

G OPEN ACCESS

Citation: Li L, Verma M, Wang B, Nakashima Y, Nagahara H, Kawasaki R (2023) Automated grading system of retinal arterio-venous crossing patterns: A deep learning approach replicating ophthalmologist's diagnostic process of arteriolosclerosis. PLOS Digit Health 2(1): e0000174. https://doi.org/10.1371/journal. pdig.0000174

Editor: Dukyong Yoon, Yonsei University College of Medicine, KOREA, REPUBLIC OF

Received: June 20, 2022

Accepted: November 29, 2022

Published: January 11, 2023

Copyright: © 2023 Li et al. This is an open access article distributed under the terms of the <u>Creative</u> Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All codes are available on github (https://github.com/ conscienceli/MDTNet). Image data access is possible but limited due to the range of patients' consent. Please contact the Ethics Review Board of Osaka University Hospital (https://www.med. osaka-u.ac.jp/pub/hp-crc/person_concerned/index. html) for detailed information. RESEARCH ARTICLE

Automated grading system of retinal arteriovenous crossing patterns: A deep learning approach replicating ophthalmologist's diagnostic process of arteriolosclerosis

Liangzhi Li¹, Manisha Verma¹, Bowen Wang¹, Yuta Nakashima¹, Hajime Nagahara¹, Ryo Kawasaki²*

1 Institute for Datability Science (IDS), Osaka University, Osaka, Japan, 2 Graduate School of Medicine, Osaka University, Osaka, Japan

* ryo.kawasaki@ophthal.med.osaka-u.ac.jp

Abstract

The morphological feature of retinal arterio-venous crossing patterns is a valuable source of cardiovascular risk stratification as it directly captures vascular health. Although Scheie's classification, which was proposed in 1953, has been used to grade the severity of arteriolosclerosis as diagnostic criteria, it is not widely used in clinical settings as mastering this grading is challenging as it requires vast experience. In this paper, we propose a deep learning approach to replicate a diagnostic process of ophthalmologists while providing a checkpoint to secure explainability to understand the grading process. The proposed pipeline is threefold to replicate a diagnostic process of ophthalmologists. First, we adopt segmentation and classification models to automatically obtain vessels in a retinal image with the corresponding artery/vein labels and find candidate arterio-venous crossing points. Second, we use a classification model to validate the true crossing point. At last, the grade of severity for the vessel crossings is classified. To better address the problem of label ambiguity and imbalanced label distribution, we propose a new model, named multi-diagnosis team network (MDTNet), in which the sub-models with different structures or different loss functions provide different decisions. MDTNet unifies these diverse theories to give the final decision with high accuracy. Our automated grading pipeline was able to validate crossing points with precision and recall of 96.3% and 96.3%, respectively. Among correctly detected crossing points, the kappa value for the agreement between the grading by a retina specialist and the estimated score was 0.85, with an accuracy of 0.92. The numerical results demonstrate that our method can achieve a good performance in both arterio-venous crossing validation and severity grading tasks following the diagnostic process of ophthalmologists. By the proposed models, we could build a pipeline reproducing ophthalmologists' diagnostic process without requiring subjective feature extractions. The code is available (https://github.com/ conscienceli/MDTNet).

Funding: This work was supported by Council for Science, Technology and Innovation (CSTI) crossministerial Strategic Innovation Promotion Program (SIP) "Innovative AI Hospital System" (Funding Agency: National Institute of Biomedical Innovation, Health and Nutrition (NIBIOHN)) (LL, MV, HN, RK). This work was also supported by JSPS KAKENHI Grant Number 19K10662 (RK), 20K23343 (LL), and 21K17764 (LL). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Author Summary

Assessment of arterio-venous crossing points in retinal images provides rich cues for quick screening of arteriosclerosis and even for classifying them into different severity grades. Considering the ever-increasing demand for ophthalmologic examination, computer-aided diagnosis (CAD) is extremely helpful for quick screening. However, retinal image analysis for CAD is a challenging task due to the high complexity of the vessel system and huge visual differences among retinal images. To address the aforementioned problems, we propose a whole pipeline for an automatic method for severity grading of artery hardening. Our method can find and validate possible arterio-venous crossing points, for which the severity grade is predicted. We also design a new model, MDTNet, which uses the focal loss to address the problem of data ambiguity and unbalance. Therefore, we believe that this research contributes to the advancement of research in machine learning on retinal images.

