n}$, that is as close as possible to each low-dimensional projection $U_m X_m$, $m = 1 \dots M$. It leads to problem (2) with $\gamma = 0$. This relaxation is solved using an eigenvalue decomposition. Chen et al. [18] have proposed GMCCA, in which some graph-based prior knowledge on S , when available, can be incorporated to increase the clustering performance. This is done by ensuring smoothness of S on this graph. By doing so, graph-regularized CCA problem can be posed as (with L the Laplacian matrix):

$$\min_{(U_m)_{m=1 \dots M}} \sum_{m=1}^M \|U_m X_m - S\|_2^2 + \gamma S L S^T \text{ s.t. } S S^T = I_d. \quad (2)$$

The solution S^* has columns equal to the k leading eigenvectors of $\sum_{m=1}^M X_m^T (X_m X_m^T) X_m - \gamma L$ [17]. For large datasets, this method involving the eigendecomposition does not always scale well.

2.2. Probabilistic CCA

The CCA solution of eq. (1) can be obtained from a graphical model [15], where the views come from a common source Z . Let us define a prior distribution on this latent space $p(z)$, and the conditional probability (decoders) for each views as $p_{\theta_m}(x_1|z)$, and are given as¹:

$$\begin{aligned} p(z) &\sim \mathcal{N}(0, I_d). \\ p_{\theta_m}(x_m|z) &\sim \mathcal{N}(W_m \cdot z + \mu_m, \Psi_m). \end{aligned} \quad (3)$$

with $\mu_m \in \mathbb{R}^{d_m}$, $W_m \in \mathbb{R}^{d_m \times d}$ and $\Psi_m \in \mathbb{R}^{d_m \times d_m} \succ 0$. We collect the parameters in θ_m as $\theta_m = (W_m, \mu_m, \Psi_m)$, $m = 1, 2$. The optimal solution to the probabilistic two-views CCA (θ_1^*, θ_2^*) is computed by maximizing data log-likelihood². From the Bayes theorem, the probabilities (defining encoders), $P_{\theta_1^*}(z|x_1^i)$ and $P_{\theta_2^*}(z|x_2^i)$ are well defined. Their expectation is exactly the optimal projection (up to an operator $M_m \in \mathbb{R}^{d \times d}$ defined in [15]): $\mathbb{E}_{z \sim p} P(z|x_m^i) = M_m^T U_m^* x_m^i$, $m = 1, 2$. In this framework, CCA has a natural multiview extension to $M > 2$.

¹Notation $p(\cdot) \sim q$, means $p = q$.

² $\log p_{\theta}(X_1, X_2) = \sum_{i=1}^n \sum_{m=1}^2 \log \int_{\mathbb{R}^d} p_{\theta_m}(x_m^i|z) p(z) dz$

We will use such an extension, while incorporating graph regularization like in [18]. Yet, solving a problem such as eq. (3) (i.e., a multi-dimensional probability distribution) is often intractable because it requires maximization of the log-likelihood and thus to integrate over all the latent spaces. A variational approach solves this issue.

2.3. Variational bound and graph autoencoder

Kingma et al. [20] have shown that introducing parametric distributions $q_{\eta}(z|x)$, with parameters η , of some untractable distribution $p(z|x)$, one can build a lower bound of log-likelihood called evidence lower bound objective (ELBO)³:

$$\begin{aligned} \log p_{\theta}(X_1, X_2) &\geq \sum_{i=1}^n \mathbb{E}_{z \sim q_{\eta}(\cdot|x_1^i, x_2^i)} [\log(p_{\theta}(x_1^i, x_2^i|z))] \\ &\quad - \sum_{i=1}^n D_{KL}(q_{\eta}(\cdot|x_1^i, x_2^i) \| p). \end{aligned} \quad (4)$$

ELBO is easier to approximate than the data log-likelihood so we maximise this lower bound with respect to both θ and η . The first term of ELBO ensures a correct data reconstruction due to encoded latent representations. The second term acts as a regularizer ensuring that posterior distribution q_{η} for each multiview element remains coherent in the latent space. It corresponds to the loss function of the variational autoencoders, used in most existing variational CCA methods.

