



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

RET-CLIP: A Retinal Image Foundation Model Pre-trained with Clinical Diagnostic Reports

Jiawei Du¹, Jia Guo¹, Weihang Zhang¹, Shengzhu Yang¹, Hanruo Liu^{1,2,3},
Huiqi Li^{1*}, and Ningli Wang⁴

¹ Beijing Institute of Technology, Beijing, China

² Beijing Tongren Hospital, Beijing, China

³ Capital Medical University, Beijing, China

⁴ Beijing Institute of Ophthalmology, Beijing, China
{jiaweidu, huiqili}@bit.edu.cn

Abstract. The Vision-Language Foundation model is increasingly investigated in the fields of computer vision and natural language processing, yet its exploration in ophthalmology and broader medical applications remains limited. The challenge is the lack of labeled data for the training of foundation model. To handle this issue, a CLIP-style retinal image foundation model is developed in this paper. Our foundation model, RET-CLIP, is specifically trained on a dataset of 193,865 patients to extract general features of color fundus photographs (CFPs), employing a tripartite optimization strategy to focus on left eye, right eye, and patient level to reflect real-world clinical scenarios. Extensive experiments demonstrate that RET-CLIP outperforms existing benchmarks across eight diverse datasets spanning four critical diagnostic categories: diabetic retinopathy, glaucoma, multiple disease diagnosis, and multi-label classification of multiple diseases, which demonstrate the performance and generality of our foundation model. The source code and pre-trained model are available at <https://github.com/sStonemason/RET-CLIP>.

Keywords: Vision-Language Pre-training · Foundation Model · Retinal Fundus Image.

1 Introduction

Foundation models trained on large-scale, multi-task datasets are now becoming increasingly popular and have achieved success in the fields of computer vision and natural language processing. Foundation models excel in generalization in feature extraction, offering significant potential for addressing the complex challenges of clinical applications. However, the development of medical foundation models is still in its nascent phase, primarily hindered by the lack of high-quality data and concerns around patient privacy. Although initial efforts have been

* Corresponding author.

made [23, 19, 24, 9, 12, 6, 11], the effectiveness of these models, particularly in analyzing retina fundus images, has yet to meet expectations, underscoring the urgent need for focused advancements in this area.

In the clinical diagnosis and treatment of ocular diseases, medical imaging, such as color fundus photography (CFP), and the detailed image interpretations and diagnostic reports written by professional ophthalmologists are indispensable. This makes the clinics of ophthalmology inherently rich in image-text multi-modality data, which holds significant potential for enhancing clinical applications. RETFound [25] is a foundation model for retinal images based on self-supervised learning. However, it solely utilizes image data and overlooks the equally vast amount of clinical diagnostic text. To address this limitation, CLIP [17], a powerful vision-language self-supervised paradigm, is widely explored in foundation models. By aligning the information of image and text in a shared representation space using a large corpus of image-text pairs, CLIP-style models can understand and associate visual content with natural language information. This results in feature representations with stronger generalization capabilities. Many studies focus on training vision-text models in the medical field [23, 19, 9, 22, 18, 7, 20, 2]. PMC-CLIP [9] collects image-description pairs from large amount of scientific documents and trains a CLIP-style model based on them. FLAIR [18] is a pre-trained vision-language model designed to understand retinal fundus images. The textual data utilized in such research often comes from captions in medical papers or through the manual annotation of simple labels. However, clinical diagnostic reports, rich in valuable textual information, remain underutilized in this context.

Moreover, the conventional approaches often involve treating CFPs of individual eyes as separate entities during model training. This necessitates the extraction of information corresponding to each eye from the original clinical diagnostic reports, which may not always clearly differentiate between left and right eyes. The manual processing involved in this procedure requires specialized knowledge and could introduce errors and increase costs significantly due to the potential for human-induced noise. Conversely, considering both eyes of a patient together provides a more holistic and clinically meaningful approach in clinical scenarios.

