

MBZUAI

Digital.Commons@MBZUAI

---

Computer Vision Faculty Publications

Scholarly Works

---

2-3-2022

## SubOmiEmbed: Self-supervised representation learning of multi-omics data for cancer type classification

Sayed Hashim

*Mohamed Bin Zayed University Of Artificial Intelligence*

Muhammad Ali

*Mohamed Bin Zayed University Of Artificial Intelligence*

Karthik Nandakumar

*Mohamed Bin Zayed University Of Artificial Intelligence*

Mohammad Yaqub

*Mohamed Bin Zayed University Of Artificial Intelligence*

Follow this and additional works at: <https://dclibrary.mbzuai.ac.ae/cvfp>



Part of the [Artificial Intelligence and Robotics Commons](#), [Biomedical Engineering and Bioengineering Commons](#), and the [Medicine and Health Sciences Commons](#)

Preprint: arXiv

Archived with thanks to arXiv

Preprint License: CC BY 4.0

Uploaded 25 March 2022

---

### Recommended Citation

S. Hashim, M. Ali, K. Nandakumar, and M. Yaqub, "SubOmiEmbed: Self-supervised representation learning of multi-omics data for cancer type classification", 2022. arXiv:2202.01672

This Article is brought to you for free and open access by the Scholarly Works at Digital.Commons@MBZUAI. It has been accepted for inclusion in Computer Vision Faculty Publications by an authorized administrator of Digital.Commons@MBZUAI. For more information, please contact [libraryservices@mbzuai.ac.ae](mailto:libraryservices@mbzuai.ac.ae).

---

# SubOmiEmbed: Self-supervised Representation Learning of Multi-omics Data for Cancer Type Classification

---

Sayed Hashim    Muhammad Ali    Karthik Nandakumar    Mohammad Yaqub

Mohamed Bin Zayed University of Artificial Intelligence, UAE

{sayed.hashim, muhammad.ali, karthik.nandakumar, mohammad.yaqub}@mbzuai.ac.ae

## Abstract

For personalized medicines, very crucial intrinsic information is present in high dimensional omics data which is difficult to capture due to the large number of molecular features and small number of available samples. Different types of omics data show various aspects of samples. Integration and analysis of multi-omics data give us a broad view of tumours, which can improve clinical decision making. Omics data, mainly DNA methylation and gene expression profiles are usually high dimensional data with a lot of molecular features. In recent years, variational autoencoders (VAE) [13] have been extensively used in embedding image and text data into lower dimensional latent spaces. In our project, we extend the idea of using a VAE model for low dimensional latent space extraction with the self-supervised learning technique of feature subsetting. With VAEs, the key idea is to make the model learn meaningful representations from different types of omics data, which could then be used for downstream tasks such as cancer type classification. The main goals are to overcome the curse of dimensionality and integrate methylation and expression data to combine information about different aspects of same tissue samples, and hopefully extract biologically relevant features. Our extension involves training encoder and decoder to reconstruct the data from just a subset of it. By doing this, we force the model to encode most important information in the latent representation. We also added an identity to the subsets so that the model knows which subset is being fed into it during training and testing. We experimented with our approach and found that SubOmiEmbed produces comparable results to the baseline OmiEmbed [25] with a much smaller network and by using just a subset of the data. This work can be improved to integrate mutation-based genomic data as well.

## 1 Introduction

The word omics refers to areas of studies in biological sciences which end with omics, such as genomics, proteomics, transcriptomics, or metabolomics [21]. The advent of Next Generation Sequencing (NGS) has brought about huge amounts of multi-omics data. Different types of omics data show various aspects of samples. Omics data, mainly DNA methylation and gene expression profiles are usually high dimensional data with a lot of molecular features from less samples. This makes it hard to extract insights from them. This is called 'the curse of dimensionality' in machine learning.

Omics profiling usually tells us the differences associated with a disease from an omics point of view. Such profiling can be used as markers of the disease, in extracting insights about the biological pathways of the disease as well as about treatment options [11]. Moreover, integration and analysis of multi-omics data give us a broader view of tumours, which can improve clinical decision making,

as each type of profiling reveals particular characteristics of the tumours. Therefore, it is safe to say that extracting information through intelligent computational tools can go a long way in improving patient outcomes.

Huge number of molecular features of less number of samples makes it challenging to build computational tools to extract insights from genome-wide high dimensional multi-omics data. As most of the omics profiling nowadays is done genome-wide, the dimensionality of such data is pretty large. For instance, RNA-Seq gene expression profiling comprises of more than 60,000 values for a single sample, while DNA methylation profiling contains more than 485,000 values for a sample. At the same time, the number of samples available in a multi-omics dataset is usually less, because of the time and resources involved in collecting samples from patients. Curse of dimensionality, can cause huge overfitting of a model and make samples harder to group.

In this project, we use a deep embedding network that learns to represent multi-omics data types in a latent space with lower dimension. We implemented the baseline called OmiEmbed [25] by Zhang et al. that uses a variational autoencoder for low dimensional latent space extraction by integrating genome-wide miRNA & gene expression as well as DNA methylation profiles. The key idea here is to make the model learn meaningful representations from different types of omics data, which could then be used for downstream tasks such as cancer type classification. The main goals here are to overcome the curse of dimensionality in multi-omics data, combine information about different aspects of samples, and hopefully extract biologically relevant features.

