

# Dynamic Multimodal Information Bottleneck for Multimodality Classification

Yingying Fang<sup>1\*</sup>, Shuang Wu<sup>3\*</sup>, Sheng Zhang<sup>1</sup>, Chaoyan Huang<sup>4</sup>, Tiejong Zeng<sup>4</sup>, Xiaodan Xing<sup>2</sup>,  
Simon Walsh<sup>1</sup>, Guang Yang<sup>1,2†</sup>

<sup>1</sup>National Heart and Lung Institute, Imperial College London, London, SW7 2AZ, UK

<sup>2</sup>Bioengineering Department and Imperial-X, Imperial College London, London, W12 7SL, UK

<sup>3</sup>Black Sesame Technologies, Fusionopolis, Singapore

<sup>4</sup>The Chinese University of Hong Kong, Shatin, Hong Kong

y.fang; sheng.zhang; x.xing; s.walsh; g.yang@imperial.ac.uk

wushuang@outlook.sg cyhuang; zeng@math.cuhk.edu.hk

## Abstract

*Effectively leveraging multimodal data such as various images, laboratory tests and clinical information is becoming increasingly attractive in a variety of AI-based medical diagnosis and prognosis tasks. Most existing multi-modal techniques only focus on enhancing their performance by leveraging the differences or shared features from various modalities and fusing feature across different modalities. These approaches are generally not optimal for clinical settings, which pose the additional challenges of limited training data, as well as being rife with redundant data or noisy modality channels, leading to subpar performance. To address this gap, we study the robustness of existing methods to data redundancy and noise and propose a generalized dynamic multimodal information bottleneck framework for attaining a robust fused feature representation. Specifically, our information bottleneck module serves to filter out the task-irrelevant information and noises in the fused feature, and we further introduce a sufficiency loss to prevent dropping of task-relevant information, thus explicitly preserving the sufficiency of prediction information in the distilled feature. We validate our model on an in-house and a public COVID19 dataset for mortality prediction as well as two public biomedical datasets for diagnostic tasks. Extensive experiments show that our method surpasses the state-of-the-art and is significantly more robust, being the only method to remain performance when large-scale noisy channels exist. Our code is publicly available at <https://github.com/ayanglab/DMIB>.*

\*Equal Contribution

†Corresponding Author

## 1. Introduction

Medical practitioners utilize various sources of data such as electronic health records, laboratory tests, genetic information and medical imaging modalities such as Computerized Tomography (CT), Magnetic Resonance Imaging (MRI), Positron Emission Tomography (PET) etc. for medical diagnosis and prognosis. Historically, partly due to domain gaps and specialization, clinicians primarily operated on each data modality in silo, drawing upon the distinctive features from an individual source to perform diagnosis and prognosis. Nonetheless, it is undeniable that integrating multi-modal data effectively would be beneficial for diagnosis and prognosis tasks, by not only incorporating extra guidance but also possibly providing novel insights, enabled by a more holistic understanding of the entirety.

The advent of deep learning has sparked several lines of works centered on healthcare applications [46], among which AI-based medical image analysis [34] has enjoyed prominent success, with deep learning models achieving performance on par with or even surpassing radiologists on some tasks. Likewise, deep learning has also demonstrated immense potential for analysing electronic health records [32] and genetic information [48]. Yet, despite impressive performance on individual modalities, developing techniques to effectively leverage multiple modalities remains challenging [3, 7, 21, 26].

The general multimodal learning approach is to separately train a model for each modality to extract a modality-specific feature vector, and subsequently fusing these individual features to obtain a multimodal feature representation. The fused feature is then propagated to downstream task modules to perform supervised learning. Common approaches for feature fusion include concatenation, attention-weighted, common subspace projection, graph-

based, and transformer-based methods [1,9].

While experimenting with these different fusion schemes in the context of clinical applications, we observe that they generally suffer from significant drops in model performance. Clinical applications typically have limited training data, and the difficulty is compounded by the fact that clinical data is often subject to various forms of noise such as missing data, inaccurate clinical records, and subjective biases in patients' self-assessment. Existing methods generally suffer from a susceptibility to overfitting, and learning to extract task-relevant information is often suboptimal when training data is scarce. An even more glaring weakness is the sensitivity and low robustness towards noisy modality channels, where most existing models suffer from significant performance drops.

To address these issues, we propose a Dynamic Multimodal Information Bottleneck (DMIB) framework, drawing inspirations from mutual information theory [29,40], information bottleneck [12,38,41]. Specifically, our DMIB consists of the following key components: i) an information bottleneck module along with dropout regularization and masking of modalities to remediate feature redundancy and model overfitting; ii) an explicit supervision to maximize the task relevant information in the final fused feature. We conduct comprehensive experiments for multimodality classification tasks across four datasets. Our method not only achieves state-of-the-art performance, but also demonstrates remarkable robustness, retaining similar performance when a modality consists of pure noise.

To summarize, our key contributions are: i) We design an information bottleneck module together with a mutual information inspired sufficiency loss which can be applied to arbitrary multimodal classification tasks as a plug-and-play module. Our fusion strategy dynamically filters out noise, maximizes the inclusion of relevant information and eliminates the need for heuristic or greedy feature selection approaches that were often employed in previous studies, leading to a superior performance. ii) To our knowledge, we are the first to conduct principled experiments to study the performance of fusion methods under different levels of noisy and redundant modalities. iii) Our method demonstrates outstanding robustness by retaining prediction performance even when noise and redundancy are introduced, making it particularly suitable for clinical settings where datasets tend to be small and biomarkers are unclear.

## 2. Related Works

### 2.1. Multimodal Learning for Clinical Applications

In recent years, there has been growing interests and efforts in developing multimodal fusion for clinical applications [1,3,9,26]. Compared to the early fusion cannot address the domain gaps between modalities [24] and

late fusion strategies which overlook the cross-modality information [33], intermediate fusion methods which perform fusion at the feature level [14,30] has attracted much investigation on various techniques for extracting multimodal feature learning in deep neural techniques. These methods could be broadly categorised into concatenation-based [2,24,45], attention-based [10,19,30], projection into aligned subspaces [13,39,44,50], graph-based [8,20,49], and recent Transformer-based [6,25,43]. While concatenation methods are parameter-free and straightforward, they might not identify the intricate relationships between diverse modalities. Attention-based methods learns importance scores for intermediate multi-modality features, which allows for modeling the complex relationships between different modalities. However, introducing the attention mechanisms, along with additional network parameters increases the propensity to overfitting. Recently, a major trend in multimodal learning is to leverage transformer modules to perform cross-attention across the various modalities to obtain a fused feature representation. While transformer based multimodal learning deliver state-of-the-art performance for large scale datasets, they tend to be lackluster when training data are scarce [43]. Besides, graph models are also gaining traction in multimodal tasks with their unique ability to leverage the relations between the data from various modalities.

