
A Variational Information Bottleneck Approach to Multi-Omics Data Integration

Changhee Lee
University of California, Los Angeles

Mihaela van der Schaar
University of Cambridge
University of California, Los Angeles
Alan Turing Institute

Abstract

Integration of data from multiple omics techniques is becoming increasingly important in biomedical research. Due to non-uniformity and technical limitations in omics platforms, such integrative analyses on multiple omics, which we refer to as views, involve learning from incomplete observations with various view-missing patterns. This is challenging because i) complex interactions within and across observed views need to be properly addressed for optimal predictive power and ii) observations with various view-missing patterns need to be flexibly integrated. To address such challenges, we propose a deep variational information bottleneck (IB) approach for *incomplete multi-view* observations. Our method applies the IB framework on marginal and joint representations of the observed views to focus on intra-view and inter-view interactions that are relevant for the target. Most importantly, by modeling the joint representations as a product of marginal representations, we can efficiently learn from observed views with various view-missing patterns. Experiments on real-world datasets show that our method consistently achieves gain from data integration and outperforms state-of-the-art benchmarks.

1 Introduction

Technological advances in high-throughput biology enable integrative analyses that use information across

multiple omics layers – including genomics, epigenomics, transcriptomics, proteomics, and metabolomics – to deliver more comprehensive understanding in biological systems [1, 2]. Unfortunately, due to limitations of experimental designs or compositions from different data sources (e.g., TCGA¹), integrated samples commonly have one or more entirely missing omics with various missing patterns. Learning from such incomplete observations is challenging. Discarding samples with missing omics greatly reduces sample sizes (especially when integrating many omics layers) [3] and simple mean imputation can seriously distort the marginal and joint distribution of the data [4].

In this paper, we model multi-omics data integration as learning from *incomplete multi-view observations* where we refer to observations from each omics data as *views* (e.g., DNA copy number and mRNA expressions). However, direct application of the existing multi-view learning methods does not address the key challenges of handling missing views in integrating multi-omics data. This is because these methods are typically designed for complete-view observations assuming that all the views are available for every sample [5]. Therefore, we set our goal to develop a model that not only learns complex intra-view and inter-view interactions that are relevant for target tasks but also flexibly integrates observed views regardless of their view-missing patterns in a single framework.

Contribution. To this goal, we propose a deep variational information bottleneck (IB) approach for incomplete multi-view observations, which we refer to as DeepIMV. Our method consists of four network components: a set of *view-specific encoders*, a set of *view-specific predictors*, a *product-of-experts* (PoE) module, and a *multi-view predictor*. More specifically, for flexible integration of the observed views regardless of the view-missing patterns, we model the *joint* representations as a PoE over the marginal representations, which are further utilized by the multi-view predictor.

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s).

¹<http://cancergenome.nih.gov/>

Thus, the joint representations combine both common and complementary information across the observed views. The entire network is trained under the IB principle [6, 7] which encourages the marginal and joint representations to focus on intra-view and inter-view interactions that are relevant to the target, respectively. Experiments on real-world datasets show that our method consistently achieves gain from data integration and significantly outperforms state-of-the-art benchmarks with respect to measures of prediction performance.

2 Related Works

2.1 Multi-View Learning

Complete Multi-View Observations. To utilize information across multiple views, a variety of methods have been proposed in recent years. Canonical correlation analysis (CCA) [8] and its kernel or deep learning extensions [9, 10, 11] are representative methods which aim to learn a common (latent) space that is shared between the two views such that their canonical correlation is maximized in a purely unsupervised fashion. The learned representations can then be used for supervised learning as a downstream task. Meanwhile, some existing methods [12, 13, 14] have focused on fully utilizing the label information when seeking for the common (latent) space that are not only shared between the views but also discriminative for the target. Although these methods have shown promising performance in different applications [5], they are designed for complete-view observations.

Incomplete Multi-View Observations. Most of the existing methods require multi-stage training – that is, constructing complete multi-view observations based on imputation methods [15, 16] and then training a multi-view model – or relying on auxiliary inference steps to generate the missing views [17, 18, 19]. There are a few methods that flexibly handle incomplete multi-view observations. Authors in [20] introduced a generative model that utilizes latent factorization enabling cross-view generation without requiring multi-stage training regimes or additional inference steps. Matrix factorization was extended in [21] to learn low-dimensional representations that capture the joint aspects across different views. However, these methods are trained in a purely unsupervised fashion. Thus, while information relevant to the reconstruction of views will be well-captured in learned representations, information relevant to the target task may be lost.

Our work is most closely related to CPM-Nets [22]. Both methods aim to find representations in a common space for supervised learning of predicting the target. A

notable distinction from CPM-Nets is how we integrate incomplete multi-view observations. In CPM-Nets, the authors directly learn the mapping from latent representations to the original views without utilizing any encoder structure.² Instead, for all the training samples, the corresponding latent representations are randomly initialized and, then, iteratively updated to minimize the reconstruction and classification losses. We conjecture that this approach has two limitations: First, the method relies on reconstruction to find latent representations of the testing samples, which makes it difficult to capture task-relevant information. Second, random initialization means that there are no inherent relations across the latent representations of the training samples. Thus, all training samples must be updated at the same time, which increases the memory burden during training. In contrast, in DeepIMV, we utilize posterior factorization for flexible integration of observed views and focus latent representations to be task-relevant, thereby effectively capturing the complementary information for predicting the target. In addition, CPM-Nets is designed only for classification tasks; thus, extension to regression tasks can lose the advantages of cluster-friendly representations.

2.2 Information Bottleneck

The information bottleneck (IB) principle [6, 17, 23] is an information-theoretic approach that defines the intuitive idea of what a *task-relevant representation* is in terms of the fundamental trade-off between having a concise representation and one that provides good prediction power. The IB framework [7] has been recently studied to address multi-view problems [24, 25]. Authors in [24] combined marginal representations from each view using an auxiliary layer into joint representation on which the IB principle is applied. However, this work cannot handle missing views both during training and testing. In [25], the IB principle was extended to the two-view unsupervised setting to learn robust representations that are common to both views. Thus, this is not applicable to multi-omics data integration where the goal is to flexibly integrate incomplete views (i.e., observations from different omics layers), that often contain both common and complementary information, in a supervised fashion.

