

Multi-disease Detection in Retinal Images Guided by Disease Causal Estimation

Jianyang Xie¹, Xiuju Chen², Yitian Zhao³, Yanda Meng⁵, He Zhao¹, Anh Nguyen⁴, Xiaoxin Li²(✉), and Yalin Zheng¹(✉)

¹Eye and Vision Sciences Department, University of Liverpool, Liverpool, UK

²Xiamen Eye Center, Xiamen University, China

³Ningbo Institute of Materials Technology and Engineering, CAS, China

⁴Computer Science Department, University of Liverpool, Liverpool, UK

⁵Computer Science Department, University of Exeter, Exeter, UK

yzheng@liverpool.ac.uk

Abstract. There have been significant advancements in analyzing retinal images for the diagnosis of eye diseases and other systemic conditions. However, a key challenge is multi-disease detection, particularly in addressing the demands of real-world applications where a patient may have more than one condition. To address this challenge, this study introduces a novel end-to-end approach to multi-disease detection using retinal images guided by disease causal estimation. This model leverages disease-specific features, integrating disease causal relationships and interactions between image features and disease conditions. Specifically, 1) the interactions between disease and image features are captured by cross-attention in a transformer decoder. 2) The causal relationships among diseases are automatically estimated as the directed acyclic graph (DAG) based on the dataset itself and are utilized to regularize disease-specific feature learning with disease causal interaction. 3) A novel retinal multi-disease dataset of 500 patients, including six lesion labels, was generated for evaluation purposes. Compared with other methods, the proposed approach not only achieves multi-disease diagnosis with high performance but also provides a method to estimate the causal relationships among diseases. We evaluated our method on two retinal datasets: a public colour fundus photography and an in-house fundus fluorescein angiography (FFA). The results show that the proposed method outperforms other state-of-the-art multi-label models. Our FFA database and code have been released ¹.

Keywords: Retinal image · Multi-disease detection · Disease causal estimation · Disease-specific features · Disease-feature interaction.

1 Introduction

Retinal images are essential in computer-aided diagnosis to evaluate conditions such as glaucoma [5], cataracts [26], and age-related macular degeneration [13].

¹ <https://github.com/davelailai/multi-disease-detection-guided-by-causal-estimation.git>

Deep learning has shown remarkable potential for specific disease detection [10,25,6,21] with the assumption that each instance corresponds to only a specific disease. However, the complexity of retinal diseases in the real world, often involving multiple conditions in a single patient [8], makes accurate diagnosis difficult. Therefore, the demand for deep learning techniques to handle retinal multi-disease detection is growing evident in clinical settings.

In recent decades, multi-label classification of retinal images has gained significant attention. This work falls within the category of label-specific feature modeling, which has emerged as a mainstream. In these methods, the label co-occurrence or the interaction between label and image features is utilized to guide label-specific feature generation. For example, in the label co-occurrence modeling [3,11,22], they leveraged dependencies between labels as prior knowledge, employing graph-based methods to guide model inference. In the transformer-based methods [12,16,23,17], the interactions between label embeddings and image spatial features were captured to learn effective label-specific features.

While the aforementioned label-specific methods have yielded promising results, they encounter several limitations. On the one hand, these methods often consider only one aspect of correlation; for instance, the label co-occurrence method overlooks the feature-disease interaction, while transformer-based methods miss the label co-occurrence. On the other hand, in label co-occurrence methods [3,11,22], the statistical graph constructed from a small dataset may introduce a frequency-bias issue, potentially leading to overfitting [22,14,19]. Additionally, this label co-occurrence fails to capture the causal relationships between labels [20]. These causal associations were proved in [2], where the authors stated that Type-2 diabetes mellitus might cause senile cataracts, glaucoma, and disorders of the optic nerve and visual pathways.

Inspired by these insights, we propose a multi-disease detection in retinal images guided by causal estimation. Our approach integrates both disease causal relationships and interactions between image features and disease conditions. More specifically, cross-attention in transformer-based models is employed to capture the interaction between retinal image features and disease conditions. Furthermore, a causal matrix among the diseases is learned as a directed acyclic graph (DAG) and serves as an auxiliary task, facilitating the exploration of causal relationships between diseases, and then regularizing the disease-specific feature learning with disease causal interaction. In this case, the disease-specific features in our method are more informative and discriminative, thereby enhancing detection performance. Additionally, the estimated causal matrix provides insights to explore causal relationships among diseases.

