

Major Temporal Arcade Segmentation through Hyperparameter Fine-Tuning in an Attention-Based U-Net

Hiram Efraín Orocio-García, Ivan Cruz-Aceves*, Arturo Hernández-Aguirre

Centro de Investigación en Matemáticas,
Mexico

{hiram.orocio, ivan.cruz, artha}@cimat.mx

Abstract. This study presents an innovative approach based on a U-Net architecture for the automatic segmentation of the major temporal arcade (MTA) in fundus images. The strategy consists of attention and segmentation procedures. Vascular structure segmentation using the popular VGG-16 deep learning model and transfer learning using the U-Net architecture constitute the initial phase. An attention module is added in the second stage to capture the MTA shape and remove any vascular structures other than MTA. The effectiveness of the suggested method is evaluated in terms of F1 score, sensitivity, specificity, and accuracy, and compared with various state-of-the-art vascular segmentation methods. A data augmentation approach was applied only to the training set. An expert categorized MTA images from the popular DRIVE database, which included 40 fundus images, to determine the vascular structure. 1377 training patches were obtained by dividing each training image into 81 patches and applying data augmentation to ensure a reliable training procedure. A comparative study was conducted, showing how the best metrics were achieved, as well as comparing them with other methods for segmenting arterial structures. The best results were obtained with an accuracy of 0.9923 and an F1 of 0.7700 using the test set of fundus images. The numerical results of the suggested method on the automatic MTA segmentation problem outperform those of the comparative approaches.

Keywords. Deep learning, image segmentation, major temporal arcade, u-net attention architecture, parameter tuning.

1 Introduction

Biomedical imaging plays a crucial role in modern medicine, offering clinicians the ability

to visualize the internal structures of the human body. Different imaging modalities provide detailed information about specific tissue types, including bones, muscles, and nerves. These images are essential for diagnosis, treatment planning, and monitoring the progression of diseases. Within these modalities, retinal fundus imaging stands out as a key tool for assessing ocular conditions such as retinal detachment, as well as systemic disorders like hypertension, congenital vascular malformations, and trauma-related injuries.

Moreover, fundus images support general health evaluations because the retina is the only location in the body where blood vessels can be directly observed. This characteristic enables a non-invasive assessment of vascular health, offering valuable clues about systemic conditions that influence overall patient well-being.

A major difficulty in extracting or segmenting the Major Temporal Arcade (MTA) from retinal fundus images arises from several factors: low contrast between vessels and background, the narrow diameter of the vessels, and uneven illumination.

Automatic vessel segmentation typically involves two main stages: vessel enhancement (or detection) and binary segmentation. Cutting-edge approaches often rely on thresholding techniques, with intraclass thresholding [16] being the most common. Their main variations occur during the detection stage, where images are filtered either spatially or in the frequency domain to compensate for illumination inconsistencies and highlight vessel-like structures.

To extract vascular structures, Eiho and Qian [7] proposed a mathematical morphology filter based on a top-hat operator with a disk-shaped structuring element. Qian et al. [18] extended this idea by incorporating a multiscale top-hat operator to detect vessels of different sizes. Frangi et al. [8] introduced a multiscale vesselness filter derived from the eigenvalues of the Hessian matrix, later refined in subsequent studies [21]. Chaudhuri et al. [4] developed the Gaussian Matched Filter, which was later adapted for multiscale detection using neural networks [5]. Additionally, Nguyen et al. [13] proposed a multiscale linear filter for retinal vessel detection. In the frequency domain, Gabor filters—augmented for multiscale analysis [19]—have been widely adopted for automatic vessel detection [2].

Deep learning architectures, particularly U-Net [20], have significantly advanced vessel segmentation in multiple medical imaging modalities, including coronary arteries in X-ray angiograms, retinal fundus images [22], phase-contrast tomography [25], and 3D vascular lumen reconstruction [12]. Despite these advancements, MTA segmentation remains relatively unexplored. Oloumi et al. [14, 15] used the Hough transform to model the MTA in diabetic retinopathy, along with single-scale and multiscale Gabor filters for segmentation. Alvarado et al. [1] implemented U-Net combined with second-order splines to segment the MTA. Soto et al. [23] integrated multiscale linear filtering, thresholding, and length filtering, employing Jacobi polynomials for numerical modeling. Nevertheless, these methods (including Gabor filters, U-Net, and multiscale filtering) have difficulty consistently isolating the MTA from thinner vessels in terms of length and width.