To alleviate the above issues, we have the following contributions in this paper: Firstly, we propose a vision-language foundation model for CFPs, named RET-CLIP, which we believe is the first attempt to leverage clinical diagnostic reports to build a retinal foundation model, enriching the model’s visual encoding capabilities with practicality and authenticity. The diagnostic reports in Chinese are included, extending the linguistic versatility of the research domain beyond English. Secondly, a novel strategy is proposed to decouple the information of left and right eyes in diagnostic reports, which is a simple yet effective paradigm for building a retinal foundation model. In practical scenarios, diagnostic reports are usually patient-level, mixing information from both eyes, which brings a big challenge for directly using CLIP to build foundation models. The proposed monocular and patient-level contrastive learning approach can

handle this challenge in the ophthalmology domain. Lastly, our model achieves state-of-the-art performance across diverse tasks and datasets, confirming the effectiveness of the proposed training strategy.

2 Method

2.1 Data Collection and Preprocessing

Dataset acquisition. We collected a dataset of retina fundus binocular images-text triplets (RET-Clinical) at the patient level for RET-CLIP. The dataset includes a total of 193,865 samples from Beijing Tongren Hospital, Beijing, China. Each patient’s triplet includes two CFPs for left and right eyes, alongside a clinical diagnostic report.

Data preprocessing and augmentation. For the CFPs, all of them are resized to 512×512 . The augmentation includes random crop followed by resizing to 224×224 , random horizontal flipping, color jitter, and image normalization. For diagnostic reports, we focus on correcting typos and consecutive punctuation errors caused by human input, restoring abbreviations to their full expressions, unifying mixed Chinese and English expressions into Chinese to align with our text encoder’s language capabilities, and ensuring the text is coherent and grammatically correct by manual scrutiny. It’s important to highlight that the preprocessing of text data only involves basic text standardization mentioned above, avoiding the need for advanced clinical knowledge or modifications that may alter the original content or meaning.

2.2 Model Architecture

As shown in Figure 1, we trained a Visual-Language model called RET-CLIP under the CLIP paradigm using our constructed binocular images-text triplets. RET-CLIP consists of a visual encoder and a text encoder, which extract image features from CFPs and text features from clinical diagnostic reports, respectively. During pre-training, image-text contrastive learning is performed at the monocular and patient level jointly. Patient level examines data features from a holistic patient perspective, effectively leveraging the information in raw data while minimizing the interference of manual preprocessing in the pre-training phase. Concurrently, the binocular level guides the model towards acquiring finer-grained features than the patient level. Combined together, these methodologies can improve RET-CLIP’s performance.

Given a mini-batch containing N binocular images-text triplets (i.e., N patients), $\mathcal{D} = \{(\mathcal{I}_1^l, \mathcal{I}_1^r, \mathcal{T}_1), \dots, (\mathcal{I}_N^l, \mathcal{I}_N^r, \mathcal{T}_N)\}$, where \mathcal{I}_i^l , \mathcal{I}_i^r and \mathcal{T}_i represents the CFP of left eye, the CFP of right eye and the diagnostic report of the i th patient, respectively. The visual encoder takes \mathcal{I}_i^l and \mathcal{I}_i^r as input, while the text encoder is fed with \mathcal{T}_i .

Visual encoder. The left and right $(\mathcal{I}^l, \mathcal{I}^r)$ CFPs for a patient are encoded to the embedding dimension of d using a ViT-based [5] encoder $\Phi_v(\cdot)$ respectively:

$$\mathbf{V}^l = \Phi_v(\mathcal{I}^l) \in \mathbb{R}^d, \mathbf{V}^r = \Phi_v(\mathcal{I}^r) \in \mathbb{R}^d. \quad (1)$$