3 Incomplete Multi-View Problem

The presence of missing views remains an inevitable and prevalent problem in multi-omics data integration. To

²The “encoding networks” in [22] denote networks that reconstruct original views from latent representations; there are no such networks that map original views to latent representations (in our context, encoders).

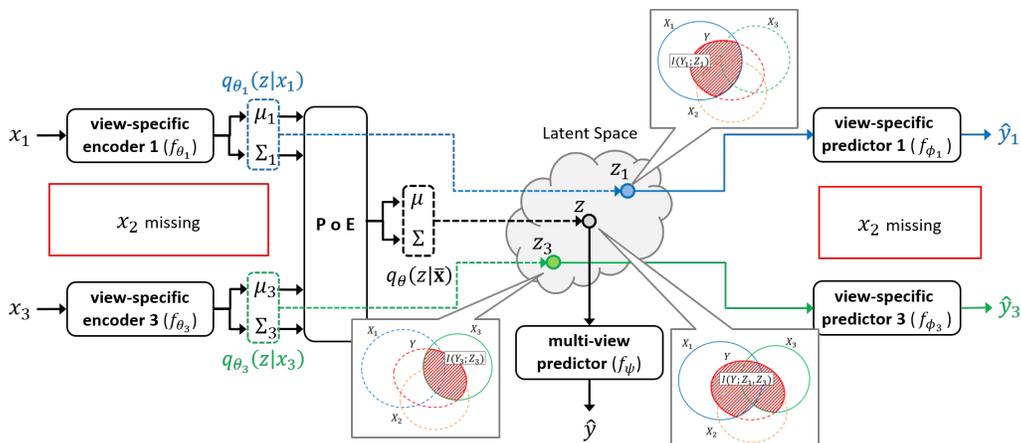


Figure 1: An illustration of the proposed network architecture with $V = 3$ views. For an illustration purpose, we assume that the current sample has the second view missing, i.e., $x_2 = \emptyset$. Here, dotted lines correspond to drawing samples in the latent representation from the respective distributions, and green and blue colored lines indicate marginal representations and the corresponding predictions.

address this, we start off by defining such an integrative analysis as an *incomplete multi-view problem* where some of the views may be missing with arbitrary view-missing patterns.

Notation. Let $X_v \in \mathcal{X}_v$ be a random variable for the d_v -dimensional input feature from the v -th view and $Y \in \mathcal{Y}$ be a random variable for the output label. We say the v -th view is missing if $x_v = \emptyset$ and is available if $x_v \in \mathcal{X}_v$ where x_v is a realization of X_v . To incorporate with arbitrary missing views, we denote a set of observed views by $\mathcal{V} = \{v : x_v \neq \emptyset\} \subseteq [V]$ which we refer to as a *view-missing pattern*. Here, $[V]$ is the complete set of available views. Then, we can finally define a random variable for a *multi-view observation* as $\bar{\mathbf{X}} = (X_v : v \in \mathcal{V})$; we say the observation is *incomplete* when $\mathcal{V} \subset [V]$ and *complete* when $\mathcal{V} = [V]$. Throughout the paper, we will often use lower-case letters to denote realizations of a random variable.

Definition 1. (Incomplete Multi-View Problem) Suppose that we have a training dataset $\mathcal{D} = \{\bar{\mathbf{x}}^n, y^n\}_{n=1}^N$ which contains one or more sample(s) with missing views, i.e., $\mathcal{V}^n \subset [V]$ for some n . Then, we define an *incomplete multi-view problem* as a supervised learning problem of predicting target y – that is, C -classification for $\mathcal{Y} = [C]$ and regression for $\mathcal{Y} = \mathbb{R}$ – for a new multi-view observation $\bar{\mathbf{x}}$ with an arbitrary view-missing pattern.³

Solving an incomplete multi-view problem involves two main challenges: First, we want to learn representations in a *common space* that leverages both *marginal* and *joint* aspects of the observed views for prediction. Second, the learned representations must flexibly integrate

incomplete observations with various view-missing patterns in a unified framework.

4 Method: DeepIMV

To address such challenges, we propose a deep variational information bottleneck approach, which we refer to as DeepIMV⁴, that consists of four network components as illustrated in Figure 1:

- a set of V *view-specific encoders*, parameterized by $\theta = (\theta_1, \dots, \theta_V)$, each of which stochastically maps observations from each individual view into a common latent space;
- a *product-of-experts* (PoE) module that integrates the marginal latent representations into a joint latent representation in the common space;
- a *multi-view predictor*, parameterized by ψ , which provides label predictions based on the joint representations; and
- a set of V *view-specific predictors*, parameterized by $\phi = (\phi_1, \dots, \phi_V)$, each of which provides label predictions based on the marginal latent representations determined by the corresponding view-specific encoders.

Throughout the remainder of the paper, we will often use f to denote a deterministic mapping and use q to denote a stochastic mapping.

4.1 Toward Task-Relevant Representations

Finding *task-relevant representations* in a common latent space that contains both *marginal* and *joint*

³Note that a similar definition has been introduced in [22] for classification tasks.

⁴Source code is available at <https://github.com/ch18856/DeepIMV>.

aspects of the observations is crucial for solving the incomplete multi-view problem. To this goal, we apply the IB principle [7, 23] since it aims at learning task-relevant representations by discarding as much information about the input as possible that is irrelevant to the target task and thereby encouraging the predictor to be robust to overfitting.