While the choice of architecture is central to performance, the effectiveness of a deep learning model is also highly dependent on the training strategy and the selection of hyperparameter choices. Fine-tuning these parameters not only improves convergence but also leads to more robust and generalizable segmentations.

Therefore, in addition to evaluating the proposed attention-based U-Net, this work presents a systematic analysis of hyperparameter fine-tuning,

highlighting its impact on segmentation quality. We focus on the effect of different learning rates, batch sizes, and training schedules (freezing and unfreezing the encoder) on the accuracy, sensitivity, specificity, F1-score, and inference time of the model. Our experiments demonstrate that appropriate fine-tuning substantially enhances the performance of the model, providing insights that are valuable for applying deep learning segmentation methods in settings with limited training data.

Longitudinal monitoring of patients represents another relevant clinical scenario where the method can offer advantages. Consistent and repeatable segmentation across multiple visits would enable clinicians to track gradual vascular changes that may not be visually appreciable, such as localized narrowing, shape deformation, or early indications of microvascular damage. Reliable temporal comparisons may support earlier interventions or more precise evaluation of treatment responses. Although the results indicate strong performance on the curated dataset, several practical considerations may affect real-world applicability. Clinical imaging workflows often present substantial variability in patient demographics, disease conditions, and acquisition settings. Models trained on limited anatomical diversity may underperform when encountering retinal structures affected by edema, hemorrhages, or irregular vessel morphology, where the shape of the arcade deviates from typical patterns. Broader validation on datasets with a wider spectrum of clinical presentations would be required to confirm stability under these conditions, as well as adaptations to maintain robustness when facing non-standard vascular configurations.

2 Materials and Methods

The following section outlines the core principles of the U-Net architecture, a widely adopted convolutional neural network for medical image segmentation, together with the integration of VGG16 as its backbone. It also provides a description of the database containing manual delineations of the Major Temporal Arcade in retinal fundus images used throughout this study.

2.1 Database of Major Temporal Arcade

For the purposes of this research, the publicly available DRIVE database of retinal fundus images was employed. This dataset is commonly used as a benchmark for evaluating algorithms designed for automatic retinal vessel segmentation. In this work, the DRIVE images were repurposed to focus specifically on segmenting the Major Temporal Arcade [24], while discarding other vessel-like structures not belonging to the MTA. The dataset contains 35 RGB images, each with a resolution of 576×576 pixels. The corresponding ground-truth annotations are provided as binary images, where low-intensity values correspond to the background and high-intensity values denote the manually delineated MTA.

2.2 U-Net Deep Learning Architecture

Deep learning, a subfield of artificial intelligence, employs deep neural networks to model and address highly complex tasks by mimicking human decision-making processes. It has shown outstanding performance across numerous domains, particularly in image classification. Within biomedical imaging, deep learning has reshaped how images are interpreted and processed, enabling major improvements in diagnosis, treatment planning, and scientific research.

Among the various deep learning applications in biomedical imaging, segmentation is the most relevant to our work. Using convolutional neural networks (CNNs), it is possible to perform highly accurate segmentation of organs and tissues in medical images, offering valuable support for clinical decision-making and planning.

One of the key strengths of U-Net lies in its ability to handle large images while preserving spatial information, which is essential for segmentation tasks. The architecture is composed of two main components: the encoder and the decoder. The encoder reduces the spatial dimensions through successive convolutional operations and downsampling steps, capturing essential features such as edges and textures. In parallel, skip connections play a vital role by preserving spatial structure and enabling the decoder to reconstruct

detailed segmentation maps by linking encoder outputs to corresponding decoder layers.

The U-Net architecture highlights critical operations including max pooling, transposed convolutions (or up-convolutions), and standard convolutions.