We also explore the field of Self-Supervised Learning (SSL) in this project. SSL has been successfully used in recent years to learn meaningful representations of data in natural language processing [15, 17] and image domains [5, 7]. It is usually done by taking advantage of temporal, semantic, spatial or structural relationships in the data. But such methods are less effective due to the absence of such relationships in tabular data. Augmentation methods like color transformation, affine and projective transformations etc. are domain specific, and not suited to tabular setting. Due to the above reasons, the domain of SSL has not been explored much in tabular setting compared to other domains.

In this project, we divide the features into subsets and train the model to reconstruct all features from a single subset of features, thus making the model learn from multiple views of the dataset. This is similar to feature bagging in ensemble learning or cropping in image domain. Training the model to reconstruct the data from a subset of its features results in the encoder learning better representations. During testing, representations from multiple subsets are aggregated. Additionally, we explore adding an identity for the subset while training for downstream tasks. This helps the model identify the subset of features that is fed into it. Thus the model learns the corresponding view during training and uses this information during testing.

## 2 Literature review

The success of representation learning ability of DNNs (deep neural networks) in computer vision and natural language processing motivated researchers to apply DNNs in various downstream tasks using high-dimensional omics data. These tasks include classification of diseases, predicting age and others. For classification and discrimination of breast tumour samples from benign samples, Danaee et al. [6] proposed a new model used for cancer detection where they utilized gene expression data by using autoencoder with stacked denoising. In another work, Lyu and Haque [14] used convolutional neural network for tumour type classification. For this purpose, initially images were obtained from RNA-Seq high dimensional data which achieved an accuracy of 95.59%.

Graph convolution neural network (GCN) along with relation network (RN) were proposed by Rhee et al. [19] to classify subtypes of breast tumours. In this method, gene expression data along with protein-protein interaction networks were used. Zhang et al. proposed OmiVAE [26] (variational autoencoder based generative model strategy) for tumour type classification and it achieved an accuracy of 97.49% from 33 tumour categories and normal class using gene expression as well as DNA methylation data from GDC dataset.

Apart from classification tasks, scientist also tried to use deep neural networks for various other downstream tasks using high dimensional multi-omics data. For the prediction of survival for 20 different tumour types a multi-modal approach was proposed by Cheerla and Gevaert [4] using GDC dataset. Along with multi-omics data, they used additional clinical information & histopathology

whole slide images (WSIs) and achieved an overall C-index of 0.78. To further explore in this direction, an innovative deep neural framework called SAUCIE was proposed by Amodio et al. [1]. In this work, they used single cell gene expression data to perform imputation, clustering, batch correction, denoising, and visualisation. The state-of-the-art in this domain is OmiEmbed [25] (to the best of our understanding), which we implemented as baseline.

Corrupting the data by adding noise is a common SSL approach in tabular data [22]. In this approach, an autoencoder is trained to represent corrupted examples of data on a latent space, after which it is trained to reconstruct the uncorrupted data from this latent representation. Thus, the model is trained to learn representations that are robust to the noise in the input. This type of autoencoder is called a denoising autoencoder. A drawback of this approach is that it makes the assumption that all features are equally important and treats them equally. But, adding noise to uninformative features may not help in this method.

Another work called VIME (Value Imputation and Mask Estimation) [24] uses SSL on tabular data by training a denoising autoencoder on an additional pretext task of predicting mask vectors from corrupted data on top of reconstructing the data. To predict which features have been masked in mask estimation task, a classifier connected to the latent representation layer is trained on corrupted data. But, this approach still relies on noisy data in the input as in the case of a standard denoising autoencoder. Moreover, training a classifier on an imbalanced binary mask for large-dimensional data might not make the model learn meaningful representations.

A recent approach named SubTab [20] (Subsetting features of Tabular data) reduces the reliance on noise by changing the task from learning from tabular data to a multi-view representation learning problem through subsetting the features, i.e., dividing the input features into subsets. They show that reconstructing the input data from the subset of its features instead of corrupted data in can make the autoencoder learn better latent representations. At test time, the joint representation is calculated as the aggregate of latent representations of each subset, which they call collaborative inference.

An approach named TabNet [2] uses sequential attention to decide the features to use at each decision step. It performs unsupervised representation learning through masked SSL. This provides a better encoder model for the supervised learning task. In a recent work on continual learning on multiple tasks [23], authors show that task identity can be used to avoid catastrophic forgetting when the model is trained sequentially on multiple tasks. We take inspiration from this work and add subset identity in our framework.

### 3 Methods and Materials

#### 3.1 Dataset

For our experiments, we used The Cancer Genome Atlas (TCGA) pan-cancer multi-omics dataset [10]. An overview of the dataset is given in Table 3.1. It is one of the most popular multi omics dataset. It consists of omics data as well as phenotypic information of patients. We used 3 different omics data from the TCGA dataset, namely, DNA methylation, miRNA stem loop expression as well as gene expression. They are 485,577, 1881 and 60,483 dimensional respectively. The dataset contains samples of 33 different tumour types, as well as of normal tissues.

