

Multi-Modal Multi-Instance Learning for Retinal Disease Recognition

Xirong Li
MoE Key Lab of DEKE,
Renmin University of China
China
xirong@ruc.edu.cn

Yang Zhou
Vistel AI Lab,
Beijing Visionary Intelligence Ltd.
China
yang.zhou@vistel.cn


Jie Wang
Hailan Lin
AIMC Lab, School of Information,
Renmin University of China
China

Jianchun Zhao
Dayong Ding
Vistel AI Lab,
Beijing Visionary Intelligence Ltd.
China

Weihong Yu
Youxin Chen
Dept. of Ophthalmology, Peking
Union Medical College Hospital
China



(a) A multi-modal sample acquired during an eye examination, consisting of a color fundus photo (CFP) and an array of 12 OCT B-scan images.



(b) Activation maps produced by our proposed model, showing abnormal regions in the CFP and abnormal B-scans in the OCT array.

Figure 1: Illustrating multi-modal retinal disease recognition, with (a) input and (b) output. Diseases: ERM, ME.

ABSTRACT

This paper attacks an emerging challenge of multi-modal retinal disease recognition. Given a multi-modal case consisting of a color

fundus photo (CFP) and an array of OCT B-scan images acquired during an eye examination, we aim to build a deep neural network that recognizes multiple vision-threatening diseases for the given case. As the diagnostic efficacy of CFP and OCT is disease-dependent, the network’s ability of being both selective and interpretable is important. Moreover, as both data acquisition and manual labeling are extremely expensive in the medical domain, the network has to be relatively lightweight for learning from a limited set of labeled multi-modal samples. Prior art on retinal disease recognition focuses either on a single disease or on a single modality, leaving multi-modal fusion largely underexplored. We

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '21, October 20–24, 2021, Virtual Event, China

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8651-7/21/10...\$15.00

<https://doi.org/10.1145/3474085.3475418>

propose in this paper Multi-Modal Multi-Instance Learning (MM-MIL) for selectively fusing CFP and OCT modalities. Its lightweight architecture (as compared to current multi-head attention modules) makes it suited for learning from relatively small-sized datasets. For an effective use of MM-MIL, we propose to generate a pseudo sequence of CFPs by over sampling a given CFP. The benefits of this tactic include well balancing instances across modalities, increasing the resolution of the CFP input, and finding out regions of the CFP most relevant with respect to the final diagnosis. Extensive experiments on a real-world dataset consisting of 1,206 multi-modal cases from 1,193 eyes of 836 subjects demonstrate the viability of the proposed model.

CCS CONCEPTS

• **Computing methodologies** → **Visual content-based indexing and retrieval.**

KEYWORDS

Retinal disease recognition, multi-modal retinal imaging, multi-modal feature fusion, deep multi-instance learning

ACM Reference Format:

Xirong Li, Yang Zhou, Jie Wang, Hailan Lin, Jianchun Zhao, Dayong Ding, Weihong Yu, and Youxin Chen. 2021. Multi-Modal Multi-Instance Learning for Retinal Disease Recognition. In *Proceedings of the 29th ACM International Conference on Multimedia (MM '21), October 20–24, 2021, Virtual Event, China*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3474085.3475418>

1 INTRODUCTION

Multi-modal imaging is a routine measure in the screening and diagnosis of many diseases. Consider retinal disease diagnosis for instance. Retinal diseases manifest themselves in various ways in retina – a ten-layered structure covering over 200 degrees of eye fundus. Accurate diagnoses depend on accurate localization and visualization of the pathological alterations in this subtle and complex structure. Images of distinct modalities are acquired during an eye examination, including a color fundus photograph (CFP) and an array of optical coherence tomography (OCT) B-scan images, see Fig. 1. OCT B-scans provide cross-sectional information of the retina and choroid, while CFPs show en face information of the fundus. Due to such different imaging mechanisms, using a specific modality alone is often inadequate to spot alterations for all diseases. For dryAMD at its early stage, drusenoid pigment epithelial detachment can be observed in OCT images, even when its counterpart remains invisible in a CFP. Similarly for Macular Edema (ME), its main symptom is retinal thickening. The resultant alteration in CFP is hard exudate (not obvious in the early stage) and that in OCT B-scan is cystoid macular edema (noticeable as the layered structure shown in B-scan changes tangibly). Hence, OCT B-scans play a crucial role in recognizing ME, dryAMD and wetAMD in particular at their early stage. Meanwhile, as typical alterations by DR include microaneurysm, retinal hemorrhage, and venous beading, all are vascular lesions, and thus more observable in CFP than in OCT B-scans, making CFP more suitable for DR recognition. For EpiRetinal Membrane (ERM) and Pathological Myopia (PM), both modalities matter. In clinical practice, ophthalmologists use both modalities as a standard method for fundus checks. Hence,

multi-modal deep learning that effectively exploits a limited set of labeled multi-modal samples is crucial for artificial intelligence (AI) assisted disease diagnosis.

Previous efforts on retinal disease recognition are mostly based on a single modality, let it be a CFP [15], an OCT image [23] or a sequence of OCT images [7]. They are thus not directly applicable to deal with the multi-modal input. An initial attempt on end-to-end multi-modal learning is by Wang *et al.* [21], where authors present MM-CNN, a two-stream CNN that takes a CFP and a single OCT image as a paired input. While MM-CNN is shown to be superior to its single-modal counterparts for categorizing subclasses of age-related macular degeneration (AMD), the OCT image has to be manually selected by a technician that operates the OCT device. We argue that this requirement limits the practical use of MM-CNN, as the technician is unlikely to be sufficiently trained to select the most appropriate B-scan images for all diseases. Moreover, by using only the chosen image with the other B-scans ignored, the OCT modality is underexplored.