Define \mathcal{Z} be the common latent space. We consider the *marginal representation* $Z_v \in \mathcal{Z}$ be a stochastic encoding of X_v , i.e., $q_{\theta_v}(Z_v|X_v)$, which is defined by the v -th view-specific encoder for $v \in \mathcal{V}$. Similarly, we consider the *joint representation* $Z \in \mathcal{Z}$ be a stochastic encoding of $\bar{\mathbf{X}}$, which is defined by the encoder block $q_{\theta}(Z|\bar{\mathbf{X}})$ that combines outputs of the V view-specific encoders. Then, given a latent representation $z \in \mathcal{Z}$ drawn from $q_{\theta}(z|\bar{\mathbf{x}})$, the multi-view predictor estimates the target label Y based on the distribution defined as $q_{\psi}(Y|z)$. To learn *joint* aspects of the observed views for predicting the target Y , we apply the IB principle on Z based on the following loss:

$$\begin{aligned} \mathcal{L}_{\text{IB-J}}^{\theta, \psi}(\bar{\mathbf{X}}, Y) &= -I(Y; Z) + \beta I(\bar{\mathbf{X}}; Z) \\ &\approx \mathbb{E}_{\bar{\mathbf{x}}, y \sim p(\bar{\mathbf{x}}, y)} \mathbb{E}_{z \sim q_{\theta}(z|\bar{\mathbf{x}})} \left[-\log q_{\psi}(y|z) \right] \\ &\quad + \beta \mathbb{E}_{\bar{\mathbf{x}} \sim p(\bar{\mathbf{x}})} \left[KL(q_{\theta}(Z|\bar{\mathbf{x}}) \| q(Z)) \right] \end{aligned} \quad (1)$$

where $\beta \geq 0$ is a coefficient chosen to balance between the two information quantities. Here, $KL(p(Z) \| r(Z))$ denotes the Kullback-Leibler (KL) divergence between the two distributions $p(Z)$ and $r(Z)$. Detailed derivations can be found in the Supplementary Material.

4.2 Product-of-Experts for Incomplete Views

One question is outstanding. How should we design the joint representation such that it integrates the marginal representations of the observed views with arbitrary view-missing patterns? Our solution is to use a product-of-experts (PoE) that factorizes the joint posterior $q_{\theta}(Z|\bar{\mathbf{X}})$ into a *product* of the marginal posteriors $q_{\theta_v}(Z|X_v)$ for $v \in \mathcal{V}$. Formally, the joint posterior can be defined as the following [20]:

$$p(z|\bar{\mathbf{x}}) \approx C \cdot p(z) \prod_{v \in \mathcal{V}} q_{\theta_v}(z|x_v) \stackrel{\text{def}}{=} q_{\theta}(z|\bar{\mathbf{x}}) \quad (2)$$

where C is a normalizing constant.

Factorizing the joint posterior using a PoE is promising for solving an incomplete multi-view problem compared to that using a mixture-of-experts (MoE), i.e. $q_{\theta}(z|x_1, \dots, x_V) = \sum_{v=1}^V \alpha_v \cdot q_{\theta_v}(z|x_v)$: First, by employing PoE, we can simply ignore missing views when finding the joint representation Z of an input $\bar{\mathbf{X}}$ regardless of its view-missing patterns. Hence, we can fully

utilize samples with incomplete views without discarding them (or applying view-completion in advance) for training [19], and can avoid auxiliary inference steps of generating the missing views [17, 18] during both training and testing. Second, aside from its flexibility, a PoE (i.e., the joint posterior) can produce a much sharper distribution than the individual experts (i.e., the marginal posteriors) allowing for them to be specialized in a particular aspect of the target task [26]. This is a desirable property for multi-omics data integration where individual views often contain view-specific (complementary) information or uneven amount of information about the target.

Computing the Joint Posterior. Suppose that we utilize marginal posteriors of the form $q_{\theta_v}(z|x_v) = \mathcal{N}(z|\mu_v, \Sigma_v)$ where $(\mu_v, \Sigma_v) = f_{\theta_v}(x_v)$ and use the prior of the form $p(z) = \mathcal{N}(z|\mu_0, \Sigma_0)$ (typically, a spherical Gaussian). Because a product of Gaussian is itself Gaussian [27], we can derive the joint posterior as $q_{\theta}(z|\bar{\mathbf{x}}) = \mathcal{N}(z|\mu, \Sigma)$ where $\mu = (\mu_0 \Sigma_0^{-1} + \sum_{v \in \mathcal{V}} \mu_v \Sigma_v^{-1}) (\Sigma_0^{-1} + \sum_{v \in \mathcal{V}} \Sigma_v^{-1})^{-1}$ and $\Sigma = (\Sigma_0^{-1} + \sum_{v \in \mathcal{V}} \Sigma_v^{-1})^{-1}$. Hence, we can efficiently compute the joint posterior of the incomplete multi-view observations in terms of the available marginal posteriors.

4.3 Building View-Specific Expertise

However, training a PoE appears to be difficult, requiring artificial sub-sampling of the observed views [20] or applying variants of contrastive divergence [26] in order to ensure that the individual views are learnt faithfully. To mitigate such an issue, we instead introduce a set of V view-specific predictors and apply IB principle on the marginal representations in order to allow each encoder to build view-specific expertise for predicting the target.

Formally, for each view, given the latent representation $z_v \in \mathcal{Z}$ drawn from $q_{\theta_v}(z_v|x_v)$, the view-specific predictor estimates Y_v , which is a random variable for the target that can be solely described by the corresponding view X_v , based on the distribution defined as $q_{\phi_v}(Y_v|z_v)$. Then, we apply the IB principle on the marginal representations to capture view-specific aspects of the observed views for predicting the target based on the following loss:

$$\begin{aligned} \mathcal{L}_{\text{IB-M}}^{\theta_v, \phi_v}(X_v, Y_v) &= -I(Y_v; Z_v) + \beta_v I(X_v; Z_v) \\ &\approx \mathbb{E}_{x_v, y \sim p(x_v, y)} \mathbb{E}_{z \sim q_{\theta_v}(z|x_v)} \left[-\log q_{\phi_v}(y|z) \right] \\ &\quad + \beta_v \mathbb{E}_{x_v \sim p(x_v)} \left[KL(q_{\theta_v}(Z_v|x_v) \| q(Z_v)) \right] \end{aligned} \quad (3)$$

where $\beta_v \geq 0$ is a balancing coefficient. Minimizing (3) encourages Z_v to become a minimal sufficient statistics of X_v for Y_v [28], enforcing the marginal representations