Dataset		TCGA	
Domain		Pan-cancer	
Tumour types		33 + 1(normal) = 34	
Omics data type	Gene expression	DNA methylation	miRNA expression
No of features	60,483	485,577	1881
No of samples	11,538	9736	11,020

Table 1: An overview of TCGA pan-cancer dataset

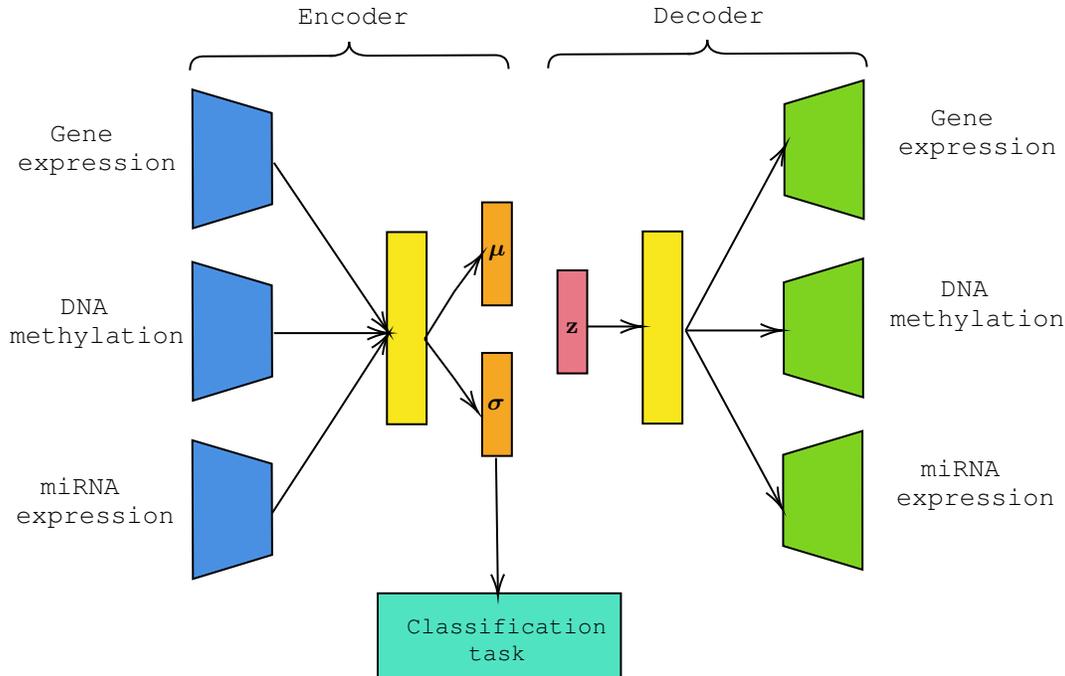


Figure 1: SubOmiEmbed Architecture: It consists of an encoder that learns meaningful representations of the inputs and a decoder that reconstructs the data from the latent representation.  $\mathbf{z}$ ,  $\sigma$  and  $\mu$  represent the latent vector, standard deviation vector and mean vector respectively.  $\mathbf{z}$  is sampled from  $\sigma$  and  $\mu$  using the reparameterisation trick shown in Equation 2

### 3.2 Data preprocessing

We downloaded harmonised data of 3 different omics from UCSC Xena data portal [8]. RNA-Seq gene expression dataset contains 60,483 identifiers, one for each gene. Gene expression level is quantified as the  $\log_2$  transformation of fragments per kilobase of transcript per million mapped reads (FPKM) value. In the case of miRNA expression dataset, the expression levels of miRNA stem-loop features were quantified as the  $\log_2$  transformation of reads per million mapped reads (RPM) value.

DNA methylation dataset consists of beta values for each CpG site and represents the ratio of methylated to total array intensity for the corresponding CpG site [12]. Lower values of beta imply lower levels of methylation and vice-versa. The missing beta values in DNA methylation dataset were imputed using the mean of beta values of the respective CpG sites. All three datasets were rescaled so that the vector of each sample has unit norm, independent of the sample's distribution.

### 3.3 Architecture

SubOmiEmbed comprises of a deep embedding network and a downstream classification network, as shown in Figure 1. The deep embedding network learns low-dimensional latent space representation from high-dimensional multi-omics data. This latent space is then used for the downstream task of cancer type classification. Fully connected layers are used in the encoder, decoder and downstream network.

The deep embedding module of SubOmiEmbed has a variational autoencoder (VAE) [26] as its backbone. VAE is a generative network that can effectively represent high-dimensional data in a lower dimension. We make the assumption that a latent vector  $\mathbf{z}^i \in \mathbb{R}^q$  can be used to represent each sample  $\mathbf{x}^i \in \mathbb{R}^d$  in the dataset  $\mathcal{D}$ , where  $q \gg d$ . To generate samples, each latent vector is sampled from a prior distribution  $p_\theta(\mathbf{z})$ , after which the conditional distribution  $p_\theta(\mathbf{x} | \mathbf{z})$  generates multi-omics data for each sample. Here,  $\theta$  is the set of decoder parameters that can be learned. Due to the difficulty in estimating the true posterior  $p_\theta(\mathbf{z} | \mathbf{x})$ , the variational distribution  $q_\phi(\mathbf{z} | \mathbf{x})$  is used to approximate  $p_\theta(\mathbf{z} | \mathbf{x})$ . Here,  $\phi$  is the set of encoder parameters that can be learned. Therefore, by

maximizing the variational lower bound given below, the VAE model can be optimised:

$$\mathbb{E}_{\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{x})} \log p_\theta(\mathbf{x}|\mathbf{z}) - D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x})\|p_\theta(\mathbf{z})) \quad (1)$$

Here,  $D_{\text{KL}}$  is the Kullback–Leibler (KL) divergence between 2 probability distributions [9].

