

Quantification of Corneal Neovascularization by Segmentation of Blood Vessels in Slit-Lamp Microscopy Images

Bruno Domingues de Oliveira
Federal University of ABC
São Bernardo do Campo, São Paulo,
Brazil

Tais H Wakamatsu
Federal University of São Paulo,
São Paulo, Brazil

Daniel Cavalcante de Oliveira
Federal University of ABC
São Bernardo do Campo, São Paulo,
Brazil

John A. Sims
Federal University of ABC
São Bernardo do Campo, São Paulo,
Brazil
ORCID: 0000-0002-6513-6860

Ugo Ibusuki
Federal University of ABC
São Bernardo do Campo, São Paulo,
Brazil
ORCID: 0000-0003-2478-4551

Abstract:

The cornea can form blood vessels as a result of inflammatory disease. To assess the evolution of vascularization and assist clinical decision making, ophthalmologists can assess slit lamp microscopy images, but this is qualitative and subjective, and more detailed analysis requires a time-consuming manual process. In this work, we propose automated image segmentation algorithms using classical methods and deep learning for 20 images, manually annotated by specialists. Artificial intelligence models, known as convolutional neural networks, based on the U-Net architecture, were trained and tested in three experiments and results compared with those from B-COSFIRE, a traditional image processing approach. The results with U-Net were the most promising, achieving Sørensen-Dice Similarity values of 48%.

Keywords — Corneal neovascularization, convolutional neural networks, ophthalmology, image segmentation, slit lamp microscopy.

I. INTRODUCTION

Stevens-Johnson Syndrome (SJS) is a rare dermatological disease, whose symptoms include continuous fever, generalized rashes (changes in skin tone and/or skin texture) and lesions on the genital, oral and ocular mucosa [1]. The ocular sequelae of the syndrome range from photophobia and discomfort to the development of corneal ulcers, which should be followed up by specialists.

In ophthalmological procedures, doctors monitor the condition of patients who have blood vessel formation in the cornea, a transparent tissue with a fundamental role in the formation of vision. This condition is known as corneal neovascularization (CNV), a possibly severe manifestation of SJS in the eyes.

The treatment of CNV may involve anti-inflammatory drugs or anti-vascular endothelial growth factor (anti-VEGF) agents, such as topical eye drops [2]. Anti-VEGF agents have the function of inhibiting protein signals that stimulate overexpression of blood vessel formation (angiogenesis). The gold standard for image monitoring for the detection of blood vessels in the eyes is fluorescein angiography [3], a tracer dye method that guarantees rapid examination, higher image quality and diagnostic accuracy, as it promotes greater contrast between the vessels in the cornea and its posterior tissues. However, it is an invasive and intravenous exam, with

risks of adverse reactions to the dye, such as nausea and anaphylactic shock [4].

Alternatively, images can be acquired with a slit-lamp microscope, which show superficial structures of the eye, such as the cornea (fig. 1). To monitor effectiveness of treatments to stop disease progression, the physician periodically evaluates the images and, as a measure of disease, counts the number of vascularized pixels, manually marking all blood vessels in the corneal region. This enables quantitative analysis of the advance or regression of CNV [5].

However, unlike retinal fundus images, the vessels of interest tend to have low contrast against the high variability of iris color background, pigmented tissue posterior to the cornea. Further, disease progression can provide partial or complete iris opacity, further confounding the segmentation. Additional challenges are specular reflections of the lamp on the surface of the eye and short depth of focus, meaning that a part of the cornea is always out of focus, either the central part, closest to the camera lens, or the edges, further from the camera.

In our literature review, we found an extensive number of studies for segmentation of blood vessels in the area of ophthalmology, but mainly focused on fundus images of the retina. This is due to the relationship between retinal structures with diagnosis and prognosis of cardiovascular diseases and diabetes [6]. However, we found no publications describing the use of techniques to aid segmentation of vessels in the cornea, except for a methodology for semi-automatic segmentation of mouse corneas [7].

In the clinical routine, assessment of patient condition is performed using the previously described subjective comparison of slit microscopy images from two or more consultations. Manual segmentation by an ophthalmologist is laborious, expensive and with high intra and inter-operator variability. The use of automated segmentation methods to quantify the vessels can reduce costs and operator subjectivity, however automated segmentation is a challenging task.