### 2.2. Mutual Information

Mutual information is a measure of the statistical dependency between random variables [11]. The core principle of deep learning approaches is to automatically learn how to extract optimal features from data, instead of manually crafting features from heuristic guidance. In this light, mutual information lies at the heart of deep learning, since it could be invoked to quantify the dependency between the learned feature representation and desired network output.

**Mutual Information Estimation** [22] proposes the InfoMax principle which seeks to maximize the mutual information between feature and model output. This serves as a general prescription which is generally pertinent to neural networks and there have been many works in the recent years that explore optimal ways for mutual information estimation. One line of approach is to employ an additional neural network for estimation [4,17] while another approach seeks variational bounds for mutual information [31]. However, mutual information remains a notoriously difficult problem due to the curse of dimensionality [29]. In other words, the amount of data samples to accurately bound mutual information scales exponentially with the dimensionality of the feature. To tackle this, researchers have proposed alternatives to mutual information, such as the Wasserstein dependency measure in [29] and the variational distillation loss in [36].

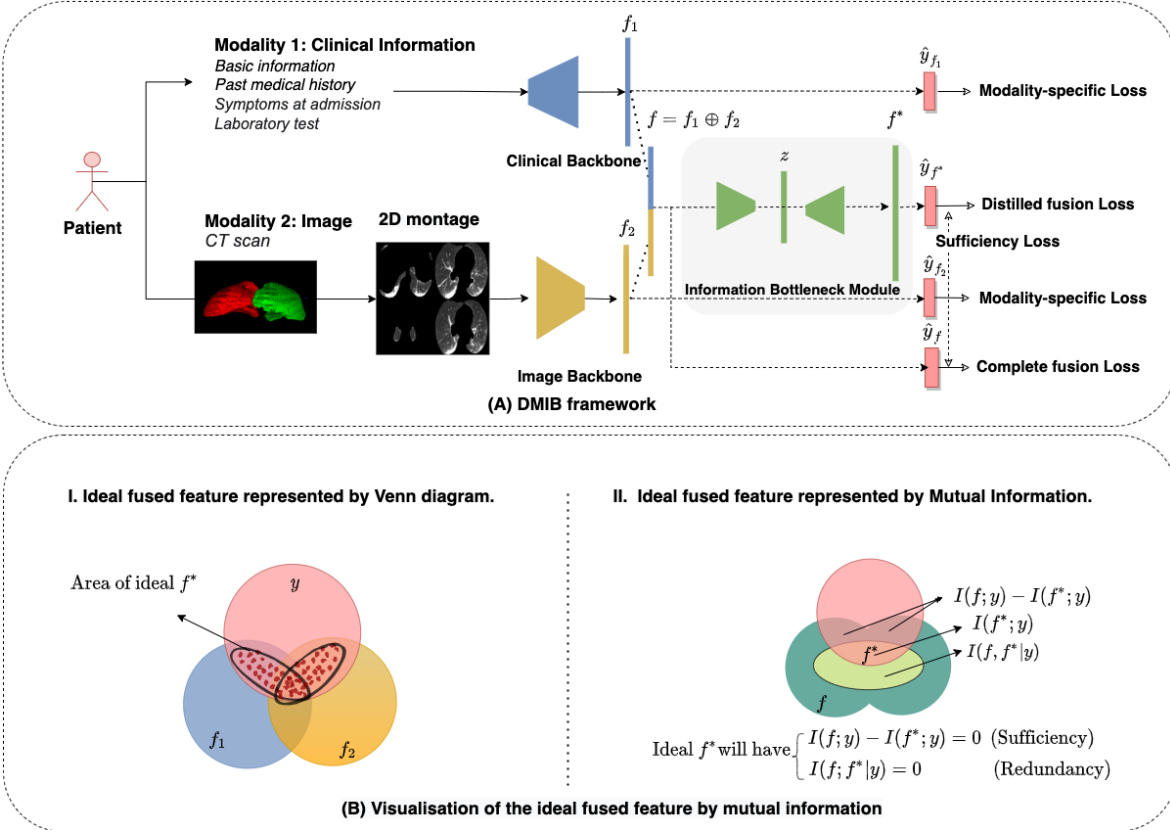


Figure 1. (A) Illustration of our Dynamic Multimodal Information Bottleneck framework for a lung disease prognosis task: a) Separate backbones are used to first extract features from each modality. b) These features are concatenated to form a preliminary fused feature  $f$ . c)  $f$  is distilled via an information bottleneck module, arriving at a final fused feature  $f^*$ . d) We formulate a sufficiency loss to preserve task-relevant information in  $f^*$ . (B) The mutual information between various entities can be visualized with Venn diagrams.

**Information Bottleneck** Closely related to mutual information is the concept of information bottleneck. Introduced in [37], the authors further advance this principle to put forward an explanation of the learning process in deep neural networks [38]. Given raw data inputs, the notion of a bottleneck in the model effectuates an information distillation process, which retains only useful information, while discarding irrelevant information and superfluous noises. This is analogous to how humans learn and master knowledge and abstractions. The information bottleneck method is highly pertinent for improving model generalization and robustness and reducing overfitting [5, 40]. This can be inherently understood from the intuition that a generalisable model should learn feature abstractions that capture the essence of the task, instead of memorizing instance-specific characteristics in the training set.

### 3. Methodology

An overview of our DMIB framework is found in Figure 1 (A). We employ separate backbones to extract fea-

tures from different modalities. Supervision is applied for each backbone to guarantee the intermediate features have sufficient prediction information<sup>1</sup>. Subsequently, we concatenate the extracted features for each modality to form an initial fused feature  $f$ , which preserves complete information from each modality. We introduce an information bottleneck module to perform information distillation, arriving at our final feature  $f^*$ . To ensure sufficiency of task information in  $f^*$ , we introduce a sufficiency loss to ensure that no task-relevant information are being discarded.

To begin, we introduce our overall supervision objective displayed in Figure 1 (B) which motivates the design of our DMIB framework<sup>2</sup>. Subsequently, we cover the details of our information bottleneck module and our sufficiency loss.

<sup>1</sup>To prevent the model predictions from being dominated by the modality with larger dimensions, we also enforce dimension equality by upsampling low dimensional modality features.

<sup>2</sup>An overview of relevant mutual information definitions is given in the supplementary material.

