

Incomplete Modality Disentangled Representation for Ophthalmic Disease Grading and Diagnosis

Chengzhi Liu^{1,2*}, Zile Huang^{1,2*}, Zhe Chen², Feilong Tang^{2†}, Yu Tian⁴, Zhongxing Xu²,
Zihong Luo², Yalin Zheng², Yanda Meng^{1,2†}

¹ Department of Computer Science, University of Exeter, UK

² Department of Eye and Vision Sciences, University of Liverpool, UK

⁴ Center For AI And Data Science For Integrated Diagnostics, University of Pennsylvania, USA
Y.M.Meng@exeter.ac.uk

Abstract

Ophthalmologists typically require multimodal data sources to improve diagnostic accuracy in clinical decisions. However, due to medical device shortages, low-quality data and data privacy concerns, missing data modalities are common in real-world scenarios. Existing deep learning methods tend to address it by learning an implicit latent subspace representation for different modality combinations. We identify two significant limitations of these methods: (1) implicit representation constraints that hinder the model’s ability to capture modality-specific information and (2) modality heterogeneity, causing distribution gaps and redundancy in feature representations. To address these, we propose an Incomplete Modality Disentangled Representation (IMDR) strategy, which disentangles features into explicit independent modal-common and modal-specific features by guidance of mutual information, distilling informative knowledge and enabling it to reconstruct valuable missing semantics and produce robust multimodal representations. Furthermore, we introduce a joint proxy learning module that assists IMDR in eliminating intra-modality redundancy by exploiting the extracted proxies from each class. Experiments on four ophthalmology multimodal datasets demonstrate that the proposed IMDR outperforms the state-of-the-art methods significantly.

Introduction

Retinal Fundus Imaging and Optical Coherence Tomography are widely used 2D and 3D imaging techniques for detecting ophthalmic diseases. Recent methods have integrated multiple modalities to enhance diagnostic accuracy (Watanabe et al. 2022; Wang et al. 2023b; Peng et al. 2024b; Zou et al. 2023; Tang et al. 2024a; Trinh et al. 2024; Xiong et al. 2024; Li et al. 2024b; Wang et al. 2023d; Duan et al. 2024; Qu et al. 2023; Meng et al. 2024; Peng et al. 2023). Despite the benefits of more comprehensive diagnostic insights compared to single-modality approaches, factors such as low-quality data, lack of equipment, and data privacy concerns in real-world practice settings can lead to missing modalities (Warner et al. 2024). Developing a reliable method for

ophthalmic disease diagnosis that effectively addresses incomplete modalities remains challenging.

Existing methods for handling missing-modality problems can be classified into two categories: (1) Naive Modality Generation (Wang, Cui, and Li 2023; Chen et al. 2024) reconstructs missing modalities from the available ones. However, controlling and generating medical image quality is complex due to a lack of clinical knowledge background, thus introducing noise and suffering from computational burden. (2) Latent Subspace Methods (Liu et al. 2023a; Wang et al. 2023a; Liu et al. 2023b; Wang et al. 2023c; Shi et al. 2024; Li et al. 2024a; Tang et al. 2024b; Yang et al. 2024) directly project inputs with different modality combinations into a deterministic embedding, ensuring consistency between the features and logits of the student and teacher models, as illustrated in Fig. 1 (a). Despite the success of these methods, they suffer from poor scalability and limited feature representations during knowledge distillation and do not intrinsically investigate the integration among various modalities.

Under the missing-modality setting, we identify a sub-optimal challenge in multimodal distillation: Although the teacher model generates more valuable feature representations than the student model, these representations are redundant and lack modality-specific information, resulting in ineffective distillation. The reasons for this challenge are twofold: **(i) Modality Heterogeneity:** The discrepancies in the inherent properties, statistical distribution, and structural characteristics of data across different modalities lead to a feature distribution gap and information redundancy during multimodal joint learning, which is confirmed by the empirical observation in (Liang et al. 2022; Udandarao, Gupta, and Albanie 2023; Hu et al. 2024). These discrepancies are further exacerbated under the missing-modality setting, leading to degraded model performance. **(ii) Implicit Representation Constraint:** As shown in Fig. 1 (c), different input combinations from the same class are forced to learn the embeddings in the same direction of latent space, reducing the diversity of features. The lack of feature diversity results in sub-optimal distillation, affecting the model performance.

To this end, we propose an **Incomplete Modality Disentangled Representation (IMDR)** strategy for ophthalmic disease diagnosis to disentangle features into ex-

*These authors contributed equally.

†Corresponding author: Feilong Tang; Yanda Meng
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

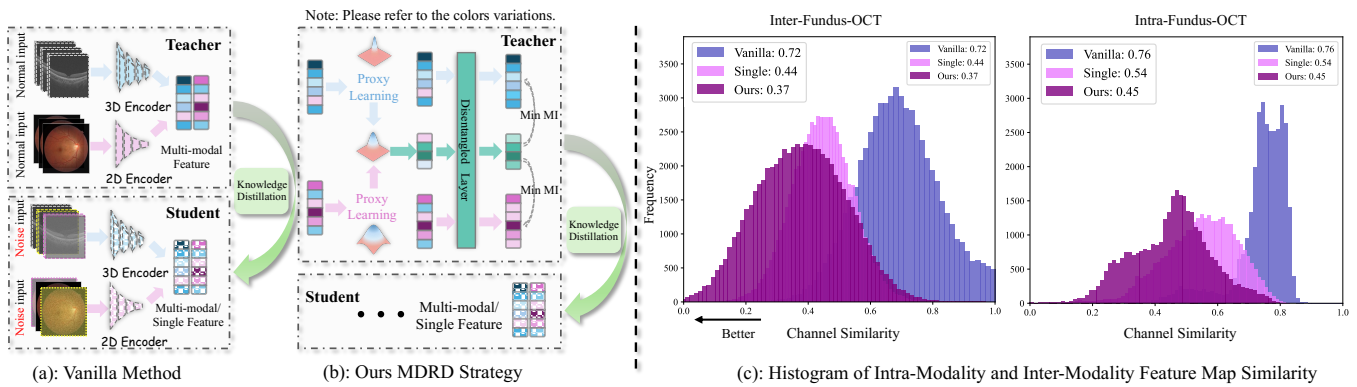


Figure 1: (a) Vanilla latent subspace methods. (b) Our proposed IMDR strategy effectively decouples multimodal data by employing explicit constraints to minimize mutual information in a Disentangle Extraction layer, guided by a joint distribution. (c) Illustration of intra-modality and inter-modality inter-channel distances between encoder feature maps. “Single” is the model that trains the encoders of each modality independently, providing the ideal feature diversity without inter-modality interference. “A” denotes the histogram mean. Lower inter-channel similarity means higher diversity. More details in the Appendix B.

explicit independent modal-shared and modal-specific features by the guidance of the modality joint distribution, distilling informative knowledge from the teacher model to the student network and enabling it to reconstruct valuable missing semantics and produce robust multimodal representations, as shown in Fig. 1 (b). Specifically, IMDR constructs probabilistic modality-specific representations obeying Gaussian distribution, allowing their combination into a joint distribution that estimates the latent space for each modality input. Modal-shared information is extracted by sampling from this distribution. The Disentangle Extraction (DE) layer, using an attention mechanism, minimizes mutual information (MI) between modality mean representations, ensuring the features are effectively disentangled into independent modal-shared and modality-specific components, thus distilling informative knowledge to the student network.

