

MULTIMODAL FUSION USING SPARSE CCA FOR BREAST CANCER SURVIVAL PREDICTION

Vaishnavi Subramanian¹, Tanveer Syeda-Mahmood², Minh N. Do^{1,3}

¹ Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, USA

² IBM Research, Almaden Research Center, San Jose, USA

³ College of Engineering and Computer Science, VinUniversity, Hanoi, Vietnam

ABSTRACT

Effective understanding of a disease such as cancer requires fusing multiple sources of information captured across physical scales by multimodal data. In this work, we propose a novel feature embedding module that derives from canonical correlation analyses to account for intra-modality and inter-modality correlations. Experiments on simulated and real data demonstrate how our proposed module can learn well-correlated multi-dimensional embeddings. These embeddings perform competitively on one-year survival classification of TCGA-BRCA breast cancer patients, yielding average F1 scores up to 58.69% under 5-fold cross-validation.

1. INTRODUCTION

In a complex disease such as cancer, the interactions between the tumor and host can exist at the molecular, cellular, tissue, and organism levels. Thus, evidence for the disease and its evolution may be present in multiple modalities across scale such as clinical, genomic, molecular, pathological, and radiological imaging. An improved disease understanding requires bridging scales of observation through multimodal fusion and is the focus of recent cancer research in survival prediction [1–3] and disease understanding [4, 5]. The majority of multimodal techniques focus on representation, alignment, and fusion of modalities. These include deep learning methods [1–3] and linear analyses [4–6].

Our focus in multimodal fusion is to identify the shared (common) information present across modalities for obtaining a robust characterization of the underlying problem. Accurate quantification of the shared information should account for the correlations within and across modalities that capture the underlying dependencies. Of the different fusion techniques, the classical formulation of canonical correlation analysis (CCA) [6] has been useful in discovering cross-modality correlations by identifying highly-correlated features from two modalities as a set of canonical variates. When applied to cancer, we have earlier shown that CCA enables the discovery of genotype-phenotype associations [4, 5]. Although the classical CCA formulation can model the cor-

relations across modalities, it does not explicitly capture the correlations and dependencies of features within each modality. To account for these intra-modality correlations, group and graph structures can be incorporated [7, 8] within a sparse CCA framework (SCCA) [9] to yield higher real data correlations. Further, while CCA-based feature fusion has been applied for downstream prediction tasks in image recognition [10] and autism [11], the utility of the learnt canonical variates is yet to be investigated under cancer settings and with sparse variants of CCA.

In this work, we capture intra-modality correlations through modality-specific graph representations and inter-modality correlations through the CCA objective to develop a linear feature embedding module. We propose an iterative scheme to generate projections of the two modalities' features onto multi-dimensional well-correlated spaces which can be used for downstream prediction tasks such as cancer survival, as outlined in Fig. 1.

2. METHOD

In this section, we review a graph-structured variant of sparse CCA and present how to generate multi-dimensional embeddings using an iterative approach. Finally, we present our proposed two-modality feature embedding.

2.1. Graph-based SCCA (1-GCCA)

The CCA formulation [6] considers matched observations of n samples $\mathbf{X} \in \mathbb{R}^{p \times n}$ and $\mathbf{Y} \in \mathbb{R}^{q \times n}$ from two views. The goal of CCA is to identify embedding directions $\mathbf{u} \in \mathbb{R}^p$ and $\mathbf{v} \in \mathbb{R}^q$ to maximize the correlation coefficient, between resulting 1-dimensional embeddings $\mathbf{u}^T \mathbf{X}$ and $\mathbf{v}^T \mathbf{Y}$:

$$\rho^* = \max_{\mathbf{u}, \mathbf{v}} \mathbf{u}^T \mathbf{X} \mathbf{Y}^T \mathbf{v}, \text{ s.t. } \|\mathbf{X}^T \mathbf{u}\|_2 = \|\mathbf{Y}^T \mathbf{v}\|_2 = 1.$$

Sparse CCA (SCCA) [9] further imposes sparsity in the entries of the embedding directions \mathbf{u} and \mathbf{v} as additional constraints. The different features of \mathbf{X} and \mathbf{Y} often demonstrate intra-modality correlations/dependencies which can