### 3.1. Overall supervision

Given  $f^*$  an encoded fused feature from the direct fused feature  $f$ , we would like the ideal fused feature representation to contain sufficient predictive information in  $f$  while discarding all redundancy and noise. The information contained in  $f^*$ , is given by the mutual information between  $f^*$  and  $f$ :

$$I(f; f^*) = \mathbb{E}_{p(f, f^*)} \left[ \log \frac{p(f, f^*)}{p(f)p(f^*)} \right]. \quad (1)$$

It can be further decomposed into two components by the chain rule of mutual information [12] (given in the supplementary material) as:

$$I(f; f^*) = I(f; f^*|y) + I(f^*; y). \quad (2)$$

$I(f; f^*|y)$  quantifies the amount of task-irrelevant information in  $f^*$  and  $I(f^*; y)$  quantifies the predictive information in  $f^*$  for task  $y$ . The goal is to simultaneously minimize redundancy  $I(f; f^*|y)$  and maximize task relevance  $I(f^*; y)$ , which can be combined into the information bottleneck objective [37]:

$$\min_{f^*} \mathcal{L}_{IB} = I(f; f^*|y) - \gamma I(f^*; y). \quad (3)$$

However, estimating for mutual information in high dimensions is intractable in general [27, 31] due to the curse of dimensionality [29], *i.e.* accurate estimation of the expected information in empirical distributions requires a sample size that scales exponentially with the dimension of the data. In light of this, we do not explicitly optimize for Eqn. (3) via mutual information estimation. Instead, we delegate information distillation to our information bottleneck module in Section 3.2. For maximizing the task-relevant information  $I(f^*; y)$ , we formulate a tractable sufficiency loss in Section 3.3. The joint framework consisting of the information bottleneck module and the sufficiency loss enables  $f^*$  to converge to an ideal feature with maximal relevance and minimal redundancy.

### 3.2. Information Bottleneck Module

Given  $n$  features  $f_1, \dots, f_n$  extracted from  $n$  modalities, we first generate an initial fused feature  $f$  by direct concatenation with a masking

$$f = \bigoplus_{i=1}^n m_i f_i,$$

where  $f \in \mathbb{R}^N$  and  $m_i$  denotes a masking coefficient. Specifically, at each training iteration, we uniformly sample a random number  $u \sim U([0, 1])$ , and the masking coefficients are formally given by  $m_i = \mathbb{1}_{[\frac{i-1}{2n}, \frac{i}{2n}]}(u)$ . In other words, the model can only access the full features across all

modalities for half of the training iterations, while for the other half, a single modality  $f_i$  is masked out. The motivation of masking individual modalities is to allow the model to function in the absence of data-streams from a modality, thus improving robustness [15]. Subsequently, our information bottleneck module comprises of two linear projection layers along with dropouts and ReLU. The operations may be summarized as:

$$f \xrightarrow[\text{Dropouts + ReLU}]{\text{Linear Projection}} z \xrightarrow[\text{Dropouts + ReLU}]{\text{Linear Projection}} f^*. \quad (4)$$

$z$  is of dimension  $p < n$  while the final feature representation  $f^*$  has  $n$  dimensions. The purpose of re-projecting  $z$  back to the same dimensions as  $f$  is in order to perform more effective feature-level supervision of the distilled feature and the initial feature to learn a more predictive feature  $f^*$ , which is introduced in the next section. It is observed that the feature-level supervision is more effective when  $f^*$  and the initial feature  $f$  is aligned with equal dimensions.

### 3.3. Sufficiency Loss

In this section, we introduce our sufficiency loss which serves as a feature-level supervision to maximize the task-relevant information  $I(f^*; y)$  in Eqn. (3). Since  $f^*$  is a feature extracted from  $f$ , the information contained in  $f^*$  cannot exceed that of  $f$ , and  $I(f^*; y) \leq I(f; y)$ . Maximizing  $I(f^*; y)$  is therefore equivalent to:

$$\min I(f; y) - I(f^*; y). \quad (5)$$

To solve this optimization problem, we make use of the following proposition [23, 36] (proof can be found in the supplementary material):

$$KL[p(y|f)||p(y|f^*)] = 0 \implies I(y; f) - I(y; f^*) = 0, \quad (6)$$

where  $KL$  denotes the Kullback-Leibler divergence. Therefore, instead of having to estimate the mutual information in Eqn. (5), we formulate our sufficiency loss as

$$\mathcal{L}_{\text{Sufficiency}} = KL[p_{\theta}(y|f)||p(y|f^*)]. \quad (7)$$

### 3.4. Overall objective

The overall loss function in our DMIB framework is:

$$\mathcal{L} = \mathcal{L}_f(y, \hat{y}_f) + \alpha \sum_i \mathcal{L}_{\text{modality}}(y, \hat{y}_{f_i}) + \underbrace{\mathcal{L}_{f^*}(y, \hat{y}_{f^*}) + \beta \mathcal{L}_{\text{sufficiency}}(\hat{y}_f, \hat{y}_{f^*})}_{\text{Supervision of Information Bottleneck Module}} \quad (8)$$

where  $\mathcal{L}_f$  and  $\mathcal{L}_{f^*}$  are respective losses to supervise the direct fused expression datafeature and distilled fused features, and  $\mathcal{L}_{\text{modality}}$  supervises each modality backbone. Here,  $\alpha$  and  $\beta$  are model hyperparameters to control the roles of modal-specific supervision and feature-level supervision, respectively.  $\hat{y}_f$  and  $\hat{y}_{f^*}$  are the classifier results from the observed fused feature and distilled fused feature.

## 4. Experiments

### 4.1. Datasets

We trained our proposed DMIB for various multimodal medical classification tasks on four datasets: i) Our in-house **ITAC dataset**, which includes HRCT and 18 clinical features of 566 COVID19 inpatients, is utilized for predicting the 10-day mortality rate of COVID19 inpatients; ii) the public **iCTCF dataset** [28]<sup>3</sup>, which includes HRCT scans and up to 81 clinical features, is used to predict the morbidity outcomes of 751 COVID19 patients and the detection of 751 COVID19 patients from 529 non-COVID19 patients respectively; iii) **BRCA dataset** for diagnosis of breast carcinoma PAM50 subtypes (mRNA expression data, DNA methylation data, and miRNA expression data) of 875 patients [42]; iv) **ROSMAP dataset** for diagnosis of Alzheimer’s Disease from 351 patients [42]<sup>4</sup>. A summary of the various datasets is found in Table 1. It is worth mentioning that ITAC and iCTCF are collected from different countries and have no overlap in enrolled patients.

### 4.2. Implementation details

**ITAC and iCTCF.** For data preprocessing, we crop the lung regions from the complete 3D scan and resize each slice to the dimensions of  $350 \times 350$  pixels, preserving its original depth. Subsequently, we generate 2D montages for each patient from their 3D HRCT scans. Each montage is composed of 4 randomly chosen axial slices, with each slice originating from one of the 4 equally divided regions of the scan, placed on a  $2 \times 2$  grid. The utilization of 2D montages enables us to generate diverse montages for each scan without duplication. This approach can effectively mitigate the challenges of data scarcity and data imbalance commonly encountered in medical image datasets. It offers an advantage over using the entire scan as input, as well as compared to single-slice input, by preserving an adequate amount of predictive information in the input. The number of generated montages for each class is summarized in Table 1.

For training and testing the proposed model, we conduct patient-level data splitting to prevent information leakage. We reserve 20% of patients from each class for testing and using the remaining patients for model training with five-fold cross-validation. Patients are divided into five folds with no overlaps using stratified sampling, with each having the same patient distribution. Each of this fold serves as a validation set, and we train five models from scratch on the remaining four folds. From the five models, the model attaining the top AUC on its validation set is chosen as the final testing model. For testing, we generate 10 montages for each patient and take the median prediction over the 10 montages as the final result.