To eliminate intra-modality redundancy and obtain a robust joint distribution, we introduce a Joint Proxy Learning (JPL) module that utilizes multiple sets of learnable proxies for each modality. Each set corresponds to a latent space distribution aligned with its ground truth label, effectively capturing discriminative information. The module reduces the overlap in feature representations between classes by maximizing the similarities between encoded single-modality features and positive proxies while minimizing those with negative proxies. Consequently, a joint distribution derived from proxies replaces the original joint distribution as the guidance, which becomes more refined and less redundant, enhancing the robustness of the model.

We evaluate the proposed method on intra and inter-modality incompleteness conditions across four ophthalmology multimodal datasets, where our approach achieves state-of-the-art performances. The contributions of our work are summarized as:

- We propose an Incomplete Modality Disentangled Representation (IMDR) strategy that disentangles features into explicit independent modal-shared and modal-specific features and distills informative knowledge from

the teacher model to the student network.

- We propose an auxiliary Joint Proxy Learning (JPL) module to eliminate intra-modality redundancy to obtain a robust joint distribution.
- Experiments on four ophthalmology multimodal datasets demonstrate the effectiveness of the IMDR strategy.

Related Work

Incomplete Multimodal Learning. Recent methods in incomplete multimodal learning emphasize data reconstruction and latent subspace methods. Traditional data reconstruction methods utilize Generative Adversarial Networks (Jue et al. 2019; Liu et al. 2021) to enable models to simulate complete datasets but incur high computational costs. To address this, advanced techniques (Li, Li, and Hu 2023; Miao et al. 2023; Xiong et al. 2023; Poudel et al. 2024; Hu et al. 2025; Xu et al. 2024) reduce computational overhead. Additionally, statistical models have been introduced to align distributions of reconstructed and available data (Zou et al. 2023; Wang, Cui, and Li 2023). Latent subspace methods map different data modalities into a shared latent space, enabling the learning of intra-modality relationships (Ebrahimi et al. 2023; Sun et al. 2024; Wang, Cui, and Li 2023). This work employs explicit constraints to remove redundant information and capture discriminative features, hence relaxing the implicit representation constraint.

Decoupled Multimodal Representation. Decoupled representation techniques can be categorized into Modality-Specific Learning and Training-Inference Decoupling, which are vital for optimizing performance across diverse data types and operational phases. Modality-Specific Learning ensures that the features of each data type are processed in a way that maintains their unique characteristics while being effectively integrated into a unified model (Tripathi et al. 2024; Li, Wang, and Cui 2023; Cong et al. 2024; Zhao et al. 2024). Training-Inference Decoupling involves using different models during the training phase compared

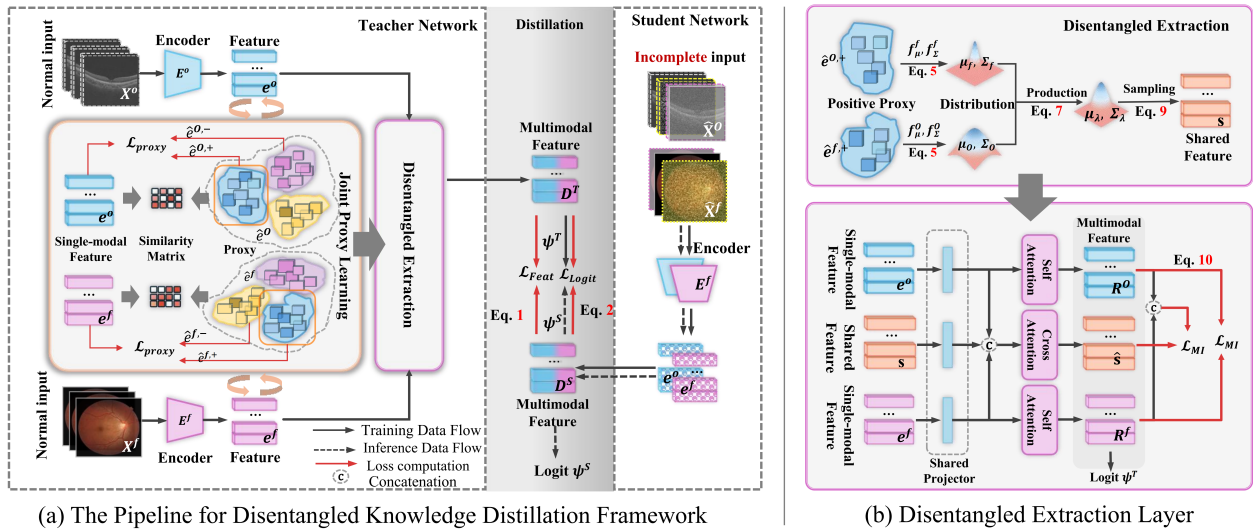


Figure 2: Overview of our proposed framework. (a): We train a teacher model using complete modality data, followed by co-training with a student model on incomplete inputs for knowledge distillation. The distillation is supervised by feature loss \mathcal{L}_{Feat} and logit loss \mathcal{L}_{Logit} . During the training of the teacher model, the encoder outputs the single-modality feature e^f and e^o . We build a set of proxies for a modality, with each set representing a class. Positive proxies are selected by a similarity matrix between \hat{e} and e . All proxies are optimized through the proxy loss \mathcal{L}_{prox} . Consequently, $\hat{e}^{f,+}$ and $\hat{e}^{o,+}$, together with features e^f and e^o are then passed to the IMDR. (b): Details for IMDR strategy. We estimate the distributions of $\hat{e}^{f,+}$ and $\hat{e}^{o,+}$, then combine them using Eq. 7 to obtain the joint distribution $\mathcal{P}(\hat{e}|x^f, x^o)$. The modality-shared feature s is sampled from this distribution. This feature s guides the decoupling via an attention layer, supervised by the loss \mathcal{L}_{MI} to minimize the mutual information between extracted shared features \hat{s} and specific features ($\mathcal{R}^f, \mathcal{R}^o$), as well as between \mathcal{R}^f and \mathcal{R}^o .

to the inference phase to better adapt to varying task demands (Zhang et al. 2024a; Tang et al. 2023; Yang et al. 2023; Peng et al. 2024a). Although these single-modality-based methods achieved remarkable improvements in feature extraction, they do not explicitly remove redundant information from inter-modality or intra-modality features. This work utilizes joint distribution across modalities to further guide disentangling multimodal data while preserving diversity in modality shared and specific representations.