Algorithm 1 Pseudo-code for training DeepIMV

Input: $\mathcal{D} = \{\bar{\mathbf{x}}^n, \mathbf{y}^n\}_{n=1}^N$, learning rate η , mini-batch size n_{mb} , coefficients $\alpha, \beta, \{\beta_v\}_{v=1}^V$, and μ_0, Σ_0
Output: DeepIMV parameters (θ, ϕ, ψ)
Initialize parameters (θ, ϕ, ψ) via Xavier Initializer
repeat
 Sample a mini-batch: $\{\bar{\mathbf{x}}^n, \mathbf{y}^n\}_{n=1}^{n_{mb}} \sim \mathcal{D}$
 for $n = 1, \dots, n_{mb}$ **do**
 for $v \in \mathcal{V}^n$ **do**
 $\mu_v^n, \Sigma_v^n \leftarrow f_{\theta_v}(x_v^n)$
 $z_v^n \leftarrow \mu_v^n + \epsilon \cdot \Sigma_v^n$ where $\epsilon \sim \mathcal{N}(0, 1)$
 $\hat{y}_v^n \leftarrow f_{\phi_v}(z_v^n)$
 end for
 $\mu^n, \Sigma^n \leftarrow PoE((\mu_v^n, \Sigma_v^n)_{v \in \mathcal{V}^n})$
 $z^n \leftarrow \mu^n + \epsilon \cdot \Sigma^n$ where $\epsilon \sim \mathcal{N}(0, 1)$
 $\hat{y}^n \leftarrow f_{\psi}(z^n)$
 end for
 $\hat{\mathcal{L}}_{IB-J}^{\theta, \psi} = \frac{1}{n_{mb}} \sum_{n=1}^{n_{mb}} (-\sum_{c=1}^C y^n[c] \log(\hat{y}^n[c]) + \beta \cdot KL(\mathcal{N}(\mu^n, \Sigma^n) || \mathcal{N}(\mu_0, \Sigma_0)))$
 $\hat{\mathcal{L}}_{IB-M}^{\theta, \phi} = \frac{1}{n_{mb}} \sum_{n=1}^{n_{mb}} \sum_{v \in \mathcal{V}^n} (-\sum_{c=1}^C y^n[c] \log(\hat{y}_v^n[c]) + \beta_v \cdot KL(\mathcal{N}(\mu_v^n, \Sigma_v^n) || \mathcal{N}(\mu_0, \Sigma_0)))$
 $(\theta, \phi, \psi) \leftarrow (\theta, \phi, \psi) - \eta \nabla_{(\theta, \phi, \psi)} (\hat{\mathcal{L}}_{IB-J}^{\theta, \psi} + \alpha \hat{\mathcal{L}}_{IB-M}^{\theta, \phi})$
until convergence

to learn the view-specific (possibly complementary) aspects of the target, which eventually ease the training of PoE.

4.4 Training

We train the overall network – the view-specific encoder-predictor pairs (θ, ϕ) and the final predictor (ψ) – by minimizing a combination of the marginal and joint IB losses:

$$\mathcal{L}_{\text{Total}}^{\theta, \phi, \psi}(\bar{\mathbf{X}}, Y) = \mathcal{L}_{\text{IB-J}}^{\theta, \psi}(\bar{\mathbf{X}}, Y) + \alpha \sum_{v \in \mathcal{V}} \mathcal{L}_{\text{IB-M}}^{\theta_v, \phi_v}(X_v, Y_v) \quad (4)$$

where $\alpha \geq 0$ is a hyper-parameter that trades off between the two losses. The pseudo-code of DeepIMV is provided in Algorithm 1.⁵

5 Experiments

Throughout the experiments, we evaluate different multi-view learning methods on two real-world multi-omics datasets that were collected by the Cancer Genome Atlas (TCGA)⁶ and by the Cancer Cell Line Encyclopedia (CCLE) [29] in the context of integrating multi-omics observations for predicting 1-year mortality and drug sensitivity of cancer cells, respectively.

Benchmarks. We compare DeepIMV with 2 baselines and 6 state-of-the-art multi-view learning methods:

⁵Here, we assume a C -classification task and slightly abuse notation; we will write y to denote a one-hot vector and write $y[c]$ to denote the c -th element of y .

⁶<https://www.cancer.gov/tcga>

The baselines include a pre-integration method that simply concatenates observations from multiple views (denoted as **Base1**) and a post-integration method that integrates predictions of individual predictors trained on each view as an ensemble (denoted as **Base2**). The multi-view learning methods that assume complete multi-view observations include **GCCA** [30], **DCCA** [10], and **DCCAE** [11], and the methods that flexibly integrate incomplete multi-view observations include **MVAE** [20], **CPM-Nets** [22], and **MOFA** [21]. Table 1 summarizes the key characteristics. It is worth highlighting that i) for the methods that cannot handle incomplete multi-view observations, we use mean values (and further utilize the reconstructed inputs of MVAE in Tables 3 and 6) to impute missing views of training and testing samples, and ii) for the methods that do not utilize label information, we train a multi-layer perceptron (MLP) based on the learned representations as a downstream task.

Table 1: Comparison table.

| Methods | Task Oriented | Incomplete Views |
|----------|---------------|------------------|
| Base1 | ✓ | ✗ |
| Base2 | ✓ | ✗ |
| GCCA | ✗ | ✗ |
| DCCA | ✗ | ✗ |
| DCCAE | ✗ | ✗ |
| MVAE | ✗ | ✓ |
| CPM-Nets | ✓ | ✓ |
| MOFA | ✗ | ✓ |
| DeepIMV | ✓ | ✓ |

To focus our experiments on the integrative analysis and to overcome “curse-of-dimensionality” in the high-dimensional multi-omics data, we extracted low-dimensional representations (i.e., 100 features) using the kernel-PCA (with polynomial kernels) on each view [31] to train all the multi-view learning methods except for MOFA, which is well-known for capturing sparse factors across multiple views. Please see the Supplementary Material for more details.

