

PRETI: Patient-Aware Retinal Foundation Model via Metadata-Guided Representation Learning

Yeonkyung Lee¹, Woojung Han¹, Youngjun Jun¹,
Hyeonmin Kim², Jungkyung Cho², and Seong Jae Hwang¹

¹ Yonsei University, Seoul, South Korea

{yeonkyung.lee, dnwjddl, youngjun, seongjae}@yonsei.ac.kr

² Mediwhale, Seoul, South Korea

{luke.kim, kyla.cho}@mediwhale.com

Abstract. Retinal foundation models have significantly advanced retinal image analysis by leveraging self-supervised learning to reduce dependence on labeled data while achieving strong generalization. Many recent approaches enhance retinal image understanding using report supervision, but obtaining clinical reports is often costly and challenging. In contrast, metadata (e.g., age, gender) is widely available and serves as a valuable resource for analyzing disease progression. To effectively incorporate patient-specific information, we propose PRETI, a retinal foundation model that integrates metadata-aware learning with robust self-supervised representation learning. We introduce Learnable Metadata Embedding (LME), which dynamically refines metadata representations. Additionally, we construct patient-level data pairs, associating images from the same individual to improve robustness against non-clinical variations. To further optimize retinal image representation, we propose Retina-Aware Adaptive Masking (RAAM), a strategy that selectively applies masking within the retinal region and dynamically adjusts the masking ratio during training. PRETI captures both global structures and fine-grained pathological details, resulting in superior diagnostic performance. Extensive experiments demonstrate that PRETI achieves state-of-the-art results across diverse diseases and biomarker predictions using in-house and public data, indicating the importance of metadata-guided foundation models in retinal disease analysis. Our code and pretrained model are available at <https://github.com/MICV-yonsei/PRETI>

Keywords: Foundation Model · Masked Image Modeling · Metadata

1 Introduction

Foundation models have recently gained attention for their remarkable adaptability across various tasks [5, 24, 30]. By training on extensive unlabeled data, they capture core patterns and relationships that can be fine-tuned with minimal effort for various downstream applications [11, 38]. While their applications in areas like computer vision and natural language processing have shown promising results [5, 21, 30], their use in medical contexts remains relatively unexplored.

A key motivation for developing medical foundation models lies in the scarcity of annotated medical data, as labeling clinical images requires expert knowledge and is both time-consuming and expensive. Recently, medical foundation models have shown the potential of using large-scale unlabeled data [26, 40]. Especially in ophthalmology, retinal foundation models are being studied due to the discovery that clinical features associated with ocular diseases can be observed in color fundus photography (CFP) [7, 35]. RETFound [40] demonstrated how image-based pre-training can serve as a strong foundation for retinal fundus analysis. Expanding on this, foundation models such as FLAIR [31], RETCLIP [13], and UrFound [36] incorporate report data, enhancing the multimodal understanding of retinal diseases. While these approaches have significantly advanced retinal image analysis through report supervision, expert descriptions of patients such as clinical diagnostic reports are difficult to obtain, thereby reducing the amount of available data. On the other hand, among various types of metadata, age and gender serve as minimal patient information that can be easily obtained. Moreover, they exert a critical influence on understanding how ocular pathologies develop, progress, and respond to treatment [22, 33]. Thus, we incorporate metadata into foundation models to enable them to capture patient-specific variations.

Despite the significance of metadata, research on foundation models leveraging metadata has been relatively scarce [4, 14, 17]. These approaches treat metadata as static embeddings [4] or incorporate it through a contrastive loss criterion [14, 17]. However, relying on metadata in such a static manner fails to fully leverage metadata for capturing complex image-metadata relationships. Instead, we introduce the *Learnable Metadata Embedding (LME)*. LME dynamically integrates metadata into the learning process, allowing it to adapt and evolve alongside visual features. In contrast to static embeddings or fixed constraints, LME enables a more flexible representation of retinal data. Specifically, we focus on the two key metadata variables (age and gender) which are fundamental to understanding retinal disease patterns [3, 28].

Enhancing representation learning through metadata is a notable step forward. However, retinal foundation models still face challenges from imaging variations due to device differences, lighting, and inter-eye anatomical distinctions. To address these issues, we construct *patient-level data pairs*, grouping images from the same individual, even when captured with different cameras. Additionally, considering the clinical symmetry between the left and right eyes in ocular diseases [27, 32, 34], we treat them as equivalent when forming patient-level pairs. This strategy accounts for non-disease-related variations, enabling the model to learn robust representations. To effectively utilize both metadata and patient-level pairs, we build upon a self-supervised modeling framework based on Siamese Masked Autoencoder (SiamMAE) [15], an extension of MAE [16]. This architecture suits our patient-level pairs by disrupting one image and restoring it using information from the other, learning visual correspondence between both eyes.