Introduction

Retina provides a window to directly visualize vascular structure *in vivo*, and ophthalmologic examination has been regarded as an important routine for detecting not only eye diseases but also ocular manifestations of cardiovascular diseases or their accumulated risks [1]. Among these detectable retinal vascular signs, arteriolosclerosis is critical yet asymptomatic, of which diagnosis requires detailed retinal observation. It is not widely conducted in the modern medical practice as it depends on mostly subjective qualitative observations, and most importantly, it requires vast experiences.

Assessment of arterio-venous crossing points in retinal images provides rich cues for screening arteriosclerosis and for evaluating accumulated cardiovascular risks. Typically, arterio-venous crossing points are classified into severity grades [2]. The assessment is based on some diagnostic criteria, for example, Scheie's classification [3], as shown in Fig 1(b)-1(e). The grades are described as follows: (i) *none* (no anomaly observed); (ii) *mild* (slight shrink in the caliber at venular edges); (iii) *moderate* (narrowed caliber at a single venular edge); and (iv) *severe* (narrowed caliber at both venular edges).

However, human graders are subjective and usually with different levels of experience, and there has been a criticism of the low reproducibility of severity grading, which makes grading results from human graders unreliable for clinical practice, screening, and clinical trials [4]. Also, considering the ever-increasing demand for ophthalmologic examination, computer-aided diagnosis (CAD) is extremely helpful for quick screening. Yet, retinal image analysis for CAD is a challenging task due to the high complexity of the vessel system and huge visual differences among retinal images.

In fact, most researchers in this area have been focusing on preliminary tasks, such as vessel segmentation [5-7], artery/vein classification [8-10], etc. A few works address higher-level tasks [4, 11], mostly on top of vessel segmentation, such as vessel width measurement, vessel-to-vessel ratio calculation, etc. However, they usually struggle in actual diagnoses: Firstly, vessel segmentation in retinal images *per se* is a challenging task. The vessel maps in Fig 1(c)-1(e), which are produced by the state-of-the-art segmentation model [12], cannot capture such deformation. This may imply that deformation is too minor to be captured by segmentation models, although such kind of segmentation-based approach is a typical solution for automatic severity grading. Secondly, the existing methods detect arterio-venous crossing points by applying some morphological operators to vessel maps [13]. This approach may not be



Fig 1. Typical examples of our prediction targets. Images in the first and second rows are raw retinal patches and automatically-generated vessel maps with manually-annotated artery/vein labels, respectively. Red represents arteries while blue represents veins. (a) is false crossing (the vein runs above the artery), while (b)–(e) are for *none, mild, moderate*, and *severe* grades, respectively. Note that even the state-of-the-art segmentation techniques cannot capture caliber narrowing, therefore, the arterioloscleroses are not very obvious in the vessel maps.

https://doi.org/10.1371/journal.pdig.0000174.g001

accurate enough to find crossing points that satisfy diagnostic requirements. For example, we can only use crossing points at which the artery is above the vein for diagnosis, and Fig 1(a) is not a diagnostic crossing point since the artery goes below the vein.

Instead of fully relying on segmentation results, we propose a multi-stage approach, in which segmentation results are used only for finding crossing point candidates, and actual prediction of the severity grade is conducted for an image patch around each crossing point after validating if the crossing point is an actual and informative one. To the best of our knowledge, this is the first work proposing a fully-automatic methodology aiming at grading arteriolo-sclerosis through the joint detection and analysis of retinal crossings.