Methodology

Problem Formulation

We denote $X = \{x_i, y_i\}_{i=1}^N$ as a multimodal dataset that has N samples. Each x_i consists of M inputs from different modalities as $x_i = \{x_i^m\}_{m=1}^M$ and $y_i \in \{1, 2, \dots, C\}$, where M is the number of modalities; C is the number of categories. The modality encoder E , learns from the input image x to produce a single-modal representation e . The incomplete modality is denoted as \hat{X} , and we define two cases of incomplete modalities to simulate the natural and holistic challenges in real-world scenarios: (i) **intra-modality incompleteness**, referring to impaired or noisy data within a specific modality. (ii) **inter-modality incompleteness**, where some modalities are entirely missing, such as the frame-level features in the OCT layers. We use the scenario of intra-modality incompleteness as an example for clarity of illustration, as shown in Fig. 2.

Overall Framework

Training and Inference. We first train the teacher model on complete-modality data X , and co-train with a student model with incomplete inputs \hat{X} for knowledge distillation. Our goal is to model both intra- and inter-modality relations to create more informative multimodal features D^T for distillation, enabling student model to reconstruct information in features D^S during inference and obtain more accurate logits ψ_S for ophthalmic disease grading and diagnosis across any degree of incomplete modalities.

Teacher Model Training. To eliminate intra-modality redundancy, we establish a set of proxies \hat{e} to capture discriminative information relevant to the label y for each modality. We optimize the parametric proxy \hat{e} to align the distribution of the data x . By calculating the similarity between sets of e and \hat{e} , we can identify a positive set of proxies \hat{e}^+ . For better stability, we predict the mean μ and covariance matrix Σ of proxies distribution by the distribution predictor f_μ and f_Σ , rather than modeling $\mathcal{P}(\hat{e}|x)$. The joint distribution $\mathcal{P}(z|x^f, x^o)$ is then obtained by multiplying the distributions of the positive proxies, and we randomly sample the modality-shared feature s from this joint distribution. Utilizing the modality-shared feature s as guidance, we disentangle multimodal data into independent modality-shared representations \hat{s} and modality-specific representations \mathcal{R}^f and \mathcal{R}^o for fundus and OCT within a Disentangle Extraction layer. Specifically, a parameter-shared projector first obtain the query, key and value of s, e^f , and e^o , respectively. Con-

sequently, modality-shared feature \hat{s} is extracted by a plain cross-attention, which is detailed in Appendix D. \mathcal{R}^f and \mathcal{R}^O are extracted by self-attention module. A constraint is set to minimize the mutual information between the mean representations of \hat{s} , \mathcal{R}^f and \mathcal{R}^O . Finally, \hat{s} , \mathcal{R}^f and \mathcal{R}^O are concatenated as multimodal feature \mathcal{D}^T to obtain the predicted logits ψ_T through a task-specific classifier.

Co-training Distillation. We employ feature-based and logit-based distillation to ensure consistency between the teacher model and the student model, enabling the student model to handle various scenarios of incomplete modality.

In feature-based distillation, for each class c , a softmax loss is applied to the multimodal features to generate soft targets, denoted as $\Gamma_c^T = \text{softmax}(D_c^T/\tau)$, where τ is the temperature parameter that controls the softening of the distribution. The class deterministic feature Γ_c^S of the student model is obtained similarly. The Kullback-Leibler (KL) divergence is then utilized to align the class deterministic feature of the student with those of the teacher, defined as:

$$\mathcal{L}_{\text{Feat}} = \sum_{c=1}^C \Gamma_c^T \log \left(\frac{\Gamma_c^T}{\Gamma_c^S} \right). \quad (1)$$

In logit-based distillation, given the logits $\psi^T = \{\psi_c^T\}_{c=1}^C$ from the teacher model and ψ^S from the student model, we enforce consistency between their cosine similarity matrices to effectively transfer inter-class relationship information. The similarity matrix for the teacher model is computed as $Q_c^T = \text{softmax}(\text{sim}(\psi_c^T, \psi^T)/\tau)$, and similarly for the student model, Q_c^S . The alignment between these similarity matrices is achieved through the KL divergence loss:

$$\mathcal{L}_{\text{Logit}} = \sum_{c=1}^C Q_c^T \log \left(\frac{Q_c^T}{Q_c^S} \right). \quad (2)$$

The comprehensive loss function integrates the classification loss and the distillation losses.

$$\mathcal{L}_{\text{Distill}} = \mathcal{L}_{\text{CE}} + \mathcal{L}_{\text{Feat}} + \mathcal{L}_{\text{Logit}}, \quad (3)$$

where \mathcal{L}_{CE} is a cross-entropy loss for ophthalmic disease diagnosis task, allowing the same input to be associated with multiple classes.

Distribution-guided Disentangling Strategy

We introduce an IMDR strategy to effectively capture both modality-shared and modality-specific representations across input combinations from different modalities while eliminating inter-modality redundancy.

We consider building probabilistic embeddings, i.e. $e \sim \mathcal{P}(e|x)$, to achieve a more adaptable representation space for modality-specific representations. Specifically, we define probabilistic embeddings e following a spherical Gaussian distribution as common. Since computing the posterior distribution is intractable, we variationally approximate the posterior distribution $\mathcal{P}(e|x)$ as:

$$\mathcal{P}(e|x) \approx q_\theta(e|x) = \mathcal{N}(e; \mu, \Sigma), \quad (4)$$

where $q_\theta(e|x)$ is a variation approximation, and the mean μ and covariance matrix Σ of the Gaussian distribution are

predicted by distribution predictors f_μ and f_Σ . Unlike existing methods (Guo et al. 2017; Wei et al. 2022) that estimate μ and Σ for the logits after pooling, we directly estimate μ and Σ for the feature e . For each modality, the parameters μ and Σ are defined as:

$$\mu = f_\mu(\theta_\mu, e), \quad \log(\Sigma) = f_\Sigma(\theta_\Sigma, e), \quad (5)$$

where θ_μ and θ_Σ are the parameters for the distribution predictors $f_\mu(\cdot)$ and $f_\Sigma(\cdot)$, respectively. To enhance stability, we directly predict $\log(\Sigma)$ instead of Σ . Both $f_\mu(\cdot)$ and $f_\Sigma(\cdot)$ are implemented using a simple Multilayer Perceptron with batch normalization, introducing minimal additional parameters.