Multi-modal CNNs that accept a sequence of images as input have been extensively studied in the context of video action classification [3, 11, 25]. To exploit the spatio-temporal information in video frames, a 3D-CNN consisting of stacked 3D convolutional blocks is commonly used, either alone [11] or in a two-stream manner where another 3D-CNN runs in parallel to process the sequence of optical flow images [3, 25]. Note that the good performance of 3D-CNNs is subject to the availability of very large-scale training data [3]. This is however difficult to be fulfilled in the medical domain where both data acquisition and manual labeling are expensive. Moreover, as the sequence is processed as a whole, 3D-CNNs lack an explicit mechanism to interpret the contribution per frame, a wanted property for AI-assisted disease diagnosis. Therefore, multi-modal CNNs with a relatively lightweight architecture and interpretability is in demand for retinal disease recognition.

For automated screening of Retinopathy of Prematurity (ROP) given multiple CFPs acquired per case, Li *et al.* [16] adapt instance-attention based deep multiple instance learning (MIL) [17]. The given CFPs, treated as a bag of instances, are fed in parallel into a shared 2D-CNN for feature extraction. Such instance-level features are weighed in terms of the corresponding attended weights and summed up for case-level classification. The shared 2D backbone substantially reduces the amount of trainable parameters, whilst the attention mechanism naturally explains the contribution of each CFP to the final prediction. Again, as [16] works on a single modality, how to exploit MIL for the multi-modal scenario remains open. In fact, because the CFP is clearly outnumbered by the OCT images in a given input, extending MIL by treating the CFP and OCT images as individual instances in a bag is problematic.

Inspired by the success of deep MIL for the ROP classification task, we propose in this paper Multi-Modal Multi-Instance Learning (MM-MIL) for multi-label retinal disease recognition. The overall architecture, as illustrated in Fig. 2, follows the classical two-stream framework [3], with MM-MIL as a novel module for multi-modal feature fusion. To conquer the data imbalance issue and to improve the resolution of the CFP channel, over sampling on CFP is performed to form a pseudo sequence of CFPs. In sum, our contributions are as follows:

- We propose MM-MIL, a new module for multi-modal feature fusion. It inherits the interpretability of the instance-attention based MIL. Its lightweight architecture (as compared to current multi-head attention modules [20]) makes it suited for learning from relatively small-sized datasets.
- For an effective use of MM-MIL, we propose to generate a pseudo sequence of CFPs by over sampling a given CFP. The benefits of this tactic are multifold, including well balancing the samples between the CFP and OCT modalities, increasing the resolution of the CFP input, and finding out regions of the CFP most relevant with respect to the final prediction.
- Extensive experiments on a real-world dataset collected from an outpatient clinic justify the superiority of the proposed method against the state-of-the-art. Compared to the best single-modal baseline (OCT-MIL with an overall AP of 0.7748) and the best multi-modal baseline (MM-CNN++ with AP of 0.8172), our model scores the highest AP of 0.8539.

2 RELATED WORK

Since the seminal work by Google [10], which shows the initial success of deep CNNs for CFP-based diabetic retinopathy (DR) screening, many deep learning based methods have been proposed for eye disease recognition [2, 6, 8, 9, 13, 16, 24]. The majority of the methods target at a single disease, *e.g.*, DR [8], Pathological Myopia (PM) [6], Glaucoma [24] or AMD [14], and make predictions based on a single-modal input, either a CFP or an OCT image.

Although scarce, there are emerging efforts on recognizing multiple diseases from a CFP [15] or from a sequence of OCT B-scan images [7]. Li *et al.* [15] train an ensemble of SeResNext-50 to detect twelve major retinal diseases. To process OCT volume data, De Fauw *et al.* [7] first train a 3D-Unet to segment layers of B-scan images, and then feed the segmentation results into a 3D-DenseNet for multiple disease diagnosis and referral recommendation. Again, the above methods are single-modal.

Probably due to the lack of public multi-modal data, we see few work on multi-modal retinal disease recognition. In a pilot study [21], Wang *et al.* develop a two-stream CNN termed MM-CNN for multi-modal AMD categorization. Given a CFP and a manually chosen OCT B-scan image as a paired input, MM-CNN uses a ResNet-18 to extract features from the CFP and another ResNet-18 to extract features from the OCT image. The CFP and OCT feature vectors are simply concatenated and fed into a regular classification block (a linear layer followed by softmax) for categorizing the input as normal, dryAMD, or wetAMD. Despite its encouraging performance against the single-modal alternatives, we argue that MM-CNN has the following deficiencies. First, manual selection of one B-scan introduces extra workload on the technician. It also affects the model performance when there is discrepancy between the B-scan selection strategy used in the test stage and that used for training, see our experiments. Moreover, using merely one image per OCT scan makes the OCT modality largely underexplored. Lastly, fusion by feature concatenation lacks the ability of selectively exploiting the multi-modal information for different diseases.

Technically, we are inspired by the successful use of instance-attention based deep multiple instance learning (MIL) [17] for ROP

classification [16]. In order to capture different zones of the pre-mature retina, an ROP examination typically collects multiple CFPs. By viewing each CFP as an instance, Li *et al.* [16] formulate ROP classification as an MIL problem. As their study is single-modal, how to exploit MIL for the multi-modal scenario is untouched. Straightforward solutions such as combining MM-CNN with MIL, *e.g.*, substituting MIL for the OCT branch in MM-CNN, are inadequate for multi-modal retinal disease recognition.

The state-of-the-art on multi-modal representation, *e.g.*, cross-attention [22] and co-attention [5], is essentially multi-head self-attention (MHA) used in Transformer-like architectures [20]. Compared to MHA, our proposed MM-MIL performs better for the new task yet with much less parameters. Moreover, while the interpretability of Transformers is a research problem on its own [4], MM-MIL has an explicit mechanism for model interpretability.