Our objective is to develop methodologies for automated segmentation of corneal vessels from slit-lamp microscopy images for patients with SJS, providing a pixel count corresponding to blood vessels and allowing quantification of the proportion that they occupy in the cornea. This will enable evaluation of CNV progression/regression in patients with SJS when treated with topical eye drops.



Fig. 1. Slit-lamp microscopy image of the cornea (Source: authors).

II. MATERIALS

A. Acquisition of images

The SJS database contains 20 images captured from the outer superficial portion of the eyes of patients with SJS, who have advanced CNV and severe ocular opacity. Each image is a high-resolution color image, on average 1800×1800 pixels, from different ophthalmological evaluations. Images for analysis were acquired by the Department of Ophthalmology, UNIFESP (Federal University of São Paulo), São Paulo, Brazil, with a slit-lamp microscope (TOPCON SL-D7 DC3 8MP) according to the following imaging protocol: 16X zoom, illumination with 20 aperture slit light, full corneal framing excluding the eyelid, without a light filter and peripheral light focus (central illumination in the pupillary portion) in order to enhance the vessels in the peripheral region of the cornea.

The image set was manually segmented using ImageJ [8] by specialist ophthalmologists at the hospital with experience in identifying blood vessels in advanced CNV caused by SJS. Labels showing the vascularized pixels were superimposed on the original image rather than containing a separate layer with the gold standard segmentation (fig. 2a). An initial task was to recover a binary mask from these labels. Figure 2b shows the manually-segmented binary map of the vessels extracted from the corresponding image, fig. 2a.

This research was approved by the Ethics and Research Committee of UNIFESP, with Certificate of Presentation of Ethical Appreciation number




Fig. 2. (a) Manual segmentation provided by UNIFESP, performed in ImageJ [8]; (b) Generated gold standard binary segmentation.

28211320.0.0000.5505, and all patients included in the study signed a Free and Informed Consent Term.

B. Computational resources

Training convolutional neural networks requires intense computational power, due to the high number of trainable parameters. To reduce training time, it is common to use a GPU (Graphics Processing Unit) with high RAM memory to load the model and images, so as not to overload the CPU. Computational demand also varies according to the model architecture, the memory size of the processed images and the model's hyperparameters. In this study, we used the Google Colab Pro platform, providing execution on a Tesla P100-PCIe GPU with 16 GB RAM and 100 GB storage on Google Drive.

III. METHODOLOGY

Medical image manipulation and processing is a well-established and developed field, with solutions readily available for similar problems to those proposed by this work [9]. Image segmentation, in general, aims to establish which pixels in a region of an image belong to the foreground, so isolating it from the background. Thus, the goal of image segmentation in our study is to compute a binary map, separating pixels corresponding to vessels in the cornea from those belonging to background structures, such as the iris.

Many commonly-used and easily-applied techniques can provide satisfactory results, such as thresholding, where pixel levels above a threshold value are considered to be figure and below this level are considered to be background. They are usually tested first in projects and used only if they present satisfactory results, as they generally do not require high computational power, and can avoid the need for heavy algorithms, connection to a server and GPU processing. However, in more complex cases, they may not be accurate enough, even after employing adaptations to make them more effective. It may be necessary to use more sophisticated algorithms to solve our proposed problem, considering that it must be able to deal with the limitations presented in NVC segmentation, such as colored and blurred images, non-uniform lighting, and low contrast.

Our chosen baseline for classical segmentation was B-COSFIRE (Combination of Shifted Filter Responses) version 1.4.0.0 written in MATLAB R14 [10], which uses Gabor filters in various scales and orientations [11]. The original article contains an example of blood vessel segmentation in retinal fundus images from the DRIVE database, which was adjusted to segment vasculature from our SJS images. While we consider it to be a classical approach (does not use neural networks), it is an algorithm of much greater complexity than thresholding.