While MAE is a self-supervised learning method that learns image features by reconstructing images from randomly masked inputs, its standard random masking strategy does not align well with retinal images, where important struc-

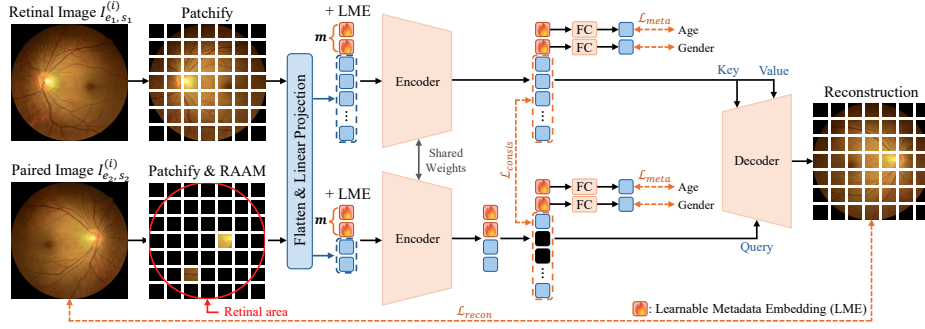


Fig. 1. The architecture of PRETI, a foundation model for retinal image representation learning, illustrates the integration of color fundus photography (CFP) and metadata through Learnable Metadata Embedding (LME). The model processes paired retinal images using Patient-level Pairs and employs shared encoders with dynamic metadata adaptation (age and gender). It also incorporates the Retina-Aware Adaptive Masking (RAAM) strategy for selective masking and reconstructing images via a decoder to enhance representation learning for robust clinical generalization.

tures are confined to the eye region. To address this, we implement *Retina-Aware Adaptive Masking (RAAM)*, a technique designed to align with the specific properties of retinal image data. RAAM selectively applies masking within the retinal region and dynamically adjusts the masking ratio during training. This refinement enhances the ability of the model to capture both global structures and fine details, making it better suited for retinal image analysis.

Based on patient-level pairs incorporating LME and RAAM, we confidently propose PRETI, a foundation model that focuses on CFP and metadata for comprehensive assessment. Through extensive experiments, we validate its state-of-the-art performance across diverse disease and biomarker prediction tasks.

2 Methods

Our foundation model, PRETI, integrates visual and metadata information for retinal image representation learning (Fig. 1). We first introduce patient-level data pairs, followed by the retina-aware adaptive masking strategy (RAAM), which is designed specifically for retinal images. Then, we present the learnable metadata embedding (LME) for capturing the complex relationship between images and metadata, along with the objectives introduced for training.

2.1 Patient-level Data Pairs

We construct patient-level data by grouping CFPs for each individual, designing each pair to consist of two images from the same patient, potentially spanning both left and right eyes and different scanners. When pairs include left and right eye images, we leverage the natural symmetry and biological correlation to

extract intrinsic retinal characteristics specific to the individual [27, 32, 34]. To account for real-world imaging variability, we incorporate images from multiple scanners rather than restricting pairs to a single device, ensuring that the learned representations generalize across diverse acquisition settings. To maximize the number of pairs, we generate all possible combinations of a patient’s retinal images, capturing variations in eye laterality (l : left, r : right) and scanner type \mathcal{S} . For each patient i , we define the image set as $\mathcal{P}_i = \{I_{e,s}^{(i)} | e \in \{l, r\}, s \in \mathcal{S}\}$, where $I_{e,s}^{(i)}$ represents a e -side CFP of patient i acquired by scanner s from the scanner set \mathcal{S} . Each pair is then formed as $(I_{e_1,s_1}^{(i)}, I_{e_2,s_2}^{(i)})$ for some $e_1, e_2 \in \{l, r\}$ and $s_1, s_2 \in \mathcal{S}$, including cases where $s_1 = s_2$ if multiple images from the same scanner are available. This approach integrates both patient-specific retinal information and inter-scanner variability for effective model training.