<sup>3</sup>Available at <https://ngdc.cncb.ac.cn/ictcf/>

<sup>4</sup>iii) and iv) are available at <https://github.com/txWang/MOGONET>

We employ DensetNet-121 [18] as the backbone for the CT montage modality to obtain an image feature  $f_1 \in \mathbb{R}^{1024}$ . For the clinical data, we perform mean imputation for missing records followed by data normalization. We employ four linear layers to obtain a clinical feature  $f_2 \in \mathbb{R}^{1024}$ . The dimension of  $z$  in the IB module is set to 1024 and loss function weights are set to  $\alpha = 1, \beta = 10$ . We use the cross-entropy loss for  $\mathcal{L}_f, \mathcal{L}_{f^*}, \mathcal{L}_{\text{modality}}$ . All training are done with the Adam optimizer (initial learning rate of  $1e-6$  and a linear decay rate of  $1e-2$ ), a batch size of 8, and trained for 70 epochs on a single RTX3090.

**ROSMAP and BRCA.** For fair comparison, we adopt similar experimental settings as [14], and we reenacted their experiments using their open-sourced codes. We set the dimension of  $z$  to 1024 for the bottleneck.

### 4.3. Evaluation setting

**ITAC.** To evaluate the robustness of different fusion strategies in handling noises and redundant information, we consider multiple experimental settings, starting with the most informative and complete clinical variables, and gradually expanding to those with larger degrees of redundancy and noise. Specifically, we fix five settings: i) 1 clinical variable (age) which is the most informative feature. ii) 4 clinical variables (i) + oxygen saturation, platelets, measured saturation oxygen, which are also deemed as relevant by domain experts. iii) 7 clinical variables (ii) + respiratory rate, pO<sub>2</sub>, and D-Dimer, which are informative but have a considerable 30% of the data are missing. iv) (iii) + reported symptoms and health records, which are more prone to missing data and subjective biases in the symptoms description, and deemed to be of secondary importance. v) 1 random integer between 0 and 100 that serves as a noise input replacing age. In the supplementary material, we further include the details of the clinical data for our ITAC dataset.

We benchmark DMIB against single modality models (CT backbone and clinical backbone in our method), as well as various multimodal fusion schemes including direct concatenation (Concatenation [24]), fusion via channelwise attention (Attention [10]), fusion via transformer-based cross-attention (Transformer [25]), fusion via dynamic weighing of each modality channel (Dynamic [14]), fusion via projection to a common subspace (Fusion: Subspace [50]) and graph-based fusion (Graph [49]). We adopt the Area Under Receiver Operating Characteristic Curve (AUC) as the primary evaluation metric, with accuracy, sensitivity, specificity, and the weighted Youden indices  $JW_{0.5}$  and  $JW_{0.6}$  as auxiliary metrics. Results are reported in Table 2.

**iCTCF.** To eliminate the manual selection of optimal clinical data combinations for improved fusion performance, we extend our evaluation on this dataset to assess the robustness of the proposed model for both diagnostic and prognostic tasks. We integrate the CT scans with all 81 available

Table 1. Summary of datasets

Dataset	Modality	Task Description	Enrolled Patients	Generated montages
ITAC	HRCT scans + up to 18 clinical features	Prognosis for COVID19 mortality in 10 Days	Deceased: 257 / Cured: 309	Deceased: 3084 / Cured: 3090
iCTCF	HRCT scans + up to 81 clinical features	Prognosis for COVID19 morbidity outcome	Severe symptoms: 202 / Mild symptoms: 549	Severe symptoms: 606 / Mild symptoms: 549
		Diagnosis for COVID19 patients	PCR positive: 751 / PCR negative: 529	PCR positive: 3755 / PCR negative: 3174
BRCA	mRNA, DNA methylation, miRNA	Diagnosis for breast carcinoma PAM50 subtype	Normal: 115 / Basal: 131 / Her2: 46 / LumA: 436 / LumB: 147	/
ROSMAP	mRNA, DNA methylation, miRNA	Diagnosis for Alzheimer’s Disease	Normal: 169 / AD: 182	/

Table 2. Performance of various methods on the ITAC dataset. Bold denotes the clinical settings where each methods achieves best AUC.

	Test							Test					
	AUC	Accuracy	Sensitivity	Specificity	JW <sub>0.5</sub>	JW <sub>0.6</sub>		AUC	Accuracy	Sensitivity	Specificity	JW <sub>0.5</sub>	JW <sub>0.6</sub>
	Single modality							Fusion: Concatenation [24]					
clinical 1	0.76	0.69	0.75	0.63	0.69	0.70	clinical 1	0.77	0.71	0.74	0.68	0.71	0.72
<b>clinical 4</b>	<b>0.82</b>	0.75	0.77	0.73	0.75	0.75	clinical 4	0.78	0.69	0.65	0.72	0.69	0.68
<b>clinical 7</b>	<b>0.82</b>	0.75	0.74	0.76	0.75	0.75	<b>clinical 7</b>	<b>0.83</b>	0.76	0.72	0.78	0.75	0.74
clinical 18	0.81	0.73	0.71	0.76	0.74	0.73	clinical 18	0.80	0.74	0.62	0.85	0.74	0.71
CT only	0.77	0.72	0.57	0.85	0.71	0.68	pure noise	0.74	0.66	0.72	0.62	0.67	0.68
	Fusion: Attention [10]							Fusion: Transformer [25]					
<b>clinical 1</b>	<b>0.81</b>	0.76	0.72	0.78	0.75	0.74	clinical 1	0.66	0.65	0.54	0.74	0.64	0.62
<b>clinical 4</b>	<b>0.81</b>	0.74	0.77	0.72	0.75	0.75	clinical 4	0.64	0.57	0.51	0.62	0.62	0.55
clinical 7	0.68	0.62	0.51	0.71	0.61	0.59	<b>clinical 7</b>	<b>0.73</b>	0.67	0.66	0.68	0.67	0.67
clinical 18	0.76	0.67	0.65	0.69	0.67	0.66	<b>clinical 18</b>	<b>0.73</b>	0.66	0.65	0.67	0.66	0.66
pure noise	0.65	0.62	0.51	0.72	0.61	0.59	pure noise	0.65	0.64	0.54	0.73	0.64	0.62
	Fusion: Dynamic [14]							Fusion: Graph [49]					
clinical 1	0.82	0.76	0.58	0.91	0.75	0.71	clinical 1	0.73	0.72	0.67	0.77	0.72	0.71
clinical 4	0.83	0.70	0.85	0.58	0.72	0.74	clinical 4	0.69	0.68	0.62	0.74	0.68	0.67
<b>clinical 7</b>	<b>0.85</b>	0.76	0.77	0.74	0.76	0.76	<b>clinical 7</b>	<b>0.76</b>	0.75	0.75	0.78	0.77	0.76
clinical 18	0.83	0.74	0.88	0.63	0.75	0.78	clinical 18	0.75	0.75	0.74	0.76	0.75	0.75
pure noise	0.75	0.66	0.66	0.67	0.67	0.66	pure noise	0.68	0.67	0.60	0.74	0.67	0.66
	Fusion: Subspace [50]							Fusion: DMIB					
<b>clinical 1</b>	<b>0.84</b>	0.75	0.69	0.79	0.74	0.73	clinical 1	0.85	0.76	0.82	0.72	0.77	0.78
<b>clinical 4</b>	<b>0.84</b>	0.76	0.75	0.76	0.76	0.75	clinical 4	0.85	0.78	0.78	0.77	0.78	0.78
<b>clinical 7</b>	<b>0.84</b>	0.78	0.77	0.78	0.78	0.77	<b>clinical 7</b>	<b>0.86</b>	0.80	0.75	0.83	0.79	0.78
clinical 18	0.83	0.76	0.57	0.92	0.75	0.71	<b>clinical 18</b>	<b>0.86</b>	0.76	0.77	0.76	0.77	0.77
pure noise	0.71	0.63	0.55	0.69	0.62	0.61	pure noise	0.76	0.72	0.66	0.77	0.72	0.70