Implementations details and sensitivity analyses on the hyper-parameters (α, β) of DeepIMV and details on the benchmarks can be found in the Supplementary Material. Throughout the experiments, all the results are reported based on 10 and 100 random 64/16/20 train/validation/test splits for the TCGA dataset and the CCLE dataset, respectively.

5.1 Results: TCGA Dataset

Dataset Description. We analyze 1-year mortality based on the comprehensive observations from multiple omics on 7,295 cancer cell lines (i.e. samples). The

Table 2: Comparison of the AUROC performance (mean \pm 95%-CI) for (*complete*) the multi-view learning methods that are trained with only complete multi-view samples and (*incomplete*) those that are trained with both complete and incomplete multi-view samples. The values are reported by varying the number of observed views from $|\mathcal{V}^n| = 1$ to $|\mathcal{V}^n| = 4$ during testing.

| Methods | 1 View | | 2 Views | | 3 Views | | 4 Views | |
|----------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|
| | complete | incomplete | complete | incomplete | complete | incomplete | complete | incomplete |
| Base1 | 0.660 \pm 0.04 | 0.675 \pm 0.02 | 0.722 \pm 0.03 | 0.739 \pm 0.02 | 0.750 \pm 0.02 | 0.765 \pm 0.02 | 0.766 \pm 0.02 | 0.781 \pm 0.01 |
| Base2 | 0.711\pm0.02 | 0.717 \pm 0.02 | 0.746 \pm 0.01 | 0.766 \pm 0.00 | 0.767 \pm 0.02 | 0.775 \pm 0.01 | 0.783\pm0.02 | 0.790 \pm 0.01 |
| GCCA | 0.680 \pm 0.02 | 0.650 \pm 0.03 | 0.737 \pm 0.02 | 0.737 \pm 0.03 | 0.764 \pm 0.01 | 0.769 \pm 0.02 | 0.783\pm0.01 | 0.792 \pm 0.01 |
| DCCA | 0.702 \pm 0.01 | 0.638 \pm 0.03 | 0.745 \pm 0.03 | 0.761 \pm 0.02 | 0.758 \pm 0.02 | 0.775 \pm 0.01 | 0.776 \pm 0.02 | 0.784 \pm 0.01 |
| DCCAE | 0.623 \pm 0.04 | 0.605 \pm 0.04 | 0.747 \pm 0.03 | 0.763 \pm 0.01 | 0.774 \pm 0.02 | 0.775 \pm 0.01 | 0.776 \pm 0.02 | 0.778 \pm 0.02 |
| MVAE | 0.592 \pm 0.05 | 0.589 \pm 0.04 | 0.677 \pm 0.02 | 0.674 \pm 0.02 | 0.731 \pm 0.02 | 0.730 \pm 0.01 | 0.774 \pm 0.01 | 0.781 \pm 0.01 |
| CPM-Nets | 0.700 \pm 0.02 | 0.709 \pm 0.01 | 0.748 \pm 0.02 | 0.761 \pm 0.02 | 0.766 \pm 0.01 | 0.771 \pm 0.01 | 0.781 \pm 0.01 | 0.788 \pm 0.01 |
| MOFA | 0.681 \pm 0.03 | 0.646 \pm 0.01 | 0.732 \pm 0.01 | 0.734 \pm 0.01 | 0.756 \pm 0.01 | 0.764 \pm 0.02 | 0.781 \pm 0.02 | 0.785 \pm 0.02 |
| DeepIMV | 0.701 \pm 0.02 | 0.724\pm0.02 | 0.757\pm0.02 | 0.772\pm0.01 | 0.776\pm0.01 | 0.791\pm0.01 | 0.783\pm0.01 | 0.801\pm0.01 |

data consists of observations from 4 distinct views on each cell line across 3 different omics layers: (**View 1**) mRNA expressions, (**View 2**) DNA methylation, (**View 3**) microRNA expressions, and (**View 4**) reverse phase protein array. Among 7,295 samples, 3,282 samples have incomplete multi-view observations with various view-missing patterns: average missing rates were 0.10, 0.24, 0.13, and 0.74 for View 1, View 2, View 3, and View 4, respectively.

5.1.1 Multi-Omics Data Integration

We start off by exploring the benefit of integrating multi-omics data, specifically augmenting samples with incomplete multi-view observations. To this goal, we compare predictions of different multi-view learning methods in terms of area under the receiver operating characteristics (AUROC) in Table 2 by augmenting $N_I = 3,282$ samples with incomplete views on the already available $N_C = 3,210$ samples with complete views for training. For evaluating the performance, we artificially created missing views for hold-out testing samples by varying the number of observed views from $|\mathcal{V}^n| = 1$ to $|\mathcal{V}^n| = 4$.

There are a couple of things to be highlighted from Table 2: First, DeepIMV better integrates samples with incomplete views as the performance improvement were most significant outperforming the benchmarks regardless of the number of observed views. Second, even when trained only with complete-view samples, our method better handles different view missing patterns during testing as it provided the highest performance (except for 1 View) with partially observed views. Third, MVAE and MOFA sacrifice their discriminative power since the latent representations focus on retaining the information of the input for view generation (reconstruction), which results in discarding the task-relevant discriminative information.