It is possible to account geometric structure, by using the variational autoencoder extension proposed by Kipf et al. [16] for link prediction on graphs. In their singleview framework, data resides on the nodes of a graph with the weight matrix $A \in [0, 1]^{n \times n}$. We denote $\mathcal{V}(x_1^i)$ the neighborhood of x_1^i (including x_1^i) up to a certain distance in the graph. Then the graph-aware ELBO loss function is defined as :

$$\begin{aligned} \mathcal{L}_{ELBO} &= \sum_{i,j=1}^n \mathbb{E}_{z^i \sim q_{\eta}(\cdot|\mathcal{V}(x_1^i), A)} [\log(p_{\theta}(A_{i,j}|z^i, z^j))] \\ &\quad - \sum_{i=1}^n D_{KL}(q_{\eta}(\cdot|\mathcal{V}(x_1^i), A) \| p). \end{aligned} \quad (5)$$

The cross entropy term allows for graph reconstruction [16]. The posterior probability q_{η} can be now parametrized as a graph neural network. We will use a similar approach here, and extend it for multiple views.

3. VARIATIONAL GRAPH M-CCA

We now present the proposed method that estimate a probabilistic Multiview-CCA. Starting from eq. (3), our first contribution is to simply extend this framework for $m > 2$. To do so, we model the M -view data given the latent space

³Notation $z \sim p$, means z follows distribution p .

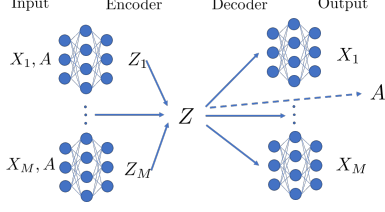


Fig. 1. Representation of MVGCCA. All the views are encoded to their own latent space Z_m using the common graph. They are merged to form a common view Z . Finally, Z is tailored to decode all the views and original graph.

vector $z \in \mathbb{R}^d \sim p = \mathcal{N}(0, I_d)$ using a parametric decoder based on a multilayer perceptron (denoted **MLP**). Hence, for $m = 1, 2, \dots, M$, we have $p_{\theta_m}(x_m|z) \sim \mathcal{N}(\mathbf{W}_m^{\mu \text{dec}} \mathbf{MLP}_m(z), \mathbf{\Psi}_m)$.

Next, we take into account a prior graph structure on the latent space of z , inspired by [16]: a link decoder distribution $p_g(a = 1|z, z') \sim \mathcal{B}(1, \ell(z^T z'))$ is introduced, parametrized by a Bernoulli law \mathcal{B} , and $\ell(\cdot)$ is the logistic sigmoid function. Here $a \in [0, 1]$ is a weight space vector and $z' \in \mathbb{R}^d$ is second latent space vector.

With the hypothesis of independence of all the views and of the links in the graph, we have the factorized decoder distribution $p_{\theta}((x_m^i)_{m=1}^M|z) = \prod_{m=1}^M p_{\theta}(x_m^i|z)$, graph decoder distribution⁴ $p_{\theta}(A|Z) = \prod_{j=1}^n p_g(A_{i,j}|z^i, z^j)$, and the parametric decoder distribution $q_{\eta}(z|\mathcal{V}(x_m^i)_{m=1}^M, A) = \prod_{m=1}^M q_{\eta_m}(z|\mathcal{V}(x_m^i), A)$. The ELBO takes the form⁵:

$$\begin{aligned} \mathcal{L}_{ELBO} = & \sum_{i,j=1}^n \mathbb{E}_{z^i \sim q_{\eta}(\cdot|\mathcal{V}(x_m^i)_{m=1}^M, A)} \log p_g(A_{i,j}|z^i, z^j) \\ & - \sum_{i,j=1}^n \mathbb{E}_{z^j \sim q_{\eta}(\cdot|\mathcal{V}(x_m^j)_{m=1}^M, A)} \log p_g(A_{i,j}|z^i, z^j) \\ & + \sum_{i=1}^n \sum_{m=1}^M \mathbb{E}_{z \sim q_{\eta}(\cdot|\mathcal{V}(x_m^i)_{m=1}^M, A)} \log p_{\theta}(x_m^i|z) \\ & - \sum_{i=1}^n D_{KL}(q_{\eta}(\cdot|\mathcal{V}(x_m^i)_{m=1}^M, A) \| p). \end{aligned} \quad (6)$$

The last term involves views encoders:

$$\begin{aligned} q_{\eta_m}(\cdot|\mathcal{V}(x_m^i), A) &= \mathcal{N}(\mu_m^{\text{enc}}(\cdot, i), \text{diag}(\sigma_m^{\text{enc}}(\cdot, i))). \\ \mu_m^{\text{enc}} &= \mathbf{W}_m^{\mu \text{enc}} \mathbf{Krylov}_m(X_m, A). \\ \log \sigma_m^{\text{enc}} &= \mathbf{W}_m^{\sigma \text{enc}} \mathbf{Krylov}_m(X_m, A). \end{aligned} \quad (7)$$