To extract modality-shared information, we leverage the joint distribution $\mathcal{P}(e|x^f, x^O)$ across multiple modalities. This joint posterior is derived using the product-of-experts (Hinton 2002), where individual distributions are combined multiplicatively to form a cohesive model. The joint distribution $\mathcal{P}(e|x^f, x^O)$ can be formulated to:

$$\mathcal{P}(e|x^f, x^O) \propto \mathcal{P}(e)\mathcal{P}(e|x^f)\mathcal{P}(e|x^O), \quad (6)$$

where $\mathcal{P}(e|x^f)$ and $\mathcal{P}(e|x^O)$ could be approximated by Gaussian distributions $\mathcal{N}(e; \mu^f, \Sigma^f)$ and $\mathcal{N}(e; \mu^O, \Sigma^O)$, as detailed in Eq. 4. Since we define that the probabilistic representation e follows a spherical Gaussian distribution, as proved in (Zhang et al. 2024b), we can assume the prior distribution $\mathcal{P}(e)$ is also a spherical Gaussian $\mathcal{N}(e; \mu_\Delta, \Sigma_\Delta)$. Hence it can be shown that the product of Gaussian distributions is also a Gaussian distribution:

$$\begin{aligned} \Sigma_\lambda &= (\Sigma_\Delta^{-1} + \sum_{m \in \{f, O\}} \Sigma_m^{-1})^{-1}, \\ \mu_\lambda &= (\mu_\Delta \Sigma_\Delta^{-1} + \sum_{m \in \{f, O\}} \mu_m \Sigma_m^{-1}) \Sigma_\lambda^{-1}, \end{aligned} \quad (7)$$

which yields the parameters $\mu_\lambda, \Sigma_\lambda$ for joint distribution formulated as a Gaussian distribution:

$$\mathcal{P}(e|x^f, x^O) = \mathcal{N}(e; \mu_\lambda, \Sigma_\lambda), \quad (8)$$

hence, the modality-shared feature s becomes a stochastic embedding sampled from the Gaussian distribution $\mathcal{N}(e; \mu_\lambda, \Sigma_\lambda)$. To enable differentiable sampling, we employ the reparameterization trick (Kingma and Welling 2013). In general, we sample noise from $\mathcal{N}(0, \mathbf{I})$ and construct the embedding s by reparameterizing, instead of directly sampling from $\mathcal{N}(e; \mu_\lambda, \Sigma_\lambda)$:

$$s = \mu_\lambda + \epsilon \Sigma_\lambda, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}). \quad (9)$$

We further disentangle multimodal data into ideally independent modal-shared representations \hat{s} and modal-specific representations $\mathcal{R}^f, \mathcal{R}^O$ through an attention layer. Then, we minimize the mutual information (MI) between modal-shared and modal-specific representations to preserve modality-specific information via a defined loss function:

$$\mathcal{L}_{\text{MI}} = I(\hat{s}, \tilde{\mathcal{R}}) + I(\mathcal{R}^f, \mathcal{R}^O), \quad (10)$$

where $I(\cdot, \cdot)$ represents the MI that measures the dependence between two variables, and $\tilde{\mathcal{R}}$ is the concatenation of \mathcal{R}^f and \mathcal{R}^O as complete modal-specific information. Since MI is generally intractable, we employ CLUB (Cheng et al. 2020) to implement \mathcal{L}_{MI} . Detailed computation procedures are provided in the Appendix C.

Joint Proxy Learning Module

We introduce a JPL module to mitigate redundancy in modality-shared features from interfering with the disentanglement process. For each modality, the module directly approximates the distribution $\mathcal{P}(e|x)$ with the distribution of a parametric proxy \hat{e} , with distribution $\mathcal{P}(\hat{e})$ represented by a set of N_p proxies for each class, denoted as $\mathcal{P} = \{\mathcal{N}(\hat{e}; \mu_c, \Sigma_c)\}_{c=1}^C$. Each proxy is designed to capture discriminative information for the ground truth label y , approximating the conditional probability distribution $\mathcal{P}(\hat{e}|y) = \mathcal{N}(\hat{e}; \mu_y, \Sigma_y)$. The feature representations e are then expected to align with the proxies \hat{e} with the same label y . Consequently, the objective for the variation approximation in Eq. 4 is adjusted as:

$$\mathcal{P}(e|x) = \mathcal{P}(e|x, y) \approx \mathcal{P}(\hat{e}|y). \quad (11)$$

To achieve this, we maximize the similarity between $\mathcal{P}(\hat{e})$ and distributions of features e . As a result, we just need to optimize the parametric proxies \hat{e} . The objective of approximating $\mathcal{P}(e|x, y)$ with proxies $\hat{e}|y$ in Eq. 11 can be achieved by a proxy loss. For each modality m , the loss pulls the feature e^m closer to positive proxies $\hat{e}^{m,+}$ while pushing them away from negative ones $\hat{e}^{m,-}$, formulated as:

$$\mathcal{L}_{\text{Prox}} = \frac{1}{M} \sum_{m=1}^M \left(-\text{Sim}(\hat{e}^{m,+}, e^m) + \frac{1}{C-1} \sum_{n=1}^{C-1} \text{Sim}(\hat{e}_n^{m,-}, e_n^m) \right). \quad (12)$$

During training, positive proxies \hat{e}^+ are selected based on their matching labels with the input data. During inference, this is determined by calculating a similarity matrix and selecting the set with the highest mean similarity.

Now that we can obtain $\hat{\mu}$ and $\hat{\Sigma}$ using positive proxies as inputs to the distribution predictor. We reformulated the distribution prediction as:

$$\hat{\mu} = f_{\mu}(\theta_{\mu}, \hat{e}^+), \quad \log(\hat{\Sigma}) = f_{\Sigma}(\theta_{\Sigma}, \hat{e}^+). \quad (13)$$

Hence, with $\mathcal{P}(\hat{e})$ being a spherical Gaussian $\mathcal{N}(\hat{e}; \mu_{\Delta}, \Sigma_{\Delta})$ the μ_{λ} and Σ_{λ} for joint distribution defined in Eq. 7 can be reformulated to:

$$\begin{aligned} \hat{\Sigma}_{\lambda} &= (\Sigma_{\Delta}^{-1} + \sum_{m \in \{f, O\}} \hat{\Sigma}_m^{-1})^{-1}, \\ \hat{\mu}_{\lambda} &= (\mu_{\Delta} \Sigma_{\Delta}^{-1} + \sum_{m \in \{f, O\}} \hat{\mu}_m \hat{\Sigma}_m^{-1}) \hat{\Sigma}_{\lambda}^{-1}, \end{aligned} \quad (14)$$

which leads to the joint distribution $\mathcal{P}(\hat{e}|x^f, x^O)$:

$$\mathcal{P}(\hat{e}|x^f, x^O) = \mathcal{N}(\hat{e}; \hat{\mu}_{\lambda}, \hat{\Sigma}_{\lambda}). \quad (15)$$

The overall learning objective for the teacher model is:

$$\mathcal{L}_{\text{Teacher}} = \mathcal{L}_{\text{CE}} + \omega_1 \mathcal{L}_{\text{MI}} + \omega_2 \mathcal{L}_{\text{Prox}}. \quad (16)$$

where ω_1 and ω_2 are weights that control the contribution of the mutual information and proxy losses, respectively.

Experiments

Datasets and Evaluation Metrics

We evaluate the proposed framework using four publicly available multimodal datasets: **GAMMA** dataset (Wu et al.

2023) and three subsets from Harvard-30k (Luo et al. 2024), including **Harvard-30k AMD**, **Harvard-30k DR**, and **Harvard-30k Glaucoma**, covering Age-related Macular Degeneration (AMD), Diabetic Retinopathy (DR), and Glaucoma. GAMMA contains cases with three-tier grading, with OCT and fundus images sized at $256 \times 512 \times 992$ and 1956×1934 , respectively. Harvard-30k subsets are annotated with four-tier AMD grading and two-tier for glaucoma and DR, with image dimensions of 448×448 for fundus and $200 \times 256 \times 256$ for OCT (where 200 denotes the number of OCT slices). Detailed descriptions are in Appendix E.