2.2 Retina-Aware Adaptive Masking Strategy

To account for the distinct properties of retinal images, where the key features primarily reside within the retinal region, we propose a Retina-Aware Adaptive Masking (RAAM). This adaptive masking strategy selectively applies masking to the retinal region and dynamically adjusts the masking ratio during training.

(1) Region-Specific Retinal Masking. Traditional MAE training [16] applies random masking across the entire image. However, in retinal images, the non-retinal regions (e.g., background) do not contain meaningful information for learning. To ensure that the model focuses on critical retinal structures within the circular retinal image area, we introduce a Region-Specific Masking approach, where the unmasked patches are selectively chosen from within the retinal area (a red circle in Fig. 1). A patch is considered part of the retinal region if more than 50% of its pixels fall within the retinal mask. To implement this, we generate a retinal region mask using the automated retinal image analysis tool, Automorph [39].

(2) Cosine Decay-Based Masking Ratio Scheduling. In retinal image analysis, capturing both global structural patterns and fine-grained details is essential. To achieve this, we introduce a dynamic masking schedule that progressively reduces the masking ratio during training. Starting with a high masking ratio r_0 , the model is encouraged to prioritize learning broad structural features, minimizing early reliance on localized information [8]. As training progresses, the masking ratio decreases towards the final masking ratio r_T , allowing the model to focus on fine-grained retinal structures while maintaining strong global representations. To ensure stable learning and a smooth transition, we adopt a Cosine Decay-Based Scheduling [25], which keeps the masking ratio stable early on and reduces it more rapidly as training progresses. The masking ratio r_t follows a cosine decay function: $r_t = 0.5(1 - \cos(\pi t/T))(r_T - r_0) + r_0$, where T is the total number of training epochs.

2.3 Learnable Metadata Embedding

To dynamically integrate metadata into the learning process, we introduce the Learnable Metadata Embedding (LME), which enables metadata to adapt and evolve alongside visual features. We utilize the patient’s age and gender as the sole metadata, keeping the approach lightweight and focused, and introduce learnable meta embeddings, m_{age} and m_{gender} , for each attribute. Formally, we define a set of learnable embeddings $m \in \{m_{\text{age}}, m_{\text{gender}}\}$ and prepend them to the image patch embeddings, ensuring seamless integration with visual features. After the forward pass, the model produces an output embedding for each prompt, which is optimized to align with its corresponding ground-truth metadata labels (y_{age} or y_{gender}). Concretely, we optimize these embeddings through a meta loss function $\mathcal{L}_{\text{meta}}$, as outlined in Eq. (1) of the following section. By minimizing $\mathcal{L}_{\text{meta}}$, each prompt specializes in representing its corresponding clinical attribute: m_{age} adapts to capture age-related variations, and m_{gender} learns to reflect gender-specific patterns in CFPs.

2.4 Joint Optimization Objectives

To train our foundation model, we combine multiple objectives to capture both visual features and metadata. These include reconstruction loss for robust feature extraction, consistency loss to maintain stable representations across patient variations, and meta loss to guide the learnable embeddings for age and gender.

Reconstruction Loss. We include a reconstruction loss, following the objective used in Masked Autoencoders [16], to guide the encoder toward learning discriminative visual features. Given the necessity of capturing fine-grained anatomical structures in retinal images, we further incorporate perceptual loss [18] to make the model restore subtle details more effectively.

Consistency Loss. To further ensure stable learning and robust integration across patient-level variations, we use a feature consistency loss based on cosine similarity: $\mathcal{L}_{\text{consis}} = \frac{1}{N} \sum_{n=1}^N (1 - \hat{z}_{e_1, s_1}^n \cdot \hat{z}_{e_2, s_2}^n)$, where \hat{z}_{e_1, s_1}^n are the normalized feature embeddings of an unmasked patch and \hat{z}_{e_2, s_2}^n corresponds to the embedding of its paired counterpart in the opposite spatial location. Here, N denotes the total number of unmasked patch pairs across all images in a batch.

Meta Loss. To incorporate clinically relevant metadata into the learning process, we introduce a meta loss that ensures the model effectively captures both continuous and categorical patient attributes:

$$\mathcal{L}_{\text{meta}} = \text{RMSE}(\hat{y}_{\text{age}}, y_{\text{age}}) + \text{CE}(\hat{y}_{\text{gender}}, y_{\text{gender}}), \quad (1)$$

where RMSE is the root-mean-square error, applied to enforce accurate age prediction and CE is the cross-entropy loss used for gender classification.