We choose in this model to parametrize q_{η} by the posterior distribution of a **Krylov** Graph convolutional neural networks (GCN) [21]. The reason is that GCN are efficient to extract features information of a node considering its neighborhood [22, 23]. They have been widely used in node classification, node clustering, and other graph related tasks. Currently,

⁴with $Z \in \mathbb{R}^{d \times n}$ with $Z(:, i) \sim q_{\eta}(\cdot|X, A)$

⁵log in $\log \sigma_m$ is applied element-wise on matrix σ_m .

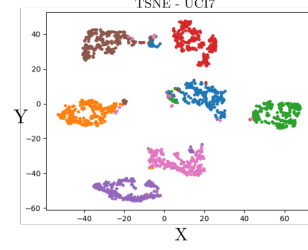


Fig. 2. t-SNE visualisation in 2D of the latent space ($d = 3$) for the uci7 dataset. Each color represents a different class.

many methods have similar performance [24], and since the choice is not critical, we simply use a truncated Krylov GCN architecture [21] which has been proven to have good properties when stacking across graph layers.

Finally, the trainable parameters of the models are the weights $\mathbf{\Psi}_m$, $\mathbf{W}_m^{\mu \text{dec}}$, $\mathbf{W}_m^{\sigma \text{enc}}$, $\mathbf{W}_m^{\mu \text{enc}}$, the parameters of the multilayer perceptron \mathbf{MLP}_m and of the Krylov GCN layers **Krylov** _{m} . The training can be done with the ELBO as loss function, using a suitable optimization method. The latent representations for each view x_m^i is $\mathbb{E}_{z \sim q_{\eta_m}(\cdot|\mathcal{V}(x_m^i), A)} = \mu_m^{\text{enc}}(\cdot, i)$. And the common latent representation is given by

$$\mathbb{E}_{z \sim q_{\eta}(\cdot|\mathcal{V}(x_m^i)_{m=1}^M, A)} = \left[\sum_{m=1}^M \frac{\mu_m^{\text{enc}}(k, i)}{\sigma_m^{\text{enc}2}(k, i)} / \sum_{m=1}^M \frac{1}{\sigma_m^{\text{enc}2}(k, i)} \right]_{k=1}^d.$$

This common representation is used for the experiments.

4. DATASETS

4.1. UCI handwritten datasets

UCI Handwritten digits dataset⁶ is a multiview dataset of $n = 2000$ small images representing digits. Each image has a label from 0 to 9 (200 elements for each) and 6 views of varying dimensions: $d_1 = 76$, $d_2 = 216$, $d_3 = 64$, $d_4 = 240$, $d_5 = 47$ and $d_6 = 6$; all views corresponding to specific transformations of the original image. Clustering and classification tasks are performed on this dataset (uci10) and on a partial version (uci7) where classes 0, 5 and 6 have been removed.

4.2. Twitter friend recommendation

A multiview dataset⁷ based on tweets from Twitter has been proposed in [5]. It consists of multiview representations of messages of users. There are $n = 102327$ users. Each user has 6 1000-dimensional views: EgoTweets, MentionTweets, FriendTweets, FollowersTweets, FriendNetwork, and FollowerNetwork. A task of friend recommendation is performed as follows. The followed accounts are known for each user; given a highly followed account and a part of their followers, the goal is to determine, for each other users, whether or not

⁶archive.ics.uci.edu/ml/datasets/Multiple+Features

⁷http://www.cs.jhu.edu/~mdredze/data/

Dataset	uci7			uci10			Recommendation		
Metric	Acc.	ARI	ARI2	Acc.	ARI	ARI2	Prec.	Recall	MRR
PCA	0.87	0.55	-	0.74	0.41	-	0.1511	0.0795	0.3450
GPCA	0.95	0.74	0.79	0.90	0.63	0.62	0.1578	0.0831	0.3649
MCCA	0.89	0.66	-	0.79	0.59	-	0.0815	0.0429	0.2225
GMCCA	0.96	0.84	0.83	0.92	0.73	0.76	0.2290	0.1206	0.4471
MVGCCA	0.97	0.79	0.85	0.95	0.69	0.76	0.1951	0.0665	0.4230

Table 1. Results of experiments on the different datasets and tasks; see text for comments. MRR is mean reciprocal rank.

he will follow this account after March 2015. For this task, a graph based on the Twitter dataset is built as in [18] with the views Egoweets, FollowersTweets, and FriendNetwork.