clinical variables, to showcase the superiority of DMIB in fusing potentially redundant and noisy data without requiring manual selection. Furthermore, we perform fusion with single variables, ‘Age’ and ‘pure noise’ (consisting of random age and temperature), as reference points for assessing fusion performance with all clinical information.

**BRCA & ROSMAP.** To further evaluate the generalization of DMIB to other modalities, we also performed experiments on the BRCA and ROSMAP dataset consisting of multimodal genomic data. We benchmark against Fusion: Dynamic [14], the state-of-the-art method which employs fusion strategy consisting of model-specific supervision, attention, sparsity constraints and dynamic assignment of confidence to each modality. Following [14], we report AUC, WeightedF1 and MACroF1 and F1 in Table 5.

#### 4.4. Results

**ITAC.** As reported in Table 2, our DMIB outperforms all competing methods across all settings for our ITAC dataset. As expected, we observe that inclusion of more clinical variables often fails to improve prognostic performance in many existing methods. In fact, the inclusion of more modalities in attention-based, transformer-based and graph-

based fusion might even lead to weaker performance than single modalities. Such fluctuations across different experimental settings reflect a lack of robustness in the models.

By inspecting gradients for each variable, we observe that age is predominant, followed by GOT, O2 saturation, PCR, Glucose, Platelets of secondary importance, in agreement with expert knowledge. This is reflected in our method, where our model attains peak performance when fusing the image modality with 7 clinical variables. Furthermore, our DMIB maintains peak performance upon adding more noisy and redundant clinical features, unlike most existing methods which show declining trend. Another noteworthy point is when fusing the CT modality with a random noise input, our method retains its performance, whereas all existing fusion methods except concatenation method experience considerable performance drops. Again, this demonstrates the robustness of our approach and its ability to filter out noises. Overall, our method is more feasible and reliable for real-world clinical applications, when noise and redundant information are often present, and there is no prior guidance for which features are informative.

We discuss some insights for the various fusion schemes. Concatenation [24] fails to surpass the clinical modality for

Table 3. Performance of various methods on the iCTCF dataset for morbidity prediction. Bold denotes the best AUC in each fusion setting.

	Test							Test					
	AUC	Accuracy	Sensitivity	Specificity	JW <sub>0.5</sub>	JW <sub>0.6</sub>		AUC	Accuracy	Sensitivity	Specificity	JW <sub>0.5</sub>	JW <sub>0.6</sub>
	Single modality							Fusion: Concatenation [24]					
clinical I	0.72	0.62	0.71	0.59	0.65	0.66	clinical I	0.71	0.63	0.67	0.62	0.64	0.65
All clinical	0.74	0.69	0.67	0.70	0.68	0.68	All clinical	0.77	0.72	0.65	0.75	0.70	0.69
CT only	0.71	0.68	0.63	0.70	0.67	0.66	<b>pure noise</b>	<b>0.71</b>	0.65	0.67	0.64	0.66	0.66
	Fusion: Attention [10]							Fusion: Transformer [25]					
clinical I	0.72	0.65	0.63	0.65	0.64	0.64	clinical I	0.59	0.74	0.06	0.99	0.53	0.43
All clinical	0.77	0.75	0.65	0.79	0.72	0.70	All clinical	0.67	0.67	0.59	0.70	0.65	0.63
pure noise	0.56	0.69	0.31	0.83	0.57	0.52	pure noise	0.65	0.38	0.90	0.18	0.54	0.61
	Fusion: Dynamic [14]							Fusion: Graph [49]					
clinical I	0.78	0.75	0.16	0.96	0.56	0.48	clinical I	0.71	0.70	0.67	0.72	0.70	0.69
All clinical	0.78	0.50	0.90	0.35	0.62	0.68	All clinical	0.74	0.72	0.67	0.77	0.72	0.71
pure noise	0.67	0.73	0.00	1.00	0.50	0.40	pure noise	0.62	0.64	0.69	0.60	0.65	0.65
	Fusion: Subspace [50]							Fusion: DMIB					
clinical I	0.78	0.63	0.88	0.54	0.71	0.75	<b>clinical I</b>	<b>0.79</b>	0.70	0.76	0.68	0.72	0.73
All clinical	0.80	0.68	0.76	0.64	0.70	0.72	<b>All clinical</b>	<b>0.82</b>	0.73	0.71	0.74	0.72	0.72
<b>pure noise</b>	<b>0.71</b>	0.58	0.75	0.52	0.63	0.66	<b>pure noise</b>	<b>0.71</b>	0.63	0.69	0.62	0.65	0.66

Table 4. Performance of various methods on the iCTCF dataset for COVID19 diagnosis. Bold denotes the best AUC in each fusion setting.

	Test							Test					
	AUC	Accuracy	Sensitivity	Specificity	JW <sub>0.5</sub>	JW <sub>0.6</sub>		AUC	Accuracy	Sensitivity	Specificity	JW <sub>0.5</sub>	JW <sub>0.6</sub>
	Single modality							Fusion: Concatenation [24]					
clinical I	0.51	0.51	0.55	0.46	0.51	0.52	<b>clinical I</b>	<b>0.81</b>	0.74	0.72	0.76	0.74	0.73
All clinical	0.80	0.74	0.77	0.70	0.73	0.74	All clinical	0.85	0.78	0.85	0.68	0.76	0.78
CT only	0.80	0.72	0.85	0.55	0.70	0.73	<b>pure noise</b>	<b>0.81</b>	0.73	0.81	0.61	0.71	0.73
	Fusion: Attention [10]							Fusion: Transformer [25]					
clinical I	0.70	0.67	0.74	0.58	0.66	0.68	clinical I	0.62	0.63	0.86	0.29	0.58	0.63
All clinical	0.80	0.75	0.80	0.69	0.74	0.76	All clinical	0.66	0.64	0.68	0.59	0.63	0.64
pure noise	0.50	0.49	0.61	0.33	0.47	0.50	pure noise	0.64	0.59	0.60	0.59	0.59	0.59
	Fusion: Dynamic [14]							Fusion: Graph [49]					
clinical I	0.80	0.74	0.85	0.59	0.72	0.74	clinical I	0.74	0.73	0.74	0.73	0.74	0.74
All clinical	0.86	0.76	0.86	0.62	0.74	0.76	All clinical	0.83	0.83	0.84	0.83	0.84	0.84
pure noise	0.80	0.74	0.82	0.62	0.72	0.74	pure noise	0.73	0.73	0.72	0.74	0.73	0.73
	Fusion: Subspace [50]							Fusion: DMIB					
clinical I	0.79	0.70	0.80	0.56	0.68	0.71	clinical I	0.80	0.73	0.86	0.53	0.70	0.73
All clinical	0.86	0.77	0.79	0.74	0.77	0.77	<b>All clinical</b>	<b>0.89</b>	0.82	0.78	0.87	0.83	0.82
pure noise	0.80	0.72	0.95	0.40	0.67	0.73	pure noise	0.80	0.73	0.74	0.71	0.73	0.73