To ensure reliable results, each dataset underwent five-fold cross-validation, and the model is assessed using four key metrics: Accuracy (ACC), F1 score (F1), Area Under the Curve (AUC), and Specificity (Spec), effectively predicting the severity and type of various ophthalmic diseases.

Comparison with State-of-the-art Methods

Effectiveness in Complete-modality Fusion. We compare IMDR with six representative complete-modality fusion methods using CNN and Transformer architectures, as shown in Table 1. ResNet50 2D and 3D serve as CNN backbones, while Swin-Transformer and UNETR are employed as Transformer backbones. **(1) B-IF**, a baseline utilizing an intermediate multimodal fusion method; **(2) HFS-IL**, which employs a hybrid fusion strategy that combines intermediate and late fusion; **(3) B-EF** (Hua et al. 2020), an early fusion strategy; **(4) CR-AF** (Zheng et al. 2023), incorporating a multimodal cross-attention fusion approach; **(5) M²LC** (Woo et al. 2018), integrating both channel attention and spatial attention for fusion; **(6) Eye-Most** (Zou et al. 2024), an evidence fusion model based on the inverse gamma prior distribution. The Complete-Modality Fusion section of Table 1 shows that the proposed IMDR model consistently outperforms other methods across all datasets and backbones evaluated. For example, on the Harvard-30k AMD dataset, our IMDR model outperforms the state-of-the-art Eye-Most model by over +3.58% in accuracy and exceeds the baseline model (B-IF) by +6.33% in specificity when compared to the Transformer baseline architecture.

Robustness to Inter-modality Incompleteness. In the Inter-Modality Missing section of Table 1, we evaluate our model by comparing the performance of the single backbone and other models. All models show a performance decline when a modality is missing, with the most pronounced drop occurring in the absence of fundus images. This highlights the critical role of complementary information from heterogeneous modalities. However, IMDR demonstrates greater robustness, retaining strong performance when the OCT modality is missing and surpassing other models when the crucial fundus modality is absent. This robustness is due to IMDR’s capability to disentangle multimodal features, distill informative knowledge to the student network, and reconstruct missing semantics for robust representations.

As shown in Fig. 3, we compare GradCAM (Selvaraju et al. 2017) heatmaps of IMDR with other models under missing OCT modality on AMD and Glaucoma datasets. Unlike baseline and other models, which miss critical re-

Method	Modality		GAMMA			Harvard-30k AMD			Harvard-30k DR			Harvard-30k Glaucoma		
	OCT	Fundus	ACC	AUC	F1	ACC	AUC	F1	ACC	AUC	F1	ACC	AUC	F1
Inter-Modality Missing with ResNet-50 Backbone														
2D-Resnet50 Backbone		✓	0.7050	0.7896	0.6409	0.7312	0.7535	0.7223	0.7381	0.7915	0.7047	0.7312	0.7535	0.7223
3D-Resnet50 Backbone	✓		0.6860	0.7453	0.6210	0.6517	0.6998	0.6963	0.7073	0.6994	0.6197	0.6572	0.6989	0.7098
B-IF + distill		✓	0.7153	0.8005	0.7233	0.7235	0.7198	0.7003	0.7362	0.6750	0.6968	0.7339	0.7661	0.7247
B-IF + distill	✓		0.6891	0.7859	0.6844	0.6957	0.7014	0.6745	0.6905	0.6525	0.6793	0.6964	0.6895	0.6718
M ² LC + distill		✓	0.7607	0.7561	0.7548	0.7324	0.7267	0.7380	0.7304	0.6789	0.7459	0.7278	0.7023	0.7111
M ² LC + distill	✓		0.7205	0.7566	0.7275	0.6897	0.7223	0.6506	0.6720	0.6505	0.6433	0.6770	0.7122	0.6560
IMDR (Ours)	✓	✓	0.7900	0.7655	0.7433	0.7517	0.8048	0.7659	0.7619	0.7907	0.7218	0.7544	0.7847	0.7512
IMDR (Ours)	✓		0.7483	0.7912	0.8097	0.7062	0.7269	0.7190	0.7262	0.7469	0.7290	0.7116	0.7507	0.7037
Complete-Modality Fusion with ResNet-50 Backbone														
B-IF	✓	✓	0.7099	0.8610	0.6691	0.7167	0.8122	0.6901	0.7355	0.7544	0.7396	0.7366	0.7937	0.7289
B-EF	✓	✓	0.6837	0.8024	0.6433	0.6933	0.7912	0.6064	0.7498	0.7413	0.7717	0.7419	0.7752	0.7222
HFS-IL	✓	✓	0.7431	0.8347	0.6800	0.7500	0.8174	0.7104	0.7557	0.7531	0.7647	0.7334	0.8096	0.7281
CR-AF	✓	✓	0.6703	0.8123	0.6777	0.7300	0.8137	0.6773	0.7657	0.7835	0.7627	0.7392	0.7748	0.7309
M ² LC	✓	✓	0.7200	0.8609	0.6804	0.7400	0.8456	0.7073	0.7736	0.8362	0.7601	0.7473	0.7745	0.7354
Eye-Most	✓	✓	0.8600	0.8493	0.8135	0.7555	0.8230	0.7100	0.7500	0.8034	0.7342	0.7433	0.7540	0.7201
IMDR (Ours)	✓	✓	0.8700	0.8320	0.8366	0.7645	0.8310	0.7357	0.7803	0.8534	0.7823	0.7650	0.7711	0.7437
Complete-Modality Fusion with ResNet-101 Backbone														
B-IF	✓	✓	0.7100	0.8856	0.7076	0.7317	0.8382	0.7125	0.7636	0.7795	0.7561	0.7339	0.7332	0.7211
B-EF	✓	✓	0.7326	0.8106	0.7005	0.7217	0.8181	0.7196	0.7083	0.6119	0.6884	0.7415	0.7878	0.7681
CR-AF	✓	✓	0.7425	0.8077	0.7562	0.7436	0.8274	0.7348	0.7489	0.7592	0.7581	0.7324	0.7593	0.7420
M ² LC	✓	✓	0.7860	0.7800	0.6823	0.7493	0.8239	0.7120	0.7521	0.7968	0.7439	0.7498	0.7645	0.7423
Eye-Most	✓	✓	0.8200	0.8321	0.7488	0.7592	0.8334	0.7002	0.7688	0.8259	0.7588	0.7210	0.7200	0.7342
IMDR (Ours)	✓	✓	0.8775	0.8495	0.7958	0.7950	0.8509	0.7252	0.7857	0.8500	0.7704	0.7731	0.7898	0.7890

Table 1: Quantitative Results for Missing-Modality Tasks on GAMMA, Harvard-30k AMD, DR, and Glaucoma Datasets.




Figure 3: Comparative visualization of attention maps under the inter-modality incompleteness setting: The first row is AMD dataset and the second row is Glaucoma dataset.

gions, IMDR accurately captures essential features, such as optic nerve head information for Glaucoma and hemorrhage features for AMD, due to its effectiveness in capturing discriminative information relevant to the target.