In addition, the proposed method still outperformed the benchmarks that do not handle incomplete views, when we replaced the mean imputation method with an

Table 3: Comparison of the AUROC performance (mean \pm 95%-CI) with different view completion methods on the TCGA dataset with $|\mathcal{V}^n| = 3$.

| Methods | Mean Impt. | MVAE Impt. |
|----------|----------------------------------|------------------|
| Base1 | 0.765 \pm 0.02 | 0.771 \pm 0.01 |
| Base2 | 0.775 \pm 0.01 | 0.784 \pm 0.01 |
| GCCA | 0.769 \pm 0.02 | 0.774 \pm 0.01 |
| DCCA | 0.775 \pm 0.01 | 0.784 \pm 0.02 |
| DCCAE | 0.775 \pm 0.01 | 0.773 \pm 0.02 |
| MVAE | 0.730 \pm 0.01 | – |
| CPM-Nets | 0.771 \pm 0.01 | – |
| MOFA | 0.764 \pm 0.02 | – |
| DeepIMV | 0.791\pm0.01 | – |

advanced multi-view imputation method (i.e., MVAE) as shown in Table 3. (More results can be found in the Supplementary Material.)

Visualization. We visually compare the principle component analysis (PCA) projections of the latent representations of the proposed method with different combinations of observed views in Figure 2. Here, DeepIMV is trained with both complete and incomplete multi-view samples. The PCA projections of latent representations become more discriminative and representative of the target labels as DeepIMV incorporates more views during testing. This can be highlighted by how the PCA projections of a representative cell line labeled with 1-year mortality, that are marked by star dots, move toward the opposite direction of the class boundary as the proposed method collects more views from $\mathcal{V}^n = \{2\}$ (which provides the highest $I(Y_v; Z_v)$; see Table 5) to $\mathcal{V}^n = \{1, 2, 3, 4\}$. Also, DeepIMV was able to achieve very similar representations of complete views without incorporating all the observed views. In particular, the PCA projections of latent representations without incorporating View 4 (which provided the smallest $I(Y_v; Z_v)$; see Table 5), i.e., $\mathcal{V} = \{1, 2, 3\}$, were almost the same with those of latent representations with $\mathcal{V} = \{1, 2, 3, 4\}$. This highlights a potential role of DeepIMV in designing multi-omics experiments

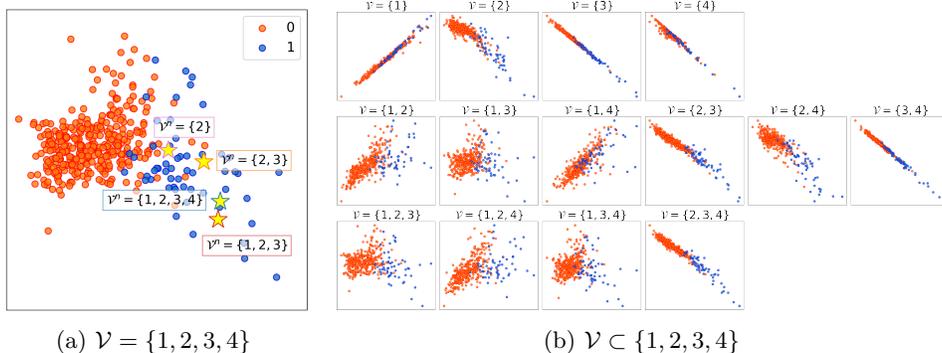


Figure 2: Comparison of PCA projections of latent representations from DeepIMV trained with both complete and incomplete multi-view samples for the TCGA dataset. Here, all possible combinations of different view-missing patterns $\mathcal{V} \subseteq [4]$ are illustrated.

Table 4: Ablation study results. AUROC performances (mean \pm 95%-CI) are reported for DeepIMV variations trained with both complete and incomplete multi-view samples on the TCGA dataset.

| Methods | 1 View | 2 Views | 3 Views | 4 Views |
|-----------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|
| MoE | 0.629 \pm 0.03 | 0.691 \pm 0.02 | 0.736 \pm 0.02 | 0.768 \pm 0.01 |
| MoE with marginal IBs | 0.712 \pm 0.01 | 0.766 \pm 0.01 | 0.786 \pm 0.01 | 0.790 \pm 0.01 |
| PoE | 0.655 \pm 0.04 | 0.719 \pm 0.03 | 0.755 \pm 0.03 | 0.783 \pm 0.02 |
| PoE with marginal IBs | 0.724\pm0.02 | 0.772\pm0.01 | 0.791\pm0.01 | 0.801\pm0.01 |

by providing advice which omics layers should not be measured for cost-efficient predictions.

Table 5: Information quantities on the TCGA dataset.

| $I(Y_1; Z_1)$ | $I(Y_2; Z_2)$ | $I(Y_3; Z_3)$ | $I(Y_4; Z_4)$ | $I(Y; Z)$ |
|---------------|---------------|---------------|---------------|-----------|
| 0.319 | 0.506 | 0.487 | 0.157 | 0.562 |

5.1.2 Ablation Study

In Table 4, we study the effect of i) utilizing PoE over MoE, and ii) introducing view-specific predictors and the marginal IB losses in (3) in order to provide additional insight into the source of gain. As discussed in Section 4.2, a PoE allows for encoders to specialize in analyzing their corresponding views and to build different expertise. Furthermore, introducing view-specific predictors and marginal IB losses encourage the view-specific encoders to focus on (possibly complementary) task-relevant information of each view and, thus, to ease the training associated with the PoE factorization. We observe that these components clearly contributes to DeepIMV’s performance improvement. Here, we use the same experimental setting in Section 5.1.1.

5.2 Results: CCLE Dataset

Dataset Description. We analyze sensitivities of heterogeneous cell lines to 4 different drugs – that are,

Irinotecan, *Panobinostat*, *Lapatinib*, and *PLX4720* – based on the multiple omics observations on 504 cancer cell lines (i.e. samples). Drug response was converted to a binary label by dividing cell lines into quartiles ranked by ActArea; the top 25% were assigned to the “sensitive” class and the rest were assigned to the “non-sensitive” class. The data consists of observations from 6 distinct views on each cell line across 5 different omics layers: (**View 1**) DNA copy number, (**View 2**) DNA methylation, (**View 3**) mRNA expressions, (**View 4**) microRNA expressions, (**View 5**) reverse phase protein array, and (**View 6**) metabolites.