Each of the three types of input data was first encoded into respective vectors. These vectors were then concatenated together before feeding them into the hidden layers in the encoder. To produce the standard deviation vector  $\sigma$  and the mean vector  $\mu$ , the output vector from hidden layers was connected to 2 bottleneck layers. The variational distribution  $q_\phi(\mathbf{z}|\mathbf{x})$  is defined by these vectors, which is also the Gaussian distribution  $\mathcal{N}(\mu, \sigma)$  of the latent variable  $\mathbf{z}$  given the input  $\mathbf{x}$ . The reparameterization trick as shown in equation 2, is used because sampling  $\mathbf{z}$  from the distribution is not differentiable and is unsuitable for backpropagation.

$$\mathbf{z} = \mu + \sigma \epsilon \quad (2)$$

Here,  $\epsilon$  is a random variable which is sampled from the standard Gaussian distribution  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ . To reconstruct the input data  $\mathbf{x}'$ , latent variable  $\mathbf{z}$  is then passed to a symmetrical decoder.

### 3.4 Training Strategy

The loss function in the model comprises of two major parts, the embedding loss and classification loss. Let's denote the input data  $\mathbf{x}$  of  $i$ th omics type as  $\mathbf{x}_i$  and the reconstructed data as  $\mathbf{x}'_i$ . Now the embedding loss can be formulated as shown below.

$$\mathcal{L}_{\text{embed}} = \frac{1}{N} \sum_{i=1}^N \text{MSE}(\mathbf{x}_i, \mathbf{x}'_i) + D_{\text{KL}}(\mathcal{N}(\mu, \sigma)\|\mathcal{N}(\mathbf{0}, \mathbf{I})) \quad (3)$$

Here, the 1st part of the loss is the mean squared error loss between input and reconstructed omics data, while the 2nd part represents the KL divergence between the learnt probability distribution and a standard normal distribution.

Loss function for the downstream classification task is formulated as follows

$$\mathcal{L}_{\text{classification}} = \text{CE}(y, y') \quad (4)$$

Here,  $y$  is the label,  $y'$  the prediction and CE the cross entropy loss.

Additionally, the joint loss is defined as

$$\mathcal{L}_{\text{joint}} = \mathcal{L}_{\text{embed}} + \mathcal{L}_{\text{classification}} \quad (5)$$

Training is divided into three phases. The 1st phase is unsupervised where the loss minimized is the embedding loss and the encoder and decoder parameters are learnt. Labels are not needed in this phase and phase 1 could be used individually for visualisation or reducing dimensions. In phase 2, the layers of embedding module are kept frozen while the layers of downstream network are trained. Here, only the classification loss is minimized. In the third phase, the overall loss function as shown in Equation 5 is minimized. In this last phase, the whole network consisting of both the embedding layers and the classification network layers are trained. This way, the network learns representations that are useful for the downstream classification task.

### 3.5 Feature subsetting

In this setting, we divided the features into subsets and trained encoder and decoder on each of the subsets. This is shown in Figure 2. Suppose we divide the  $i$ th omics-type dataset  $\mathbf{x}_i$  into  $M$  subsets. For each subset  $\mathbf{x}_i^{(j)}$ , we reconstruct the entire dataset and the reconstruction is denoted by  $\mathbf{x}_i^{(j) \prime}$ . The embedding loss for  $j$ th subset can be computed as shown below

$$\mathcal{L}_{\text{embed}}^{(j)} = \frac{1}{N} \sum_{i=1}^N \text{MSE}(\mathbf{x}_i^{(j)}, \mathbf{x}_i^{(j) \prime}) + D_{\text{KL}}(\mathcal{N}(\mu^{(j)}, \sigma^{(j)})\|\mathcal{N}(\mathbf{0}, \mathbf{I})) \quad (6)$$

The joint loss is given as

$$\mathcal{L}_{\text{joint}} = \frac{1}{M} \sum_{j=1}^M \mathcal{L}_{\text{embed}}^{(j)} + \mathcal{L}_{\text{classification}} \quad (7)$$

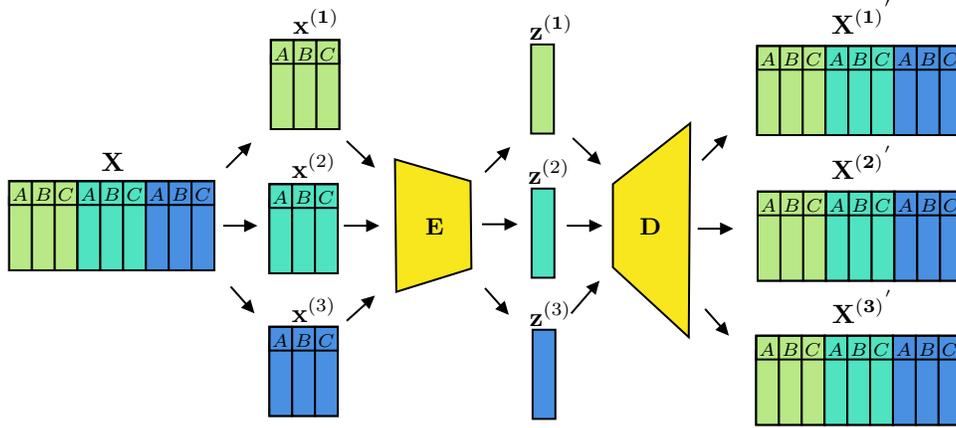


Figure 2: Illustration of feature subsetting during training. Here, the dataset  $X$  is divided into 3 subsets,  $x^{(1)}$ ,  $x^{(2)}$  and  $x^{(3)}$ . All omics-type A, B and C are proportionately divided into 3 subsets. For instance,  $x^{(1)}$  contains  $\frac{1}{3}$ rd of each type A, B and C. Each subset is fed to the same encoder which produces latent representations  $z^{(1)}$ ,  $z^{(2)}$  and  $z^{(3)}$  for the corresponding subsets. These representations are then fed into the decoder. The decoder reconstructs the whole training set  $X$ . The reconstruction from  $z^{(1)}$  is named  $X^{(1)'}$  and so on.