Robustness to Intra-modality Incompleteness. We estimate the robustness of intra-modality incompleteness by comparing performance across four datasets under different data missing rates, as shown in Fig. 4. To simulate intra-modality data loss, we introduce Gaussian noise at varying levels to represent different degrees of information degradation. As noise increases, all models show significant performance declines, highlighting the impact of single-modality data loss on the stability of multimodal representations. However, IMDR exhibits superior robustness, particularly under high information loss, maintaining stable performance compared to other models. This underscores the effectiveness of our proxy learning in capturing class relations, enabling student network to reconstruct missing information.

Exp	C-Feat	C-Dist	JPL	DE	ACC	AUC	F1
I	✓				0.633	0.625	0.677
II		✓			0.654	0.643	0.623
III		✓		✓	0.681	0.755	0.712
IV		✓	✓		0.725	0.672	0.680
V		✓	✓	✓	0.752	0.805	0.766

Table 2: Ablation study with missing OCT modality.

Ablation Studies

Effectiveness of each component. To validate the effectiveness of each proposed component, we conduct five ablation experiments on the AMD dataset with the OCT modality missing, as shown in Table 2. In Experiments I and II, C-Dist outperforms C-Feat, indicating that distribution-based methods effectively capture and integrate cross-modal representations. Experiment III shows that introducing our DE layer into the joint distribution, which preserves independent modality-shared and modality-shared features, increases accuracy by +2.7%. Experiment IV, which incorporates our JPL module to eliminate redundancy, resulted in a +7.1% performance boost. Experiment V combines the DE layer and the JPL module, further improving performance by a significant +9.8% over Experiment II. This confirms that integrating our proposed components substantially enhances the robustness of the student model in the Ophthalmic Disease Diagnosis, demonstrating that the JPL and DE modules can work harmoniously together, with JPL capturing discriminative features and DE disentangling features through the joint distribution.

To illustrate the effects of different modules, we perform a t-SNE visualization on 600 randomly selected samples from the AMD test set under missing OCT modality conditions. As shown in Fig. 5 (a), the Baseline model struggles to extract distinct features from incomplete data. However, with the integration of JPL or DE modules, as shown in Fig. 5

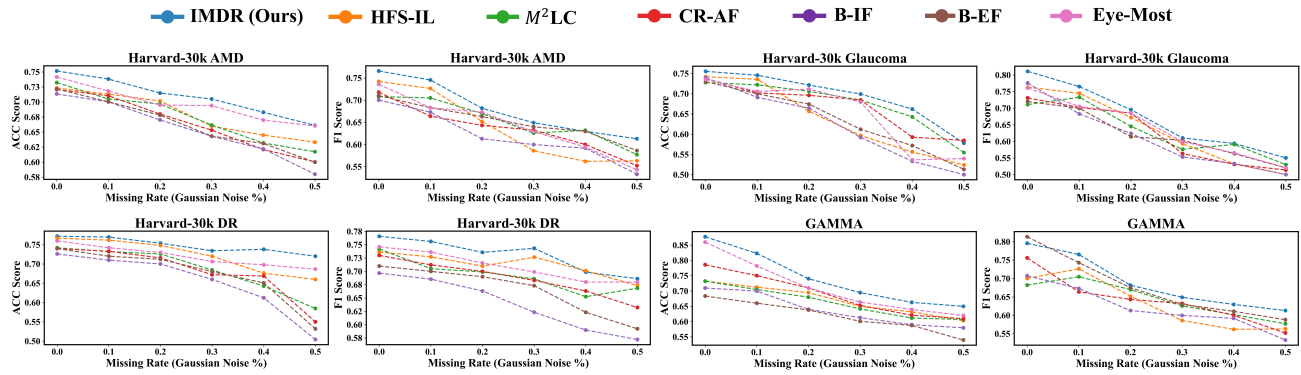


Figure 4: The comparison of performance across various missing rates under intra-modality incompleteness.

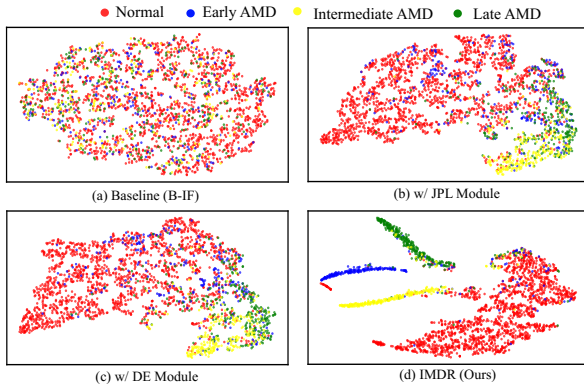


Figure 5: Ablation Study of visualization under condition of missing OCT modality on Harvard-30k AMD test set.

(b) and (c), features for Intermediate and Late AMD become more distinguishable, though some overlap remains. The model incorporating both JPL and DE modules achieves clear feature clustering within the same class and distinct separation between different classes, demonstrating the effectiveness of our disentangled representation strategy.

Comparison with other distribution estimations. To validate the effectiveness of Joint Proxy Learning module in reducing intra-modality redundancy and enhancing the joint distribution, we assess the distribution estimation capabilities of various methods (Table 3), as shown in Table 3. Compared to PE (Shi and Jain 2019) and PCME (Chun et al. 2021), IMDR demonstrates a clear performance advantage, confirming the ability of the Joint Proxy Learning module to capture informative features, guide feature disentangling, and effectively handle severe modality incompleteness.

Sensitivity Analysis of Hyperparameters. We conduct a sensitivity analysis of key parameters within IMDR, as shown in Fig. 6. The results indicate an initial increase of r enhances model performance, then further increases lead to a decline due to noise from oversampling, which diminishes the overall effectiveness of our model. Furthermore, as the τ coefficients increase, model performance initially improves but ultimately declines under higher parameter settings.

Method	ACC	AUC	F1
C-Dist	0.706	0.743	0.687
PE	0.717	0.770	0.706
PCME	0.735	0.789	0.713
IMDR (Ours)	0.752	0.804	0.766

Table 3: Comparison with other distribution methods. C-Dist: directly estimates the distribution of features from each modality. PE: uses a fully connected layer to estimate the feature vector’s distribution. PCME: incorporates attention modules to aggregate information from the feature map.




Figure 6: Ablation study of hyperparameters under the condition of missing OCT modality on Harvard-30k AMD test set. r : the ratio r of number of proxy N_p to the number of samples in the training set N . τ : distillation temperature.

Conclusions

This paper identifies two critical limitations of existing methods: implicit representation constraints that limit capturing modality-specific information and modality heterogeneity, leading to feature distribution gaps and redundancy. To overcome it, Incomplete Modality Disentangled Representation strategy disentangles features into distinct modal-shared and modal-specific components guided by mutual information. This enables student network to reconstruct missing semantics and produce robust multimodal representations. Additionally, we introduce a Joint Proxy Learning module to eliminate intra-modality redundancy by leveraging class-specific proxies. Experiments on four ophthalmology multimodal datasets demonstrate that IMDR significantly outperforms state-of-the-art methods.