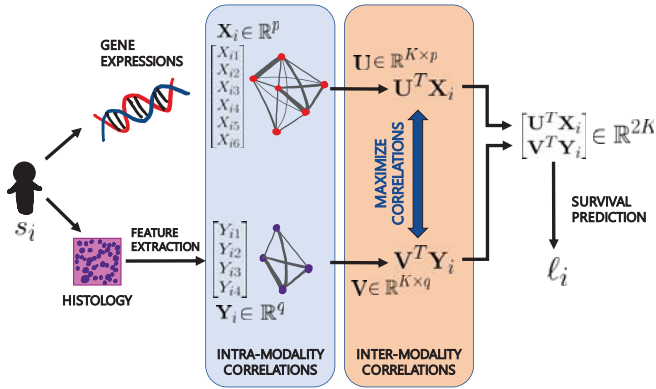


Fig. 1: Overview: We make use of intra-modality and inter-modality correlations to learn embedding matrices \mathbf{U} and \mathbf{V} which project to well-correlated spaces. The projections are used for downstream prediction tasks.

be estimated using sample covariance matrices and represented as underlying graphs for each modality. It is then additionally desired that the entries of embedding directions \mathbf{u} and \mathbf{v} weigh well-connected features on the graph similarly, such that underlying properties captured by these well-connected features are highlighted. This is enforced in the 1-dimensional graph-based CCA (1-GCCA) formulation:

$$\max_{\mathbf{u}, \mathbf{v}} \mathbf{u}^T \mathbf{X} \mathbf{Y}^T \mathbf{v} \text{ s.t. } \|\mathbf{u}^T \mathbf{X}\|_2 \leq 1, \|\mathbf{u}\|_1 \leq c_1, \mathbf{u}^T \mathbf{L}_1 \mathbf{u} \leq c_0, \\ \|\mathbf{v}^T \mathbf{Y}\|_2 \leq 1, \|\mathbf{v}\|_1 \leq d_1, \mathbf{v}^T \mathbf{L}_2 \mathbf{v} \leq d_0,$$

where c_0, c_1, d_0, d_1 are constants, \mathbf{L}_1 and \mathbf{L}_2 are the graph Laplacian matrices corresponding to the two respective modalities' underlying graphs.

This bi-convex problem can be solved to a local optimum using alternate optimization as shown in Algorithm 1. Algorithm 1 takes as input the correlation matrices $\Sigma_x = \mathbf{X} \mathbf{X}^T \in \mathbb{R}^{p \times p}$, $\Sigma_y = \mathbf{Y} \mathbf{Y}^T \in \mathbb{R}^{q \times q}$, cross-correlation matrix $\Sigma_{xy} = \mathbf{X} \mathbf{Y}^T \in \mathbb{R}^{p \times q}$, and graph Laplacians $\mathbf{L}_1 \in \mathbb{R}^{p \times p}$ and $\mathbf{L}_2 \in \mathbb{R}^{q \times q}$, and returns embedding vectors \mathbf{u} and \mathbf{v} .

2.2. Multi-dimensional Embeddings (K-GCCA)

The 1-GCCA algorithm provides a way to identify embedding directions \mathbf{u} and \mathbf{v} which maximize the correlations of the 1-dimensional embeddings $\mathbf{u}^T \mathbf{X}$ and $\mathbf{v}^T \mathbf{Y}$. However, 1-dimensional embeddings are often too restrictive and result in information loss. We extend the idea of 1-GCCA to identify K embedding directions $\mathbf{U} = \{\mathbf{u}^1 \dots \mathbf{u}^K\} \in \mathbb{R}^{p \times K}$, $\mathbf{V} = \{\mathbf{v}^1 \dots \mathbf{v}^K\} \in \mathbb{R}^{q \times K}$ using the Hotelling's deflation scheme, similar to K -factor penalized matrix decomposition [9]. We obtain 1-dimensional embeddings repeatedly, subtracting the contribution of already found embedding directions using projections followed by normalization. The algorithm for multi-dimensional graph-based CCA (K -GCCA) is described in Algorithm 2.