3 PROPOSED METHOD

Based on multi-modal fundus imaging comprised of OCT and color fundus photography, we aim for automated categorization of the fundus condition of a specific eye. We use x to indicate an eye. Its examination by a multi-modal fundus camera produces a color fundus photograph (CFP), denoted as x_c , and an array of n OCT B-scan images, denoted as $\mathbf{x}_o = \{x_{o,1}, \dots, x_{o,n}\}$. Suppose there are m distinct categories to be considered. We aim to build a deep neural network G that accepts the multi-modal input (x_c, \mathbf{x}_o) and produces a probabilistic vector p accordingly,

$$p := G(x_c, \mathbf{x}_o), \quad (1)$$

where p_i denotes the probability of the eye belonging to class i , $i = 1, \dots, m$. In what follows, we describe the overall architecture of G in Section 3.1, followed by MM-MIL in Section 3.2 and model interpretation in Section 3.3.

3.1 Multi-Modal Deep Classification Network

The overall architecture is illustrated in Fig. 2. Our proposed network conceptually consists of three blocks, *i.e.*, 1) instance-level feature extraction, 2) multi-modal feature fusion, and 3) case-level classification.

1) Instance-level feature extraction. In order to match the number of the instances in the OCT modality, we perform over sampling on the original CFP to produce a sequence of n (sub-)images, denoted as $\{x_{c,1}, \dots, x_{c,n}\}$. Specifically, the multi-modal fundus camera used in our experiments ran in a radial scan mode, producing $n = 12$ B-scans per case. Therefore, in the test stage, we crop the four corners and the center of the original CFP with a fixed-sized window. With horizontal flip on the big image and the five sub-images, the over sampling operation eventually generates 12 images in total for the CFP modality.

Due to the noticeable difference in the visual appearance of CFP and OCT, we use two 2D-CNNs to extract instance-level features for the two modalities, respectively. The backbone of both 2D-CNNs is ResNet-50 [12], initialized using ImageNet-pretrained models. Per OCT instance $x_{o,i}$, we use $F_{o,i}$ to denote the feature maps of size $w \times h \times d$ extracted by the OCT 2D-CNN. In a similar vein, we define $F_{c,i}$ for a CFP instance $x_{c,i}$. By applying spatial-wise global average pooling (SW-GAP) on the feature maps, we obtain n d -dimensional

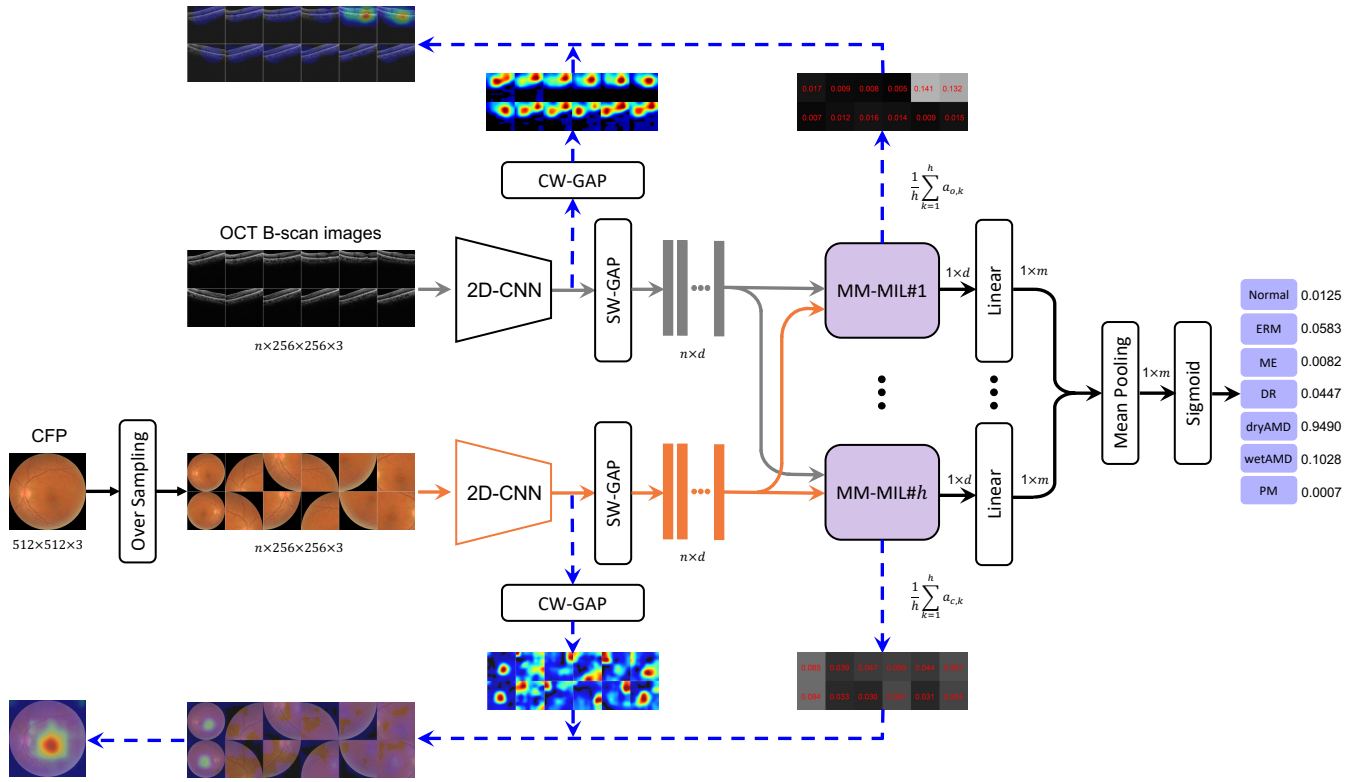


Figure 2: Proposed multi-modal retinal disease classification network in its inference mode. Given a multi-modal case with a CFP and a sequence of n OCT B-scan images, the network performs over sampling on the given CFP to generate a sequence of n CFPs. The two sequences are then fed in parallel into two 2D-CNNs to extract feature maps of size $w \times h \times d$ per instance. Per modality, spatial-wise global average pooling (SW-GAP) is applied upon the feature maps to obtain a d -dimensional feature vector for each instance. The proposed Multi-Modal Multiple Instance Learning (MM-MIL) module with h heads, see Fig. 3, aggregates the $2n$ instance-level features into h case-level features, each of which is converted into category-wise decision scores by a linear layer. Final probabilistic prediction is obtained by mean pooling plus sigmoid activation. For model interpretation, the feature maps per instance are converted to a $w \times h$ activation map by channel-wise global average pooling (CW-GAP). Multiplied by the corresponding instance-attention weight from MM-MIL, the activation maps are used to visualize the contribution of the OCT images and regions of the CFP to the prediction. Best viewed on screen.

instance-level features per modality, denoted as $\{f_{o,1}, \dots, f_{o,n}\}$ and $\{f_{c,1}, \dots, f_{c,n}\}$, respectively.