References

- Chen, Q.; Chen, X.; Song, H.; Xiong, Z.; Yuille, A.; Wei, C.; and Zhou, Z. 2024. Towards generalizable tumor synthesis. In *CVPR*.
- Cheng, P.; Hao, W.; Dai, S.; Liu, J.; Gan, Z.; and Carin, L. 2020. Club: A contrastive log-ratio upper bound of mutual information. In *ICML*.
- Chun, S.; Oh, S. J.; De Rezende, R. S.; Kalantidis, Y.; and Larlus, D. 2021. Probabilistic embeddings for cross-modal retrieval. In *CVPR*.
- Cong, C.; Xuan, S.; Liu, S.; Zhang, S.; Pagnucco, M.; and Song, Y. 2024. Decoupled optimisation for long-tailed visual recognition. In *AAAI*.
- Duan, M.; Qu, L.; Yang, Z.; Wang, M.; Zhang, C.; and Song, Z. 2024. Towards Arbitrary-Scale Histopathology Image Super-resolution: An Efficient Dual-branch Framework via Implicit Self-texture Enhancement. *arXiv preprint arXiv:2401.15613*.
- Ebrahimi, S.; Arik, S. O.; Dong, Y.; and Pfister, T. 2023. Lanistr: Multimodal learning from structured and unstructured data. *arXiv preprint arXiv:2305.16556*.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On calibration of modern neural networks. In *ICML*.
- Hinton, G. E. 2002. Training products of experts by minimizing contrastive divergence. *Neural computation*.
- Hu, M.; Xia, P.; Wang, L.; Yan, S.; Tang, F.; Xu, Z.; Luo, Y.; Song, K.; Leitner, J.; Cheng, X.; et al. 2025. Ophnet: A large-scale video benchmark for ophthalmic surgical workflow understanding. In *ECCV*.
- Hu, M.; Yuan, K.; Shen, Y.; Tang, F.; Xu, X.; Zhou, L.; Li, W.; Chen, Y.; Xu, Z.; Peng, Z.; et al. 2024. OphCLIP: Hierarchical Retrieval-Augmented Learning for Ophthalmic Surgical Video-Language Pretraining. *arXiv preprint arXiv:2411.15421*.
- Hua, C.-H.; Kim, K.; Huynh-The, T.; You, J. I.; Yu, S.-Y.; Le-Tien, T.; Bae, S.-H.; and Lee, S. 2020. Convolutional network with twofold feature augmentation for diabetic retinopathy recognition from multi-modal images. *JBHI*.
- Jue, J.; Jason, H.; Neelam, T.; Andreas, R.; Sean, B. L.; Joseph, D. O.; and Harini, V. 2019. Integrating cross-modality hallucinated MRI with CT to aid mediastinal lung tumor segmentation. In *MICCAI*.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Li, M.; Yang, D.; Zhao, X.; Wang, S.; Wang, Y.; Yang, K.; Sun, M.; Kou, D.; Qian, Z.; and Zhang, L. 2024a. Correlation-Decoupled Knowledge Distillation for Multimodal Sentiment Analysis with Incomplete Modalities. In *CVPR*.
- Li, W.; Xiong, X.; Xia, P.; Ju, L.; and Ge, Z. 2024b. TP-DRSeg: improving diabetic retinopathy lesion segmentation with explicit text-prompts assisted SAM. In *MICCAI*.
- Li, X.; Li, C.; and Hu, J. 2023. Missing-Modal Face Anti-Spoofing: A Benchmark. In *2023 8th International Conference on Image, Vision and Computing*. IEEE.
- Li, Y.; Wang, Y.; and Cui, Z. 2023. Decoupled multimodal distilling for emotion recognition. In *CVPR*.
- Liang, V. W.; Zhang, Y.; Kwon, Y.; Yeung, S.; and Zou, J. Y. 2022. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. *NeurIPS*.
- Liu, A.; Tan, Z.; Wan, J.; Liang, Y.; Lei, Z.; Guo, G.; and Li, S. Z. 2021. Face anti-spoofing via adversarial cross-modality translation. *IEEE Transactions on Information Forensics and Security*.
- Liu, H.; Wei, D.; Lu, D.; Sun, J.; Wang, L.; and Zheng, Y. 2023a. M3AE: multimodal representation learning for brain tumor segmentation with missing modalities. In *AAAI*.
- Liu, Z.; Wei, J.; Li, R.; and Zhou, J. 2023b. SFusion: Self-attention based n-to-one multimodal fusion block. In *MIC-CAI*.
- Luo, Y.; Tian, Y.; Shi, M.; Elze, T.; and Wang, M. 2024. Eye Fairness: A Large-Scale 3D Imaging Dataset for Equitable Eye Diseases Screening and Fair Identity Scaling.
- Meng, Y.; Yang, Z.; Duan, M.; Shi, Y.; and Song, Z. 2024. Continuous K-space Recovery Network with Image Guidance for Fast MRI Reconstruction. *arXiv preprint arXiv:2411.11282*.
- Miao, Y.; Zhu, G.; Liu, T.; Zhang, W.; Wen, Y.; and Zhou, M. 2023. Multimodal Sentiment Analysis based on Supervised Contrastive Learning and Cross-modal Translation under Modalities Missing. In *2023 IEEE 14th International Symposium on Parallel Architectures, Algorithms and Programming*. IEEE.
- Peng, Z.; Wang, G.; Xie, L.; Jiang, D.; Shen, W.; and Tian, Q. 2023. Usage: A unified seed area generation paradigm for weakly supervised semantic segmentation. In *ICCV*.
- Peng, Z.; Xu, Z.; Zeng, Z.; Xie, L.; Tian, Q.; and Shen, W. 2024a. Parameter efficient fine-tuning via cross block orchestration for segment anything model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3743–3752.
- Peng, Z.; Xu, Z.; Zeng, Z.; Yang, X.; and Shen, W. 2024b. Sam-parser: Fine-tuning sam efficiently by parameter space reconstruction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4515–4523.
- Poudel, P.; Shrestha, P.; Amgain, S.; Shrestha, Y. R.; Gyawali, P.; and Bhattarai, B. 2024. CAR-MFL: Cross-Modal Augmentation by Retrieval for Multimodal Federated Learning with Missing Modalities. *arXiv preprint arXiv:2407.08648*.
- Qu, L.; Ma, Y.; Yang, Z.; Wang, M.; and Song, Z. 2023. Ope-nal: An efficient deep active learning framework for open-set pathology image classification. In *MICCAI*, 3–13. Springer.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*.
- Shi, J.; Shang, C.; Sun, Z.; Yu, L.; Yang, X.; and Yan, Z. 2024. PASSION: Towards Effective Incomplete Multi-Modal Medical Image Segmentation with Imbalanced Missing Rates. *arXiv preprint arXiv:2407.14796*.