In recent decades, the areas of artificial intelligence and its subfield of deep learning have contributed to medical and biological imaging, offering possibilities for solving more challenging problems. Artificial intelligence can be defined as the field of study that seeks to automate intellectual tasks normally performed by humans [12]. One of its sub-areas is supervised machine learning, which introduces a new paradigm of algorithmic development. An algorithm is given data and required responses as inputs, and then trained to formulate rules that dictate decision-making. The outputs of the sample inputs should be close to their expected responses, in order to be generalized to instances never seen by the algorithm [12].

Dense neural networks are applicable to a myriad of academic and industrial problems, but in the context of computer vision convolutional neural networks (convnets) are commonly used. While dense networks are trained to learn global patterns in dense layers, convnets learn local patterns in their layers, translation invariance (for example, to recognize an edge, regardless of the angle represented in the image) and with spatial hierarchy (primary layers learn simple patterns such as edges and curves, and with the depth of the network these learned patterns are gradually more complex) [12].

In machine learning, however, it is customary to notice the inverse proportionality between the model's degree of complexity and its interpretability. Neural networks are capable of formulating complex solutions, at the cost of a high level of abstraction, that is, low interpretability. Nevertheless, studies such as [13] and [14] seek to recover how the layers of the network evolve hierarchically in their ability to extract complex patterns.

Some neural network architecture was specifically developed to solve segmentation problems in medical and biomedical fields, and highlight the importance of techniques aimed at implementing encoder-decoder and autoencoder [15], with emphasis on U-Net [17] in 2D images (such as tissue analysis) and V-Net [16] in 3D images.

U-Net is an encoder-decoder type network: (i) the encoder consists of convolutional layers to extract contextual features from the image, followed by down-sampling, to reduce dimensionality and resolution, similar to other networks; (ii) the decoder is symmetrical to the encoder and performs up-sampling operations of feature maps (intermediate representations of the already processed input data), providing an increase in dimensionality, concatenating them to the feature maps of the primitive layers from the beginning of the network. This concatenation operation, illustrated by the gray arrows in fig. 3, introduces the concept of concatenated skip connections, which guarantee that the network incorporates the contextual information captured in the primary layers to the deep layers, and which reduce the loss of information after many operations. These long-range connections allow for greater retrieval of visual information and gradient flow when updating weights.

The down-sampling of images in the encoder stage is mediated by an operation called max pooling, which can be visualized as a sliding window that extracts the maximum




Fig. 3. U-Net architecture [17].

value of each image channel within a small window (2×2 pixels) with a certain pixel field. This process has two advantages: (i) down-sampling drastically reduces the number of trainable network parameters, since the next layer receives feature maps of reduced size; (ii) down-sampling guarantees

hierarchical learning of spatial features, otherwise the network would need to recognize unreduced global patterns of the image. In turn, the decoder part, through up-sampling, enables the precise location of pixels [17]. These mechanisms seek to avoid the need for large data sets for convergence of weights, i.e., learn to recognize patterns. Additionally, our implementation of U-Net used batch normalization, which was not included in the original architecture. The following paragraphs describe the segmentation experiments which were performed.

A. Data augmentation

The U-Net architecture, combined with extensive application of data augmentation, enabled the model to achieve unprecedented performances in medical image segmentation [17]. Data augmentation describes techniques to increase the amount of training data with moderately modified copies. In problems with little-annotated data, common in the medical field, data augmentation ensures that the model is invariant to elastic deformations, rotation, translation and perspective, providing robustness in the generalization of properties found in the training set. In addition to performing data augmentation, the U-Net architecture avoids some challenges common in deep learning with low quantities of data. In our data augmentation strategy, each image was rotated by multiples of 15° , resulting in 24 rotated images, and a total set of 480 images.

B. Batch size

One of the neural network's hyperparameters that demands the most computational power is the batch size, or the number of images that are given as input to the network at each iteration. For batches with many images, the changes in weights oscillate less than with a smaller batch. However, as batch size increases, it becomes more difficult to load data into memory. The advantage of having a smaller batch size is that learning does not need the entire dataset at each iteration. As neural network systems are prone to overfitting, the objective is to balance the batch size tradeoff in order to pass several batches through the network, each batch being a "noisy" representation of the entire data set, and generating a "tug and pull" dynamic in the updating of weights. This dynamic helps to prevent the neural network from overfitting the training set, providing better performance on the test set.