2) Multi-modal feature fusion. As shown in Fig. 2, the $2n$ instance-level features are simultaneously fed into our proposed MM-MIL module, which aggregates these features into a case-level feature \tilde{f}_{mm} of the same dimension d , expressed as

$$\tilde{f}_{mm} := \text{MM-MIL}(\{f_{c,1}, \dots, f_{c,n}, f_{o,1}, \dots, f_{o,n}\}). \quad (2)$$

3) Case-level classification. We use a linear layer to convert \tilde{f}_{mm} into categorize-wise decision scores. Since a specific eye might have multiple diseases, we formulate the retinal disease recognition task as a multi-label classification problem. As such, the sigmoid activation is adopted. Accordingly, the generic formula given in Eq. 1 can now be realized as

$$p := \text{sigmoid}(\text{Linear}_{d \times m}(\tilde{f}_{mm})). \quad (3)$$

When a multi-head MM-MIL is used to produce multiple \tilde{f}_{mm} , we let each \tilde{f}_{mm} go through a distinct linear layer and use mean pooling

in advance to the sigmoid layer, see Fig. 2. A standard BCE loss between p and case-level labels is computed for model training.

Next, we detail the procedure of computing the case-level modal-fused feature \tilde{f}_{mm} .

3.2 MM-MIL for Multi-Modal Feature Fusion

To make the paper self-contained, we outline the single-modality MIL block previously used for CFP-based ROP classification [16]. Given $\{f_i\}$ as n feature vectors extract from n CFP instances, instance-attention weights $\{a_i\}$ is obtained by stacking the features and feeding them into the following feedforward network,

$$\{a_i\} := \text{softmax}(\text{Linear}_{128 \times 1}(\tanh(\text{Linear}_{d \times 128}(\{f_i\}))))). \quad (4)$$

Accordingly, a case-level feature \tilde{f} is obtained as a weighted sum of the instance-level features, i.e.,

$$\tilde{f} := \sum_{i=1}^n a_i f_i. \quad (5)$$

In the multi-modal context, as the CFP features $\{f_{c,i}\}$ and OCT features $\{f_{o,i}\}$ are extracted by distinct 2D-CNNs, they are not comparable by definition. In order to make them additive, we introduce a cross-modal projection block, implemented as a modality-specific linear layer followed by Layer Normalization [1], see Fig. 3. Consequently, the CFP features are transformed into new features of the same size $\{\hat{f}_{c,i}\}$, *i.e.*,

$$\hat{f}_{c,i} := \text{LayerNorm}(\text{Linear}_{d \times d}(f_{c,i})). \quad (6)$$

In a similar manner we obtain the new OCT features $\{\hat{f}_{o,i}\}$.

By stacking $\{\hat{f}_{c,i}\}$ and $\{\hat{f}_{o,i}\}$ and substituting them for $\{f_i\}$ in Eq. 4, we obtain instance-attention weights for the $2n$ multi-modal instances, denoted by $\{a_{c,1}, \dots, a_{c,n}, a_{o,1}, \dots, a_{o,n}\}$. By definition, we have $\sum_{i=1}^n a_{c,i} + \sum_{i=1}^n a_{o,i} = 1$, where the first term indicates the importance of the CFP modality and the second term indicates the importance of the OCT modality. The case-level modal-fused feature is computed by putting the corresponding terms into Eq. 5:

$$\bar{f}_{mm} := \sum_{i=1}^n a_{c,i} f_{c,i} + \sum_{i=1}^n a_{o,i} f_{o,i}. \quad (7)$$

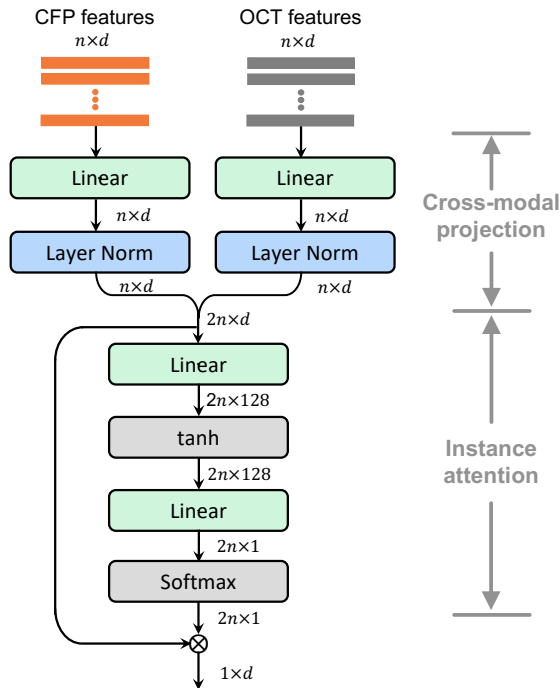


Figure 3: Proposed Multi-Modal Multiple Instance Learning (MM-MIL) module for multi-modal feature fusion. Instance-level CFP / OCT features are aggregated with instance-attention weights into a case-level feature vector. A multi-head version of MM-MIL can be obtained with ease by using multiple instance-attention blocks.

A conceptual diagram of MM-MIL (with a single head) is shown in Fig. 3. By using multiple instance-attention blocks, a multi-head

MM-MIL can be derived with ease. As Table 1 shows, when compared to multi-head self-attention (MHA), the proposed MM-MIL has much fewer parameters.