Although originally developed for medical image segmentation, U-Net has been widely adopted for numerous segmentation challenges. It is particularly recognized for its symmetric design and its strong performance even when only limited training data are available. The architecture revolves around two stages: the encoder, or contraction path, and the decoder, or expansion path. The encoder extracts meaningful visual features through convolutional layers and downsampling, reducing spatial resolution while capturing informative patterns. The decoder then reconstructs the segmentation output by leveraging the encoded features, progressively restoring the resolution through upsampling and convolutional layers. Throughout this process, detailed spatial information is reinforced through skip connections, ensuring that the shapes and structures within the image are accurately preserved.

2.3 Hyperparameter Fine-Tuning

Fine-tuning hyperparameters such as learning rate, batch size, and training schedule is critical for achieving optimal performance in deep learning models, particularly when training data is limited. Recent reviews have emphasized that hyperparameter tuning significantly affects model generalization and performance trade-offs, even in applied biomedical contexts [10, 11]. Studies in small-dataset segmentation tasks also highlight systematic hyperparameter selection as an effective strategy to improve accuracy and reliability [3]. Specifically for deep learning in computer vision, fine-tuning learning rates and batch sizes has been shown to greatly impact convergence speed and segmentation quality, especially when adapting pre-trained models to new domains [17, 9]. Importantly, tuning strategies that includes freezing vs. unfreezing layers, adjusting learning rates gradually, or exploring batch sizes in low-data regimes can lead to both improved

performance and reduced overfitting risk. To assess the impact of hyperparameter choices on the segmentation performance, we systematically varied the learning rate (1^{-3} , 1^{-4} , 1^{-5}), the batch size (1, 2, 4, 8) and the number of epochs allocated to the frozen and unfrozen phases. For each configuration, we trained the model twice, first with the encoder frozen, and then with all layers unfrozen. After every epoch, we recorded training loss, validation loss, F1-score, accuracy, sensitivity, specificity, and inference time per image. This experimental design enabled a direct comparison of how fine-tuning influences both segmentation quality and training efficiency.

A structured workflow was followed to analyze the effect of hyperparameters on segmentation quality.

The process begins by selecting a combination of learning rate, and batch size. Once the configuration is defined, the encoder is kept frozen during an initial training phase, allowing the decoder to adapt to the segmentation task while retaining the high-level representations learned by VGG16. After this stage, all layers are unfrozen and the model is trained again, enabling the encoder to specialize toward the retinal domain. For each configuration, the model is trained for a predefined number of epochs in both phases. Training loss, validation loss, accuracy, sensitivity, specificity, F1-score, and inference time are recorded after every epoch.

These metrics allow a direct comparison of how each hyperparameter combination influences convergence behavior, generalization, and computational efficiency. The workflow not only facilitates performance assessment but also reveals the trade-offs associated with batch size and learning rate choices.

3 Proposed U-Net Architecture based Attention Module

To segment the MTA in retinal fundus images (576×576 pixels), a U-Net model equipped with VGG16 as a pre-trained backbone and an integrated attention mechanism is employed. VGG16 provides robust feature extraction, enabling the U-Net architecture to capture both coarse and fine vascular characteristics. The model effectively

handles common challenges such as image noise, vessel overlap, and illumination inconsistencies.

Using ground-truth annotations during training, the network distinguishes the Major Temporal Arcade from surrounding vessels, improving segmentation performance. Figure 1 presents the architecture of the attention-augmented U-Net, where the attention modules reinforce the model's capacity to emphasize relevant structures and suppress non-essential information.

The attention mechanism enables neural networks to prioritize meaningful regions within an image. It assigns varying weights to feature representations, highlighting crucial structures while reducing the influence of background noise. This is particularly advantageous in retinal vessel segmentation, where vessels differ in width, contrast, and orientation. The mechanism computes relevance by comparing queries and keys via a dot-product operation, followed by a normalization step using the softmax function. Attention is mathematically defined as:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (1)$$

where d_k denotes the dimensionality of the keys, Q corresponds to the query matrix, K to the key matrix, and V to the value matrix. The softmax operation ensures that the resulting weights sum to one, accentuating the most relevant features while diminishing the influence of less informative components. Integrating this mechanism into the segmentation pipeline allows the model to prioritize the Major Temporal Arcade by reducing interference from adjacent vessels.

To manage dense and overlapping vasculature, the initial stage uses a patch-based segmentation strategy, where the U-Net processes the retinal fundus image in smaller patches to segment all visible vessels. The reconstructed mask includes both the MTA and surrounding veins. This approach enhances segmentation accuracy and computational efficiency by reducing spatial variability, preserving fine details, and supporting parallel operations. The attention-augmented U-Net then applies an additional refinement step.