Table 5. Performance of multimodal methods on BRCA and ROSMAP. † denotes the result showing significant difference to that of DMIB with  $p < 1e - 4$  by paired T test.

Method	BRCA			ROSMAP		
	ACC	WeightedF1	MacroF1	ACC	F1	AUC†
Fusion: Dynamic	<b>87.1+0.5†</b>	<b>87.4+0.6†</b>	<b>83.5+0.9†</b>	81.7+1.5†	82.3+1.5†	90.0+1.2†
Fusion: DMIB	86.0+0.7	86.0+0.8	81.6+0.9	<b>84.9+1.8</b>	<b>85.3+1.7</b>	<b>91.6+0.7</b>

1 and 4 clinical variables. This is probably due to the image feature dimensions being much larger than the clinical feature dimensions, resulting in the dominant influence of one modality. Channel-wise attention fusion [10] only boosts performance when the modality data is complete and informative. It is particularly sensitive to noise, and drops significantly upon inclusion of noisy and redundant data. Transformer-based fusion also delivers very poor performance, never surpassing single-modality results. This is consistent with observations that attention-based mech-

anisms and transformers typically require a larger corpus of training data for effective training [43]. Dynamic fusion [14], subspace-based fusion [50] and graph-based fusion [25] demonstrated better consistency and solid fusion performance. However, they are prone to be affected by noisy modalities. For subspace projection-based fusion, replacing the clinical variable with noise leads to considerable performance dips. Intuitively, aligning normal inputs to random noise could remove predictive information in the normal channels. For dynamic fusion, learning to weigh the reliability of features from a modality does not generalize well to test samples at an instance level, leading to a drop when including unreliable clinical inputs.

**iCTCF.** Furthermore, we demonstrate the exceptional fusion performance of DMIB when combined with all clinical data in both prognosis and diagnosis tasks, along with its resilience to noise input, as evidenced in Table 3 and Table 4. Notably, when compared to using CT only, adding

Table 6. Ablation studies on ITAC and BRCA datasets

No.	Ablation Setting					ITAC				BRCA		
	$\mathcal{L}_f$	IB	$\mathcal{L}_{f^*}$	$\mathcal{L}_{\text{sufficiency}}$	$\mathcal{L}_{\text{modality}}$	AUC	Accuracy	Sensitivity	Specificity	ACC	WeightedF1	MacroF1
1.	✓	-	-	-	-	0.83	0.78	0.65	0.90	81.4+0.6 <sup>†</sup>	81.1+0.7 <sup>†</sup>	76.2+1.0 <sup>†</sup>
2.	-	✓	✓	-	✓	0.83	0.77	0.74	0.79	81.0+1.0 <sup>†</sup>	80.2+1.2 <sup>†</sup>	75.3+1.2 <sup>†</sup>
3.	✓	✓	✓	-	✓	0.83	0.75	0.69	0.79	81.1+0.8 <sup>†</sup>	80.4+0.8 <sup>†</sup>	75.4+1.0 <sup>†</sup>
4.	✓	✓	✓	✓	-	0.85	0.77	0.75	0.78	83.9+1.5 <sup>†</sup>	83.8+1.7 <sup>†</sup>	79.7+1.8 <sup>†</sup>
5.	✓	✓	-	✓	✓	0.84	0.78	0.65	0.88	86.1+0.5 <sup>†</sup>	86.2+0.6 <sup>†</sup>	81.8+0.7 <sup>†</sup>
6.	✓	✓	✓	✓	✓	0.86	0.76	0.77	0.76	86.0+0.7 <sup>†</sup>	86.0+0.8 <sup>†</sup>	81.6+0.9 <sup>†</sup>

Table 7. Experimenting with different image backbones and unchanged clinical backbone on ITAC (with all clinical variables). Bond fonts denote the best result for each backbone.

Image Backbone	Experiment	AUC	ACC	Sens	Spec	JW5	JW6
ResNet18 [16]	CT Only	0.77	0.71	0.68	0.74	0.71	0.7
	Concat	0.77	0.72	0.7	0.75	0.73	0.72
	Subspace	0.79	0.73	0.73	0.73	0.73	0.73
	<b>Proposed</b>	<b>0.83</b>	<b>0.77</b>	<b>0.77</b>	<b>0.77</b>	<b>0.77</b>	<b>0.77</b>
ResNet50 [16]	CT Only	0.71	0.69	0.75	0.63	0.69	0.7
	Concat	0.75	0.67	0.74	0.62	0.68	0.69
	Subspace	0.77	0.72	0.75	0.69	<b>0.72</b>	<b>0.73</b>
	<b>Proposed</b>	<b>0.81</b>	<b>0.72</b>	<b>0.74</b>	<b>0.71</b>	<b>0.72</b>	<b>0.73</b>
EfficientNet_b0 [35]	CT Only	0.71	0.62	0.58	0.65	0.62	0.61
	Concat	0.76	0.66	0.72	0.62	0.67	0.68
	Subspace	0.74	0.65	0.71	0.60	0.66	0.67
	<b>Proposed</b>	<b>0.82</b>	<b>0.73</b>	<b>0.74</b>	<b>0.72</b>	<b>0.73</b>	<b>0.73</b>
PoolFormer_v2_tiny [47]	CT Only	0.69	0.62	0.62	0.63	0.62	0.62
	Concat	0.73	0.67	0.66	0.68	0.67	0.67
	Subspace	0.77	0.71	0.71	0.72	0.71	0.71
	<b>Proposed</b>	<b>0.79</b>	<b>0.73</b>	<b>0.74</b>	<b>0.72</b>	<b>0.73</b>	<b>0.73</b>

the singular feature ‘Age’, has a significant impact on predicting COVID-19 patient outcomes (Table 3). However, it shows only a marginal improvement in the COVID-19 diagnosis task (Table 4). As such, for COVID19 diagnosis, age can be deemed as a redundant feature and indeed. We observe similar performance in iCTCF when combining the CT modality with the noise feature. In both tables, our method and the concatenation method demonstrate the highest level of robustness, with our method achieving the best fusion performance.