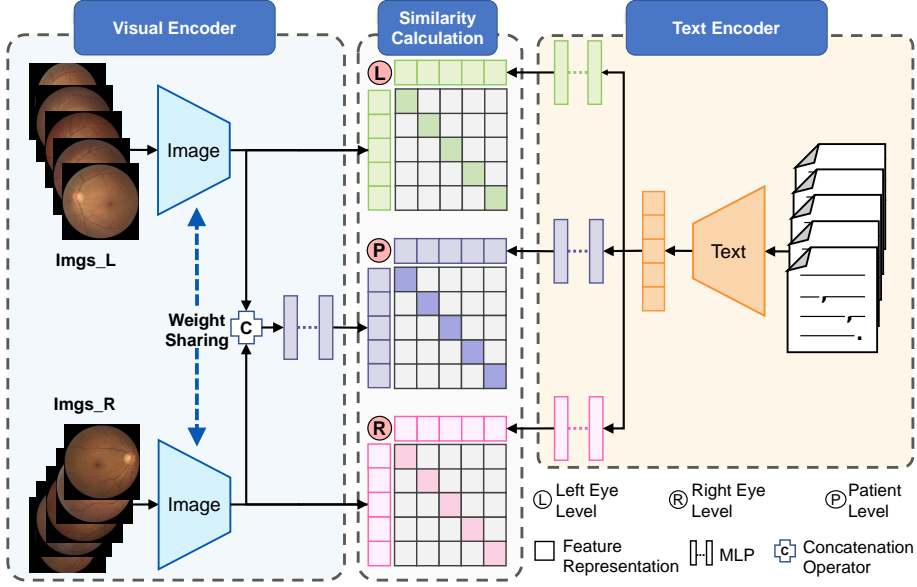


Fig. 1. Overview of the RET-CLIP foundation model.

where \mathbf{V}^l and \mathbf{V}^r represent the image features of the left and right eye, respectively. Next, concatenation and a simple Multilayer Perceptron (MLP) $F_v(\cdot)$ are employed to merge the image features of left and right eyes to derive comprehensive patient-level image features:

$$\mathbf{V}^p = F_v(\mathbf{V}^l \oplus \mathbf{V}^r) \in \mathbb{R}^d, \quad (2)$$

where \oplus denotes concatenation.

Text encoder. For a given patient’s diagnostic report \mathcal{T} , a BERT-based [4] encoder $\Phi_t(\cdot)$ is implemented to encode the clinical descriptions with a text token of length l :

$$\mathbf{T} = \Phi_t(\mathcal{T}) \in \mathbb{R}^{l \times d}, \quad \mathbf{T}_0 \in \mathbb{R}^d, \quad (3)$$

where \mathbf{T} denotes the sentence embedding, \mathbf{T}_0 denotes the embedding for [CLS] token. We then implement three stacked two-layer nonlinear MLPs $F_l(\cdot)$, $F_r(\cdot)$, $F_p(\cdot)$ to decouple \mathbf{T}_0 into textual features representing the left eye, right eye, and patient level, termed as \mathbf{T}^l , \mathbf{T}^r , and \mathbf{T}^p , respectively:

$$\mathbf{T}^l = F_l(\mathbf{T}_0), \quad \mathbf{T}^r = F_r(\mathbf{T}_0), \quad \mathbf{T}^p = F_p(\mathbf{T}_0), \quad \mathbf{T}^l, \mathbf{T}^r, \mathbf{T}^p \in \mathbb{R}^d. \quad (4)$$

2.3 Training Objective

For the provided mini-batch, termed as \mathcal{D} , the extracted feature set \mathcal{F} , which is $\{(\mathbf{V}_1^l, \mathbf{V}_1^r, \mathbf{V}_1^p, \mathbf{T}_1^l, \mathbf{T}_1^r, \mathbf{T}_1^p), \dots, (\mathbf{V}_N^l, \mathbf{V}_N^r, \mathbf{V}_N^p, \mathbf{T}_N^l, \mathbf{T}_N^r, \mathbf{T}_N^p)\}$, is

then divided into three subsets: $\mathcal{F}^l = \{(\mathbf{V}_1^l, \mathbf{T}_1^l), \dots, (\mathbf{V}_N^l, \mathbf{T}_N^l)\}$, $\mathcal{F}^r = \{(\mathbf{V}_1^r, \mathbf{T}_1^r), \dots, (\mathbf{V}_N^r, \mathbf{T}_N^r)\}$, and $\mathcal{F}^p = \{(\mathbf{V}_1^p, \mathbf{T}_1^p), \dots, (\mathbf{V}_N^p, \mathbf{T}_N^p)\}$, corresponding to left eye, right eye, and patient level, respectively. The image and text features of the same patient in each subset are positive samples of each other, while the rest are negative samples. The cosine similarity matrix is calculated on each subset.