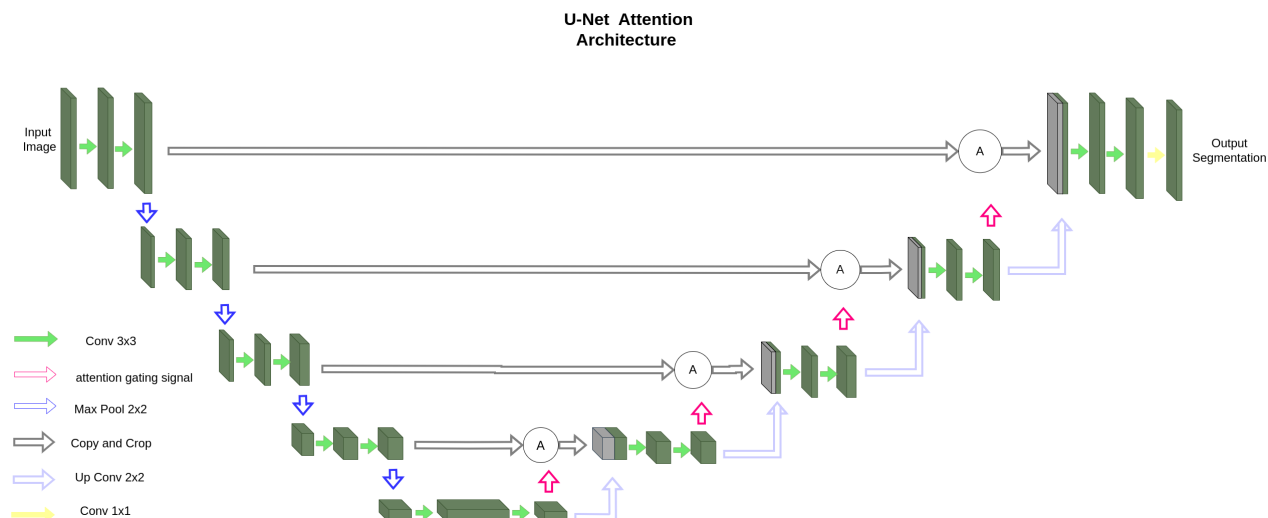


Fig. 1. Proposed U-Net attention architecture

A second model U-Net with an attention module built upon the VGG16 backbone—further suppresses noise, irrelevant structures, and illumination variations to isolate the MTA. This refinement stage minimizes false positives, prevents over-segmentation, and adapts to variations in vessel caliber and overall image quality by leveraging ground-truth masks that contain only the MTA. The resulting segmentation is substantially more precise, enabling improved anatomical assessment for clinical analysis and diagnostic applications.

4 Experimental Results

The following section presents the results achieved in segmenting the MTA using the proposed approach and compares these outcomes with those obtained from several traditional techniques.

Five specialized vessel enhancement techniques were applied to segment the Major Temporal Arcade and assess the performance of the proposed methodology. Eiho and Qian [7] introduced a morphology-based spatial filter using the Top-Hat operator, where the filter size served as the primary parameter. After testing values in the interval [5,15] on the training set, a size of nine pixels provided the best F1-score. Nguyen et al. [13]

developed a multiscale spatial filter tailored for linear structures to enhance blood vessels, tuning the scale parameter within the interval [7,17]. By refining the search process, the most effective range was determined to be [15:2:19], resulting in the highest F1-score.

Frangi et al. [8] proposed a Hessian-based vesselness filter that measures vessel likelihood using the Hessian matrix eigenvalues. Selecting an appropriate sigma range was essential to capture vessels of different widths; after evaluating values with a step of 0.5, the optimal range was identified as [1:0.5:7], which maximized the F1-score. Chaudhuri et al. [4] introduced a Gaussian-matched filter in which sigma determines the target vessel width.

Expanding the search to the interval [1.5,3] and adjusting stride values led to an optimal sigma of 1.75 based on F1-score performance. Lastly, the single-scale Gabor filter, operating in the frequency domain, relied on width, length, and the number of orientations. The width parameter was particularly important for detecting vessels of varying thickness, and after testing values from [4.5,7], the highest F1-score was obtained with a width of 7 pixels.