Algorithm 1: 1-Graph-based SCCA (1-GCCA)

Input: Correlation matrices $\Sigma_x, \Sigma_y, \Sigma_{xy}$, graph

Laplacian matrices $\mathbf{L}_1, \mathbf{L}_2$

Parameters: $\Theta = (\alpha_1, \beta_1, \lambda_1, \alpha_2, \beta_2, \lambda_2)$

Output: Embedding directions \mathbf{u} and \mathbf{v}

Initialize $\mathbf{u}_{(0)} = \mathbf{1}_p/p, \mathbf{v}_{(0)} = \mathbf{1}_q/q, i = 1$

while not converged do

$\mathbf{D}_u \leftarrow \text{diag}(1/|\mathbf{u}_{(i-1)}|), \mathbf{D}_v \leftarrow \text{diag}(1/|\mathbf{v}_{(i-1)}|)$

$\mathbf{u}_{(i)} \leftarrow (\alpha_1 \Sigma_x + \beta_1 \mathbf{D}_u + \lambda_1 \mathbf{L}_1)^{-1} \Sigma_{xy} \mathbf{v}_{(i-1)}$

$\mathbf{v}_{(i)} \leftarrow (\alpha_2 \Sigma_y + \beta_2 \mathbf{D}_v + \lambda_2 \mathbf{L}_2)^{-1} \Sigma_{xy}^T \mathbf{u}_{(i)}$

$i \leftarrow i + 1$

end

$\mathbf{u} \leftarrow \frac{\mathbf{u}_{(i-1)}}{\mathbf{u}_{(i-1)}^T \Sigma_x \mathbf{u}_{(i-1)}}, \mathbf{v} \leftarrow \frac{\mathbf{v}_{(i-1)}}{\mathbf{v}_{(i-1)}^T \Sigma_y \mathbf{v}_{(i-1)}}.$

Algorithm 2: K -Graph-based SCCA (K-GCCA)

Input: Data matrices $\mathbf{X} \in \mathbb{R}^{p \times n}, \mathbf{Y} \in \mathbb{R}^{q \times n}$, graph

Laplacian matrices $\mathbf{L}_1, \mathbf{L}_2$

Parameters: $K, \Theta = (\alpha_1, \beta_1, \lambda_1, \alpha_2, \beta_2, \lambda_2)$

Output: Embedding direction matrices

$\mathbf{U} = \{\mathbf{u}^1 \dots \mathbf{u}^K\}, \mathbf{V} = \{\mathbf{v}^1 \dots \mathbf{v}^K\}$

$\Sigma_x = \mathbf{X} \mathbf{X}^T, \Sigma_y = \mathbf{Y} \mathbf{Y}^T, \Sigma_{xy} = \mathbf{X} \mathbf{Y}^T$

for $k \in 1 \dots K$ **do**

$\mathbf{u}^k, \mathbf{v}^k \leftarrow \text{GCCA}_{\Theta}(\Sigma_x, \Sigma_y, \Sigma_{xy}, \mathbf{L}_1, \mathbf{L}_2)$

$\Sigma_{xy} \leftarrow \Sigma_{xy} - \frac{\langle \Sigma_{xy}, \mathbf{u}^k (\mathbf{v}^k)^T \rangle}{\|\mathbf{u}^k (\mathbf{v}^k)^T\|_2} \cdot \mathbf{u}^k (\mathbf{v}^k)^T$

$\Sigma_{xy} \leftarrow \Sigma_{xy} / \|\Sigma_{xy}\|_F$

end

2.3. Feature Embedding Module

We propose a feature embedding module that utilizes the K -GCCA algorithm. The module generates embedding matrices \mathbf{U} and \mathbf{V} , from which embeddings $\mathbf{X}_e \in \mathbb{R}^{K \times n}, \mathbf{Y}_e \in \mathbb{R}^{K \times n}$ of data matrices \mathbf{X} and \mathbf{Y} are generated as linear combinations $\mathbf{X}_e = \mathbf{U}^T \mathbf{X}, \mathbf{Y}_e = \mathbf{V}^T \mathbf{Y}$ and the final embedding of the two modalities is obtained as $\mathbf{Z}_e = \begin{bmatrix} \mathbf{X}_e \\ \mathbf{Y}_e \end{bmatrix} \in \mathbb{R}^{2K \times n}$.