Another issue in our severity grading task, which is very common in medical imaging, is the imbalanced label distribution. Most patients in our dataset have the slightest signs (*none* and *mild*) of arteriolosclerosis while only a few patients suffer from the *severe* grades of artery hardening. Also, the boundaries among different severity labels are not always obvious, making accurate diagnosis challenging.

Inspired by the concept of the multidisciplinary team [14], which strives to make a comprehensive assessment of a patient, we propose a multi-diagnosis team network (MDTNet) in this paper to address the imbalanced label distribution and label ambiguity problems at the same time. MDTNet can combine the features from multiple classification models with different structures or different loss functions. Some of the underlying models in MDTNet use the class-balanced focal loss [15] to handle hard or rare samples, of which the original version requires hyperparameter tuning, while MDTNet can utilize the advantage of the focal loss without tuning its hyperparameters.

Our main contribution is two-fold: (i) We propose a whole pipeline for an automatic method for severity grading of artery hardening. Our method can find and validate possible arterio-venous crossing points, for which the severity grade is predicted. (ii) We design a new model, MDTNet, which uses the focal loss to address the problem of data ambiguity and unbalance.

Dataset

Ethics statement

This study was performed in accordance with the World Medical Association Declaration of Helsinki. Patients gave written informed consent to participate and the study protocol was approved by the institutional review board of the Osaka University Hospital.

We built a vessel crossing point dataset extracted from our retinal image database of the Ohasama study, a cohort to study cardiovascular diseases risk, where we could utilize 1, 440 images in the size of 5, 184×3 , 456 pixels, which are captured by the CR-2 AF Digital Non-Mydriatic Retinal Camera (Canon, Tokyo) between 2013 and 2017 as JPEG files. This database includes the medical data of 684 people, which are with an average age of 64.5 (standard deviation: 6.1). The ratio between female and male is 65.2% : 34.8% and 47.6% of all participants have hypertension. Details of the study profile were published elsewhere [16].

To find crossing points in these images (Fig 2(a)-2(d)), we used a segmentation model ([12]) to get vessel maps. We then classified each pixel on extracted vessels into artery/vein using [17]. We combine the vessel segmentation and classification results to find crossing points because classification results, which are more beneficial for crossing point detection, tend to have more errors while segmented vessel maps are more accurate. Therefore, we refine the classification results based on the vessel maps. A classic approach then finds crossing points in these refined artery/vein maps. Specifically, we find the artery pixels neighbouring vein pixels and check whether it is a crossing point or not using the skeletonized vessel map. The points marked in yellow in Fig 2 are detected crossing point candidates. Note that for cup zones as indicated by a pink circle and dot in Fig 2, we exclude candidates because the vessel system in this area is with high complexity and thus segmentation and classification are not reliable. Image patches are of size 150 × 150, centered at the crossing point candidates. Consequently, we detected 4, 240 crossing points and extracted corresponding image patches, centered at these crossing points.

Each image patch was carefully reviewed by a highly experienced ophthalmologist. Due to the errors in vessel segmentation and artery/vein classification, the detected crossing points may not be actual or informative. Therefore, the specialist first annotated each image patch with a label on its validity, *i.e.*, if the image patch contains an actual and informative crossing point (*true*) or not (*false*). The numbers of true and false crossing points are 2, 507 and 1, 733,



https://doi.org/10.1371/journal.pdig.0000174.g002

respectively. For each true crossing point, the specialist gave its severity label in $C = \{none, mild, moderate, severe\}$. The numbers of image patches with respective labels are 1, 177, 816, 457, and 57. In both tasks, the datasets will be divided into training, validation, and test sets following a ratio of 8:1:1. As an examinee may have multiple retinal images, it is important to strictly put them into one same subset to prevent the training data contamination.