Latent representations,  $z^1, z^2, \dots, z^M$  are aggregated by mean (default), max, min or sum as shown in 3. The aggregated representation is then used for the classification task. The classification loss is defined in Equation 4.

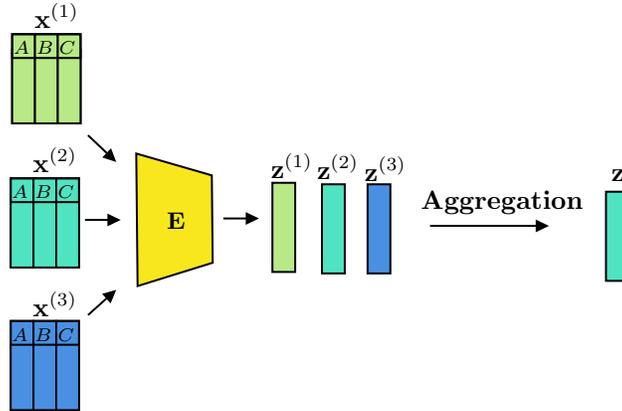


Figure 3: Illustration of aggregation of representations at test time

### 3.6 Adding subset identity

To help the model identify which subset is fed into it, we add a one-hot encoded subset identity. Once the latent representation  $z^j$  is produced by the encoder, it is concatenated with the one-hot encoded subset identity to produce  $z^{j'}$  as shown in Figure 4. This concatenated vector is then passed onto the decoder and downstream layers.

In the setting where subset identity is used, the aggregation of representation at test time is not required. Therefore, the downstream layer output from any one of the subsets is used as the model output. We noticed that aggregating the downstream outputs produced by representations from multiple subsets performed similar to the setting where representation from any one of the subsets is used to produce downstream output. This shows the power of adding subset identity.

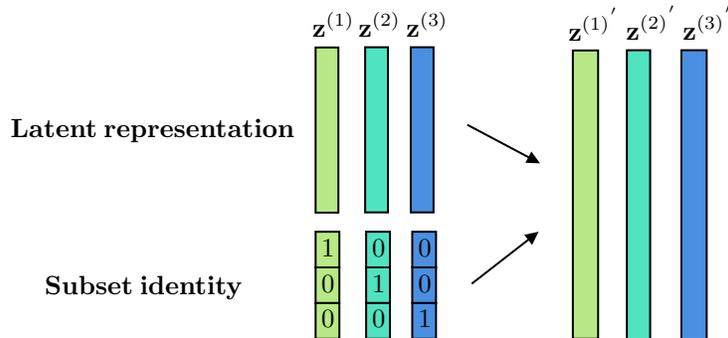


Figure 4: Illustration of adding subset identity. Latent representations  $\mathbf{z}^{(1)}$ ,  $\mathbf{z}^{(2)}$  and  $\mathbf{z}^{(3)}$  are concatenated with their corresponding subset identity vectors to produce  $\mathbf{z}^{(1)'}$ ,  $\mathbf{z}^{(2)'}$  and  $\mathbf{z}^{(3)'}$ . These vectors are then input to decoder and downstream layers.

## 4 Experiments and results

### 4.1 Implementation details

Our codebase (<https://github.com/hashimsayed0/OmiEmbed>) is built upon the codebase of the baseline that we implemented [25] (<https://github.com/zhangxiaoyu11/OmiEmbed>). The framework is built on PyTorch [16]. It supports multiple omics modes, backbone networks, reconstruction losses and latent space dimensions. The model was tuned on each of these hyperparameters. The following sections describe the experiments in detail.

### 4.2 Baseline experiments

#### 4.2.1 Handling missing values

The DNA methylation dataset had missing values. We tried various approaches to handle them and evaluated the model’s performance on cancer type classification. Zero & mean imputations as well as dropping features containing missing values were the three approaches attempted. Mean imputation provided best performance as shown in Table 2. Therefore, mean imputed DNA methylation dataset was used for further experiments.

Imputation method	Accuracy	Precision	Recall	F1	AUC
Dropping features	0.9512	0.9387	0.9065	0.9190	0.9976
Zero imputation	0.9517	0.9421	0.9256	0.9305	<b>0.9987</b>
<b>Mean imputation</b>	<b>0.9589</b>	<b>0.9470</b>	<b>0.9552</b>	<b>0.9503</b>	0.9981

Table 2: Effect of various imputation methods on cancer type classification performance with only DNA methylation dataset

#### 4.2.2 Omics mode

For this set of experiments, we tried multiple combinations of omics data types. As shown in Table 3, using all three data types showed best performance. This can be because each type of data reveals some kind of information about the tumour, which eventually helps the model classify tumours. Also, note that miRNA expression dataset with under 2000 features alone gives almost 95% accuracy in the task, showing the predictive capacity of omics data. Figure 5 shows how training evolves through the 3 phases discussed earlier.