All loss terms are combined into a weighted sum, where λ_{recon} , λ_{consis} , λ_{meta} are fixed scaling factors: $\mathcal{L} = \lambda_{\text{recon}}\mathcal{L}_{\text{recon}} + \lambda_{\text{consis}}\mathcal{L}_{\text{consis}} + \lambda_{\text{meta}}\mathcal{L}_{\text{meta}}$.

Table 1. Comparison of various methods on disease prediction and biomarker prediction on an in-house dataset. The best AUROC (ROC) and AUPRC (PRC) scores are in **bold**, while the second-best scores are underlined. Our proposed method is highlighted in a pastel color.

Method	Disease Prediction						Biomarker Prediction			
	DR		Glaucoma		AMD		CAC		eGFR	
	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC
<i>train/test on in-house</i>										
RETFound (ViT-S)	0.871	0.586	0.531	0.508	0.488	0.499	0.783	0.623	0.614	0.598
PRETI (ViT-S)	0.977	0.866	0.717	0.582	0.751	0.653	0.898	0.732	0.687	0.642
<i>test on in-house</i>										
MAE (ViT-B) [16]	0.965	0.818	0.697	0.590	<u>0.729</u>	0.616	0.805	0.620	0.504	0.497
RETFound (ViT-L) [40]	0.969	0.838	<u>0.747</u>	0.587	0.695	0.608	0.890	0.696	0.582	0.566
UrFound (ViT-B) [36]	<u>0.971</u>	0.910	0.691	<u>0.591</u>	0.700	<u>0.625</u>	0.860	0.689	0.572	0.556
RET-CLIP (ViT-B) [13]	0.857	0.579	0.642	0.547	0.516	0.504	0.795	0.630	0.522	0.543
PRETI (ViT-B)	0.982	<u>0.885</u>	0.777	0.622	0.772	0.659	0.899	0.734	0.699	0.643

3 Experiments

We train our foundation model PRETI on a large-scale dataset and then evaluate it through downstream classification tasks, with minimal additional training.

Datasets. PRETI was trained on an in-house dataset comprising 1,017,549 CFPs from 292,006 patients across six medical institutions. This dataset includes non-public data and UK Biobank, which was internally processed for consistency; we thus refer to the entire dataset as in-house. For downstream tasks, we evaluate disease prediction for Diabetic Retinopathy (DR), Glaucoma, and Age-Related Macular Degeneration (AMD), as well as biomarker prediction such as Coronary Artery Calcium (CAC) score and estimated Glomerular Filtration Rate (eGFR) using an in-house dataset. These biomarkers were selected due to their known associations with retinal features and strong correlation with demographic metadata such as age and sex. These downstream datasets contain 1,110 to 6,861 CFPs per disease or biomarker, separate from pre-training data, with a 70:15:15 train:val:test split. Additionally, we conduct disease prediction tasks using public datasets (IDRID [29], APTOS2019 [20], MESSIDOR-2 [9], GF [1], PAPILA [23], Retina [19], and JSIEC [6]).

Training Details. The images are resized to 224×224 and augmented using random crop and flip, with different augmentation ratios applied to each eye. In our pre-training setup, we use a Vision Transformer (ViT) architecture [12], with ViT-S or ViT-B for the encoder and ViT-S for the decoder. The encoder is initialized with ImageNet-1K pre-trained weights [10]. We train with a batch size of 512 per GPU on 8 H100 GPUs for 300 epochs (15 warm-up). The learning rate is $5e-5$. We set the loss function weights as $\lambda_{\text{recon}} = 1.4$, $\lambda_{\text{meta}} = 0.2$, and $\lambda_{\text{consis}} = 0.4$ by default. In reconstruction loss, the original MAE loss weight is 1.0, and the perceptual loss weight is 0.4. We also use RAAM with making ratios

Table 2. Performance comparison across different datasets and foundation models on public datasets.