Severity grading pipeline

Our method forms a pipeline with three main modules, *i.e.*, preprocessing, patch validation, and severity grade prediction. The whole pipeline is shown in Fig 2.

Preprocessing

Steps (a)–(d) in the figure are preprocessing, in which the same processes as our dataset construction are applied to get image patches of 150×150 pixels with crossing point candidates.

Crossing point validation

Both crossing point validation and severity grading are classification problems, whereas validation is easier because the label distribution is more balanced and the differences between real and false crossing points are more obvious. We find that commonly used classification models, such as [18–20], work well for our validation task (refer to *Experiments and Results* Section).

Severity grade prediction

The severity grade prediction task is much more challenging: Firstly, the label distribution is highly biased. For example, samples with the *none* label account for 68% of the total samples, while ones with the *severe* label only take up 3%. Secondly, the difference among samples with different labels may not be clear enough. Even medical doctors may make diverse decisions on a single image patch.

For such classification tasks with ambiguous or imbalanced classes, the focal loss [15] has been used, which makes a model more aware of hard samples than easy ones. The focal loss introduces a hyperparameter γ , on which a model's performance depends significantly. Tuning this hyperparameter is extremely important yet computationally expensive [21]. A greater γ may make the model focus too much on hard samples, spoiling the accuracy of other samples, while a smaller γ may decrease its ability to classify hard samples.

We propose a multi-diagnosis team network (MDTNet) to address the aforementioned problems in severity grade prediction. As shown in Fig 3, MDTNet consists of three modules, *i.e.*, a base module, a focal module, and a fusion module.

The base and focal modules have multiple sub-models, and all of them take the same image patch as input. The difference between the sub-models in the base and focal modules is the losses: Ones in the base module adopt the cross entropy (CE) loss while ones in the focal module use the focal loss. These sub-models are trained independently with respective losses. The fusion module concatenates all features (*i.e.*, the outputs of the second last layers of the sub-models) into a single vector, which is then fed into two fully-connected layers to make the final prediction.

The focal loss is originally designed for object detection [15], defined as

$$L(y,t) = -\sum_{l} t_{l} (1-y_{l})^{\gamma} \log y_{l}, \qquad (1)$$

where *t* is the one-hot representation of label and *y* is the softmax output from a model (t_l and y_l are the *l*-th entries of *t* and *y*); γ is a hyperparameter to weight hard examples. The focal loss



https://doi.org/10.1371/journal.pdig.0000174.g003

reduces to the CE loss when $\gamma = 0$, and a larger γ weights more on hard examples. One possible criticism of the focal loss is its sensitivity to γ . We therefore propose to ensemble sub-models with different γ 's. The hypothesis behind this choice is that different γ 's may rely on different cues for prediction and aggregating respective features may help in improving the final decision. This is embodied in the focal module. The same idea can also be applied to different network architectures, embodied in the base module. These sub-models thus provide diagnostic features that may complement each other.

To cope with the imbalanced class distribution, we adopt class weighting [22, 23]. We multiply weight $\alpha_l = \ln N_l / \ln N$ to each term (*i.e.* different *l*'s) in the CE/focal loss, where N and N_l are the numbers of all samples and of samples with the label corresponding to the *l*-th entry of *t*. We pre-train the sub-models using their own classifiers and losses, and then freeze their weights to train the additional two fully-connected layers for the final decision.

Data augmentation

We adopt extensive data augmentation. During the training process, the input images have 50% chance of getting each operator in Fig 4. Among them, $(b \sim h)$ are used for shape modification, changing the locations and the shapes of the attention areas of the deep learning models; $(i \sim k)$ are to provide variety on imaging quality by blurring or adding random noises; (l) represents sensor characteristics of color (hue and saturation).