Table 1: The number of trainable parameters in distinct attention modules. The number of depth and heads in Multi-Head self-Attention (MHA) is denoted by d and h , respectively. MM-MIL $\times h$ stands for MM-MIL with h heads. Even compared with a mini-version of MHA with $d = 1$ and $h = 1$, the MM-MIL series has substantially fewer parameters.

Attention module	#Parameters
MHA($d = 1, h = 1$)	50.35 M
MHA($d = 4, h = 4$)	201.41 M
MM-MIL $\times 1$	8.67 M
MM-MIL $\times 2$	8.93 M
MM-MIL $\times 4$	9.45 M
MM-MIL $\times 8$	10.50 M

3.3 Model Interpretability

As aforementioned, the contribution of the individual instances is subject to their attention weights $\{a_{c,1}, \dots, a_{c,n}, a_{o,1}, \dots, a_{o,n}\}$ calculated by MM-MIL. Meanwhile, the feature maps of each instance reflect how it responds spatially. Therefore, combining the two factors enables the interpretability of our multi-modal classification network in multiple aspects. In particular, for the OCT modality, sorting the B-scans by their attention weights in descending order helps select the most abnormal B-scan for further examination. For the CFP modality, as each instance corresponds to a specific region of the original CFP, overlaying an attention-weight based heatmap on top of the CFP visualizes which region is the most relevant w.r.t. the final prediction. Moreover, comparing the accumulated attention weights per modality reveals which modality the model truly counts on for recognizing specific diseases.

For a specific instance, say $x_{c,i}$, we compute its attention-weighted activation map $A_{c,i}$ as

$$A_{c,i} := a_{c,i} \cdot \text{CW-GAP}(F_{c,i}), \quad (8)$$

where CW-GAP indicates channel-wise global average pooling that compresses $F_{c,i}$ into a $w \times h$ feature map. For an MM-MIL with h heads, $a_{c,i}$ is computed by $\frac{1}{h} \sum_{k=1}^h a_{c,k,i}$, where $a_{c,k,i}$ is the attention weight of $x_{c,i}$ produced by MM-MIL# k . By default, MM-MIL refers to a four-head MM-MIL, *i.e.*, MM-MIL $\times 4$, unless otherwise stated.

4 EVALUATION

4.1 Experimental Setup

Data collection. Due to the lack of public data for multi-modal retinal disease recognition, we built a new dataset as follows. Multi-modal cases were collected using a Topcon Maestro-1 (Topcon Corp., Japan) at the outpatient clinic, the Department of Ophthalmology in a state hospital from July 2020 to October 2020. The radial scan mode of the Topcon fundus camera was used, allowing us to acquire a CFP and 12 OCT B-scans simultaneously in a single examination, see Fig. 4. A panel of four retinal experts was formed and asked to categorize each case with respect to dozens of pre-specified

diseases. For quality control, each case was labeled by two experts independently. In case of disagreement, a third expert was asked to make the final decision. Note that not all diseases had a reasonable amount of cases for training and evaluation. Eventually, our dataset consists of 1,206 multi-modal cases from 1,193 eyes of 836 subjects, where each case has been labeled either as *normal* or categories from six vision-threatening diseases including *epiretinal membrane* (ERM), *macular edema* (ME), *diabetic retinopathy* (DR), *dry age-related macular degeneration* (AMD), *wet AMD* and *pathological myopia* (PM). The number of diseases assigned to an abnormal case ranges from 1 to 4, with an average value of 1.44. Some of the diseases such as DR can co-occur in both eyes of a specific subject. Hence, to avoid any risk of data leakage, we randomly split the data collection into three disjoint subsets, *i.e.*, training, validation and test, on the basis of subjects. A person (and his / her associated samples) is exclusively assigned to one of the three subsets. The ratio of the amount of subjects in training / validation / test is 6:2:2 approximately. Table 2 summarizes basic data statistics.

Table 2: Number of subjects, eyes and cases of each category for training, validation and test in our experiments.

Category	Training			Validation			Test		
	Subjects	Eyes	Cases	Subjects	Eyes	Cases	Subjects	Eyes	Cases
<i>Normal</i>	251	358	359	87	117	118	85	126	126
<i>ERM</i>	126	144	145	44	53	54	51	57	57
<i>ME</i>	95	104	105	30	35	35	36	41	41
<i>DR</i>	74	106	108	23	37	37	27	41	41
<i>dryAMD</i>	53	64	65	19	21	21	20	26	26
<i>wetAMD</i>	36	39	40	12	14	14	14	14	14
<i>PM</i>	27	35	38	10	15	15	11	14	14
Total	485	700	710	173	238	241	178	255	255




Figure 4: A multi-modal sample acquired by a Topcon Maestro-1 fundus camera in the radial scan mode. Expert labels: ERM and DR.

Performance metrics. In a retinal disease screening program, a subject with her eye(s) predicted with higher probability of having retinal diseases shall be given higher priority in a waiting list of referrals. Hence, we report Average Precision (AP) as commonly used to assess the quality of rankings. As secondary metrics, we report AUC, Sensitivity (Sen.), Specificity (Spe.) and F1-score of Sen. and Spe. For computing Sen. and Spe., we use the default threshold of 0.5 to convert probabilistic output to hard labels.

Details of implementation. We run all experiments with PyTorch [18]. For a fair comparison between all models evaluated in this study, we use the following common setup: ResNet-50 as the backbone, with its weights initialized using an ImageNet-pretrained model, SGD as the optimizer with momentum of 0.9 and weight decay of $1e-5$. The initial learning rate is 0.01, with the OneCycle [19]

strategy to adjust the learning rate. The maximum number of training epochs is 50, except for the CFP-only ResNet which used 100 epochs due to its relatively slow convergence speed. Per training sample, we use a smaller number of 6 OCT / CFP instances for the purposes of case-level data augmentation and saving GPU memory to use bigger batch size. In particular, random down-sampling is conducted to select the OCT instances, while random crop is used to generate CFP instances for training. Every instance is resized to 256×256 . In addition, traditional low-level data augmentation is performed on instances. For a more reliable evaluation, for each model we repeat the training procedure three times and select the version that maximizes AP on the validation set.