In view of the computational limitations of having a large batch size with our high-dimensional images, we divided each rotated image from augmentation into subimages (fig. 4). The original image was resized to the nearest multiple of 336 pixels of its original value so that each resulting sub-image had dimensions of 336×336 pixels. The dimension was chosen since 336 is a multiple of 16 pixels, targeting the 4 operations of max pooling with a receptive field of size 2×2 , reducing the width and height of the original image by 16 times in the




Fig. 4. (a) Original slit-lamp microscopy image; (b) Mosaic illustrating the subimages.

encoder phase, as seen in the U-Net architecture. With this processing, mirrored for the gold standard masks, a total of 14,664 sub-images were obtained for model training, without using pre-trained weights for initialization of the network.

C. Dataset split into train, test and validation image sets

We define training, validation and test sets in the modeling process as follows. Training data allows model learning, updating the neural network weights and biases. Validation data is a separate portion of the training data, which provides the model with an unbiased assessment in order to fit the model's hyperparameters. The test data allows evaluation of the trained model in an unbiased way to estimate how generalizable the model has become. In the experiments described below, we used different splits of the data to find the best strategies, as reflected by model performance on the test set. Regarding the problem of little available data in any scenario of splitting the training and test sets, there are few images left for testing. Although the subimages resulting from the data augmentation can be used in the training set, increasing the generalization power and consistency of the model, only the original subimages (i.e., without rotation) were used for validation and testing.

D. Experiments performed

We performed three experiments with different sizes of training and test sets:

Experiment 1 (Exp1): 75% training and 25% testing (i.e., subimages for 5 of the 20 original images for testing);

Experiment 2 (Exp2): 50% training and 50% testing (i.e., subimages of 10 of the 20 original images for testing).

Experiment 3 (Exp3): use the original entire rotated images.

For Exp1 and Exp2, the batch size used was 16 images, while in Exp3 the batch size was only 1 image, because of the image weight in RAM memory. k -fold cross-validation was used for Exp1 and Exp2, which partitions the dataset into a k parts, using one part to test the model and all the remaining parts for training, later exchanging them, resulting in k trained models. The purpose of cross-validation is to check consistency of the modeling approach in different training and testing bases.

E. Segmentation metric

To compare the gold standard segmentation to the method segmentation, we used the Sørensen-Dice Similarity (SDS) measure (1).

$$SDS = \frac{2 \cdot TP}{2 \cdot TP + FP + FN} \quad (1)$$

Where TP are true positive detections: the number of pixels identified as pertaining to vessels in the gold standard that were correctly classified by our method; FP are false

positive detections: the number of pixels, not identified as pertaining to vessels in the gold standard, that were classified as vessel by our method; FN are false negative detections, the number of pixels identified as vessel in the gold standard that were identified as non-vessel by our method.

TABLE I. EXPERIMENTS PERFORMED

	<i>Batch size</i>	<i>Train set</i>	<i>Test set</i>	<i>k-fold validation</i>
Exp1	16 subimages	10998 subimages	5 images	4-fold
Exp2	16 subimages	7332 subimages	10 images	2-fold
Exp3	1 image	360 images	5 images	-

We also used the measures precision (Prec) (2) and sensitivity (Sens) (3). Prec provides a reply to the question: "What is the proportion of pixels identified by our method as vessel that are correct?" Sens provides a reply to the question: "What is the proportion of pixels that are identified as vessel by our method compared to those which should be identified as vessels?"

$$Prec = \frac{TP}{TP + FP} \quad (2)$$

$$Sens = \frac{TP}{TP + FN} \quad (3)$$

Information regarding data volumes is shown in Table 1. In all folds, 10 training epochs were performed, with one epoch consisting of the number of batches to pass the entire training set through the network. For example, in a hypothetical training of a dataset of 500 images with a batch size of 10 images, 50 batches would be needed to complete an epoch. In all experiments, the images or subimages referring to the original non-rotated images were separated from the training dataset for validation at the end of each epoch.