Experiments and results

Implementation

For sub-models in the base module, we used ResNet [18], Inception [20], and DenseNet [19]. In the focal module, DenseNet with $\gamma = 1, 2$, or 3 were used. All these models are pre-trained



Fig 4. Our data augmentation operator pool. (a) Raw image, (b) vertical flipping, (c) horizontal flipping, (d) cropping and padding, (e) scaling, (f) translating, (g) rotating, (h) sheering, (i) blurring, (j) additional noise, (k) additional frequency noise, and (l) color modification.

https://doi.org/10.1371/journal.pdig.0000174.g004

over the ImageNet dataset [24]. The fully-connected layers in the fusion module are followed by the ReLU nonlinearity. For optimization, Adam [25] was adopted with a learning rate of 0.0001. Models are trained on the training set, and the weights with the highest performance on the validation set are selected as the best models, which will be evaluated on the test set.

Performance of base models

We first evaluated the performance of the base module's sub-models for the crossing point validation and severity grade prediction tasks. For comparison, we also give the results of models without pre-training (w/o PT) and without data augmentation (w/o DA), as well as models using only the green channel (GC Only).

Models		Cross. Point Val.		Severity Grade Pred.			
	Pre.	Rec.	<i>t</i> (ms)	Acc.	Kappa	<i>t</i> (ms)	
ResNet-50	0.9427	0.9526	0.274	0.8063	0.6629	0.278	
—w/o PT	0.8646	0.6975	0.274	0.5445	0.0177	0.278	
—w/o DA	0.9531	0.8551	0.274	0.5340	0.0036	0.278	
—GC Only	0.9583	0.9154	0.273	0.7277	0.5288	0.273	
Inception v3	0.9635	0.9635	0.218	0.8534	0.7432	0.222	
—w/o PT	0.9010	0.6865	0.218	0.5183	0.0313	0.222	
—w/o DA	0.9323	0.9179	0.218	0.5393	0.0000	0.222	
—GC Only	0.9167	0.9119	0.216	0.8115	0.6771	0.216	
DenseNet-121	0.9479	0.9630	0.266	0.8795	0.7892	0.269	
—w/o PT	0.9375	0.6742	0.266	0.5288	0.0050	0.269	
—w/o DA	0.9740	0.8274	0.266	0.7225	0.4865	0.269	
—GC Only	0.9740	0.9212	0.266	0.6702	0.4406	0.267	

Table 1. Performances of base models with ablation.

https://doi.org/10.1371/journal.pdig.0000174.t001

The crossing point validation performances are shown in the left part of <u>Table 1</u>. We use two metrics, precision and recall, and the time measurement to show the timing performance. We can see that pre-training and data augmentation can improve the overall performance of the crossing point validation. The Inception model with PT and DA achieved the best recall and the second-best precision. Note that PT and DA will not change the running time of the model because they do not modify the network structure.

The right part of Table 1 gives the results of the base models on the severity grade prediction task, and Table 2 presents the performance of MDTNet and models using the focal loss. In addition to the classification accuracy, we also adopt Cohen's kappa, which can measure the agreement between the ground-truth labels and predictions. We can see that, compared with the focal loss models, the DenseNet can achieve higher overall accuracy with the CE loss. However, the combination of different models, different losses, as well as different γ values can boost the performance. MDTNet achieved the highest performance in this experiment when n = 3.

To better analyze the severity grade prediction performance, we present the confusion matrices in Fig 5. It can be seen that, with the increment of the underlying sub-models, MDTNet gains the classification ability. Fig 6 shows visual explanation of MDTNet by Grad-CAM [26]. Fig 6 (a) and 6(b) show two examples for the crossing point validation. The ground-truth labels are *false* and the predictions were also *false*, *i.e.*, these are not effective crossing points as the arteries are under the veins. The model mainly counted the red area in the second row along the vein. The model might find the vein, track it down, and reach the conclusion that it lies above the artery. Fig 6 (c) and 6(d) are for the severity grade prediction. The ground-truth labels are respectively *mild* and *moderate* and were both correctly predicted.