For the deep learning architectures, two training stages were employed. First, the VGG16 layers were frozen to prevent updates, allowing the

Table 1. Comparative and time analysis between the proposed and comparative methods for automatic MTA segmentation

Method	Sensitivity	Specificity	Accuracy	F1-score	Testing time in seconds (per image)
Mathematical Morphology [18]	0.6715	0.9481	0.9428	0.3097	1.84
Linear Multi-scale [2]	0.7800	0.9315	0.9286	0.2946	1.43
Multi-scale Hessian [21]	0.3552	0.9952	0.9829	0.4434	1.63
Gaussian Matched Filter [5]	0.4087	0.9614	0.9508	0.2412	1.46
Single-scale Gabor [15]	0.6991	0.9579	0.9530	0.3624	1.57
U-Net Architecture [20]	0.3457	0.9994	0.9869	0.5032	0.002
Deformable convolutions [6]	0.6575	0.9965	0.9901	0.7168	0.1243
Proposed U-Net Attention	0.6730	0.9985	0.9923	0.7700	0.003

Table 2. Best results by batch size and learning rate

BS	LR	F1	Sens	Spec	Acc	Best-ep	Phase	TrainLoss	ValLoss
1	1e-05	0.7225	0.6828	0.9961	0.9900	68	unfrozen	0.01241	0.04046
1	0.0001	0.7331	0.6879	0.9964	0.9905	89	unfrozen	0.00444	0.05327
1	0.001	0.7332	0.6805	0.9967	0.9906	77	unfrozen	0.00604	0.04593
2	1e-05	0.7336	0.7140	0.9954	0.9900	77	unfrozen	0.05965	0.08181
2	0.0001	0.7474	0.7210	0.9960	0.9907	72	unfrozen	0.00607	0.03705
2	0.001	0.7463	0.7084	0.9963	0.9908	81	unfrozen	0.00656	0.03940
4	1e-05	0.7390	0.7536	0.9944	0.9898	77	unfrozen	0.18098	0.18968
4	0.0001	0.7473	0.7158	0.9961	0.9907	69	unfrozen	0.00772	0.03204
4	0.001	0.7441	0.7066	0.9963	0.9907	117	unfrozen	0.00610	0.04191
8	1e-05	0.7356	0.7272	0.9952	0.9900	138	unfrozen	0.19983	0.20886
8	0.0001	0.7475	0.7230	0.9959	0.9907	96	unfrozen	0.00646	0.04097
8	0.001	0.7464	0.7178	0.9960	0.9907	123	unfrozen	0.00603	0.04053
2	0.0001	0.7700	0.6730	0.9985	0.9923	100	unfrozen	0.00550	0.03126

model to adapt to the nature of the task. After completing this stage, the VGG16 layers were unfrozen and the full network was trained jointly, enabling the architecture to specialize further and improve performance. Both models take a retinal fundus image as input for MTA segmentation, which includes general vessel extraction.

Table 1 presents the comparison of automatic MTA segmentation results between the state-of-the-art techniques, computational time, and the proposed approach. Sensitivity, specificity, accuracy, and F1-score were used as evaluation

metrics. The parameter optimization previously determined was applied to carry out the vessel enhancement procedures based on spatial and frequency filtering. The deep learning models and the Hessian-based technique achieved the best overall performance.

Figure 2 presents the confusion matrix, which summarizes the classification results of the segmentation model. This matrix provides insights into the model's performance, showing how well it distinguishes between vessel and background pixels. The high number of true negatives

suggests that the model effectively identifies non-vessel areas, while the false positive and false negative values indicate the areas where misclassifications occur.