Method	DR				Glaucoma				Multicategory			
	APTOS2019		MESSIDOR-2		GF		PAPILA		Retina		JSIEC	
	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC
<i>Report Supervision</i>												
FLAIR [31]	0.932	0.686	0.819	0.483	0.905	0.792	0.752	0.610	0.863	0.679	0.917	0.704
UrFound [36]	0.949	0.716	0.882	0.608	0.958	0.880	0.783	0.625	0.901	0.793	0.995	0.923
RET-CLIP [13]	0.951	0.748	-	-	0.958	0.889	0.853	0.754	0.942	0.871	0.999	0.972
<i>Non-report Supervision</i>												
MAE [16]	0.941	0.676	0.842	0.488	0.931	0.832	0.629	0.475	0.742	0.537	0.985	0.818
RETFound [40]	0.943	0.726	0.864	0.568	0.943	0.863	0.855	0.748	0.847	0.697	0.990	0.884
PRETI (ViT-S)	0.948	0.725	0.862	0.600	0.955	0.888	0.838	0.714	0.831	0.656	0.990	0.853
PRETI (ViT-B)	0.948	0.743	0.866	0.612	0.955	0.886	0.879	0.764	0.893	0.771	0.994	0.912

of $r_0 = 0.985$ and $r_T = 0.85$. Downstream tasks with our pre-trained encoder run for 50 epochs (batch size 16), saving the best validation model for testing.

3.1 Evaluation Results

We assess PRETI against state-of-the-art retinal foundation models, using in-house data (disease/biomarker prediction) and benchmarks (disease prediction). **Disease/Biomark Prediction.** We evaluate PRETI and baselines on internal data for disease and biomarker prediction (Table 1). For fair comparison based on training data, we also train RETFound on the same data as PRETI and evaluate its performance.

As a result, PRETI outperforms baselines in disease prediction tasks and maintains strong performance regardless of encoder size (ViT-S, B). Notably, it surpasses existing foundation models in biomarker prediction tasks, demonstrating its versatility. By effectively integrating metadata, PRETI leverages patient-specific factors, ensuring robustness across various tasks.

Benchmarks. PRETI is evaluated on disease classification benchmarks consisting of seven public datasets. Unlike report-supervision models such as FLAIR [31], UrFound [36], and RET-CLIP [13], which are trained using clinical descriptions, our model uses only images and metadata. Despite this, PRETI performs comparably to or outperforms baselines, demonstrating generalizability across diverse datasets. These evaluations shows the potential of our foundation models using only metadata, even when clinical text data is unavailable.

3.2 Ablation Study

Effect of LME & RAAM. We conduct our experiments by adding RAAM and LME to the SiamMAE on public datasets (Table 3). In most datasets, LME and RAAM each strictly improved AUROC and AUPRC, showing significant effects on MESSIDOR-2, PAPILA, Retina, and JSIEC datasets. Notably, PRETI, which

Table 3. Comparative analysis of LME and RAAM in the our PRETI model.

Method	DR				Glaucoma				Multicategory			
	APTOS2019		MESSIDOR-2		GF		PAPILA		Retina		JSIEC	
	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC
Baseline	0.935	0.692	0.846	0.495	0.940	0.844	0.813	0.690	0.832	0.675	0.976	0.788
+ RAAM	0.943	0.713	0.864	0.585	0.950	0.875	0.845	0.736	0.861	0.717	0.989	0.857
+ LME	0.943	0.724	0.851	0.534	0.953	0.883	0.864	0.740	0.871	0.718	0.989	0.875
PRETI	0.948	0.743	0.866	0.612	0.955	0.886	0.879	0.764	0.893	0.771	0.994	0.912

incorporates both RAAM and LME, achieves the highest performance due to retinal area awareness and the utilization of metadata.

Impact of LME through Attention Maps. To examine how PRETI utilizes metadata visually, we analyze attention maps. Specifically, we extract and visualize the attention maps of age and gender tokens from PRETI, along with the CLS token attention maps from both the baseline model without LME and PRETI with LME, using the APTOS2019 and GF datasets (Fig. 2). In the first row, featuring a retinal image of a patient with moderate diabetic retinopathy (DR), the attention maps of the age and gender tokens in the LME-enabled model precisely highlight structures and features critical for DR diagnosis. Notably, these metadata attention maps capture not only the macula and vessels but also drusen (yellow circles) and microaneurysms (pink circles), which the model without LME fails to detect but is important for DR [37]. This demonstrates how metadata helps the model effectively recognize key structures essential for DR prediction. Similarly, in the second row, for early glaucoma, age and gender embeddings of the LME-enabled model sharply focus on the optic disc (green circles) and the optic cup within it, both of which are critical for glaucoma diagnosis [2]. Furthermore, in both cases, the CLS token in the LME-enabled model, influenced by metadata embeddings, captures more accurate disease-relevant information for DR and glaucoma, demonstrating that metadata enhances the model’s ability to target critical diagnostic features.