#### 4.2.3 Backbone of the network

Two main backbones that we trained for the encoder decoder architecture are fully connected and convolutional networks. In the fully connected network, DNA methylation data was divided into

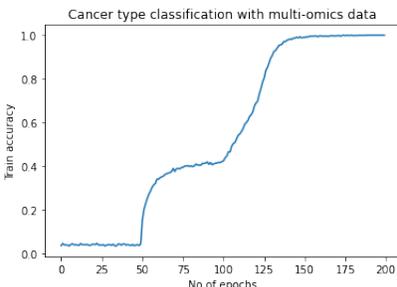


Figure 5: Downstream task train accuracy with multi-omics data (A+B+C) in 3 phases. Phase 1 and 2 are 50 epochs while phase 3 is 100 epochs.

Omics mode	Accuracy	Precision	Recall	F1	AUC
RNA-seq gene expression (A)	0.9584	0.9532	0.9283	0.9257	0.9986
DNA methylation data (B)	0.9589	0.9470	0.9552	0.9503	0.9981
miRNA stem loop expression (C)	0.9494	0.9355	0.9254	0.9286	0.9984
Multi omics (A+B)	0.9647	0.9506	0.9497	0.9497	0.9973
<b>Multi omics (A+B+C)</b>	<b>0.9709</b>	<b>0.9640</b>	<b>0.9597</b>	<b>0.9613</b>	<b>0.9988</b>

Table 3: Evaluation of OmiEmbed on cancer type classification with multiple omics type combinations

23 blocks, corresponding to the chromosome number of respective CpG sites. This reduces the number of parameters, avoids overfitting and uses less GPU memory, while also making the model learn inter-chromosome relationships in first hidden layer and intra-chromosome relationships in subsequent layers. As shown in Table 4, the extra chromosome separation step improves performance.

Backbone	Accuracy	Precision	Recall	F1	AUC
Conv-1D	0.9536	0.9345	0.9239	0.9278	0.9971
Fully connected (FC)	0.9598	0.9302	0.9388	0.9336	0.9983
<b>FC + chromosome separation</b>	<b>0.9626</b>	<b>0.9561</b>	<b>0.9488</b>	<b>0.9443</b>	<b>0.9985</b>

Table 4: Evaluation of OmiEmbed on cancer type classification with various backbones for embedding module

#### 4.2.4 Latent space dimension

This set of experiments was aimed at choosing the ideal dimension of latent space. We experimented with 64, 128 and 256. As shown in Table 5, 128 dimensional latent space gives best performance. This could be due to the fact that its large enough to embed all the essential information. The 128D latent vector was visualized using t-SNE [3] in Figure 6, which shows that the network learns to discriminate between various tumour types and the normal class. Even classes with few number of samples are separated by the model.

#### 4.2.5 Reconstruction loss

We also tried different reconstruction losses. The results are shown in Table 6. Note that binary cross entropy loss outperforms means squared error loss in all metrics except accuracy.

### 4.3 Feature subsetting experiments

#### 4.3.1 Reducing the network size

Once we applied the features subsetting approach, we noticed that the network was overfitting to the subsets. This is because we were using the same network as used in the previous experiments with just the change in input dimension. To solve this issue, we reduced the dimensions of encoder, decoder and downstream layers by encoder reduction factor, decoder reduction factor and downstream reduction factors respectively compared to the baseline network. This set of experiments were run to

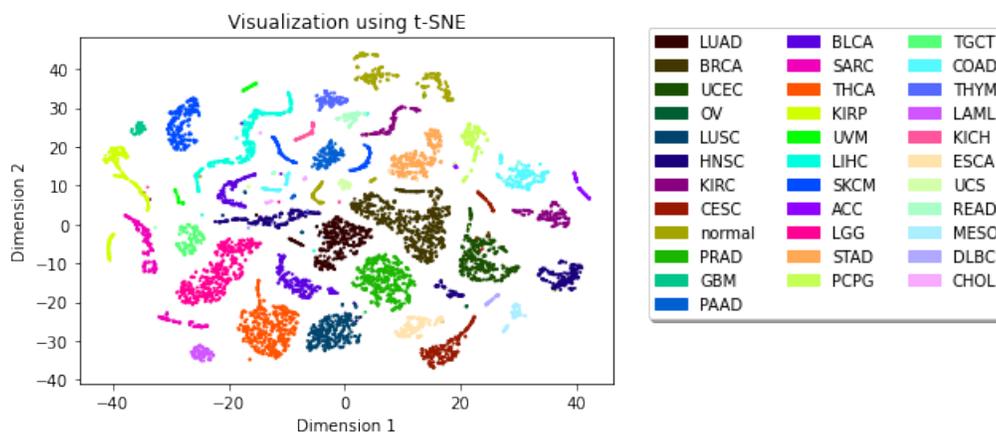


Figure 6: Visualizing the 128D latent vector learned by the end of phase 3 training using t-SNE. It can be seen that the classes are separated. Labels on the right side refer to the TCGA short codes for cancer types

Latent space dimension	Accuracy	Precision	Recall	F1	AUC
64	0.9693	0.9588	0.9578	0.9568	0.9981
256	0.9631	0.9398	<b>0.9618</b>	0.9473	0.9983
<b>128</b>	<b>0.9709</b>	<b>0.9640</b>	0.9597	<b>0.9613</b>	<b>0.9988</b>

Table 5: Evaluation of OmiEmbed on cancer type classification with various latent space dimensions

choose the optimal network size. Results are shown in Table 7. It is striking to see that with feature subsetting, a network with reduced size achieves comparable results to the baseline. Due to the success of networks with lower of sizes, the encoder, decoder and downstream layers of the network were reduced by 8, 4 and 2 in the following experiments.