2 Method

In this section, the notation of multi-disease detection based on disease-specific feature learning is introduced, and then a brief overview of the proposed method and its key modules are described in detail.




Fig. 1. Overall architecture and details of the proposed method. (a) Pipeline of the proposed model, ST , TP , L , NP are names of the disease. (b) Transformer decoder. (c) Framework of the causal estimation module.

Notations. In the context of a retinal image X containing multiple diseases, multi-disease detection aims to predict the presence of each disease. Disease-specific feature learning facilitates the extraction of unique features for each disease category, enhancing classification accuracy. Let K denote the total number of disease categories, and the corresponding labels of X are denoted as $Y = [y_1, \dots, y_K]$, where $y_k \in \{0, 1\}$ and $k = 1, \dots, K$ represents a discrete binary indicator. The disease-specific features are represented as $Q = [q_1, \dots, q_K]$, where $q_k \in \mathbb{R}^{1 \times C}$ and C is the number of feature channels. Specifically, $y_k = 1$ if the retinal image X exhibits the k_{th} disease; otherwise, $y_k = 0$. Utilizing X as input, the multi-disease detection method predicts the probabilities P of each disease category being present.

2.1 Overview of proposed method

The framework is illustrated in **Fig 1 (a)**. First, transformer-based method [12,16] was employed to acquire an initial disease-specific feature, capturing the interaction between image features and diseases. Then, a causal matrix was estimated to regulate the disease-specific feature learning process, under the assumption that features of a given disease can be modeled based on its parents' diseases using logistic regression. In this setup, disease causal relationships can be estimated from the data, and the disease-specific features with this constraint account for disease causal interactions, effectively guiding the model toward more consistent classification outcomes.

2.2 Disease-Specific Feature Learning

In our method, the disease-specific features are learned by considering both the interaction between image spatial features and diseases, as well as the causal relationships among diseases.

Disease-Feature Interaction. Our proposed method adopts a transformer-based approach [12,16], as depicted in **Fig 1. (a-b)**. Given a retinal image X , it initially undergoes processing through a feature extraction network to extract spatial features $F_0 \in \mathbb{R}^{m \times n \times c}$ reshaped as $F_0 \in \mathbb{R}^{mn \times c}$. Our approach is backbone-agnostic, allowing the utilization of any feature extraction architecture. Here, we employed ResNet50 [7]. Subsequently, to capture the interaction between image features and disease embedding, the standard transformer decoder architecture was utilized with the F_0 and disease embedding $Q_0 \in \mathbb{R}^{K \times c}$ as input. Thus, the disease-specific feature can be calculated following **Equation 1**.

$$\begin{aligned} \text{Self-Attention: } Q_1 &\leftarrow \text{MultiHeadAttn}(Q_0, Q_0, Q_0) \\ \text{Cross-Attention: } Q_1 &\leftarrow \text{MultiHeadAttn}(Q_1, F_0, F_0) \\ \text{Feed Forward Network: } Q_1 &= \text{FFN}(Q_1) \end{aligned} \quad (1)$$

where $Q_1 \in \mathbb{R}^{K \times C}$ represents the disease-specific features, incorporating disease-feature interactions. C represents the feature channel. The initial values of Q_0 are randomly initialized and set as learnable parameters.

Disease Interaction in Causal Perspective To address the unreliable statistical relationships among diseases as observed in prior works [3,11,22], and to harness the causal relationships among diseases, we propose a method to estimate disease causality and unveil underlying disease connections. Specifically, we proposed a directed acyclical graph (DAG) learning approach, assuming a causal DAG exists among initial disease-specific features Q_1 . This approach considers disease causality, where the feature q_i for disease i can be modeled according to its parents (diseases that might cause the disease i) via logistic regression based on **Equation 2**. Regularizing disease-specific features with this constraint accounts for disease-causal interaction.

$$q_i = f_i(\text{Pa}(q_i), u_i), \quad i \in K \quad (2)$$

where $\text{Pa}(q_i)$ represents the parents of q_i , and u_i are some random noise variables. f_i represents logistic regression.