Incomplete View Construction. We artificially construct incomplete multi-view observations based on the following procedures: i) randomly select $N_I = N \times R$ samples where $R \in [0, 1]$ is the rate of samples with missing views and ii) create incomplete multi-view observations by choosing one of the $(2^V - 2)$ possible view-missing patterns for each sample. (Obviously, at least one view must be observed).

5.2.1 Multi-Omics Data Integration

We explore the benefit of incorporating more samples and more views on predicting drug sensitivities of heterogeneous cell lines. To this end, we increase the set of available views from Views $\{1, 2\}$ to Views $\{1, 2, 3, 4, 5, 6\}$ and include samples that have observations from at least one of the views in this set. For example, given the set of available views is $\{1, 2\}$, we

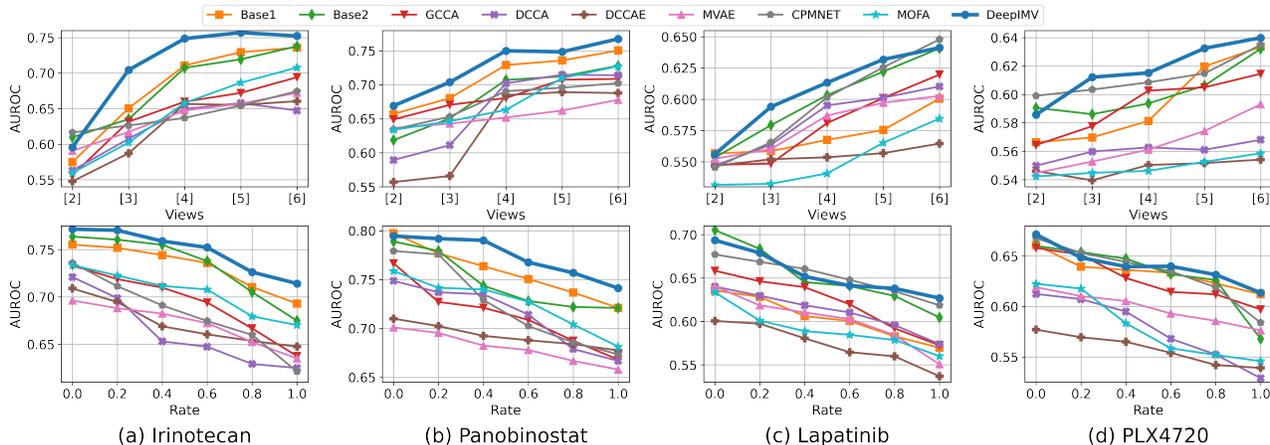


Figure 3: Average AUROC performance on four different drugs in the CCLE dataset (top) with increasing views and (bottom) with increasing view-missing rates.

integrate all the samples that satisfy $\mathcal{V}^{(n)} \cap \{1, 2\} \neq \emptyset$ for training. In the top row of Figure 3, we compare the prediction of different multi-view methods in terms of AUROC performance as we increase the set of available views from $\{1, 2\}$ to $\{1, 2, 3, 4, 5, 6\}$ with $R = 0.6$. There are a couple of things to be highlighted from this figure: First, our method provides better discriminative performance on all the tested drug sensitivity datasets (most of the times) as the number of integrated views increases. Second, the performances of DCCA and DCCAE are saturated since these methods can utilize only two views at the most, whereas GCCA provides consistently-increasing performance since it generalizes to multiple views. Third, MVAE and MOFA sacrifice their discriminative task-relevant information since the latent representations focus on retaining the information of the input for view generation (reconstruction).

Table 6: Comparison of the AUROC performance (mean \pm 95%-CI) with different view completion methods on the CCLE dataset (*Panobinostat*) with $M = 6$ and $R = 0.6$.

| Methods | Mean Impt. | MVAE Impt. |
|-----------------|----------------------------------|------------------|
| Base1 | 0.751 \pm 0.01 | 0.758 \pm 0.01 |
| Base2 | 0.728 \pm 0.02 | 0.752 \pm 0.01 |
| GCCA | 0.709 \pm 0.01 | 0.715 \pm 0.01 |
| DCCA | 0.714 \pm 0.01 | 0.717 \pm 0.01 |
| DCCAE | 0.688 \pm 0.01 | 0.697 \pm 0.01 |
| MVAE | 0.678 \pm 0.01 | – |
| CPM-Nets | 0.702 \pm 0.01 | – |
| MOFA | 0.727 \pm 0.02 | – |
| DeepIMV | 0.768\pm0.01 | – |

In addition, Table 6 shows that the proposed method still outperformed the benchmarks that do not handle incomplete views, when the mean imputation method is replaced by an advanced multi-view imputation method

(i.e., MVAE). (More results can be found in the Supplementary Material.)

5.2.2 Robustness to Missing Views

Next, we evaluate how robust the multi-view learning methods are with respect to the view-missing rate. The bottom row of Figure 3 shows the AUROC performance as the rate of samples with missing views ranges from $R = 0.0$ (all complete) to $R = 1.0$ (all incomplete) with $V = 6$. We highlight the following observations: First, our method outperforms all the benchmarks on the *Irinotecan* and *Panobinostat* datasets and provides comparable performance on the *Lapatinib* and *PLX4720* datasets to the best performing benchmark across different missing rates. Second, while other methods often fail, DeepIMV provides the most robust performance as the rate of samples with missing views increases. Third, DCCA and DCCAE show poor performance since these methods do not fully utilize the available views. Last, a similar trend can be found for MVAE and MOFA with that of the previous observation in Section 5.2.1.

6 Conclusion

Our proposed method finds intra-view and inter-view interactions that are relevant for predicting the target labels by flexibly integrating available views regardless of the view missing patterns in a unified framework. Throughout the experiments, we evaluate DeepIMV on real-world multi-omics datasets and show our method significantly outperforms existing multi-view learning methods in terms of prediction performance. In the future, further work may investigate incorporating sparsity in different omics data to address high-dimensionality.