4 Conclusion

Our PRETI, a metadata-driven retinal foundation model that leverages LME and RAAM, achieved state-of-the-art performance in disease and biomarker prediction tasks. Ultimately, PRETI serves as a patient-specific, metadata-aware foundation model, improving the generalizability by learning ocular disease-related features. Using metadata demonstrated effectiveness for retinal analysis, suggesting its potential for broader applications. While our study focused on age and gender, future work could explore additional metadata types and extend our methodology to other medical imaging domains.

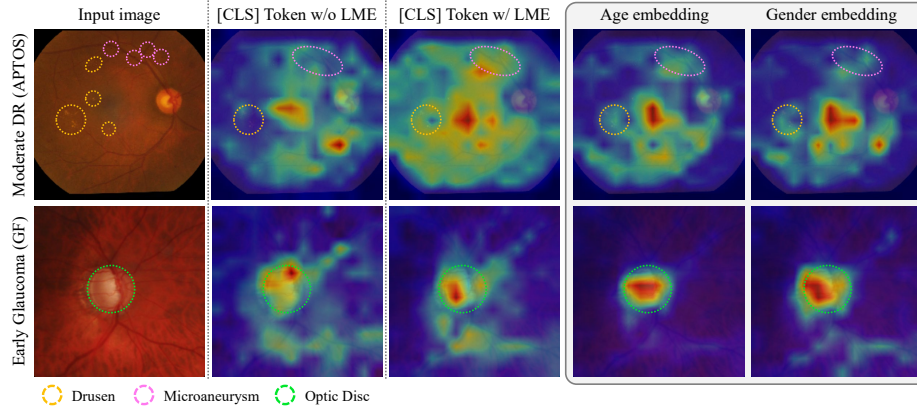


Fig. 2. Visualization of attention maps for the [CLS] token and metadata embeddings, comparing focus with and without LME. These maps highlight regions key to diagnosing eye diseases like diabetic retinopathy and glaucoma.

Acknowledgements. This work was supported in part by the IITP RS-2024-00457882 (AI Research Hub Project), IITP 2020-II201361, NRF RS-2024-00345806, NRF RS-2023-00219019, and NRF RS-2023-002620.

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

References

1. Ahn, J.M., et al.: A deep learning model for the detection of both advanced and early glaucoma using fundus photography. *PloS one* **13**(11), e0207982 (2018)
2. Almazroa, A., et al.: Optic disc and optic cup segmentation methodologies for glaucoma image detection: A survey. *Journal of Ophthalmology* **2015**, 180972 (2015). <https://doi.org/10.1155/2015/180972>
3. Bai, L., Chen, S., Gao, M., Abdelrahman, L., Ghamdi, M.A., Abdel-Mottaleb, M.: The influence of age and gender information on the diagnosis of diabetic retinopathy: Based on neural networks. *Proceedings of the 43rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* **2021**, 3514–3517 (2021). <https://doi.org/10.1109/EMBC46164.2021.9630601>
4. Cai, G., et al.: A multimodal transformer to fuse images and metadata for skin disease classification. *Vis. Comput.* **39**(7), 2781–2793 (2022)
5. Caron, M., et al.: Emerging properties in self-supervised vision transformers. In: *International conference on computer vision*. pp. 9650–9660 (2021)
6. Cen, L.P., et al.: Automatic detection of 39 fundus diseases and conditions in retinal photographs using deep neural networks. *Nature communications* **12**(1), 4828 (2021)
7. Cheung, C.Y., et al.: A deep-learning system for the assessment of cardiovascular disease risk via the measurement of retinal-vessel calibre. *Nature biomedical engineering* **5**(6), 498–508 (2021)