Metrics	DenseNet-121 (Focal Loss)				MDTNet		
	$\gamma = 1$	$\gamma = 2$	$\gamma = 5$	$\gamma = 10$	<i>n</i> = 0	<i>n</i> = 1	<i>n</i> = 3
Acc.	0.8639	0.7434	0.8639	0.7958	0.8953	0.9110	0.9162
Kappa.	0.7642	0.5685	0.7641	0.6508	0.8183	0.8453	0.8542
<i>t</i> (ms)	0.268	0.268	0.268	0.268	0.767	1.047	1.571

Table 2. Performance of MDTNet models for severity grade prediction.

https://doi.org/10.1371/journal.pdig.0000174.t002



Fig 5. Confusion matrices for three different severity grade prediction models. The recall is shown in the last row and the precision is shown in the last column. (a) MDTNet without the focal module, (b) MDTNet for n = 1, and (c) MDTNet for n = 3.

https://doi.org/10.1371/journal.pdig.0000174.g005



Fig 6. Visual explanation of prediction results. (a,b) are for the crossing point validation model and (c,d) are from the severity grade prediction model. The first row is the raw input images and the second row is the class-discriminative regions.

https://doi.org/10.1371/journal.pdig.0000174.g006

We can see the artery runs over the vein deforming the vein. Being different from the example in (a) and (b), the model looks at the crossing points and looks for possible shape deformations and their extent.

Conclusion

The paper presents a method to automatically classify the arteriolosclerosis severity from retinal images following ophthalmologists' diagnostic process. To improve the accuracy for

Table 3. The MI-CLAIM checklist.

Before paper submission			
Study design (Part 1)	Page	Notes	
The clinical problem in which the model will be employed is clearly detailed in the paper.	2-3		
The research question is clearly stated.	3		
The characteristics of the cohorts (training and test sets) are detailed in the text.	3-4		
The cohorts (training and test sets) are shown to be representative of real-world clinical settings.	3-4		
The state-of-the-art solution used as a baseline for comparison has been identified and detailed.	7-8		
Data and optimization	Page	Notes	
The origin of the data is described and the original format is detailed in the paper.	3-4		
Transformations of the data before it is applied to the proposed model are described.	6-7		
The independence between training and test sets has been proven in the paper.	4		
Details on the models that were evaluated and the code developed to select the best model are provided.	7		
Is the input data type structured or unstructured?	☑ Structured	□ Unstructured	
Model performance (Part 4)	Page	Notes	
The primary metric selected to evaluate algorithm performance, including the justification for selection, has been clearly stated.	7-8		
The primary metric selected to evaluate the clinical utility of the model, including the justification for selection, has been clearly stated.	7-8		
The performance comparison between baseline and proposed model is presented with the appropriate statistical significance.	7-9		
Model examination (Part 5)	Page	Notes	
Examination technique	7-9		
A discussion of the relevance of the examination results with respect to model/algorithm performance is presented.	7-9		
A discussion of the feasibility and significance of model interpretability at the case level if examination methods are uninterpretable is presented.	8		
A discussion of the reliability and robustness of the model as the underlying data distribution shifts is included.	4,7-9		
Reproducibility (Part 6): choose appropriate tier of transparency			
Tier 1: complete sharing of the code	\checkmark		
Tier 2: allow a third party to evaluate the code for accuracy/fairness; share the results of this evaluation			
Tier 3: release of a virtual machine (binary) for running the code on new data without sharing its details			
Tier 4: no sharing			

https://doi.org/10.1371/journal.pdig.0000174.t003

ambiguous and unbalanced samples, we design the multi-diagnosis team network (MDTNet), which can jointly consider diagnostic cues from multiple sub-models, without tuning the hyperparameter for the focal loss. Experimental results show the superiority of our method, achieving over 91% accuracy. Most importantly, the whole process can be checked to see how the grading was determined as it is designed to be a step-by-step approach replicating ophthal-mologists' diagnostic process. Therefore, the proposed method can serve as a supporting tool for experienced ophthalmologists to efficiently grade the images in a consistently reproducible manner. A quality checklist [27] for the proposed deep learning method is shown in Table 3.