4.2 Comparison with State-of-the-art

4.2.1 Baselines. We compare the following state-of-the-art:

- MM-CNN [21]: A two-stream CNN that accepts a CFP and a manually selected OCT B-scan image as a paired input. Multi-modal feature fusion is implemented by concatenating 2,048-dimensional ResNet-50 features of CFP and OCT. We train MM-CNN with B-scans selected on the basis of our expert annotations.
- OCT-MIL [16]: A very recent method that uses MIL to exploit a set of CFPs for ROP classification. We adopt its MIL module to perform classification based on the OCT B-scan sequence.
- OCT-Conv3D: A 3D-CNN is used by [13] for OCT-based retinal disease classification. However, 3D-ResNet [11] failed to converge on our training data. Hence, we try an alternative solution, which first utilizes 2D-ResNet for feature extraction. Three 3D conv. blocks are then applied on the re-shaped feature maps to exploit the spatio-temporal information. Global average pooling is performed to derive a case-level feature for classification.
- MM-CNN++: We improve MM-CNN, with its original OCT branch replaced by OCT-MIL.
- CFP: A single-modal CNN with a CFP as its input. We include this baseline for a more comprehensive evaluation.

4.2.2 Results. The performance of MM-MIL and the baselines is given in Table 3. Comparing the two modalities, they perform closely on *Normal* and *PM*. The two OCT models perform clearly better than CFP on *ERM*, *ME*, *dryAMD* and *wetAMD*. As for *DR*, its main symptoms are caused by vascular lesions, which are more vivid on color fundus photos. This explains the clear performance gap between the OCT models and the CFP model for recognizing *DR*. The result justifies the necessity of a multi-modal approach to retinal disease recognition.

MM-CNN performs closely to OCT-MIL, the best single-modal model. This result shows that the better performance of multi-modal fusion cannot be taken as granted. Table 4 shows that MM-CNN is sensitive to the choice of the B-scan chosen for classification.

Compared to OCT-MIL, MM-CNN++ improves the overall AP from 0.7710 to 0.8172 (a relative improvement of 6.0%), making it the best multi-modal baseline. The proposed MM-MIL outperforms this strong baseline, reaching the best AP of 0.8539. It is worth pointing out that none of the multi-modal baselines can beat the single-modal baselines for all diseases. For instance, MM-CNN++ is less effective than the CFP model on *DR*. By contrast, MM-MIL consistently outperforms the best of the single-modal baselines on every category. Specifically, MM-MIL×4 detects *Normal* with

Table 3: Performance of the state-of-the-art for retinal disease recognition. Neither of the multi-modal baselines, i.e., MM-CNN and MM-CNN++, can beat the best of the single-modal baselines (highlighted in lightblue cell) for all diseases. Our proposed MM-MIL×4 does so, showing its effectiveness for multi-modal feature fusion.

Model	Overall performance					AP per category						
	AP	AUC	Sen.	Spe.	F1	Normal	ERM	ME	DR	dryAMD	wetAMD	PM
<i>Single-modal:</i>												
CFP	0.6827	0.9025	0.6109	0.9500	0.7295	0.9830	0.5508	0.6405	0.8738	0.3879	0.5277	0.8152
OCT-MIL	0.7748	0.9281	0.6640	0.9495	0.7417	0.9775	0.7173	0.8451	0.7442	0.5045	0.8108	0.8245
OCT-Conv3D	0.7710	0.9283	0.6324	0.9370	0.7141	0.9756	0.7578	0.8390	0.6408	0.5049	0.8195	0.8597
<i>Multi-modal:</i>												
MM-CNN	0.7743	0.9072	0.7244	0.9522	0.8128	0.9761	0.7248	0.8020	0.7607	0.5210	0.8151	0.8200
MM-CNN++	0.8172	0.9403	0.6917	0.9612	0.7892	0.9861	0.7626	0.8429	0.8409	0.5764	0.8087	0.9030
MM-MIL×1	0.8097	0.9261	0.7311	0.9635	0.8214	0.9888	0.7576	0.7802	0.8617	0.5326	0.8325	0.9149
MM-MIL×2	0.8348	0.9424	0.7577	0.9664	0.8393	0.9933	0.7628	0.8409	0.8898	0.5584	0.8969	0.9014
MM-MIL×4	0.8539	0.9545	0.7431	0.9643	0.8291	0.9940	0.7793	0.9038	0.8939	0.6220	0.9035	0.8808
MM-MIL×8	0.8534	0.9505	0.7335	0.9720	0.8285	0.9922	0.7939	0.8633	0.9165	0.5969	0.9105	0.9006

Table 4: Performance of MM-CNN [21] given distinct B-scan selection strategies at the test stage. Manual selection is needed in order to maximize the performance of MM-CNN.

Which B-scan?	AP	AUC	Sen.	Spe.	F1
First-frame	0.7551	0.9081	0.7199	0.9498	0.8078
Middle-frame	0.7524	0.9168	0.6490	0.9495	0.7448
Last-frame	0.7502	0.9188	0.6941	0.9532	0.7904
Manual	0.7743	0.9072	0.7244	0.9522	0.8128

an AP of 0.9940 (95% confidence interval (CI), 0.9845-1.0), and the detection performance for the other classes are as follows: ERM 0.7793 (95% CI, 0.7048-0.8538), ME 0.9038 (95% CI, 0.8406-0.9669), DR 0.8939 (95% CI, 0.8281-0.9597), dryAMD 0.6220 (95% CI, 0.5028-0.7412), wetAMD 0.9035 (95% CI, 0.7964-1.0000), PM 0.8808 (95% CI, 0.7638-0.9977). MM-MIL is more selective than MM-CNN++.

4.3 Ablation Study

We now evaluate the influence of three major designs, i.e., the over sampling strategy, the number of heads in MM-MIL, and alternatives to MM-MIL, on the performance of our classification network.