In practice, after obtaining initial disease-specific feature $Q_1 \in \mathbb{R}^{K \times C}$. Let $A(\mathbf{W}) \in \mathbb{R}^{K \times K}$ represent the DAG with zeros on the diagonal, where W_{ij} denotes the causal relationship strength from disease i to disease j . Thus, the disease-specific feature can be updated by the weighted adjacency matrix $A(\mathbf{W})$ according to the **Equation 3**.

$$Q = \mathbf{W}Q_1 \quad (3)$$

Where Q represents the disease-specific feature that incorporates both the interaction between image spatial features and diseases, as well as the causal in-

teractions among diseases. In this scenario, learning the causality DAG becomes the primary challenge, which is addressed in **Section 2.3**.

2.3 Causal Relationship Estimation

Inspired by the work of Zheng et al. [24] on learning linear DAGs through continuous optimization, we adopt a similar approach for learning the DAG $\tilde{\mathbf{W}}$, as described by **Equation 4**.

$$\tilde{\mathbf{W}} \in \arg \min_{\mathbf{W}} (\lambda(L_{SEM} + \|\mathbf{W}\|_1) + R_{DAG}(\mathbf{W})) \quad (4)$$

where λ is a hyper-parameter to balance the two items, $L_{SEM} = \|Q_1 - \mathbf{W}Q_1\|_2^2$ is the score function, indicating the causal relationship that the one disease-specific feature can be modelled according to its parents. $\|\mathbf{W}\|_1$ is the $\mathbb{L}_1$ norm of $\mathbf{W}$, promotes sparsity in $\tilde{\mathbf{W}}$, $R_{DAG}(\mathbf{W})$ is the DAG regularization loss.

We propose to learn a causal DAG in a nonlinear setting, where the DAG is parameterized by a 2-layer feed-forward neural network (represented as f) with Sigmoid activation, and the Q can be represented as $Q = \mathbf{Sigm}(f(Q_1, w))$, as shown in **Figure 1(c)**. In the causal structure modeling, after passing through a linear layer with parameter size $w \in \mathbb{R}^{K \times Kd}$ (where d is the channel of DAG and set as 30) and applying a sigmoid activation function to introduce nonlinearity, Q_1 can be reconstructed as $\tilde{Q}_1 \in \mathbb{R}^{C \times d \times K}$. Subsequently, the reconstructed label-specific feature $Q \in \mathbb{R}^{C \times K}$ is obtained by applying a Conv1d operation with a kernel size of 1. Thus the L_{SEM} can be rewrite as $L_{SEM} = \|Q_1 - \mathbf{Sigm}(f(Q_1, w))\|_2^2$, and $\tilde{\mathbf{W}}$ can be calculated by optimizing **Equation 5**

$$\tilde{\mathbf{W}} \in \arg \min_{\mathbf{W}} (\lambda(\|Q_1 - \mathbf{Sigm}(f(Q_1, w))\|_2^2 + \|w\|_1) + R_{DAG}(\mathbf{W})) \quad (5)$$

with the DAG loss introduced in [1], writing $R_{DAG}(\mathbf{W})$ as **Equation 6**

$$R_{DAG}(\mathbf{W}) = -\log \det(sI - \mathbf{W} \circ \mathbf{W}) + K \log s \quad (6)$$

where $\mathbf{W} = \|w\|_2$ is the learned causal matrix, I is the identity matrix, s is a given scalar (setting as 1 here), and $\circ$ denotes the element-wise Hadamard product, K is disease categories. $\log \det$ is the log-determinant function.

2.4 Loss Function

In this paper, the model was trained end-to-end, and the overall loss function was defined as **Equation 7**:

$$Loss = L_c + \lambda(L_{SEM} + \|w\|_1) + R_{DAG}(\|w\|_2) \quad (7)$$

where L_c is simplified asymmetric loss [15] for multi-label classification.

Table 1. Details of the LID-FFA dataset, including the number of positive (negative) images for each pathology.

	L	TP	ST	SH	NP	VA	Total
Train	3371(418)	1773 (2016)	2953 (836)	1878 (1911)	1451 (2338)	2108(1681)	3789
Val	459 (87)	231 (315)	400 (146)	237 (309)	212 (334)	277 (269)	546
Test	1014 (86)	514 (586)	922 (178)	595 (505)	453 (647)	696 (404)	1100

3 Experiments

3.1 Dataset and Evaluation

Two datasets are utilized to validate the performance of the proposed method, the OIA-ODIR [9] and our in-house Lesion Intelligent Detection in Fundus Fluorescein Angiography dataset (LID-FFA).