3. EXPERIMENTS AND RESULTS

We first compare the 1-GCCA method with 1-SCCA on simulated data and breast cancer data to show that 1-GCCA learns better correlations than 1-SCCA. Next, we present correlations resulting from the multi-dimensional extensions, with $K = 100$. Finally, we employ the 100-GCCA embedding module to embed breast cancer data for one year survival prediction. All code and data is made available¹.

¹https://github.com/svaishnavi411/cca_fusion

In our experiments, we consider two different ways of defining the underlying graph structures in the algorithm: (i) using the squared correlation between features within the same modality as edge weights, and (ii) directly from prior knowledge.

3.1. Correlations on Simulated Data (1-GCCA)

Following previous works [7, 8], we generate data as follows. To construct \mathbf{u} , we generate a random fully-connected graph with all edge-weights as 1. The eigenvectors corresponding to the first l non-zero eigenvalues of the corresponding graph Laplacian are combined using randomly generated weights, and normalized to unit-norm to generate \mathbf{u} . The vector \mathbf{v} is set to be $\{10\ 3s, 10\ -1.5s, 10\ 1s, 10\ 2s, 60\ 0s\}$ and normalized to unit-norm. The data matrices \mathbf{X} and \mathbf{Y} are then generated such that columns $\mathbf{X}_i \sim \mathcal{N}(\mathbf{u}w_i, \sigma^2\mathbf{I}_{p \times p})$ and $\mathbf{Y}_i \sim \mathcal{N}(\mathbf{v}w_i, \sigma^2\mathbf{\Sigma}_v)$, where $\mathbf{\Sigma}_v[i, j] = \exp(-|v_i - v_j|)$ and $w_i \sim \mathcal{N}(0, 1)$. Data is generated for $n = 1000$ samples with $p = q = 100$, $l \in \{5, 10, 25, 50\}$ and $\sigma \in \{0.5, 0.75\}$, 25 times for each setting. Across settings and repetitions, \mathbf{v} is fixed, while \mathbf{u} and \mathbf{L}_1 change. For each setting and each repetition, the simulated data is split into 50%-10%-40% train, validation and test sets.

We compare 1-SCCA and 1-GCCA, with the best hyper-parameters chosen on the validation set. For 1-GCCA, we consider two methods of computing \mathbf{L}_1 : (i) from samples (1-GCCA), and (ii) feeding the true underlying \mathbf{L} as prior knowledge (1-GCCA-Prior). Table 1 shows the mean and standard deviation of different error metrics on the test set across different parameters ($l \in \{5, 10, 25, 50\}$ and $\sigma \in \{0.5, 0.75\}$). 1-GCCA-Prior has the advantage of the true graph and outperforms 1-SCCA and 1-GCCA across all error metrics. 1-GCCA performs better than 1-SCCA in the estimation of \mathbf{v} and correlation ρ .

Table 1: Simulated data: Overall mean and standard deviations of absolute cosine distance d_{\cos} in estimation of vectors \mathbf{u} , \mathbf{v} , absolute error in correlation ρ and relative spectral frequency $\mathbf{u}^T\mathbf{L}_1\mathbf{u}$. Lower values desired.

Error	1-SCCA	1-GCCA	1-GCCA-Prior
$d_{\cos}(\mathbf{u}, \hat{\mathbf{u}})$	44.54 ± 34.26	44.24 ± 24.31	31.95 ± 20.64
$d_{\cos}(\mathbf{v}, \hat{\mathbf{v}})$	43.89 ± 41.42	15.83 ± 20.61	11.86 ± 16.17
$ \rho - \hat{\rho} $	13.56 ± 13.37	10.19 ± 8.61	6.30 ± 6.75
$\frac{ \mathbf{u}^T\mathbf{L}_1\mathbf{u} - \hat{\mathbf{u}}^T\mathbf{L}_1\hat{\mathbf{u}} }{ \mathbf{u}^T\mathbf{L}_1\mathbf{u} }$	124.24 ± 142.92	105.74 ± 105.81	42.32 ± 60.02

3.2. Correlations on Breast Cancer Data (1-GCCA)

We work on histology imaging and gene expressions from the TCGA breast adenocarcinoma (BRCA) dataset of $n = 974$ patients to demonstrate the potential on real data.