8. Ciocarlan, A., et al.: Self-supervised learning for real-world object detection: a survey. *arXiv preprint arXiv:2410.07442* (2024)
9. Decencière, E., et al.: Feedback on a publicly distributed image database: the messidor database. *Image Analysis & Stereology* pp. 231–234 (2014)
10. Deng, J., et al.: Imagenet: A large-scale hierarchical image database. In: *Conference on computer vision and pattern recognition*. pp. 248–255. Ieee (2009)
11. Dettmers, T., et al.: Qlora: Efficient finetuning of quantized llms. *Advances in neural information processing systems* **36**, 10088–10115 (2023)
12. Dosovitskiy, A., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
13. Du, J., et al.: Ret-clip: A retinal image foundation model pre-trained with clinical diagnostic reports. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 709–719. Springer (2024)
14. Dufumier, B., et al.: Contrastive learning with continuous proxy meta-data for 3d mri classification. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 58–68. Springer (2021)
15. Gupta, A., et al.: Siamese masked autoencoders. *Advances in Neural Information Processing Systems* **36**, 40676–40693 (2023)
16. He, K., et al.: Masked autoencoders are scalable vision learners. In: *Conference on computer vision and pattern recognition*. pp. 16000–16009 (2022)
17. Holland, R., et al.: Metadata-enhanced contrastive learning from retinal optical coherence tomography images. *Medical Image Analysis* **97**, 103296 (2024)
18. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: *European conference on computer vision*. pp. 694–711. Springer (2016)
19. jr2ngb: Cataract. <https://www.kaggle.com/datasets/jr2ngb/cataractdataset> (2019), kaggle
20. Karthik, Maggie, Dane, S.: Aptos 2019 blindness detection. <https://kaggle.com/competitions/aptos2019-blindness-detection> (2019), kaggle
21. Kirillov, A., et al.: Segment anything. In: *International conference on computer vision*. pp. 4015–4026 (2023)
22. Korpole, N.R., et al.: Gender difference in ocular diseases, risk factors and management with specific reference to role of sex steroid hormones. *Journal of mid-life health* **13**(1), 20–25 (2022). https://doi.org/10.4103/jmh.jmh_28_22
23. Kovalyk, O., et al.: Papila: Dataset with fundus images and clinical data of both eyes of the same patient for glaucoma assessment. *Scientific Data* **9**(1), 291 (2022)
24. Liu, H., et al.: Visual instruction tuning. *Advances in neural information processing systems* **36**, 34892–34916 (2023)
25. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983* (2016)
26. Moor, M., et al.: Foundation models for generalist medical artificial intelligence. *Nature* **616**(7956), 259–265 (2023)
27. Nam, J.W., et al.: Clinical evaluation of unilateral open-angle glaucoma: a two-year follow-up study. *Chonnam Medical Journal* **57**(2), 144 (2021)
28. Paulus, J.K., et al.: Field synopsis of sex in clinical prediction models for cardiovascular disease. *Circulation: Cardiovascular Quality and Outcomes* **9**(2 Suppl 1), S8–S15 (2016). <https://doi.org/10.1161/CIRCOUTCOMES.115.002473>
29. Porwal, P., et al.: Indian diabetic retinopathy image dataset (idrid): a database for diabetic retinopathy screening research. *Data* **3**(3), 25 (2018)

30. Radford, A., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
31. Silva-Rodriguez, J., et al.: A foundation language-image model of the retina (flair): Encoding expert knowledge in text supervision. *Medical Image Analysis* **99**, 103357 (2025)
32. Speilburg, A.M., et al.: Symmetry of peripheral retinal nonperfusion in diabetic retinopathy by ischemic index. *Journal of optometry* **11**(4), 262–267 (2018)
33. Tillmann, A., et al.: Gender differences in retinal diseases: A review. *Clinical & Experimental Ophthalmology* **52**(3), 317–333 (2024). <https://doi.org/10.1111/ceo.14364>
34. Trivizki, O., et al.: Symmetry of macular fundus features in age-related macular degeneration. *Ophthalmology retina* **7**(8), 672–682 (2023)
35. Wagner, S.K., et al.: Insights into systemic disease through retinal imaging-based oculomics. *Translational vision science & technology* **9**(2), 6–6 (2020)
36. Yu, K., et al.: Urfound: Towards universal retinal foundation models via knowledge-guided masked modeling. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 753–762. Springer (2024)
37. Z, Y., et al.: The association between diabetic retinopathy and the prevalence of age-related macular degeneration - the kailuan eye study. *Frontiers in Public Health* **10**, 922289 (2022). <https://doi.org/10.3389/fpubh.2022.922289>
38. Zhang, R., et al.: Llama-adapter: Efficient fine-tuning of language models with zero-init attention. arXiv preprint arXiv:2303.16199 (2023)
39. Zhou, Y., et al.: Automorph: Automated retinal vascular morphology quantification via a deep learning pipeline. *Translational vision science & technology* **11**(7), 12–12 (2022)
40. Zhou, Y., et al.: A foundation model for generalizable disease detection from retinal images. *Nature* **622**(7981), 156–163 (2023)