Author Contributions

Conceptualization: Liangzhi Li, Yuta Nakashima, Hajime Nagahara, Ryo Kawasaki.

Data curation: Liangzhi Li, Bowen Wang.

Formal analysis: Liangzhi Li, Manisha Verma.

Funding acquisition: Liangzhi Li, Ryo Kawasaki.

Investigation: Liangzhi Li.

Methodology: Liangzhi Li, Manisha Verma.

Project administration: Yuta Nakashima, Hajime Nagahara, Ryo Kawasaki.

Resources: Liangzhi Li.

Software: Liangzhi Li.

Supervision: Yuta Nakashima, Hajime Nagahara, Ryo Kawasaki.

Validation: Liangzhi Li, Bowen Wang.

Visualization: Liangzhi Li.

Writing - original draft: Liangzhi Li, Manisha Verma.

Writing - review & editing: Liangzhi Li, Manisha Verma, Yuta Nakashima, Ryo Kawasaki.

References

- Chatziralli IP, Kanonidou ED, Keryttopoulos P, Dimitriadis P, Papazisis LE. The value of fundoscopy in general practice. The open ophthalmology journal. 2012; 6:4. <u>https://doi.org/10.2174/</u> 1874364101206010004 PMID: 22435081
- Hubbard LD, Brothers RJ, King WN, Clegg LX, Klein R, Cooper LS, et al. Methods for evaluation of retinal microvascular abnormalities associated with hypertension/sclerosis in the Atherosclerosis Risk in Communities Study. Ophthalmology. 1999; 106(12):2269–2280. <u>https://doi.org/10.1016/S0161-6420</u> (99)90525-0 PMID: 10599656
- Walsh JB. Hypertensive Retinopathy: Description, Classification, and Prognosis. Ophthalmology. 1982; 89(10):1127–1131. https://doi.org/10.1016/S0161-6420(82)34664-3 PMID: 7155523
- Nguyen UTV, Bhuiyan A, Park LAF, Kawasaki R, Wong TY, Wang JJ, et al. An Automated Method for Retinal Arteriovenous Nicking Quantification From Color Fundus Images. IEEE Transactions on Biomedical Engineering. 2013; 60(11):3194–3203. https://doi.org/10.1109/TBME.2013.2271035 PMID: 23807422
- Roychowdhury S, Koozekanani DD, Parhi KK. Iterative Vessel Segmentation of Fundus Images. IEEE Transactions on Biomedical Engineering. 2015; 62(7):1738–1749. <u>https://doi.org/10.1109/TBME.2015.</u> 2403295 PMID: 25700436
- 6. Kim JU, Kim HG, Ro YM. Iterative deep convolutional encoder-decoder network for medical image segmentation. In: IEEE Engineering in Medicine and Biology Society (EMBC); 2017. p. 685–688.
- Yan Z, Yang X, Cheng K. Joint Segment-Level and Pixel-Wise Losses for Deep Learning Based Retinal Vessel Segmentation. IEEE Transactions on Biomedical Engineering. 2018; 65(9):1912–1923. <u>https:// doi.org/10.1109/TBME.2018.2828137</u> PMID: 29993396
- Huang F, Dashtbozorg B, Tan T, ter Haar Romeny BM. Retinal artery/vein classification using geneticsearch feature selection. Computer Methods and Programs in Biomedicine. 2018; 161:197–207. https://doi.org/10.1016/j.cmpb.2018.04.016 PMID: 29852962
- Meyer MI, Galdran A, Costa P, Mendonça AM, Campilho A. Deep Convolutional Artery/Vein Classification of Retinal Vessels. In: Image Analysis and Recognition; 2018. p. 622–630. <u>https://doi.org/10.1007/</u> 978-3-319-93000-8_71
- Costa P, Galdran A, Meyer MI, Niemeijer M, Abràmoff M, Mendonça AM, et al. End-to-End Adversarial Retinal Image Synthesis. IEEE Transactions on Medical Imaging. 2018; 37(3):781–791. <u>https://doi.org/ 10.1109/TMI.2017.2759102 PMID: 28981409</u>
- Hatanaka Y, Muramatsu C, Hara T, Fujita H. Automatic arteriovenous crossing phenomenon detection on retinal fundus images. In: Medical Imaging 2011: Computer-Aided Diagnosis. vol. 7963; 2011. p. 79633V.
- Li L, Verma M, Nakashima Y, Nagahara H, Kawasaki R. IterNet: Retinal Image Segmentation Utilizing Structural Redundancy in Vessel Networks. In: The IEEE Winter Conference on Applications of Computer Vision; 2020. p. 3656–3665.
- 13. Cambò VBS, Cariello L, Mastronardi G. A COMBINED METHOD TO DETECT RETINAL FUNDUS FEATURES. In: IEEE European Conference on Emergent Aspects in Clinical Data Analysis; 2005.