For the subset of left eye features, we obtain the image feature matrix $\mathbf{V}^l = (\mathbf{V}_1^l, \dots, \mathbf{V}_N^l) \in \mathbb{R}^{N \times d}$ and the text feature matrix $\mathbf{T}^l = (\mathbf{T}_1^l, \dots, \mathbf{T}_N^l) \in \mathbb{R}^{N \times d}$. We measure the inter-sample similarity, termed as \mathbf{P}^{v2t} and \mathbf{P}^{t2v} , using the cosine distance $S(\cdot)$:

$$\mathbf{P}^{v2t} = S(\mathbf{V}^l, \mathbf{T}^l) \in \mathbb{R}^{N \times N}, \mathbf{P}^{t2v} = S(\mathbf{T}^l, \mathbf{V}^l) \in \mathbb{R}^{N \times N}. \quad (5)$$

Then we calculate the contrastive loss of the left eye:

$$\mathcal{L}_l = \frac{1}{2}(\text{CE}(\mathbf{P}^{v2t}, \mathbf{Y}^{v2t}) + \text{CE}(\mathbf{P}^{t2v}, \mathbf{Y}^{t2v})), \quad (6)$$

where \mathbf{Y}^{v2t} and \mathbf{Y}^{t2v} represent the one-hot labels, CE refers to InfoNCE loss [13].

Then we calculate \mathcal{L}_r and \mathcal{L}_p for right eye and patient level based on \mathcal{F}^r and \mathcal{F}^p in the same way. The final loss is the sum of the above three:

$$\mathcal{L} = \mathcal{L}_l + \mathcal{L}_r + \mathcal{L}_p. \quad (7)$$

2.4 Implementation

The vision encoder utilizes the base-sized version of the vision transformer (ViT-base) [5], while the text encoder employs the base-sized version of RoBERTa (RoBERTa-base) [10], both are initialized with the Chinese-CLIP weights [21]. AdamW is used as the optimizer. The batch size is 256, and training is performed using NVIDIA GeForce RTX 4090. The training process consists of 10 epochs, with the first 50 steps dedicated to warming up the model (from 0 to a learning rate of 3×10^{-5}).

3 Experiments

3.1 Tasks and Datasets

We focus on designing downstream evaluation experiments primarily for visual tasks. These tasks contain four main categories: diagnosis of diabetic retinopathy, glaucoma, multiple diseases, and multi-label classification of multiple diseases.

For diabetic retinopathy diagnosis, IDRID [16] and APTOS-2019 (<https://www.kaggle.com/competitions/aptos2019-blindness-detection/data>) are used. The labels for diabetic retinopathy are no, mild, moderate, severe, and proliferative retinopathy. The IDRID dataset comprises 516 images, while the APTOS dataset contains 3662 images.

For glaucoma diagnosis, PAPILA [8] (488 images in total) and Glaucoma Fundus [1] (1544 images in total) are used. They both have three categorical labels, non-glaucoma, suspected glaucoma (early glaucoma), and glaucoma (advanced glaucoma).

For multiple disease diagnosis, JSIEC [3] (1000 in total) and Retina (<https://www.kaggle.com/datasets/jr2ngb/cataractdataset>) (601 in total) are tested. JSIEC contains 39 categories of common referable fundus diseases and conditions. Retina includes labels for normal, glaucoma, cataract, and other retinal diseases.

For multi-label classification of multiple diseases, RFMID [15] and ODIR (<https://odir2019.grand-challenge.org/>) are tested. RFMID includes 3200 images with 28 categories of common referable fundus diseases and conditions. ODIR includes 10000 images (5000 patients’ paired left and right eyes) with labels of normal, diabetic retinopathy, glaucoma, cataract, age-related macular degeneration (AMD), hypertension, myopia, and other diseases.

For the IDRIR, the entire dataset is officially divided into a test set comprising 20% of the data, with the remaining 80% designated as the training set. In our experiments, we further split the training set into a training set and a validation set using a 4:1 ratio. Similarly, for the PAPLA, we follow the official partitioning method, which aligns with the approach described above. Regarding the RFMID, the official division includes distinct sets for training, validation, and testing; we adhere to this official partitioning. For all other datasets, we divide them into training, validation, and test sets using a 0.56:0.14:0.3 ratio, following RETFound’s [25] partitioning method. For all datasets, samples within each category are initially distributed based on the specified proportions before being combined to ensure consistent category distribution across the training, validation, and test sets.

When adapting to downstream tasks, the input image is mapped to a high-level feature representation by the visual encoder. A simple linear prediction head is then applied, followed by a Sigmoid or Softmax layer to achieve classification.

For each task, two adaptation methods are implemented: linear probing, training the classifier only with the encoder frozen, and fine-tuning, where both the encoder and classifier are trained. Each evaluation process consists of 50 epochs with a batch size of 16. The model weights with the best performance on the validation set are saved for testing.