OIA-ODIR [9] is the first internationally available dataset for multi-disease detection based on binocular fundus images, comprising 10,000 fundus photographs from 5,000 patients with eight types of ocular disease classifications. The dataset is divided into three subsets: training, off-site testing, and on-site testing, with 3,500, 500, and 1,000 patients, respectively. In our setting, the on-site is utilized for validation, and the off-site is utilized for testing.

LID-FFA is a multi-lesion detection dataset, consisting of 5435 fundus fluorescein angiograph (FFA) images from 500 patients. It includes six pathological features: Leakage (L), Transmission and Pooling (TP), Staining (ST), Shadowing (SH), Non-Perfusion (NP), and Vessel Abnormality (VA). Each image may contain more than one lesion pathology. Manual labeling was performed by two junior ophthalmologists at the image level, with any discrepancies resolved by a senior ophthalmologist to make the final arbitration. The FFA images were acquired by Spectralis HRA+OCT (Heidelberg Engineering, Heidelberg, Germany). In our setup, the dataset was split randomly based on patients into sets of 350 for training, 50 for validation, and 100 for testing. Further details are provided in **Table 1**. Several FFA cases are depicted in *supplementary*.

Data pre-processing. The original images were resized to a uniform image resolution of 512×512 , followed by random cropping to obtain patches of size 448×448 , and random horizontal flip.

Evaluation metrics. We use macro average precision (MAP), macro average recall (MAR), macro average F1-score (MAF1), mean average area under the receiver operating characteristic (mAUC) and mean average precision (mAP) for the evaluation. The threshold is set as 0.5 for all metrics.

3.2 Implementation Details

The public framework mmpretrain [4] was used to implement all our deep neural networks. All the experiments were run on a NVIDIA Tesla A100 with 80 G memory. The stochastic gradient descent (SGD) optimizer is applied to train

Table 2. Results of the comparison and the ablation experiments (%), text in **bold** or **red** indicates the best performance among all comparison methods, while text in **red** denotes an improvement over the corresponding baseline in the ablation experiments.

Model	OIA-ODIR					LID-FFA				
	MAP	MAR	MAF1	mAP	mAUC	MAP	MAR	MAF1	mAP	mAUC
Resnet50 [7]	17.95	92.04	29.29	47.49	74.61	80.23	83.44	81.59	76.30	66.94
AsyLoss [15]	40.26	72.16	49.62	57.06	79.97	81.44	87.42	83.76	87.22	76.94
GCN [11]	33.28	66.87	42.21	54.39	78.50	80.03	88.64	83.94	85.47	74.80
dyGCN [22]	45.51	70.10	53.70	58.09	80.34	83.15	83.99	83.46	87.47	76.63
Q2L [12]	41.62	67.01	52.73	50.13	79.68	83.62	85.70	84.59	87.69	76.73
Q2L causal	43.53	70.32	51.69	60.32	81.69	83.67	84.87	84.17	88.60	77.65
MLDecoder [16]	42.23	73.18	50.22	60.85	79.29	83.09	86.66	84.78	88.24	76.84
ML causal	42.55	80.00	50.46	59.59	82.51	83.85	86.17	84.93	86.88	78.02




Fig. 2. AUC Analysis for separate disease. The disease indication in OIA-ODIR is as follows: diabetes retinopathy (D), glaucoma (G), cataract (C), age-related macular degenerate (A), hypertension (H), myopia (M), and other abnormalities (O).

the networks. The initial learning rate is set to 0.005 with the onecycle learning rate decay policy [18]. The epoch was set as 100 with batch size as 64.

3.3 Comparison and Ablation Experiments

We compared our method with ResNet50 [7], an asymmetric loss [15] designed for unbalanced datasets, GCN-based methods [3,22] designed for label-label correlation, and Transformer-based methods [12,16] designed for label-feature interaction. To ensure a fair comparison, we replaced the feature extraction backbone of each method with ResNet50 [7], initialized with weights pretrained on the ImageNet-1k dataset. To evaluate the efficiency and generalization of our proposed method, we utilized Q2L [12] and ML-Decoder [16] as the disease-feature interaction component in **Fig 1 (a)**. The average performance is summarized in **Table 2**. Our proposed method (**Q2L causal** and **ML causal**) exhibits superior average performance compared to other methods, particularly in terms of mAUC, where it achieves the best results on both datasets. Regarding the ablation of the proposed causal estimation, when compared with the baseline (**Q2L** and **MLDecoder**), combining the causal estimation in our proposed methods




Fig. 3. Analysis of Channels in Causal Matrix.