IV. RESULTS

As for the classical processing methods used, such as channel thresholding methods, none showed adequate segmentation. In the approach with the COSFIRE algorithm, for the 20 images, an average performance of $SDS = (37 \pm 5)\%$ and maximum value of 42% was obtained, with average pixel precision ($Prec = (31 \pm 8)\%$) and sensitivity pixel mean ($Sens = (54 \pm 19)\%$). An example segmentation is shown in fig. 5c.

The U-Net model achieved a performance of $SDS = (48 \pm 2)\%$ in Exp1 (4-folds) and $SDS = (47 \pm 1)\%$ in Exp2 (2-folds). An example segmentation is shown in fig. 5d.

The results of the evolution of the model performance in each fold of the cross-validation experiments can be seen in fig. 6. As shown in Table II, the model for Exp3 achieved




Fig 5. (a) original image from SJS dataset; (b) gold standard manual segmentation; (c) segmentation by COSFIRE (image with SDS close to the mean value of 36.1%). (d) segmentation by U-Net (image with $SDS=44.7\%$ in Experiment 1 and 47.5% in Experiment 2.)

better SDS performance than the COSFIRE algorithm, whereas it did not perform as well as either of the experiments involving minibatches of sub-images.

TABLE II. MODEL RESULTS

Experiment	SDS (%)	Prec (%)	Sens (%)
Exp1	48	47	50
Exp2	47	47	49
Exp3	43	42	44
COSFIRE	37	31	54




Fig. 6. SDS results for the test sets: (a) Exp1 (2-fold cross-validation) and; (b) Exp2 (4-fold cross validation).

V. DISCUSSION

The COSFIRE algorithm shows high average sensitivity, but its high standard deviation indicates a lack of consistency across the test images. Although models for clinical decision support do not necessarily need high performance in their task, it is expected that they will act consistently across the data. While the model does not fail to recognize larger vessels in the images, it tends to increase sensitivity to the point where it makes a large number of false positive errors, resulting in reduced accuracy.

U-Net seems more robust to noise than COSFIRE, particularly in non-vascularized regions. This is reflected in the superior accuracy metric of U-Net, even with a slightly decreased sensitivity. In every case, U-Net has better SDS than COSFIRE. Furthermore, although U-Net has lower sensitivity, the standard deviation is lower than COSFIRE, reflecting the model's reliability. A comparison of the U-Net experiments shows Exp2 (reduced test set) produces similar results to Exp1 (larger test set). Similarly, we can say that reduced training data in Exp1 did not produce very different values from those in Exp2. We note that the SDS achieved in the first epoch of training was similar for both experiments, and subsequent training did not achieve a large increase in SDS.

Exp3 showed lower performance than Exp1 and Exp2. We suggest the following reasons to explain this: (a.) The batch size was a single image, due to the limitations of loading complete images into the computer memory. Training with only one sample per iteration leaves the model subject to very sudden changes in weights, decreasing convergence, as well as offering little sample diversity per batch; (b.) low-caliber blood vessels may lose relevance for a higher dimensionality of image as it undergoes dimension reductions through the network architecture.

Figure 7 shows the overlapping of a vessel prediction by U-Net and its corresponding gold standard. Note that the vessels with high and medium caliber have sufficient detection

by the algorithm, agreement denoted by the dark magenta color. However, the thinnest vessels are usually not detected by the network. Network training was performed using the dice loss function, which perhaps limits convergence of the model, resulting in poor recognition of low-caliber vessels. Based on [18], two of these limitations can be mentioned: (i) the Dice loss function cannot capture and consider the distance between two forms, making it impossible to use location information; (ii) if there are two shapes with different pixel volumes in the image, the larger shape will be favored, and the small net penalty of not having finding the smaller shape will be ignored (in this case, the narrow vessels).




Fig. 7. Vessel segmentation superposition. Pixels included in both U-Net and gold standard segmentation are shown in magenta (TP), pixels only segmented by U-Net are shown in yellow (FP), and pixels only segmented in the gold standard are shown in cyan (FN).