3.2 Comparison Methods and Evaluation Metrics

To demonstrate the superiority of our method, we compare two broad categories of models: foundation models trained on non-CFP datasets (Chinese-CLIP [21], PMC-CLIP [9], DINOv2 [14]) and models designed for CFP vision tasks (RET-Found [25], FLAIR [18]).

We use the area under the receiver operating curve (AUROC) and area under the precision-recall curve (AUPR) as the evaluation metrics. We evaluate five iterations with different random seeds for each model on each downstream dataset to calculate the mean values. We also conduct the t-test for each downstream

task to determine the significance level at which the top-performing method surpasses the others (see Supplementary Materials).

3.3 Result

RET-CLIP outperforms five comparison models across eight datasets (four categories) as introduced before, demonstrating strong generalization capabilities.

For linear probing, the results are shown in Table 1 and Table 2. RET-CLIP demonstrates superior performance on almost all datasets, which indicates that RET-CLIP has learned a rich feature representation during the pre-training phase, demonstrating the capability to capture high-quality features.

Table 1. Diabetic retinopathy and glaucoma diagnosis results for linear probing. The best results on each metric are highlighted in bold.

Models	IDRID		APTOS2019		PAPILA		Glaucoma Fundus	
	AUROC	AUPR	AUROC	AUPR	AUROC	AUPR	AUROC	AUPR
CN-CLIP [21]	0.633	0.336	0.806	0.429	0.658	0.473	0.863	0.716
PMC-CLIP [9]	0.585	0.303	0.756	0.368	0.773	0.603	0.899	0.780
DinoV2 [14]	0.748	0.463	0.783	0.432	0.740	0.556	0.891	0.746
RETFound [25]	0.665	0.368	0.745	0.370	0.620	0.511	0.899	0.773
FLAIR [18]	0.700	0.475	0.849	0.515	0.746	0.595	0.872	0.672
OURS	0.856	0.616	0.923	0.656	0.775	0.667	0.893	0.789

Table 2. Multiple disease diagnosis and multi-label classification of multiple diseases results for linear probing.

Models	JSIEC		Retina		RFMID		ODIR	
	AUROC	AUPR	AUROC	AUPR	AUROC	AUPR	AUROC	AUPR
CN-CLIP [21]	0.783	0.239	0.738	0.514	0.819	0.293	0.801	0.483
PMC-CLIP [9]	0.947	0.654	0.778	0.597	0.854	0.372	0.800	0.506
DinoV2 [14]	0.873	0.446	0.813	0.635	0.860	0.430	0.825	0.550
RETFound [25]	0.704	0.167	0.630	0.434	0.842	0.409	0.738	0.401
FLAIR [18]	0.843	0.304	0.773	0.557	0.773	0.254	0.858	0.531
OURS	0.982	0.855	0.935	0.864	0.925	0.552	0.902	0.682

For fine-tuning, as shown in Table 3 and Table 4, RET-CLIP demonstrates superior performance across nearly all tasks. This outcome substantiates RET-CLIP’s robust feature extraction and generalization capabilities. Furthermore, it suggests that RET-CLIP not only captures high-quality features but also exhibits strong adaptability, enabling effective customization for specific tasks.

It’s noteworthy that the previous foundation models designed for CFPs do not exhibit an advantage over models trained on non-CFP datasets. RETFound’s [25] image reconstruction-focused paradigm may prioritize features related to

Table 3. Diabetic retinopathy and glaucoma diagnosis results for fine-tuning.

Models	IDRID		APTOS2019		PAPILA		Glaucoma Fundus	
	AUROC	AUPR	AUROC	AUPR	AUROC	AUPR	AUROC	AUPR
CN-CLIP [21]	0.778	0.506	0.881	0.619	0.804	0.690	0.951	0.876
PMC-CLIP [9]	0.785	0.511	0.776	0.386	0.798	0.659	0.925	0.827
DinoV2 [14]	0.791	0.533	0.920	0.675	0.797	0.681	0.955	0.884
RETFound [25]	0.822	0.496	0.943	0.726	0.855	0.748	0.943	0.863
FLAIR [18]	0.795	0.529	0.932	0.686	0.752	0.610	0.905	0.792
OURS	0.863	0.630	0.951	0.748	0.853	0.754	0.958	0.889

Table 4. Multiple disease diagnosis and multi-label classification of multiple diseases results for fine-tuning.