**BRCA & ROSMAP.** As reported in Table 5, DMIB is slightly inferior for breast cancer subtype classification but significantly superior for Alzheimer’s diagnosis than [14], which proves its adaptability across different modalities.

#### 4.5. Ablation studies

We performed ablation studies on the ITAC and BRCA datasets to study the effectiveness of various key components, with results reported in Table 6. It is observed that preserving all information (No. 1) is worse than incorporating a bottleneck with sufficiency loss (No. 4, 5, 6). Furthermore, employing a bottleneck module without the sufficiency loss also fails to be useful, which is consistent with our intuitions that unconstrained discarding of information could remove predictive information (comparing No. 2, 3

to No. 6). Moreover, including modal-specific supervision  $\mathcal{L}_{\text{modality}}$  leads to improvement, suggesting that boosting the extraction of modality-level task-relevant information facilitates learning of fused feature (No. 4 and No. 6). This provides strong evidence for the effectiveness of the sufficiency loss in preserving predictive information.

We further investigate whether DMIB remains effective under different network architectures by switching the image backbones for the prognosis experiment on our ITAC dataset. As reported in Table 7, we experimented with ResNet18, ResNet50 [16], EfficientNet [35] and PoolFormer [47]. Regardless of the backbone, DMIB performs significantly better than the baselines of CT only and fusion methods such as concatenation (which retains all information) and subspace-based fusion (which extracts common information). These results strongly affirm the efficacy of our approach and its viability as a plug-and-play component for multimodal classification tasks.

## 5. Conclusion

Utilising insights from mutual information and information bottleneck theory, we have introduced a general multimodal classification approach which has achieved a state-of-the-art performance in various clinical tasks. Notably, our method holds great significance and applicability in the field of clinical practice, as it consistently delivers high performance, even when faced with limited training data and noisy input modalities. In the context of future research, our efforts will be directed towards expanding the application of information bottleneck theory to facilitate feature importance analysis and enhance the interpretability of multimodal learning.

## Acknowledgment

This study was supported in part by the ERC IMI (101005122), the H2020 (952172), the MRC (MC/PC/21013), the Royal Society (IEC\NSFC\211235), the NVIDIA Academic Hardware Grant Program, the SABER project supported by Boehringer Ingelheim Ltd, Wellcome Leap Dynamic Resilience, and the UKRI Future Leaders Fellowship (MR/V023799/1).



## References

- [1] Julián N Acosta, Guido J Falcone, Pranav Rajpurkar, and Eric J Topol. Multimodal biomedical ai. *Nature Medicine*, 28(9):1773–1784, 2022. 2
- [2] Francis Bach. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017. 2
- [3] Fatemeh Behrad and Mohammad Saniee Abadeh. An overview of deep learning methods for multimodal medical data mining. *Expert Systems with Applications*, page 117006, 2022. 1, 2
- [4] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and R Devon Hjelm. Mine: mutual information neural estimation. *arXiv preprint arXiv:1801.04062*, 2018. 2
- [5] Yonatan Belinkov, James Henderson, et al. Variational information bottleneck for effective low-resource fine-tuning. In *International Conference on Learning Representations*, 2020. 3
- [6] Gan Cai, Yu Zhu, Yue Wu, Xiaoben Jiang, Jiongyao Ye, and Dawei Yang. A multimodal transformer to fuse images and metadata for skin disease classification. *The Visual Computer*, pages 1–13, 2022. 2
- [7] Qiong Cai, Hao Wang, Zhenmin Li, and Xiao Liu. A survey on multimodal data-driven smart healthcare systems: approaches and applications. *IEEE Access*, 7:133583–133599, 2019. 1
- [8] Menglin Cao, Ming Yang, Chi Qin, Xiaofei Zhu, Yanni Chen, Jue Wang, and Tian Liu. Using deepgcn to identify the autism spectrum disorder from multi-site resting-state data. *Biomedical Signal Processing and Control*, 70:103015, 2021. 2
- [9] Can Cui, Haichun Yang, Yaohong Wang, Shilin Zhao, Zuhayr Asad, Lori A Coburn, Keith T Wilson, Bennett A Landman, and Yuankai Huo. Deep multimodal fusion of image and non-image data in disease diagnosis and prognosis: A review. *arXiv preprint arXiv:2203.15588*, 2022. 2
- [10] Hongyi Duanmu, Pauline Boning Huang, Srinidhi Brahmavar, Stephanie Lin, Thomas Ren, Jun Kong, Fusheng Wang, and Tim Q Duong. Prediction of pathological complete response to neoadjuvant chemotherapy in breast cancer using deep learning with integrative imaging, molecular and demographic data. In *International conference on medical image computing and computer-assisted intervention*, pages 242–252. Springer, 2020. 2, 5, 6, 7
- [11] Tyrone E Duncan. On the calculation of mutual information. *SIAM Journal on Applied Mathematics*, 19(1):215–220, 1970. 2
- [12] Marco Federici, Anjan Dutta, Patrick Forré, Nate Kushman, and Zeynep Akata. Learning robust representations via multi-view information bottleneck. In *8th International Conference on Learning Representations*. OpenReview. net, 2020. 2, 4
- [13] Fengjiao Gong, Yuzhou Nie, and Hongteng Xu. Gromov-wasserstein multi-modal alignment and clustering. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 603–613, 2022. 2
- [14] Zongbo Han, Fan Yang, Junzhou Huang, Changqing Zhang, and Jianhua Yao. Multimodal dynamics: Dynamical fusion for trustworthy multimodal classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20707–20717, 2022. 2, 5, 6, 7, 8
- [15] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 4
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 8
- [17] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*, 2018. 2
- [18] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 5
- [19] Grzegorz Jacenków, Alison Q O’Neil, and Sotirios A Tsafaris. Indication as prior knowledge for multimodal disease classification in chest radiographs with transformers. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2022. 2
- [20] Matthias Keicher, Hendrik Burwinkel, David Bani-Harouni, Magdalini Paschali, Tobias Czempiel, Egon Burian, Marcus R Makowski, Rickmer Braren, Nasir Navab, and Thomas Wendler. U-GAT: Multimodal graph attention network for covid-19 outcome prediction. *arXiv preprint arXiv:2108.00860*, 2021. 2
- [21] Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Foundations and recent trends in multimodal machine learning: Principles, challenges, and