The imaging data was acquired from the National Cancer Institute’s Genomic Data Commons portal. For the histology

Table 2: TCGA-BRCA: Mean and standard deviations of correlation coefficients on test set across folds. p denotes number of genes used in the correlation analysis.

p	SCCA	1-GCCA	1-GCCA-Prior
500	0.39 ± 0.11	0.50 ± 0.09	0.51 ± 0.08
800	0.41 ± 0.11	0.55 ± 0.09	0.52 ± 0.06
1000	0.46 ± 0.02	0.55 ± 0.09	0.49 ± 0.04
3000	0.36 ± 0.18	0.56 ± 0.09	0.40 ± 0.13

Table 3: TCGA-BRCA: Mean and standard deviations of the sum of correlations across the first 100 variates. p denotes number of genes used in the correlation analysis.

p	K-SCCA	K-GCCA	K-GCCA-Prior
500	19.75 ± 3.66	9.01 ± 0.81	8.84 ± 0.44
800	20.52 ± 5.49	9.42 ± 2.47	9.59 ± 3.01
1000	20.79 ± 5.87	11.66 ± 1.78	10.97 ± 1.45
3000	25.56 ± 3.73	11.85 ± 3.60	12.65 ± 1.62

images, we downloaded the nuclei segmentations from a recently published adversarial learning framework [12]. The nuclei segmentations are provided each patient in patches. We randomly selected 25 patches of size 2000 x 2000 pixels for each patient and fed the histology patch and segmentation mask to the CellProfiler tool to extract area, shape and texture properties for each nuclei and cell in the patch. Averaging these features across different patches yielded 213-dimensional imaging feature vectors for each patient.

The gene expression data was downloaded from the FireBrowse platform. We evaluated the most variant genes using the coefficient of variation (σ/μ) of the log2-transformed expression values. We selected the top 500, 800, 1000 and 3000 genes and the corresponding z-scores of the genes serve as the genomic feature vector for each patient. To extract prior-knowledge dependencies between genes we used the protein-protein interactions from the STRING database which captures biologically meaningful physical and functional interactions between proteins and assigned weights between any two interacting proteins’ corresponding genes.

We evaluate the methods with 5-fold cross validation using 544-137-292 patients in training-validation-test sets respectively. Table 2 reports results on the test set across folds. It is observed that 1-GCCA and 1-GCCA-Prior reveal higher correlations across different number of most variant genes (p). Further, as number of genes p increases, 1-GCCA is able to learn higher correlated embeddings.

3.3. Correlations on Breast Cancer Data (K-GCCA)

Next, we evaluate the K-GCCA proposed in Algorithm 2. Employing a similar approach with respect to SCCA, we obtain K-SCCA. The sum of the correlation coefficients of the first 100 directions $\{\mathbf{u}_1 \dots \mathbf{u}_{100}\}$ and $\{\mathbf{v}_1 \dots \mathbf{v}_{100}\}$ are re-

Table 4: TCGA-BRCA: Mean and standard deviations of different metrics on F1 % scores of one year survival prediction problem using single modalities, early fusion, late fusion and CCA-based fusion modules.

Method	500	800	1000	3000
Genomics	55.44 ± 1.90	58.39 ± 2.56	54.85 ± 2.80	58.36 ± 2.29
Imaging	60.92 ± 1.17	60.92 ± 1.17	60.92 ± 1.17	60.92 ± 1.17
Early Fusion	57.06 ± 5.55	58.61 ± 3.53	58.98 ± 1.01	60.97 ± 1.75
Late Fusion	53.44 ± 2.19	53.80 ± 3.20	52.02 ± 3.73	53.64 ± 4.04
100-SCCA	57.52 ± 2.91	59.09 ± 3.27	58.23 ± 2.57	56.53 ± 4.53
100-GCCA	56.36 ± 3.16	57.11 ± 3.02	57.92 ± 0.97	58.69 ± 2.16
100-GCCA-P	56.23 ± 2.23	58.52 ± 4.75	57.42 ± 1.84	57.71 ± 2.45

ported in Table 3. It is observed that K-SCCA yields higher sums across number of genes. Increasing the number of genes considered improves the sum of correlations for all methods.