Models	JSIEC		Retina		RFMID		ODIR	
	AUROC	AUPR	AUROC	AUPR	AUROC	AUPR	AUROC	AUPR
CN-CLIP [21]	0.992	0.882	0.839	0.691	0.901	0.480	0.859	0.598
PMC-CLIP [9]	0.964	0.738	0.875	0.742	0.894	0.456	0.819	0.542
DinoV2 [14]	0.996	0.918	0.893	0.771	0.914	0.547	0.867	0.621
RETFound [25]	0.990	0.884	0.847	0.697	0.889	0.489	0.850	0.620
FLAIR [18]	0.917	0.704	0.863	0.679	0.870	0.397	0.860	0.601
OURS	0.999	0.972	0.942	0.871	0.946	0.581	0.917	0.715

the rebuilding of CFP, which lack the granularity and quality needed for specific downstream tasks, hindering its broader applicability. FLAIR [18], while is a CLIP-style model, does not suit ophthalmic tasks as it uses the text provision method employed by the original CLIP [17], which is designed for natural contexts, offering limited textual insights from single labels. Moreover, its dependence on public datasets for training constrains its performance due to their limited scale and quality. In contrast, RET-CLIP leverages rich textual information from clinical reports to extract detailed features for ophthalmic tasks better, showcasing the benefits of integrating diagnostic reports into the training of medical CLIP-style models.

3.4 Ablation study

The results, as shown in Table 5, confirm the effectiveness of optimizing objectives at both monocular and patient levels. As previously discussed, the combination of the global information provided at the patient level with the finer-grained features contributed at the monocular level is essential to achieve optimal performance.

4 Conclusion

In this study, we compile a binocular images-text dataset, RET-Clinical, derived from 193,865 clinical patients, with which, we jointly optimize and pre-train a CLIP-style model, RET-CLIP, cooperating with the information of left eye,

Table 5. Results of ablation studies. Monocular-level loss refers to \mathcal{L}_l plus \mathcal{L}_r .

Monocular-level Loss Patient-level Loss	AUROC			AUPR		
	✓		✓	✓		✓
IDRID	0.863	0.860	0.847	0.63	0.623	0.619
APTOS-2019	0.951	0.945	0.941	0.748	0.737	0.759
PAPILA	0.853	0.864	0.846	0.754	0.745	0.739
Glaucoma Fundus	0.958	0.948	0.957	0.889	0.869	0.888
JSIEC	0.999	0.997	0.997	0.972	0.949	0.962
Retina	0.942	0.939	0.935	0.871	0.869	0.876
RFMID	0.946	0.924	0.940	0.581	0.573	0.578
ODIR	0.917	0.909	0.905	0.715	0.692	0.696

right eye, and patient level. RET-CLIP achieves state-of-the-art results across eight downstream tasks spanning four critical diagnostic categories. Our research narrows the existing void in ophthalmic vision-language models by integrating textual data from clinical diagnostic reports, thereby offering insights into the applicability of raw clinical texts in the wider medical domain.

Acknowledgments. This research work is supported by Beijing Natural Science Foundation (Grant No. IS23112), National Natural Science Foundation of China (NSFC) (GrantNo. 82072007), and Beijing Institute of Technology Research Fund Program for Young Scholars.