- Shi, Y.; and Jain, A. K. 2019. Probabilistic face embeddings. In *ICCV*.
- Sun, J.; Zhang, X.; Han, S.; Ruan, Y.-P.; and Li, T. 2024. RedCore: Relative Advantage Aware Cross-Modal Representation Learning for Missing Modalities with Imbalanced Missing Rates. In *AAAI*.
- Tang, F.; Trinh, M.; Duong, A.; Ly, A.; Stapleton, F.; Chen, Z.; Ge, Z.; and Razzak, I. 2024a. Discriminating retinal microvascular and neuronal differences related to migraines: Deep Learning based Cross-sectional Study. *arXiv preprint arXiv:2408.07293*.
- Tang, F.; Xu, Z.; Huang, Q.; Wang, J.; Hou, X.; Su, J.; and Liu, J. 2023. DuAT: Dual-aggregation transformer network for medical image segmentation. In *PRCV*.
- Tang, F.; Xu, Z.; Qu, Z.; Feng, W.; Jiang, X.; and Ge, Z. 2024b. Hunting Attributes: Context Prototype-Aware Learning for Weakly Supervised Semantic Segmentation. In *CVPR*.
- Trinh, M.; Tang, F.; Ly, A.; Duong, A.; Stapleton, F.; Ge, Z.; and Razzak, I. 2024. Sight for sore heads—using cnns to diagnose migraines. *Investigative Ophthalmology & Visual Science*.
- Tripathi, A.; Waqas, A.; Yilmaz, Y.; and Rasool, G. 2024. Multimodal transformer model improves survival prediction in lung cancer compared to unimodal approaches. *Cancer Research*.
- Udandarao, V.; Gupta, A.; and Albanie, S. 2023. Sus-x: Training-free name-only transfer of vision-language models. In *ICCV*.
- Wang, H.; Chen, Y.; Ma, C.; Avery, J.; Hull, L.; and Carneiro, G. 2023a. Multi-modal learning with missing modality via shared-specific feature modelling. In *CVPR*.
- Wang, L.; Dai, W.; Jin, M.; Ou, C.; and Li, X. 2023b. Fundus-enhanced disease-aware distillation model for retinal disease classification from OCT images. In *MICCAI*.
- Wang, S.; Yan, Z.; Zhang, D.; Wei, H.; Li, Z.; and Li, R. 2023c. Prototype knowledge distillation for medical segmentation with missing modality. In *ICASSP. IEEE*.
- Wang, S.; Zhu, Y.; Luo, X.; Yang, Z.; Zhang, Y.; Fu, P.; Wang, M.; Song, Z.; Li, Q.; Zhou, P.; et al. 2023d. Knowledge Extraction and Distillation from Large-Scale Image-Text Colonoscopy Records Leveraging Large Language and Vision Models. *arXiv preprint arXiv:2310.11173*.
- Wang, Y.; Cui, Z.; and Li, Y. 2023. Distribution-consistent modal recovering for incomplete multimodal learning. In *ICCV*.
- Warner, E.; Lee, J.; Hsu, W.; Syeda-Mahmood, T.; Kahn Jr, C. E.; Gevaert, O.; and Rao, A. 2024. Multimodal Machine Learning in Image-Based and Clinical Biomedicine: Survey and Prospects. *IJCV*.
- Watanabe, T.; Hiratsuka, Y.; Kita, Y.; Tamura, H.; Kawasaki, R.; Yokoyama, T.; Kawashima, M.; Nakano, T.; and Yamada, M. 2022. Combining optical coherence tomography and fundus photography to improve glaucoma screening. *Diagnostics*, 12(5): 1100.
- Wei, H.; Xie, R.; Cheng, H.; Feng, L.; An, B.; and Li, Y. 2022. Mitigating neural network overconfidence with logit normalization. In *ICML*.
- Woo, S.; Park, J.; Lee, J.-Y.; and Kweon, I. S. 2018. Cbam: Convolutional block attention module. In *ECCV*.
- Wu, J.; Fang, H.; Li, F.; Fu, H.; Lin, F.; Li, J.; Huang, Y.; Yu, Q.; Song, S.; Xu, X.; Xu, Y.; Wang, W.; Wang, L.; Lu, S.; Li, H.; Huang, S.; Lu, Z.; Ou, C.; Wei, X.; Liu, B.; Kobbi, R.; Tang, X.; Lin, L.; Zhou, Q.; Hu, Q.; Bogunović, H.; Orlando, J. I.; Zhang, X.; and Xu, Y. 2023. GAMMA challenge: Glaucoma grading from Multi-Modality images. *MIA*.
- Xiong, B.; Yang, X.; Song, Y.; Wang, Y.; and Xu, C. 2023. Client-Adaptive Cross-Model Reconstruction Network for Modality-Incomplete Multimodal Federated Learning. In *ACMMM*.
- Xiong, X.; Wu, Z.; Tan, S.; Li, W.; Tang, F.; Chen, Y.; Li, S.; Ma, J.; and Li, G. 2024. Sam2-UNET: Segment anything 2 makes strong encoder for natural and medical image segmentation. *arXiv preprint arXiv:2408.08870*.
- Xu, Z.; Tang, F.; Chen, Z.; Zhou, Z.; Wu, W.; Yang, Y.; Liang, Y.; Jiang, J.; Cai, X.; and Su, J. 2024. PolypMamba: Polyp Segmentation with Visual Mamba. In *MICCAI*. Springer.
- Yang, Y.; Chen, H.; Liu, Z.; Lyu, Y.; Zhang, B.; Wu, S.; Wang, Z.; and Ren, K. 2023. Action recognition with multi-stream motion modeling and mutual information maximization. *arXiv preprint arXiv:2306.07576*.
- Yang, Z.; Meng, Y.; Fu, K.; Wang, S.; and Song, Z. 2024. Tackling Ambiguity from Perspective of Uncertainty Inference and Affinity Diversification for Weakly Supervised Semantic Segmentation. *arXiv preprint arXiv:2404.08195*.
- Zhang, J.; Wu, S.; Gao, L.; Shen, H. T.; and Song, J. 2024a. Dept: Decoupled prompt tuning. In *CVPR*.
- Zhang, Y.; Xu, Y.; Chen, J.; Xie, F.; and Chen, H. 2024b. Prototypical information bottlenecks and disentangling for multimodal cancer survival prediction. *arXiv preprint arXiv:2401.01646*.
- Zhao, X.; Tang, F.; Wang, X.; and Xiao, J. 2024. Sfc: Shared feature calibration in weakly supervised semantic segmentation. In *AAAI*.
- Zheng, J.; Liu, H.; Feng, Y.; Xu, J.; and Zhao, L. 2023. CASF-Net: Cross-attention and cross-scale fusion network for medical image segmentation. *Computer Methods and Programs in Biomedicine*, 229: 107307.
- Zou, K.; Lin, T.; Han, Z.; Wang, M.; Yuan, X.; Chen, H.; Zhang, C.; Shen, X.; and Fu, H. 2024. Confidence-aware multi-modality learning for eye disease screening. *Medical Image Analysis*, 96: 103214.
- Zou, K.; Lin, T.; Yuan, X.; Chen, H.; Shen, X.; Wang, M.; and Fu, H. 2023. Reliable multimodality eye disease screening via mixture of student's distributions. In *MICCAI*.