5. EXPERIMENTS

All architectures and hyperparameters specified here have been fixed for all experiments unless otherwise indicated⁸. We preprocess all the views: each view is centered and normalized by its maximal absolute value. For each dataset graph adjacency matrix is rescaled with his maximal coefficient and diagonal coefficients are set to 1.

Decoders: the MLP_m used as decoders have 4 hidden layers and 1024 units in each. We parametrize Ψ_m by a single scalar: $\Psi_m = (\sigma_m^2 + 10^{-6})I$. This choice reduces complexity and improves robustness without decreasing performance.

Encoders: the Krylov GCN layers [21] encoding the mean and variance in q_η , looks up to 3 hop neighborhood; we stack 4 layers and use 1024 units in each hidden layer.

General: Batch size is set to 512. A dropout regularization of rate 0.5 is also applied after all the hidden layers. The Adam optimizer is always used for training.

5.1. UCI classification and clustering

The model was trained on the uci7 and uci10 datasets. The loss function was the ELBO from (6), trained in batches. For each batch, the graph used is the subgraph for the samples in the batch. The latent space is of dimension $d = 3$.

Once having the latent representation of the elements in the uci dataset, 10-fold K-means clustering and spectral clustering are done. The quality of the clustering is evaluated with adjusted Rand index (ARI). A 10-fold SVM classification task is also performed. In order to select the best learning rate, number of epochs, dropout rate and K-means/SVM parameters, we performed 3 runs on different parameter combinations. The results are averaged on the 10-fold SVM classification. Finally we use the best performing parameters.

Following [18], the results are compared to PCA (applied on concatenated views), graph-regularized PCA, regular MCCA and GMCCA. In order to make fair comparison in terms of hyperparameter tuning, all the experiments were done with the same protocol. The results are in Table 1.

We see that MVGCCA is competitive in both classification (Acc.) and clustering tasks. It achieves the best performance on classification. For clustering, the quality with K-means (ARI1) on both datasets is slightly lower than the top state-of-the-art, while spectral clustering (ARI2) obtains the best score. This indicates that the graph structure seems to be well integrated into the latent space.

5.2. Twitter friend recommendation

For this dataset and the recommendation task, no further hyperparameter tuning is done, values from previous experiments are used. The parameters of the methods used for comparison are extracted from [13, 18] with their best parameters. The twitter dataset is large with more than 100,000 users which would be intractable for existing methods. Hence, 2506 twitter users are randomly selected from the database as in [18] for fair comparison.

The 20 most followed accounts (over the whole dataset) are selected, For each of these, 10 users following them are chosen (at random) and the average representation from latent space is computed. The latent space is set to dimension $d = 5$. Finally, the cosine similarity is computed between this average profile and the one of the $L = 100$ closest users to these representations. If one of these 100 users actually followed the initially chose account, this is considered as a good friend’s recommendation. To assess performance, precision, recall and MRR (mean reciprocal rank) metrics are computed. The results are in Table 1 (right).

The performance of MVGCCA comparable (except for recall) to that of GMCCA, which is currently the best method for this task. We recall here that the results are assessed on a limited dataset that all methods can process. However, the full version with 100,000 users would be intractable for GMCCA, while the proposed MVGCCA scales for this dataset size.

6. CONCLUSION

We proposed MVGCCA, a novel multiview and non linear extension of CCA based on a Bayesian inference model. Our model is scalable and can take into account the available graph structural information from the data. The probabilistic nature of the model can be used for addressing missing views or link prediction tasks which we postpone for future work.