Acknowledgment

This work was supported by the National Science Foundation (NSF) (Grant Number: 1722516) and the Office of Naval Research (ONR).

References

- [1] Y. Hasin, M. Seldin, and A. Lusic. Multi-omics approaches to disease. *Genome Biology*, 18(83), 2017.
- [2] I. Subramanian, S. Verma, S. Kumar, A. Jere, and K. Anamika. Multi-omics data integration, interpretation, and its application. *Bioinformatics and Biology Insights*, 14, 2020.
- [3] N. Rappoport and R. Shamir. NEMO: Cancer subtyping by integration of partial multi-omic data. *Bioinformatics*, 35(18):3348–3356, 2019.
- [4] R. J. A. Little and D. B. Rubin. *Statistical Analysis with Missing Data, 2nd Edition*. Wiley, 2002.
- [5] Y. Li, M. Yang, and Z. Zhang. A survey of multi-view representation learning. *IEEE Transactions on Knowledge and Data Engineering*, 31:1863–1883, 2019.
- [6] N. Tishby, F. C. Pereira, and W. Bialek. The information bottleneck method. In *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing (Allerton 1999)*, 1999.
- [7] Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. Deep variational information bottleneck. In *Proceedings of the 5th International Conference on Learning Representations (ICLR 2017)*, 2017.
- [8] H. Hotelling. Relations between two sets of variates. *Biometrika*, 28(3/4):321–377, 1936.
- [9] S. Akaho. A kernel method for canonical correlation analysis. *arXiv preprint arXiv:cs/0609071*, 2006.
- [10] G. Andrew, R. Arora, J. Bilmes, and K. Livescu. Deep canonical correlation analysis. In *Proceedings of the 30th International Conference on Machine Learning (ICML 2013)*, 2013.
- [11] W. Wang, R. Arora, K. Livescu, and J. Bilmes. On deep multi-view representation learning. In *Proceedings of the 32th International Conference on Machine Learning (ICML 2015)*, 2015.
- [12] T. Diethe, D. R. Hardoon, and J. Shawe-Taylor. Multiview fisher discriminant analysis. In *NIPS Workshop on Learning from Multiple Sources*, 2008.
- [13] M. Kan, S. Shan, H. Zhang, S. Lao, and X. Chen. Multi-view discriminant analysis. In *Proceedings of the 12th European Conference on Computer Vision (ECCV 2012)*, 2012.
- [14] Kui Jia, Jiehong Lin, Mingkui Tan, and Dacheng Tao. Deep multi-view learning using neuron-wise correlation-maximizing regularizers. *IEEE Transactions on Image Processing*, 28(10), October 2019.
- [15] T. Cai, T. T. Cai, and A. Zhang. Structured matrix completion with applications to genomic data integration. *Journal of the American Statistical Association*, 111(514):621–623, 2016.
- [16] L. Tran, X. Liu, J. Zhou, and R. Jin. Missing modalities imputation via cascaded residual autoencoder. In *Proceedings of the 30th IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2017)*, 2017.
- [17] M. R. Amini, N. Usunier, and C. Goutte. Learning from multiple partially observed views – an application to multilingual text categorization. In *Proceedings of the 22nd Conference on Neural Information Processing Systems (NIPS 2009)*, 2009.
- [18] Y.-H. H. Tsai, P. P. Liang, A. Zadeh, L.-P. Morency, and R. Salakhutdinov. Learning factorized multimodal representations. In *Proceedings of the 7th International Conference on Learning Representations (ICLR 2019)*, 2019.
- [19] Y. Shi, N. Siddharth, B. Paige, and P. H.S. Torr. Variational mixture-of-experts autoencoders for multi-modal deep generative models. In *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*, 2019.
- [20] M. Wu and N. Goodman. Multimodal generative models for scalable weakly-supervised learning. In *Proceedings of the 32nd Conference on Neural Information Processing Systems (NeurIPS 2018)*, 2018.
- [21] R. Argelaguet, B. Velten, D. Arno, S. Dietrich, T. Zenz, J. Marioni, F. Buettner, W. Huber, and O. Stegle. Multi-omics factor analysis—a framework for unsupervised integration of multi-omics data sets. *Mol Syst Biol*, 14(6):3348–3356, 2018.

- [22] Changqing Zhang, Zongbo Han, Yajie Cui, Huazhu Fu, Joey T. Zhou, and Qinghua Hu. CPM-nets: Cross partial multi-view networks. *In Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*, 2019.
- [23] A. Achille and S. Soatto. Emergence of invariance and disentanglement in deep representations. *Journal of Machine Learning Research*, 19, 2018.
- [24] Q. Wang, C. Boudreau, Q. Luo, P.-N. Tan, and J. Zhou. Deep multi-view information bottleneck. *In Proceedings of the 2019 SIAM International Conference on Data Mining*, 2019.
- [25] M. Federici, A. Dutta, P. Forré, N. Kushman, and Z. Akata. Learning robust representations via multi-view information bottleneck. *In Proceedings of the 8th International Conference on Learning Representations (ICLR 2020)*, 2020.
- [26] Geoffrey E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14(8):1771–1800, 2002.
- [27] Y. Cao and D. J. Fleet. Generalized product of experts for automatic and principled fusion of gaussian process predictions. *arXiv preprint arXiv:1410.7827*, 2014.
- [28] O. Shamir, S. Sabato, and N. Tishby. Learning and generalization with the information bottleneck. *Theor. Comput. Sci.*, 411(29):2696–2711, June 2010.
- [29] J. Barretina, G. Caponigro, and N. Stransky et al. The Cancer Cell Line Encyclopedia enables predictive modelling of anticancer drug sensitivity. *Nature*, 483:603–607, 2012.
- [30] J. R. Kettenring. Canonical analysis of several sets of variables. *Biometrika*, 58(3):433–451, 1971.
- [31] Y. Shiokawa, Y. Date, and J. Kikuchi. Application of kernel principal component analysis and computational machine learning to exploration of metabolites strongly associated with diet. *Scientific Reports*, 8(3426), February 2018.