3.4. Survival Prediction on Breast Cancer (K-GCCA)

We make use of the proposed K -GCCA fusion module with $K = 100$ to generate the embedding \mathbf{Z}_e to predict one-year survival of the TCGA-BRCA patients as a binary classification problem. We feed \mathbf{Z}_e to a random forest of 100 estimators with maximum depth $d = 50$. For baselines, we compare these features to \mathbf{X} only (Genomics), \mathbf{Y} only (Imaging) and $[\mathbf{X}^T, \mathbf{Y}^T]^T$ (Early fusion). Further, we combine the predictions of genomics and imaging in a simple late fusion module (Late fusion). We further utilize the embeddings from 100-SCCA as features for the random forest (100-SCCA). As earlier, we compare the construction of graphs from data (100-GCCA) and those from prior knowledge (100-GCCA-P).

The accuracy, support-weighted F1 and support-weighted AUC scores for the same test set as before are reported in Table 4. Among all CCA-fusion methods, we observe that 100-SCCA works best for the lower number of genes ($p = \{500, 800, 1000\}$), while 100-GCCA and 100-GCCA-P work best for the largest number of genes ($p = 3000$).

4. CONCLUSION

In this work, we proposed a novel feature embedding module for multi-modality fusion with two modalities which generates well-correlated low-dimensional embeddings by taking into account intra-modality correlations. We first demonstrated the importance of accounting for intra-modality correlations in the CCA formulation. We showed that our proposed feature embedding module generates low-dimensional embeddings of the two modalities while preserving the information important for one-year survival prediction of breast cancer patients. In the future we will investigate the use of better deflation schemes for generating higher-dimensional embeddings, and conduct an extensive study across different cancers to comprehensively evaluate CCA-based fusion methods for cancer survival prediction.

5. COMPLIANCE WITH ETHICAL STANDARDS

This research study was conducted retrospectively using human subject data made available in open access by TCGA Research Network: <https://www.cancer.gov/tcga>. Ethical approval was not required as confirmed by the license attached with the open access data.

6. ACKNOWLEDGMENTS

This project has been funded by the Jump ARCHES endowment through the Health Care Engineering Systems Center and the IBM-Illinois C3SR center.

7. REFERENCES

- [1] A. Cheerla et al., “Deep learning with multimodal representation for pancancer prognosis prediction,” *Bioinformatics*, vol. 35, no. 14, pp. i446–i454, 2019.
- [2] L. A. V. Silva et al., “Pan-cancer prognosis prediction using multimodal deep learning,” in *ISBI*, 2020.
- [3] V. Subramanian et al., “Multimodal fusion of imaging and genomics for lung cancer recurrence prediction,” in *2020 ISBI*. IEEE, 2020, pp. 804–808.
- [4] V. Subramanian et al., “Correlating cellular features with gene expression using CCA,” in *ISBI*. IEEE, 2018.
- [5] V. Subramanian et al., “Integration of spatial distribution in imaging-genetics,” in *MICCAI*. Springer, 2018.
- [6] H. Hotelling, “Relations between two sets of variates,” *Biometrika*, vol. 28, no. 3/4, pp. 321–377, 1936.
- [7] X. Chen et al., “An efficient optimization algorithm for structured sparse CCA, with applications to EQTL mapping,” *Statistics in Biosciences*, 2012.
- [8] L. Du et al., “GraphNet based sparse canonical correlation analysis for brain imaging genetics,” in *International Conference on Brain Informatics and Health*. Springer, 2015.
- [9] D. M. Witten et al., “A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis,” *Biostatistics*, 2009.
- [10] Q. Sun et al., “A new method of feature fusion and its application in image recognition,” *Pattern Recognition*, vol. 38, no. 12, pp. 2437–2448, 2005.
- [11] F. Zhao et al., “Feature fusion via hierarchical supervised local CCA for diagnosis of autism spectrum disorder,” *Brain imaging and behavior*, vol. 11, no. 4, pp. 1050–1060, 2017.
- [12] L. Hou et al., “Robust histopathology image analysis: to label or to synthesize?,” in *CVPR*, 2019, pp. 8533–8542.