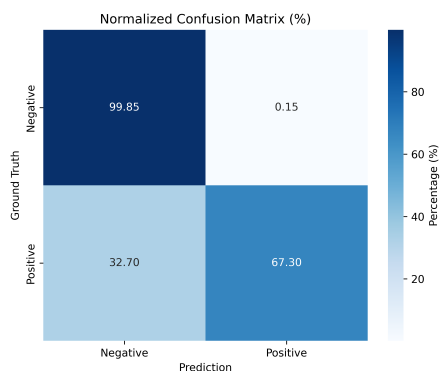


Fig. 2. Confusion matrix of the test images

In Figure 3, the results obtained for the initial segmentation of all arterial structures in the DRIVE images are presented. Additionally, the exclusive segmentation of the Major Temporal Arcade is shown, along with a comparison against the ground truth. These visualizations provide a clear assessment of the effectiveness of the proposed method in isolating the MTA while preserving relevant anatomical details. To complement these segmentation results, the analysis presented so far provides the baseline performance of the proposed model in terms of segmentation and classification.

To understand which design choices most strongly influence these outcomes, an additional experimental stage was conducted focusing on hyperparameter optimization. This second part serves to refine the training pipeline by examining how learning rate, batch size, and the allocation of frozen and unfrozen epochs affect convergence and generalization. The results complement the initial experiments by identifying the configurations that yield the most stable and accurate performance.

In Table 2, the combinations of learning rates, batch sizes, and epoch allocations considered in our experiments, the best configuration for each $(batch, lr)$ is reported with the corresponding accuracy, sensitivity, specificity, and F1-score. The systematic evaluation highlights the impact of hyperparameter tuning on segmentation




Fig. 3. Qualitative results obtained using the best-performing model identified during the hyperparameter optimization stage. First row: retinal fundus images. Second row: original images with the manual MTA outline. Third row: ground-truth MTA segmentation. Fourth row: initial segmentation of all arterial structures generated by the model. Fifth row: exclusive segmentation of the Major Temporal Arcade produced by the optimized model

performance. The highest F1-score (0.7700) was obtained using a batch size of 2 and a learning rate of $1e-4$, achieving high accuracy (0.9923) and specificity (0.9985).

These findings suggest that smaller batches tend to enhance fine-grained segmentation performance, while larger batches stabilize convergence but require more iterations to achieve optimal results. Overall, the results demonstrate that hyperparameter tuning is key to maximizing segmentation quality.

On average, the model required 0.003 seconds per image with a batch size of 1, while larger batch sizes reduced training variability but increased GPU memory demands. These results suggest that fine-tuning hyperparameters not only affects

accuracy but also the efficiency of the system, which is important for potential clinical deployment.

Figure 4 shows the training loss and validation F1-score across epochs for representative configurations. We observe that fine-tuning the encoder improves segmentation metrics after unfreezing, confirming the benefit of adapting pre-trained filters to the specific dataset.

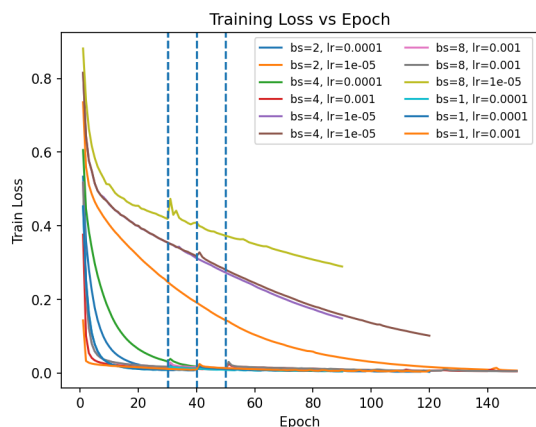


Fig. 4. Training loss curves obtained for different combinations of batch size and learning rate. The curves show the evolution of the training loss across epochs, demonstrating the convergence behavior of the models. Dashed vertical lines indicate the transition between frozen and unfrozen phases during fine-tuning

5 Conclusions

This work introduced a new deep learning strategy for automatically segmenting the MTA in retinal fundus images using a U-Net-based architecture. The proposed methodology effectively isolates the target vessel while removing non-relevant vascular structures by integrating a segmentation stage, transfer learning from the VGG16 model, and an attention module for refinement. With an accuracy of 0.9923 and an F1-score of 0.7700 on the test set, the experimental results indicate that the approach surpasses leading vascular segmentation techniques.

Additionally, the model achieves outstanding computational performance, requiring only 0.003 seconds per image on average. These findings

highlight the potential of deep learning-based architectures to deliver accurate and efficient MTA segmentation, contributing to improved retinal vascular assessment in clinical applications.