4.3.1 The influence of the over sampling strategy. As Table 5 shows, MM-MIL×1 without over sampling causes a clear performance drop, with AP decreased from 0.8097 to 0.7872. We attribute this result to the reason that without over sampling, the OCT instances become dominant during the training process. Indeed, this is also confirmed by the fact that MM-MIL×1 without over sampling performs quite closely to OCT-MIL, suggesting that the CFP modality is mostly underexplored. In addition, we try over sampling in the single-modal setting, by training CFP-MIL which uses MIL to exploit the generated sequence of CFPs. The better performance of CFP-MIL against the CFP model (0.7357 versus 0.6827 in terms of AP) further justifies the viability of the over sampling tactic.

4.3.2 Number of heads in MM-MIL. The performance of MM-MIL with distinct number of heads is presented in the last four rows of Table 3. As we have already noted, the importance of the CFP and

Table 5: Performance of MM-MIL×1 with and without over sampling.

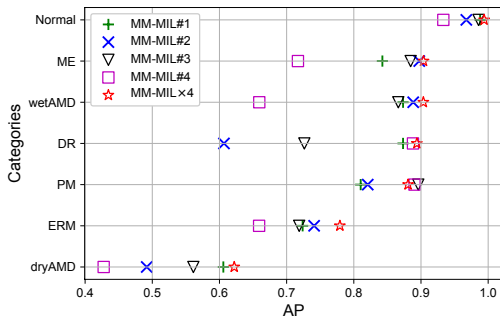
Over sampling?	AP	AUC	Sen.	Spe.	F1
Yes	0.8097	0.9261	0.7311	0.9635	0.8214
No	0.7872	0.9286	0.6729	0.9468	0.7640

OCT modalities varies over diseases. MM-MIL with more heads has a larger capacity to learn better instance attention. Therefore, better performance is observed on MM-MIL×4 and MM-MIL×8, with the latter slightly worse. Recall that we have 7 categories in total, making MM-MIL×4 sufficient for the current task.

To better understand the behavior of MM-MIL×4, we provide a zoom-in view in Fig. 5. Per-category performance of each head, denoted as MM-MIL# k , is plotted in Fig. 5(a). Recall that the k -th head produces multi-modal instance attention weights. Per case we have $(\sum_{i=1}^n a_{c,k,i}) + (\sum_{i=1}^n a_{o,k,i}) = 1$. This property means for each category, the averaged value of the accumulated attention weights per modality reflects to what extent the head is attended. Accordingly, we visualize in Fig. 5(b) each head’s per-category attention weight on the CFP modality.

A joint use of Fig. 5(a) and 5(b) allows us to interpret the behavior of a specific head. MM-MIL#4 puts more attention on CFP. Hence, it performs well on DR and PM for which this modality is suited, but is less effective for dryAMD and wetAMD, where lesions such as *drusen* appear more vividly in OCT images. By contrast, MM-MIL#2 pays much less attention to the CFP modality, making it the least effective for recognizing DR and PM.

4.3.3 MM-MIL versus MHA. Table 6 shows the performance of our multi-modal retinal disease classification network when substituting the popular MHA module for MM-MIL. The mini-version of MHA, which has one head and one depth only, decreases AP from 0.8539 to 0.8116. A wider and deeper MHA ($d = 4$ and $h = 4$) reduces the performance further. The results allow us to conclude that MHA, while being extremely popular in a wide range of tasks, is less effective than our proposed MM-MIL for multi-modal retinal disease classification.



(a) Per-category performance of each head in MM-MIL×4

	Normal	ME	wetAMD	DR	PM	ERM	dryAMD
MM-MIL#1	0.92	0.52	0.22	0.77	0.79	0.57	0.54
MM-MIL#2	0.15	0.07	0.05	0.14	0.10	0.11	0.11
MM-MIL#3	0.97	0.12	0.38	0.28	0.37	0.25	0.61
MM-MIL#4	0.67	0.66	0.56	0.84	0.59	0.66	0.71

(b) Per-category attention weights on the CFP modality

Figure 5: A zoom-in view of MM-MIL×4, showing (a) per-category performance of each of the four heads and (b) how much attention the heads pay to the CFP modality.


Table 6: Performance of our multi-modal network with distinct fusion modules, i.e., MM-MIL versus Multi-Head self-Attention (MHA). Our MM-MIL achieves better performance with much fewer parameters (see Table 1).

Fusion module	AP	AUC	Sen.	Spe.	F1
MHA (d=1, h=1)	0.8116	0.9442	0.7055	0.9668	0.8061
MHA (d=4, h=4)	0.8043	0.9310	0.6982	0.9589	0.8025
MM-MIL×4	0.8539	0.9545	0.7431	0.9643	0.8291


Fig. 6 shows some qualitative results of model prediction and instance attention-weighted activation maps, cf. Section 3.3. Consider Fig. 6(c) for instance. We observe that the activation area covers the temporal side of the optic disc (the brightest region in the CFP). This differs from the other sub-figures, neither of which has their optic disc overlapped with the activation area. Note that for PM, peripapillary atrophy is a clinical finding typically found in the area surrounding the optic disc. Recall that our model is trained exclusively on case-level labels, the result demonstrates the viability of the activation maps for model interpretability.

5 CONCLUSIONS


We develop an end-to-end deep learning approach to multi-modal retinal disease recognition. Extensive experiments on a real-world dataset support the following conclusions. As the efficacy of color



(a) Expert labels: DR → Prediction: DR (0.9542)



(b) Expert labels: wetAMD → Prediction: wetAMD (0.9967)



(c) Expert labels: PM → Prediction: PM (0.9992)

Figure 6: Some qualitative results of model prediction and visualization. Numbers in parentheses are predicted scores.

fundus photography and OCT scans is disease-dependent, the ability of being both selective and interpretable is important for multi-modal fusion. The proposed MM-MIL module possesses both properties, as demonstrated by its superior performance against the prior art. Moreover, MM-MIL has substantially fewer parameters than the prevalent Multi-Head self-Attention module, and thus can be trained on relatively small-sized data. All this makes MM-MIL attractive for AI-assisted retinal disease diagnosis.