#### 4.3.2 Aggregation method

In this set of experiments, we aggregated the representations at test time with different methods. Adding subset identity didn't require aggregation and it was also compared against the other methods. A random subset with identity concatenated with it was used as input to decoder and downstream layers in this setting. The results are shown in Table 8. From the results, it can be seen that summing up the subset is the best aggregation method to use. Perhaps this way, not much information is lost during aggregation. Another insight is that choosing a random subset with one-hot encoded identity vector performs much better than without it. Due to the efficiency of using just one random subset with identity, we use it in our following experiments.

#### 4.3.3 Number of subsets

Another set of experiments we conducted was to choose the optimal number of subsets. Results are shown in Table 9. It is surprising to see that the network with reduced dimensions works well with even 10 subsets. This shows the power of feature subsetting.

### 4.4 Comparison to other methods

The performance of OmiEmbed on cancer type classification compared to other methods are shown in Table 10. As can be seen, OmiEmbed outperforms other methods on this task.

## 5 Conclusions

In this project, we have explored the area of representation learning for multi-omics data. We began with the study of the biological relevance of the dataset, and looked for deep learning models to extract meaningful representations from the different types of omics profiling. We then identified

Reconstruction loss	Accuracy	Precision	Recall	F1	AUC
Mean absolute error (L1)	0.9620	0.9542	0.9430	0.9407	0.9977
Mean squared error	<b>0.9648</b>	0.9506	0.9374	0.9399	0.9979
<b>Binary cross entropy</b>	0.9626	<b>0.9561</b>	<b>0.9488</b>	<b>0.9443</b>	<b>0.9985</b>

Table 6: Evaluation of OmiEmbed on cancer type classification with various reconstruction losses

Reduction	Accuracy	Precision	Recall	F1	ROC AUC	Mean metric
1_1_1	0.89	0.881	0.908	0.859	0.997	0.893
1_1_2	0.922	0.925	0.894	0.879	0.998	0.914
1_4_1	0.931	0.902	0.941	0.905	0.999	0.925
1_4_2	0.91	0.892	0.92	0.888	0.998	0.907
4_1_1	<b>0.962</b>	0.93	0.909	0.909	0.999	0.934
4_1_2	0.949	0.933	0.911	0.916	0.994	0.931
4_4_1	0.954	0.937	0.939	0.929	0.999	0.943
4_4_2	0.943	0.893	0.91	0.898	0.997	0.915
8_1_1	0.945	0.927	0.935	0.928	0.999	0.936
8_1_2	0.954	0.945	<b>0.943</b>	0.94	<b>0.999</b>	<b>0.947</b>
8_4_1	0.944	<b>0.947</b>	0.907	0.907	0.999	0.933
8_4_2	0.945	0.942	0.935	<b>0.936</b>	0.999	0.941

Table 7: Evaluation of SubOmiEmbed on cancer type classification with multiple reduction factors. The dimension reduction factors are mentioned in the first column. For instance, 8\_4\_2 implies that the dimensions of encoder, decoder and downstream layers were reduced by 8, 4 and 2 respectively compared to the baseline network. Please note that the mean metric mentioned in this table as well as subsequent tables refer to the mean of accuracy, precision and recall.

the field of VAEs and its applications. After formulating the problem and loss functions, we ran extensive experiments to evaluate the learning capacity of VAEs. In our experiments, we pre-trained the embedding module on reconstruction of training data in the first phase, then kept the embedding layers frozen and trained it for the downstream task of tumour type classification in the second phase. In the last phase, we trained the model on both tasks concurrently to produce meaningful representation of the input data that is relevant for the downstream task. Various insights about the model and its hyperparameters were derived from the experiments.

After experimenting with the baseline model, we extended our framework to include the SSL technique of feature subsetting as well as adding subset identity. We first formulated the problem and then ran our experiments. There were challenges along the way, one of which being overfitting of the model to subsets. We reduced the network size to reduce overfitting in addition to giving the model information about the ordering of the subsets. With our experiments, we show that a model of a much reduced size can produce comparable results to that of the baseline model, using a single subset of the data. This shows that model learns a lot more with self supervision. This work can be extended to add noise to the data, and amending the embedding loss to include contrastive loss. Attention based techniques could also be used to integrate sequence based genomic data into the framework.

Aggregation method	Accuracy	Precision	Recall	F1	ROC AUC	Mean metric
Maximum	0.883	0.88	0.905	0.873	0.998	0.889
Mean	0.936	0.939	0.933	0.93	0.999	0.936
Minimum	0.935	0.907	0.932	0.913	0.997	0.925
Random subset	0.897	0.869	0.814	0.825	0.996	0.86
Random with identity	0.945	0.942	0.935	0.936	0.999	0.941
<b>Sum</b>	<b>0.948</b>	<b>0.951</b>	<b>0.942</b>	<b>0.942</b>	<b>0.999</b>	<b>0.947</b>

Table 8: Evaluation of SubOmiEmbed on cancer type classification with multiple aggregation methods. Using a random subset with subset identity was also tested in this set of experiments.

No of subsets	Accuracy	Precision	Recall	F1	ROC AUC	Mean metric
2	0.959	<b>0.96</b>	0.94	0.944	0.999	0.953
3	0.957	0.946	0.939	0.927	0.999	0.947
4	<b>0.963</b>	0.95	<b>0.958</b>	<b>0.952</b>	<b>0.999</b>	<b>0.957</b>
5	0.959	0.913	0.909	0.903	0.999	0.927
6	0.951	0.946	0.953	0.942	0.999	0.95
7	0.943	0.915	0.921	0.891	0.999	0.926
8	0.93	0.934	0.933	0.922	0.999	0.932
9	0.948	0.955	0.926	0.936	0.998	0.943
10	0.95	0.941	0.944	0.94	0.999	0.945
11	0.947	0.931	0.932	0.924	0.999	0.937
12	0.939	0.952	0.92	0.927	0.999	0.937

Table 9: Evaluation of SubOmiEmbed on cancer type classification with various number of subsets. Note that all experiments were run on a network with reduced size and with subset identity.