- Taylor C, Munro AJ, Glynne-Jones R, Griffith C, Trevatt P, Richards M, et al. Multidisciplinary team working in cancer: what is the evidence? The BMJ. 2010; 340:c951. <u>https://doi.org/10.1136/bmj.c951</u> PMID: 20332315
- Lin T, Goyal P, Girshick R, He K, Dollár P. Focal Loss for Dense Object Detection. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2020; 42(2):318–327. <u>https://doi.org/10.1109/TPAMI.</u> 2018.2858826 PMID: 30040631
- Inoue R, Ohkubo T, Kikuya M, Metoki H, Asayama K, Kanno A, et al. Stroke risk of blood pressure indices determined by home blood pressure measurement: the Ohasama study. Stroke. 2009; 40(8):2859–2861. https://doi.org/10.1161/STROKEAHA.108.546499 PMID: 19478224
- Li L, Verma M, Nakashima Y, Kawasaki R, Nagahara H. Joint Learning of Vessel Segmentation and Artery/Vein Classification with Post-processing. In: Medical Imaging with Deep Learning; 2020.
- He K, Zhang X, Ren S, Sun J. Deep Residual Learning for Image Recognition. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2016. p. 770–778.
- Huang G, Liu Z, Pleiss G, Van Der Maaten L, Weinberger K. Convolutional Networks with Dense Connectivity. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2019;.
- Szegedy C, Vanhoucke V, Ioffe S, Shlens J, Wojna Z. Rethinking the Inception Architecture for Computer Vision. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2016. p. 2818–2826.
- Weber M, Fürst M, Zöllner JM. Automated Focal Loss for Image based Object Detection. arXiv preprint arXiv:190409048. 2019;.
- 22. Huang C, Li Y, Loy CC, Tang X. Learning Deep Representation for Imbalanced Classification. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2016. p. 5375–5384.
- 23. Cui Y, Jia M, Lin TY, Song Y, Belongie S. Class-Balanced Loss Based on Effective Number of Samples. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2019.
- Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, et al. ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision (IJCV). 2015; 115(3):211–252. <u>https://doi.org/10.1007/s11263-015-0816-y</u>
- 25. Kingma DP, Ba J. Adam: A method for stochastic optimization. arXiv preprint arXiv:14126980. 2014;.
- Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In: International Conference on Computer Vision (ICCV); 2017. p. 618–626.
- Norgeot B, Quer G, Beaulieu-Jones BK, Torkamani A, Dias R, Gianfrancesco M, et al. Minimum information about clinical artificial intelligence modeling: the MI-CLAIM checklist. Nature medicine. 2020; 26 (9):1320–1324. https://doi.org/10.1038/s41591-020-1041-y PMID: 32908275