Acknowledgements. This research was supported by BJNSF (4202033), BJNSFC Haidian Original Innovation Joint Fund (19L2062), the Pharmaceutical Collaborative Innovation Research Project of Beijing Science and Technology Commission (Z191100007719002), and Public Computing Cloud, Renmin University of China.

REFERENCES

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450* (2016).
- [2] Philippe M Burlina, Neil Joshi, Michael Pekala, Katia D Pacheco, David E Freund, and Neil M Bressler. 2017. Automated grading of age-related macular degeneration from color fundus images using deep convolutional neural networks. *JAMA Ophthalmology* 135, 11 (2017), 1170–1176.
- [3] Joao Carreira and Andrew Zisserman. 2017. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *CVPR*.
- [4] Hila Chefer, Shir Gur, and Lior Wolf. 2021. Transformer Interpretability Beyond Attention Visualization. In *CVPR*.
- [5] Ying Cheng, Ruize Wang, Zhihao Pan, Rui Feng, and Yuejie Zhang. 2020. Look, Listen, and Attend: Co-Attention Network for Self-Supervised Audio-Visual Representation Learning. In *ACMMM*.
- [6] Siying Dai, Leiting Chen, Ting Lei, Chuan Zhou, and Yang Wen. 2020. Automatic Detection Of Pathological Myopia And High Myopia On Fundus Images. In *ICME*.
- [7] Jeffrey De Fauw, Joseph R Ledsam, Bernardino Romera-Paredes, Stanislav Nikolov, Nenad Tomasev, Sam Blackwell, Harry Askham, Xavier Glorot, Brendan O'Donoghue, Daniel Visentin, et al. 2018. Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature medicine* 24, 9 (2018), 1342–1350.
- [8] Rishab Gargeya and Theodore Leng. 2017. Automated identification of diabetic retinopathy using deep learning. *Ophthalmology* 124, 7 (2017), 962–969.
- [9] Negin Ghamsarian, Hadi Amirpourazarian, Christian Timmerer, Mario Taschwer, and Klaus Schöffmann. 2020. Relevance-Based Compression of Cataract Surgery Videos Using Convolutional Neural Networks. In *ACMMM*.
- [10] Varun Gulshan, Lily Peng, Marc Coram, Martin C Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cuadros, et al. 2016. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA* 316, 22 (2016), 2402–2410.
- [11] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. 2018. Can Spatiotemporal 3D CNNs Retrace the History of 2D CNNs and ImageNet?. In *CVPR*.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*.
- [13] Daniel S Kermany, Michael Goldbaum, Wenjia Cai, Carolina CS Valentim, Huiying Liang, Sally L Baxter, Alex McKeown, Ge Yang, Xiaokang Wu, Fangbing Yan, et al. 2018. Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell* 172, 5 (2018), 1122–1131.
- [14] Cecilia S. Lee, Doug M. Baughman, and Aaron Y. Lee. 2017. Deep Learning is Effective for Classifying Normal versus Age-Related Macular Degeneration OCT Images. *Ophthalmology Retina* 1, 4 (2017), 322–327.
- [15] Bing Li, Huan Chen, Bilei Zhang, Mingzhen Yuan, Xuemin Jin, Bo Lei, Jie Xu, Wei Gu, David Wong, Xixi He, Hao Wang, Dayong Ding, Xirong Li, Weihong Yu, and Youxin Chen. 2021. Development and evaluation of a deep learning model for the detection of multiple fundus diseases based on color fundus photography. *British Journal of Ophthalmology* (2021).
- [16] Xirong Li, Wencui Wan, Yang Zhou, Jianchun Zhao, Qijie Wei, Junbo Rong, Pengyi Zhou, Limin Xu, Lijuan Lang, Yuying Liu, Chengzhi Niu, Dayong Ding, and Xuemin Jin. 2020. Deep Multiple Instance Learning with Spatial Attention for ROP Case Classification, Instance Selection and Abnormality Localization. In *ICPR*.
- [17] Ilse Maximilian, Tomczak Jakub, and Welling Max. 2018. Attention-based Deep Multiple Instance Learning. In *ICML*.
- [18] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An imperative style, high-performance deep learning library. In *NeurIPS*.
- [19] Leslie N Smith and Nicholay Topin. 2019. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*.
- [20] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NeurIPS*.
- [21] Weisen Wang, Zhiyan Xu, Weihong Yu, Jianchun Zhao, Jingyuan Yang, Feng He, Zhikun Yang, Di Chen, Dayong Ding, Youxin Chen, and Xirong Li. 2019. Two-Stream CNN with Loose Pair Training for Multi-modal AMD Categorization. In *MICCAI*.
- [22] Xi Wei, Tianzhu Zhang, Yan Li, Yongdong Zhang, and Feng Wu. 2020. Multi-Modality Cross Attention Network for Image and Sentence Matching. In *CVPR*.
- [23] Jun Wu, Yao Zhang, Jie Wang, Jianchun Zhao, Dayong Ding, Ningjiang Chen, Lingling Wang, Xuan Chen, Chunhui Jiang, Xuan Zou, Xing Liu, Hui Xiao, Yuan Tian, Zongjiang Shang, Kaiwei Wang, Xirong Li, Gang Yang, and Jianping Fan. 2020. AttenNet: Deep Attention Based Retinal Disease Classification in OCT Images. In *MMM*.
- [24] Gang Yang, Fan Li, Dayong Ding, Jun Wu, and Jie Xu. 2021. Automatic Diagnosis of Glaucoma on Color Fundus Images Using Adaptive Mask Deep Network. In *MMM*.
- [25] Hao Ye, Zuxuan Wu, Rui-Wei Zhao, Xi Wang, Yu-Gang Jiang, and Xiangyang Xue. 2015. Evaluating Two-Stream CNN for Video Classification. In *ICMR*.