⁸Code available : <https://github.com/Yacnnn/MVGCCA>

7. REFERENCES

- [1] M. Corvellec E. Quemener, ““SIDUS”, the solution for extreme deduplication of an operating system,” *The Linux Journal*, January 2014.
- [2] Y. Yamanishi, J.-P. Vert, A. Nakaya, and M. Kanehisa, “Extraction of correlated gene clusters from multiple genomic data by generalized kernel canonical correlation analysis,” *Bioinformatics*, vol. 19, no. suppl. 1, pp. i323–i330, 07 2003.
- [3] R. Arora and K. Livescu, “Multi-view cca-based acoustic features for phonetic recognition across speakers and domains,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 7135–7139.
- [4] R. T. Collins, A. J. Lipton, H. Fujiyoshi, and T. Kanade, “Algorithms for cooperative multisensor surveillance,” *Proceedings of the IEEE*, vol. 89, no. 10, pp. 1456–1477, 2001.
- [5] A. Benton, R. Arora, and M. Dredze, “Learning multiview embeddings of twitter users,” in *Proc. Annual Meeting Assoc. Comput. Linguistics*, Aug. 7-12 2016, vol. 2.
- [6] Y. Li, Ming Yang, and Z. Zhang, “A survey of multi-view representation learning,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, pp. 1863–1883, 2019.
- [7] H. Hotelling, “Relations between two sets of variates,” *Biometrika*, vol. 28, no. 3/4, pp. 321–377, 1936.
- [8] K Chaudhuri, SM Kakade, K Livescu, and K Sridharan, “Multi-view clustering via canonical correlation analysis,” in *Proceedings of the 26th Annual International Conference on Machine Learning*. 2009, ICML ’09, p. 129–136, Association for Computing Machinery.
- [9] J.R. Kettenring, “Canonical analysis of several sets of variables,” *Biometrika*, vol. 58, no. 3, pp. 433–451, 12 1971.
- [10] S. Akaho, “A kernel method for canonical correlation analysis,” in *In Proceedings of the International Meeting of the Psychometric Society (IMPS2001)*. 2001, Springer-Verlag.
- [11] G. Andrew, R. Arora, J. Bilmes, and K. Livescu, “Deep canonical correlation analysis,” Atlanta, Georgia, USA, 17–19 Jun 2013, vol. 28 of *Proceedings of Machine Learning Research*, pp. 1247–1255, PMLR.
- [12] W. Wang, R. Arora, K. Livescu, and J. Bilmes, “On deep multi-view representation learning,” Lille, France, 07–09 Jul 2015, vol. 37 of *Proceedings of Machine Learning Research*, pp. 1083–1092, PMLR.
- [13] X. Chang, T. Xiang, and T. M. Hospedales, “Scalable and effective deep CCA via soft decorrelation,” in *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. 2018, pp. 1488–1497, IEEE Computer Society.
- [14] D. Lopez-Paz, S. Sra, A. Smola, Z. Ghahramani, and B. Schölkopf, “Randomized nonlinear component analysis,” 22–24 Jun 2014.
- [15] F. R. Bach and M. I. Jordan, “A probabilistic interpretation of canonical correlation analysis,” Tech. Rep., Department of Statistics, University of California, Berkeley, 2005.
- [16] T. N. Kipf and M. Welling, “Variational graph auto-encoders,” *arXiv:stat.ML*, vol. 1611.07308, 2016.
- [17] J. Chen, G. Wang, Y. Shen, and G. B. Giannakis, “Canonical correlation analysis of datasets with a common source graph,” *IEEE Transactions on Signal Processing*, vol. 66, no. 16, pp. 4398–4408, 2018.
- [18] J. Chen, G. Wang, and G. B. Giannakis, “Graph multi-view canonical correlation analysis,” *IEEE Transactions on Signal Processing*, vol. 67, no. 11, pp. 2826–2838, 2019.
- [19] J. Rupnik, P. Skraba, J. Shawe-Taylor, and S. Guettes, “A comparison of relaxations of multiset canonical correlation analysis and applications,” *CoRR*, vol. abs/1302.0974, 2013.
- [20] D.P. Kingma and M. Welling, “Auto-encoding variational bayes,” in *2nd International Conference on Learning Representations, ICLR*, 2014.
- [21] S Luan, M Zhao, X-W Chang, and D Precup, “Break the ceiling: Stronger multi-scale deep graph convolutional networks,” in *Advances in Neural Information Processing Systems 32*, pp. 10945–10955. 2019.
- [22] M Defferrard, X Bresson, and P Vandergheynst, “Convolutional neural networks on graphs with fast localized spectral filtering,” in *Advances in Neural Information Processing Systems 29*, pp. 3844–3852. 2016.
- [23] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” *CoRR*, vol. abs/1609.02907, 2016.
- [24] F. Wu, T. Zhang, A. H. Souza Jr., C. Fifty, T. Yu, and K. Q. Weinberger, “Simplifying graph convolutional networks,” 09–15 Jun 2019, vol. 97 of *Proceedings of Machine Learning Research*, pp. 6861–6871.