Our segmentation results could seem disappointing in terms of SDS, while detecting the larger vessels in a reasonable manner. Figure 8 shows two cases: (right) a square region, 8x8 pixels (yellow + red region), which is classified by the yellow region. The SDS in this case is $(2 \cdot 48)/(48 + 64) = 0.8571$. (left) a thin rectangular region, 4x16 pixels, which is classified by the yellow region. The SDS is $(2 \cdot 32)/(32 + 64) = 0.6667$. When SDS is used with a large square or circular region, the misclassification of a single border at the top and bottom leads to a relatively high value of SDS, given that there is a large number of true positive detections. However, with our situation of vessel segmentation, a consistent misclassification along the vessel, similar to the case of the long short rectangle in the figure, will lead to relatively lower values of SDS. While it can be seen in Figure 7 that our method did not detect the smaller vessels, we can also determine that there was misclassification along the border of most of the medium sized vessels, which could be due to mis-




Fig. 8. Intuitive example to explain relatively low values of SDS in vessel segmentation. (left) a square of pixels (8x8) is misclassified as a 6x8 rectangle, where an upper and lower row of pixels have been misclassified. (right) a long thin rectangle, 4x16 pixels, similar to the geometry of our vessels, is also misclassified by the same rows.

classification by our algorithm, or even misclassification by the ophthalmologist performing the manual segmentation.

We envisage an expansion of deep machine learning applications to the field of imaging medicine. However, we should reflect on the use of these methods in practice. One of the financial and temporal limitations of the project was the demand for computational power. Although the experiments reach a plateau in the evaluation metrics of the test bases, we would like to carry out training with more than 10 epochs to improve understanding of weight convergence in the network. Each training epoch lasted an average of 10 hours, making it impractical to carry out further experiments, one of them being the use of the leave-one-out cross-validation technique, an extreme case of k-fold in which the value of k is the same as the data sample size in the base, for example, for a base with 20 images, training would be carried out with 20 folds, leaving one of the images for testing each time.

Other improvements to our methodology could be achieved by the following: (i) Network architectures are constantly evolving and improving. While there are many other published models in recent literature, U-Net++ [19] could make a particularly interesting comparison, since it addresses other ways of including convolution layers between the encoder and decoder. (ii) Different hyperparameters could be used for training, such as changing the loss function, using different up-sampling techniques in the decoder section, and changing batch size values. (iii) More sophisticated forms of data augmentation could be used, including image transformations such as warping and varying the scale of the image in training. (iv) We note low rates of sensitivity in our segmentation method, which might be caused by poor results in the thin vessels. It might be possible to investigate this performance by using morphological operations to isolate and test segmentation of the thin vessels. Further, a more detailed statistical analysis could be performed to validate the results.

VI. CONCLUSION

In this paper, we present methodologies for segmentation of neovascularization in an image base of 20 patients with Stevens-Johnson Syndrome. Simple segmentation methods were confounded by the problems with image quality and the low contrast of the vessels against the background of the iris. COSFIRE, a sophisticated algorithm but one which does not involve machine learning, produces reasonable performance, achieving a Sorensen-Dice Similarity (SDS) of 37%. Convolutional neural networks, using the U-Net architecture, were used in three experiments: Experiment 1, a 4-fold validation, trained with minibatches of randomly selected image pieces, achieved the highest SDS of 48%. Experiment 2, a 2-fold validation, also trained with minibatches of randomly selected image pieces, achieved SDS of 47%. Experiment 3, which used single entire images for training, achieved an SDS of 43%. The U-Net architecture also produced lower standard deviation in SDS values, meaning it performed more reliably.

ACKNOWLEDGMENT

The authors thank the Ophthalmology Department of the Federal University of São Paulo (UNIFESP), Brazil, for providing the images used in this study.