An extensive experimental evaluation showed that the performance of the model can be further improved through proper fine-tuning of key hyperparameters, such as the learning rate, batch size, and number of epochs. The comparative analysis revealed that small variations in these parameters can lead to significant improvements in segmentation accuracy and F1-score, highlighting the importance of optimizing the training process rather than relying solely on architectural enhancements.

Finally, accurate segmentation of the MTA can provide valuable insights for retinal vascular characterization and may support the early detection of diseases associated with structural vascular changes, such as diabetic retinopathy or glaucoma. Future work will explore the integration of multi-task learning strategies, where segmentation is optimized together with related objectives such as vessel classification, tortuosity estimation, or artery-vein discrimination. Joint optimization may strengthen the learned representations and improve robustness in cases involving image noise, variable lighting, or atypical vascular patterns.

Additional progress may also be achieved by incorporating clinical priors into the model, including anatomical constraints, probabilistic maps of expected vessel trajectories, or disease-specific biomarkers that help guide the attention mechanism toward clinically relevant regions.

Another direction involves extending the framework to support cross-dataset generalization through domain adaptation or transfer learning techniques, reducing sensitivity to variations in imaging devices, acquisition settings, and patient demographics. Expanding the dataset with diseased cases, pediatric images, or multimodal inputs could further improve the adaptability of the model in clinical workflows. Longitudinal analysis pipelines offer an additional opportunity, enabling automatic detection of subtle vascular changes across follow-up visits to support early diagnosis and treatment monitoring.

Acknowledgments

This research has been supported by the project IxM-SECIHTI No. 3097-7185.

References

1. **Alvarado-Carrillo, D. E., Cruz-Aceves, I., Hernández-González, M. A., López-Montero, L. M. (2022).** Robust detection and modeling of the major temporal arcade in retinal fundus images. *Mathematics*, Vol. 10, pp. 1334.
2. **Ayres, F. J., Rangayyan, R. M. (2007).** Design and performance analysis of oriented feature detectors. *Journal of Electronic Imaging*, Vol. 16, pp. 023007:1–023007:12.
3. **Boyd, C., Brown, G. C., Kleinig, T. J., Mayer, W., Dawson, J., Jenkinson, M., Bezak, E. (2024).** Hyperparameter selection for dataset-constrained semantic segmentation: Practical machine learning optimization. *Journal of Applied Clinical Medical Physics*, Vol. 25, No. 12, pp. e14542. DOI: 10.1002/acm2.14542.
4. **Chaudhuri, S., Chatterjee, S., Katz, N., Nelson, M., Goldbaum, M. (1989).** Detection of blood vessels in retinal images using two-dimensional matched filters. *IEEE Transactions on Medical Imaging*, Vol. 8, pp. 263–269.
5. **Cruz-Aceves, I., Cervantes-Sanchez, F., Avila-Garcia, M. S. (2018).** A novel multiscale gaussian-matched filter using neural networks for the segmentation of x-ray coronary angiograms. *Journal of Healthcare Engineering*, Vol. 2018, pp. 5812059. DOI: 10.1155/2018/5812059.
6. **Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y. (2017).** Deformable convolutional networks. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 764–773. DOI: 10.1109/ICCV.2017.89.
7. **Eiho, S., Qian, Y. (1997).** Detection of coronary artery tree using morphological operator. *Computers in Cardiology*, Vol. 24, pp. 525–528.
8. **Frangi, A., Niessen, W., Vincken, K., Viergever, M. (1998).** Multiscale vessel enhancement filtering. *Medical Image Computing and Computer-Assisted Intervention (MICCAI'98)*, Springer, Vol. 1496, pp. 130–137.
9. **Hanifi, S., Cammarono, A., Zare-Behtash, H. (2024).** Advanced hyperparameter optimization of deep learning models for wind power prediction. *Renewable Energy*, Vol. 221, pp. 119700. DOI: 10.1016/j.renene.2023.119700.
10. **Ilemobayo, J. A., Durodola, O. (2024).** Hyperparameter tuning in machine learning: A comprehensive review. *Journal of Engineering Research and Reports*, Vol. 26, No. 6, pp. 388–395.
11. **Lenau, A., Dimiduk, D., Niezgod, S. R. (2025).** Importance of hyper-parameter optimization during training of physics-informed deep learning networks. *Integrating Materials and Manufacturing Innovation*, Vol. 14, pp. 115–135. DOI: 10.1007/s40192-025-00394-6.
12. **Materka, A., Jurek, J. (2024).** Using deep learning and b-splines to model blood vessel lumen from 3d images. *Sensors*, Vol. 24, No. 3, pp. 846. DOI: 10.3390/s24030846.
13. **Nguyen, U. T., Bhuiyan, A., Park, L. A., Ramamohanarao, K. (2013).** An effective retinal blood vessel segmentation method using multi-scale line detection. *Pattern Recognition*, Vol. 46, pp. 703–715.
14. **Oloumi, F., Rangayyan, R. M., Ells, A. L. (2012).** Parabolic modeling of the major temporal arcade in retinal fundus images. *IEEE Transactions on Instrumentation and Measurement*, Vol. 61, pp. 1825–1838.
15. **Oloumi, F., Rangayyan, R. M., Ells, A. L. (2013).** Computer-aided diagnosis of proliferative diabetic retinopathy via modeling