(**Q2L causal** and **ML causal**) yields improved performance across almost all metrics. Moreover, the AUC for each disease is illustrated in **Fig 2**, demonstrating superiority on most tasks.

3.4 Analysis of Channels in Causal Matrix

In the causal estimation section depicted in **Fig 1 (c)**, we investigated the profound impact of channel d on multi-disease detection performance by varying its values to 10, 30, 50, and 100 within the Linear layer architecture. Subsequently, we trained both the **Q2L causal** and **ML causal** models with each configuration and evaluated their performance using the mAUC. The results are presented in **Fig 3**. It can be observed that the proposed method achieves optimal performance when the channel of the causal matrix d is set to 30.

3.5 Analysis of Causal Estimation

To explore the causal matrix within our proposed method, we evaluate the causal matrix $\mathbf{W}$ learned from each dataset using the GCN-based approach [11]. More precisely, we utilize the disease co-occurrence matrix statistics obtained from the training dataset as a reference. To be a fair comparison, diagonals in the disease co-occurrence matrix are designated as zeros. Subsequently, we replace this disease co-occurrence matrix with our learned causal matrix and proceed to train the model on the corresponding datasets. For example, the causal matrix learned from OIA-ODIR was utilized to train the model on the OIA-ODIR. The results are presented in **Table 3**. The findings demonstrate that after substituting the disease co-occurrence matrix with the causal matrix learned from our method, the classification performance exhibits a significant improvement across most evaluation metrics. This result proves that the proposed causal estimation method efficiently captures the causal relationship between diseases to some extent. The visualization of the causal matrixes are generated in *supplementary*.

4 Conclusion

In this paper, we propose a method for multi-disease detection in retinal images that leverages disease-specific features and integrates disease-causal relationships. Cross-attention in the transformer model was employed to capture

Table 3. Exploration of the causal matrix, **GCN** means that the disease co-occurrence matrix with diagonals set as zero serves as the adjacency matrix (graph). **GCN+QL** (**GCN+ML**) means the causal matrix learned in **Q2L causal** (**ML causal**) is utilized.

Model	OIA-ODIR					LID-FFA				
	MAP	MAR	MAF1	mAP	mAUC	MAP	MAR	MAF1	mAP	mAUC
GCN [11]	33.28	66.87	42.21	54.39	78.50	80.03	88.64	83.94	85.47	74.80
GCN+QL	47.24	68.60	54.87	57.61	80.66	84.20	83.46	83.71	87.51	75.94
GCN+ML	40.85	67.10	49.19	55.87	79.63	83.06	87.30	84.97	87.31	77.51

interactions between images and diseases, and a causal estimation method was introduced to learn causal relationships among diseases. Experimental shows that the proposed method demonstrates improved multi-disease detection performance, and the estimated causal matrix provides insights into disease relationships. Additionally, we published a multi-lesion detection dataset in FFA.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bello, K., Aragam, B., Ravikumar, P.: Dagma: Learning dags via m-matrices and a log-determinant acyclicity characterization. *Advances in Neural Information Processing Systems* **35**, 8226–8239 (2022)
2. Chen, R., Xu, S., Ding, Y., Li, L., Huang, C., Bao, M., Li, S., Wang, Q.: Dissecting causal associations of type 2 diabetes with 111 types of ocular conditions: a mendelian randomization study. *Frontiers in Endocrinology* **14** (2023)
3. Chen, Z.M., Wei, X.S., Wang, P., Guo, Y.: Multi-label image recognition with graph convolutional networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 5177–5186 (2019)
4. Contributors, M.: Openmmlab’s pre-training toolbox and benchmark. <https://github.com/open-mmlab/mmpretrain> (2023)
5. Diaz-Pinto, A., Colomer, A., Naranjo, V., Morales, S., Xu, Y., Frangi, A.F.: Retinal image synthesis and semi-supervised learning for glaucoma assessment. *IEEE transactions on medical imaging* **38**(9), 2211–2218 (2019)
6. He, A., Li, T., Li, N., Wang, K., Fu, H.: Cabnet: Category attention block for imbalanced diabetic retinopathy grading. *IEEE Transactions on Medical Imaging* **40**(1), 143–153 (2021). <https://doi.org/10.1109/TMI.2020.3023463>
7. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
8. Ju, L., Yu, Z., Wang, L., Zhao, X., Wang, X., Bonnington, P., Ge, Z.: Hierarchical knowledge guided learning for real-world retinal disease recognition. *IEEE Transactions on Medical Imaging* (2023)
9. Li, N., Li, T., Hu, C., Wang, K., Kang, H.: A benchmark of ocular disease intelligent recognition: One shot for multi-disease detection. In: *Benchmarking, Measuring, and Optimizing: Third BenchCouncil International Symposium, Bench 2020, Virtual Event, November 15–16, 2020, Revised Selected Papers 3*. pp. 177–193. Springer (2021)