REFERENCES

- [1] Bulisani, ACP, Sanches GD, Guimarães HP, Lopes RD, Vendrame LS, Lopes AC. Síndrome de Stevens-Johnson e necrólise epidérmica tóxica em medicina intensiva. *Rev Bras Ter Intensiva, SciELO Brasil*, v. 18, p.292–297, 2006.
- [2] Dastjerdi MH, Al-Arfaj KM, Nallasamy N, Hamrah P, Jurkunas UV, Pineda R, et al. Topical bevacizumab in the treatment of corneal neovascularization: results of a prospective, open-label, noncomparative study. *Archives of ophthalmology, American Medical Association*, v. 127, n. 4, p. 381–389, 2009.
- [3] Brunner M, Romano V, Steger B, Vinciguerra R, Lawman S, Williams B et al. Imaging of corneal neovascularization: optical coherence tomography angiography and fluorescence angiography. *Invest Ophthalmol Vis Sci, ARVO*, v. 59, n. 3, p. 1263–1269, 2018.
- [4] Lira RPC, Oliveira, CL, Marques MVRB, Silva AR, Pessoa CC. Adverse reactions of fluorescein angiography: a prospective study. *Arq Bras Oftalmol, SciELO Brasil*, v. 70, p. 615–618, 2007.
- [5] Wakamatsu, TH, Takiishi AY, Hirai F, Lima AS, Gomes JAP. The efficacy of topical bevacizumab and sunitinib on corneal neovascularization in patients with Stevens-Johnson. *Invest Ophthalmol Vis Sci, ARVO*, v. 59, n. 9, p. 3329–3329, 2018.
- [6] Abdulsahib AA, Mahmoud, MA, Mohammed MA, Rasheed HH, Mostafa SA, Maashi MS. Comprehensive review of retinal blood vessel segmentation and classification techniques: intelligent solutions for green computing in medical images, current challenges, open issues, and knowledge gaps in fundus medical images. *Netw Model Anal Health Inform Bioinforms, Springer*, v. 10, n. 1, p. 1–32, 2021.
- [7] Rabiolo A, Bignami F, Rama P, Ferrari G. VesselJ: a new tool for semiautomatic measurement of corneal neovascularization. *Invest Ophthalmol Vis Sci, ARVO*, v. 56, n. 13, p. 8199–8206, 2015.
- [8] Rasband WS. ImageJ, U. S. National Institutes of Health, Bethesda, Maryland, USA, <https://imagej.nih.gov/ij/>, 1997-2018.
- [9] Dougherty G. Digital image processing for medical applications. Cambridge University Press, 2009.
- [10] Azzopardi G, Petkov N. Trainable COSFIRE filters for Keypoint Detection and Pattern Recognition. *IEEE Trans Pattern Anal.Mach.Intell.*, v. 35, n. 2, p. 490–503, 2013.
- [11] Rangayyan R, Ayres F, Oloumi F, Oloumi F, Eshghzadeh-Zanjani P. Detection of blood vessels in the retina with multiscale Gabor filters. *J. Electronic Imaging*, v. 17, p. 023018, 04 2008.
- [12] Chollet F. Deep learning with Python and Keras: The practical guide from the developer of the Keras library. MITP-Verlags GmbH & Co. KG, Bonn, 2018. [
- [13] Yosinski J, Clune J, Nguyen A, Fuchs T, Lipson H. Understanding neural networks through deep visualization. *ArXiv preprint arXiv:1506.06579*, 2015.
- [14] Zeiler M.D, Fergus R. Visualizing and understanding convolutional networks. *European conference on computer vision*, p. 818–833, 2014.
- [15] Minaee S, Boykov Y, Porikli F, Plaza A, Kehtarnavaz N, Terzopoulos D. Image segmentation using deep learning: A survey. *CoRR*, abs/2001.05566, 2020. Available at: <<https://arxiv.org/abs/2001.05566>>.
- [16] Milletari, F, Navab N, Ahmadi S. V-net: Fully convolutional neural networks for volumetric medical image segmentation. *CoRR*, abs/1606.04797, 2016. Available at: <<http://arxiv.org/abs/1606.04797>>.
- [17] Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. *CoRR*, abs/1505.04597, 2015. Available at: <<http://arxiv.org/abs/1505.04597>>.
- [18] Zhang, Y. et al. Rethinking the dice loss for deep learning lesion segmentation in medical images. *Journal of Shanghai Jiaotong University (Science), Springer*, v. 26, n.1, p.93–102, 2021.
- [19] Zhou Z, Siddiquee MMR, Tajbakhsh N, Liang, J. UNet++: A Nested U-Net Architecture for Medical Image Segmentation. *arXiv*, 2018. Available at: <<https://arxiv.org/abs/1807.10165>>