of the major temporal arcade in retinal fundus images. *Journal of Digital Imaging*, Vol. 26, pp. 1124–1130.

16. **Otsu, N. (1979)**. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 9, pp. 62–66.
17. **Probst, P., Boulesteix, A.-L., Bischl, B. (2019)**. Tunability: Importance of hyperparameters of machine learning algorithms. *Journal of Machine Learning Research*, Vol. 20, No. 53, pp. 1–32.
18. **Qian, Y., Eiho, S., Sugimoto, N., Fujita, M. (1998)**. Automatic extraction of coronary artery tree on coronary angiograms by morphological operators. *Computers in Cardiology*, Vol. 25, pp. 765–768.
19. **Rangayyan, R. M., Ayres, F. J., Oloumi, F., Eshghzadeh-Zanjani, P. (2008)**. Detection of blood vessels in the retina with multiscale gabor filters. *Journal of Electronic Imaging*, Vol. 17, pp. 023018:1–023018:7.
20. **Ronneberger, O., Fischer, P., Brox, T. (2015)**. U-net: Convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, Springer International Publishing, pp. 234–241. DOI: 10.1007/978-3-319-24574-4.
21. **Salem, N. M., Salem, S. A., Nandi, A. K. (2007)**. Segmentation of retinal blood vessels based on analysis of the hessian matrix and clustering algorithm. 15th European Signal Processing Conference (EUSIPCO 2007), Springer, pp. 428–432.
22. **Singh, L. K., Khanna, M., Thawkar, S., et al. (2024)**. Deep-learning based system for effective and automatic blood vessel segmentation from retinal fundus images. *Multimedia Tools and Applications*, Vol. 83, pp. 6005–6049. DOI: 10.1007/s11042-023-15348-3.
23. **Soto-Álvarez, J. A., Cruz-Aceves, I., Hernández-Aguirre, A., Hernández-González, M. A., López-Montero, L. M., Solorio-Meza, S. E. (2023)**. Numerical modeling of the major temporal arcade using bumdá and jacobi polynomials. *Axioms*, Vol. 12, No. 2, pp. 137. DOI: 10.3390/axioms12020137.
24. **Soto-Álvarez, J. A., Cruz-Aceves, I., Hernández-Aguirre, A., Hernández-González, M. A., López-Montero, L. M., Solorio-Meza, S. E. (2023)**. Numerical modeling of the major temporal arcade using bumdá and jacobi polynomials. *Axioms*, Vol. 12, No. 2, pp. 137. DOI: 10.3390/axioms12020137.
25. **Yagis, E., Aslani, S., Jain, Y., Zhou, Y., Rahmani, S., Brunet, J., Bellier, A., Werlein, C., Ackermann, M., Jonigk, D., Tafforeau, P., Lee, P. D., Walsh, C. L. (2024)**. Deep learning for 3d vascular segmentation in hierarchical phase contrast tomography: a case study on kidney. *Scientific Reports*, Vol. 14, pp. 27258. DOI: 10.1038/s41598-024-77582-5.

Article received on 27/08/2025; accepted on 26/11/2025.

**Corresponding author is Ivan Cruz-Aceves.*