Method	Accuracy	Precision	Recall	F1	ROC AUC
kNN	0.898	0.89	0.87	0.87	-
SVM	0.912	0.88	0.87	0.87	-
NN	0.935	0.937	0.932	-	0.986
GCN [18]	0.946	0.947	0.921	-	-
CNN [14]	0.955	<b>0.955</b>	<b>0.955</b>	0.954	-
OmiEmbed	<b>0.977</b>	-	-	<b>0.968</b>	<b>0.999</b>

Table 10: Evaluation of various methods on cancer type classification. NN refers to a simple neural network with few hidden layers. GCN [18] refers to a work that involves computing gene co-expression graph and training graph convolutional network on it. CNN [14] refers to a paper that embeds the data into 2-D images and uses a convolutional neural network for classification.

## References

- [1] M. Amodio, D. van Dijk, K. Srinivasan, W. S. Chen, H. Mohsen, K. R. Moon, A. Campbell, Y. Zhao, X. Wang, M. Venkataswamy, A. Desai, V. Ravi, P. Kumar, R. Montgomery, G. Wolf, and S. Krishnaswamy. Exploring single-cell data with deep multitasking neural networks. *bioRxiv*, 2019.
- [2] S. O. Arik and T. Pfister. Tabnet: Attentive interpretable tabular learning, 2020.
- [3] D. M. Chan, R. Rao, F. Huang, and J. F. Canny. t-sne-cuda: Gpu-accelerated t-sne and its applications to modern data, 2018.
- [4] A. Cheerla and O. Gevaert. Deep learning with multimodal representation for pancancer prognosis prediction. *Bioinformatics*, 35(14):i446–i454, 07 2019.
- [5] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. In H. D. III and A. Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR, 13–18 Jul 2020.
- [6] P. Danaee, R. Ghaeini, and D. A. Hendrix. A deep learning approach for cancer detection and relevant gene identification. *Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing*, 22:219–229, 2017.
- [7] W. Falcon and K. Cho. A framework for contrastive self-supervised learning and designing a new approach, 2020.
- [8] M. Goldman, B. Craft, M. Hastie, K. Repečka, F. McDade, A. Kamath, A. Banerjee, Y. Luo, D. Rogers, A. N. Brooks, J. Zhu, and D. Haussler. The ucsc xena platform for public and private cancer genomics data visualization and interpretation. *bioRxiv*, 2019.
- [9] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [10] R. L. Grossman, A. P. Heath, V. Ferretti, H. E. Varmus, D. R. Lowy, W. A. Kibbe, and L. M. Staudt. Toward a shared vision for cancer genomic data. *New England Journal of Medicine*, 375(12):1109–1112, 2016. PMID: 27653561.
- [11] Y. Hasin-Brumshtein, M. Seldin, and A. Lusis. Multi-omics approaches to disease. *Genome Biology*, 18, 05 2017.
- [12] N. C. Institute. Bioinformatics pipeline: Methylation liftover pipeline - gdc docs, 2021.
- [13] D. P. Kingma and M. Welling. Auto-encoding variational bayes, 2014.
- [14] B. Lyu and A. Haque. Deep learning based tumor type classification using gene expression data. In *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics, BCB '18*, page 89–96, New York, NY, USA, 2018. Association for Computing Machinery.
- [15] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space, 2013.
- [16] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019.
- [17] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.
- [18] R. Ramirez, Y.-C. Chiu, A. Hererra, M. Mostavi, J. Ramirez, Y. Chen, Y. Huang, and Y.-F. Jin. Classification of cancer types using graph convolutional neural networks. *Frontiers in Physics*, 8:203, 2020.
- [19] S. Rhee, S. Seo, and S. Kim. Hybrid approach of relation network and localized graph convolutional filtering for breast cancer subtype classification, 2018.
- [20] T. Ucar, E. Hajiramezanali, and L. Edwards. Subtab: Subsetting features of tabular data for self-supervised representation learning, 2021.
- [21] M. Vailati-Riboni, V. Palombo, and J. J. Loor. What are omics sciences? *Periparturient Diseases of Dairy Cows*, pages 1–7, 2017.

- [22] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, page 1096–1103, New York, NY, USA, 2008. Association for Computing Machinery.
- [23] J. von Oswald, C. Henning, J. Sacramento, and B. F. Grewe. Continual learning with hypernetworks, 2020.
- [24] J. Yoon, Y. Zhang, J. Jordon, and M. van der Schaar. Vime: Extending the success of self- and semi-supervised learning to tabular domain. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 11033–11043. Curran Associates, Inc., 2020.
- [25] X. Zhang, Y. Xing, K. Sun, and Y. Guo. Omiembed: A unified multi-task deep learning framework for multi-omics data. *Cancers*, 13(12):3047, Jun 2021.
- [26] X. Zhang, J. Zhang, K. Sun, X. Yang, C. Dai, and Y. Guo. Integrated multi-omics analysis using variational autoencoders: Application to pan-cancer classification. *2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, Nov 2019.