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., Kim, S., Ahn, K.S., Cho, S.H., Lee, K.B., Kim, U.S.: A deep learning model for the detection of both advanced and early glaucoma using fundus photography. *PloS one* **13**(11), e0207982 (2018)
2. Baliah, S., Maani, F.A., Sanjeev, S., Khan, M.H.: Exploring the transfer learning capabilities of clip in domain generalization for diabetic retinopathy. In: *International Workshop on Machine Learning in Medical Imaging*. pp. 444–453. Springer (2023)
3. Cen, L.P., Ji, J., Lin, J.W., Ju, S.T., Lin, H.J., Li, T.P., Wang, Y., Yang, J.F., Liu, Y.F., Tan, S., et al.: Automatic detection of 39 fundus diseases and conditions in retinal photographs using deep neural networks. *Nature communications* **12**(1), 4828 (2021)
4. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
5. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
6. Eslami, S., Meinel, C., De Melo, G.: Pubmedclip: How much does clip benefit visual question answering in the medical domain? In: *Findings of the Association for Computational Linguistics: EACL 2023*. pp. 1151–1163 (2023)
7. Huang, S.C., Shen, L., Lungren, M.P., Yeung, S.: Gloria: A multimodal global-local representation learning framework for label-efficient medical image recognition. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3942–3951 (2021)
8. Kovalyk, O., Morales-Sánchez, J., Verdú-Monedero, R., Sellés-Navarro, I., Palazón-Cabanes, A., Sancho-Gómez, J.L.: Papila: Dataset with fundus images and clinical data of both eyes of the same patient for glaucoma assessment. *Scientific Data* **9**(1), 291 (2022)
9. Lin, W., Zhao, Z., Zhang, X., Wu, C., Zhang, Y., Wang, Y., Xie, W.: Pmc-clip: Contrastive language-image pre-training using biomedical documents. *arXiv preprint arXiv:2303.07240* (2023)
10. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019)
11. Moor, M., Banerjee, O., Abad, Z.S.H., Krumholz, H.M., Leskovec, J., Topol, E.J., Rajpurkar, P.: Foundation models for generalist medical artificial intelligence. *Nature* **616**(7956), 259–265 (2023)
12. Moor, M., Huang, Q., Wu, S., Yasunaga, M., Dalmia, Y., Leskovec, J., Zakka, C., Reis, E.P., Rajpurkar, P.: Med-flamingo: a multimodal medical few-shot learner. In: *Machine Learning for Health (ML4H)*. pp. 353–367. PMLR (2023)
13. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018)
14. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193* (2023)
15. Pachade, S., Porwal, P., Thulkar, D., Kokare, M., Deshmukh, G., Sahasrabudhe, V., Giancardo, L., Quellec, G., Mériaudeau, F.: Retinal fundus multi-disease image

- dataset (rfmid): A dataset for multi-disease detection research. *Data* **6**(2), 14 (2021)
16. Porwal, P., Pachade, S., Kokare, M., Deshmukh, G., Son, J., Bae, W., Liu, L., Wang, J., Liu, X., Gao, L., et al.: Idrid: Diabetic retinopathy–segmentation and grading challenge. *Medical image analysis* **59**, 101561 (2020)
 17. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. PMLR (2021)
 18. Silva-Rodriguez, J., Chakor, H., Kobbi, R., Dolz, J., Ayed, I.B.: A foundation language-image model of the retina (flair): Encoding expert knowledge in text supervision. *arXiv preprint arXiv:2308.07898* (2023)
 19. Wang, Z., Wu, Z., Agarwal, D., Sun, J.: Medclip: Contrastive learning from unpaired medical images and text. *arXiv preprint arXiv:2210.10163* (2022)
 20. Wu, C., Zhang, X., Zhang, Y., Wang, Y., Xie, W.: Medklip: Medical knowledge enhanced language-image pre-training. *medRxiv pp. 2023–01* (2023)
 21. Yang, A., Pan, J., Lin, J., Men, R., Zhang, Y., Zhou, J., Zhou, C.: Chinese clip: Contrastive vision-language pretraining in chinese. *arXiv preprint arXiv:2211.01335* (2022)
 22. Zhang, X., Wu, C., Zhao, Z., Lin, W., Zhang, Y., Wang, Y., Xie, W.: Pmc-vqa: Visual instruction tuning for medical visual question answering. *arXiv preprint arXiv:2305.10415* (2023)
 23. Zhang, Y., Jiang, H., Miura, Y., Manning, C.D., Langlotz, C.P.: Contrastive Learning of Medical Visual Representations from Paired Images and Text (Sep 2022)
 24. Zhong, T., Zhao, W., Zhang, Y., Pan, Y., Dong, P., Jiang, Z., Kui, X., Shang, Y., Yang, L., Wei, Y., et al.: Chatradio-valuer: A chat large language model for generalizable radiology report generation based on multi-institution and multi-system data. *arXiv preprint arXiv:2310.05242* (2023)
 25. Zhou, Y., Chia, M.A., Wagner, S.K., Ayhan, M.S., Williamson, D.J., Struyven, R.R., Liu, T., Xu, M., Lozano, M.G., Woodward-Court, P., et al.: A foundation model for generalizable disease detection from retinal images. *Nature* **622**(7981), 156–163 (2023)