- open questions. *arXiv preprint arXiv:2209.03430*, 2022. [1](#)
- [22] Ralph Linsker. Self-organization in a perceptual network. *Computer*, 21(3):105–117, 1988. [2](#)
- [23] Zhenguang Liu, Runyang Feng, Haoming Chen, Shuang Wu, Yixing Gao, Yunjun Gao, and Xiang Wang. Temporal feature alignment and mutual information maximization for video-based human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11006–11016, 2022. [4](#)
- [24] Lingwei Meng, Di Dong, Liang Li, Meng Niu, Yan Bai, Meiyun Wang, Xiaoming Qiu, Yunfei Zha, and Jie Tian. A deep learning prognosis model help alert for covid-19 patients at high-risk of death: a multicenter study. *IEEE journal of biomedical and health informatics*, 24(12):3576–3584, 2020. [2](#), [5](#), [6](#), [7](#)
- [25] Satyam Mohla, Shivam Pande, Biplab Banerjee, and Subhasis Chaudhuri. Fusatnet: Dual attention based spectrospatial multimodal fusion network for hyperspectral and lidar classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 92–93, 2020. [2](#), [5](#), [6](#), [7](#)
- [26] Ghulam Muhammad, Fatima Alshehri, Fakhri Karray, Abdulmotaleb El Saddik, Mansour Alsulaiman, and Tiago H Falk. A comprehensive survey on multimodal medical signals fusion for smart healthcare systems. *Information Fusion*, 76:355–375, 2021. [1](#), [2](#)
- [27] XuanLong Nguyen, Martin J Wainwright, and Michael I Jordan. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, 2010. [4](#)
- [28] Wanshan Ning, Shijun Lei, Jingjing Yang, Yukun Cao, Peiran Jiang, Qianqian Yang, Jiao Zhang, Xiaobei Wang, Fenghua Chen, Zhi Geng, et al. Open resource of clinical data from patients with pneumonia for the prediction of covid-19 outcomes via deep learning. *Nature biomedical engineering*, 4(12):1197–1207, 2020. [5](#)
- [29] Sherjil Ozair, Corey Lynch, Yoshua Bengio, Aaron Van den Oord, Sergey Levine, and Pierre Sermanet. Wasserstein dependency measure for representation learning. *Advances in Neural Information Processing Systems*, 32, 2019. [2](#), [4](#)
- [30] Sebastian Pölsterl, Tom Nuno Wolf, and Christian Wachinger. Combining 3d image and tabular data via the dynamic affine feature map transform. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 688–698. Springer, 2021. [2](#)
- [31] Ben Poole, Sherjil Ozair, Aaron Van Den Oord, Alex Alemi, and George Tucker. On variational bounds of mutual information. In *International Conference on Machine Learning*, pages 5171–5180. PMLR, 2019. [2](#), [4](#)
- [32] Alvin Rajkumar, Eyal Oren, Kai Chen, Andrew M Dai, Nissan Hajaj, Michaela Hardt, Peter J Liu, Xiaobing Liu, Jake Marcus, Mimi Sun, et al. Scalable and accurate deep learning with electronic health records. *NPJ digital medicine*, 1(1):1–10, 2018. [1](#)
- [33] Farah E Shamout, Yiqiu Shen, Nan Wu, Aakash Kaku, Jungkyu Park, Taro Makino, Stanislaw Jastrzebski, Jan Witowski, Duo Wang, Ben Zhang, et al. An artificial intelligence system for predicting the deterioration of covid-19 patients in the emergency department. *NPJ digital medicine*, 4(1):1–11, 2021. [2](#)
- [34] Dinggang Shen, Guorong Wu, and Heung-Il Suk. Deep learning in medical image analysis. *Annual review of biomedical engineering*, 19:221, 2017. [1](#)
- [35] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. [8](#)
- [36] Xudong Tian, Zhizhong Zhang, Shaohui Lin, Yanyun Qu, Yuan Xie, and Lizhuang Ma. Farewell to mutual information: Variational distillation for cross-modal person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1522–1531, 2021. [2](#), [4](#)
- [37] Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000. [3](#), [4](#)
- [38] Naftali Tishby and Noga Zaslavsky. Deep learning and the information bottleneck principle. In *2015 IEEE information theory workshop (itw)*, pages 1–5. IEEE, 2015. [2](#), [3](#)
- [39] Tom van Sonsbeek, Xiantong Zhen, Marcel Worring, and Ling Shao. Variational knowledge distillation for disease classification in chest x-rays. In *Information Processing in Medical Imaging: 27th International Conference, IPMI 2021, Virtual Event, June 28–June 30, 2021, Proceedings 27*, pages 334–345. Springer, 2021. [2](#)
- [40] Matias Vera, Pablo Piantanida, and Leonardo Rey Vega. The role of the information bottleneck in representation learning. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 1580–1584. IEEE, 2018. [2](#), [3](#)
- [41] Zhibin Wan, Changqing Zhang, Pengfei Zhu, and Qinghua Hu. Multi-view information-bottleneck representation learning. In *Proceedings of the AAAI Con-*

- ference on Artificial Intelligence*, volume 35, pages 10085–10092, 2021. [2](#)
- [42] Tongxin Wang, Wei Shao, Zhi Huang, Haixu Tang, Jie Zhang, Zhengming Ding, and Kun Huang. Moronet: multi-omics integration via graph convolutional networks for biomedical data classification. *bioRxiv*, pages 2020–07, 2020. [5](#)
- [43] Peng Xu, Xiatian Zhu, and David A Clifton. Multi-modal learning with transformers: A survey. *arXiv preprint arXiv:2206.06488*, 2022. [2](#), [7](#)
- [44] Jiawen Yao, Xinliang Zhu, Feiyun Zhu, and Junzhou Huang. Deep correlational learning for survival prediction from multi-modality data. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11–13, 2017, Proceedings, Part II*, pages 406–414. Springer, 2017. [2](#)
- [45] Jordan Yap, William Yolland, and Philipp Tschandl. Multimodal skin lesion classification using deep learning. *Experimental dermatology*, 27(11):1261–1267, 2018. [2](#)
- [46] Kun-Hsing Yu, Andrew L Beam, and Isaac S Kohane. Artificial intelligence in healthcare. *Nature biomedical engineering*, 2(10):719–731, 2018. [1](#)
- [47] Weihao Yu, Mi Luo, Pan Zhou, Chenyang Si, Yichen Zhou, Xinchao Wang, Jiashi Feng, and Shuicheng Yan. Metaformer is actually what you need for vision. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10819–10829, 2022. [8](#)
- [48] Zexian Zeng, Chengsheng Mao, Andy Vo, Xiaoyu Li, Janna Ore Nugent, Seema A Khan, Susan E Clare, and Yuan Luo. Deep learning for cancer type classification and driver gene identification. *BMC bioinformatics*, 22(4):1–13, 2021. [1](#)
- [49] Shuai Zheng, Zhenfeng Zhu, Zhizhe Liu, Zhenyu Guo, Yang Liu, Yuchen Yang, and Yao Zhao. Multi-modal graph learning for disease prediction. *IEEE Transactions on Medical Imaging*, 41(9):2207–2216, 2022. [2](#), [5](#), [6](#), [7](#)
- [50] Jinzhao Zhou, Xingming Zhang, Ziwei Zhu, Xi-angyuan Lan, Lunkai Fu, Haoxiang Wang, and Hanchun Wen. Cohesive multi-modality feature learning and fusion for covid-19 patient severity prediction. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(5):2535–2549, 2021. [2](#), [5](#), [6](#), [7](#)