10. Li, X., Jia, M., Islam, M.T., Yu, L., Xing, L.: Self-supervised feature learning via exploiting multi-modal data for retinal disease diagnosis. *IEEE Transactions on Medical Imaging* **39**(12), 4023–4033 (2020)
11. Lin, J., Cai, Q., Lin, M.: Multi-label classification of fundus images with graph convolutional network and self-supervised learning. *IEEE Signal Processing Letters* **28**, 454–458 (2021). <https://doi.org/10.1109/LSP.2021.3057548>
12. Liu, S., Zhang, L., Yang, X., Su, H., Zhu, J.: Query2label: A simple transformer way to multi-label classification. *arXiv preprint arXiv:2107.10834* (2021)
13. Maqsood, S., Damaševičius, R., Shah, F.M., Maskeliunas, R.: Detection of macula and recognition of aged-related macular degeneration in retinal fundus images. *Computing and Informatics* **40**(5), 957–987 (2021)
14. Nie, W., Zhang, C., Song, D., Bai, Y., Xie, K., Liu, A.A.: Chest x-ray image classification: A causal perspective. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 25–35. Springer (2023)
15. Ridnik, T., Ben-Baruch, E., Zamir, N., Noy, A., Friedman, I., Protter, M., Zelnik-Manor, L.: Asymmetric loss for multi-label classification. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 82–91 (2021)
16. Ridnik, T., Sharir, G., Ben-Cohen, A., Ben-Baruch, E., Noy, A.: Ml-decoder: Scalable and versatile classification head. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 32–41 (2023)
17. Rodríguez, M.A., AlMarzouqi, H., Liatsis, P.: Multi-label retinal disease classification using transformers. *IEEE Journal of Biomedical and Health Informatics* (2022)
18. Smith, L.N., Topin, N.: Super-convergence: Very fast training of neural networks using large learning rates. In: *Artificial intelligence and machine learning for multi-domain operations applications*. vol. 11006, pp. 369–386. SPIE (2019)
19. Tan, Y., Sun, D., Shi, Y., Gao, L., Gao, Q., Lu, Y.: Bi-directional mapping for multi-label learning of label-specific features. *Applied Intelligence* pp. 1–20 (2022)
20. Tian, Y., Bai, K., Yu, X., Zhu, S.: Causal multi-label learning for image classification. *Neural Networks* **167**, 626–637 (2023)
21. Xie, J., Yi, Q., Wu, Y., Zheng, Y., Liu, Y., Macerollo, A., Fu, H., Xu, Y., Zhang, J., Behera, A., et al.: Deep segmentation of octa for evaluation and association of changes of retinal microvasculature with alzheimer’s disease and mild cognitive impairment. *British Journal of Ophthalmology* (2023)
22. Ye, J., He, J., Peng, X., Wu, W., Qiao, Y.: Attention-driven dynamic graph convolutional network for multi-label image recognition. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI* 16. pp. 649–665. Springer (2020)
23. Zhang, Y., Luo, L., Dou, Q., Heng, P.A.: Triplet attention and dual-pool contrastive learning for clinic-driven multi-label medical image classification. *Medical Image Analysis* **86**, 102772 (2023)
24. Zheng, X., Aragam, B., Ravikumar, P.K., Xing, E.P.: Dags with no tears: Continuous optimization for structure learning. *Advances in neural information processing systems* **31** (2018)
25. Zhou, Y., Wang, B., Huang, L., Cui, S., Shao, L.: A benchmark for studying diabetic retinopathy: Segmentation, grading, and transferability. *IEEE Transactions on Medical Imaging* **40**(3), 818–828 (2021). <https://doi.org/10.1109/TMI.2020.3037771>
26. Zhou, Y., Li, G., Li, H.: Automatic cataract classification using deep neural network with discrete state transition. *IEEE transactions on medical imaging* **39**(2), 436–446 (2019)