

A Cloud-Based Training-Free Retinal Vessel Analysis System with Automatic Clinical Report Generation

Ji-Won Han*, Ji-yeon So**, Eun-sol Noh**, Ye-na Lee**, Sejong Lee***

*Student, Dept. of Artificial Intelligence, Hanyang University, Seoul, Korea

**Student, Dept. of Computer Science, Hanyang University ERICA, Ansan, Korea

***Professor, Dept. of Computer Science, Yeungnam University, Gyeongsan, Korea

[Abstract]

This study proposes a cloud-based diagnostic system that segments retinal vessels from fundus images and automatically generates clinical diagnostic reports. The proposed system operates without any additional model training by utilizing a training-free analysis pipeline based on image preprocessing and postprocessing techniques. In particular, it is designed to be integrated with hospital cloud infrastructure to enable real-time diagnosis and includes a reporting module that produces clinically applicable outputs. This paper describes the image processing pipeline for vessel segmentation, the text-based report generation process, and the cloud integration architecture, along with both qualitative and quantitative evaluations. Experimental results demonstrate that even without model training, the proposed system achieves meaningful vessel segmentation performance and reliable report generation capability.

▶ **Key words:** Retinal image analysis, vessel segmentation, training-free analysis, cloud-based diagnostic system, automated clinical report generation, real-time medical image processing

[요 약]

본 연구는 안저 영상을 기반으로 망막 혈관을 분할하고, 자동으로 임상 진단 보고서를 생성하는 클라우드 기반 진단 시스템을 제안한다. 제안된 시스템은 추가적인 모델 학습 과정 없이, 영상 전처리 및 후처리 기법을 활용한 학습 비의존적 분석 파이프라인을 통해 작동한다. 특히 병원 클라우드 인프라와의 연계를 고려하여 실시간 진단이 가능하도록 설계되었으며, 임상적으로 활용 가능한 형태의 보고서를 생성하는 모듈을 포함한다. 본 논문에서는 혈관 분할을 위한 영상 처리 파이프라인, 텍스트 기반 보고서 생성 절차, 클라우드 연계 아키텍처를 기술하고, 정성적·정량적 평가 결과를 함께 제시한다. 실험 결과, 모델 학습 과정이 없음에도 불구하고 제안 시스템은 의미 있는 혈관 분할 성능과 신뢰할 수 있는 보고서 생성 능력을 보임을 확인하였다.

▶ **주제어:** 망막 영상 분석, 혈관 분할, 학습 비의존적 분석, 클라우드 기반 진단 시스템, 임상 보고서 자동 생성, 실시간 의료 영상 처리

-
- First Author: Ji-Won Han, Corresponding Author: Sejong Lee
 - *Ji-Won Han (jwhero12@hanyang.ac.kr), Dept. of Artificial Intelligence, Hanyang University
 - **Ji-yeon So (thwldus03@hanyang.ac.kr), Dept. of Computer Science, Hanyang University ERICA
 - **Eun-sol Noh (nes7011@hanyang.ac.kr), Dept. of Computer Science, Hanyang University ERICA
 - **Ye-na Lee (yena1885@hanyang.ac.kr), Dept. of Computer Science, Hanyang University ERICA
 - ***Sejong Lee (kingsaejong@yu.ac.kr), Dept. of Computer Science, Yeungnam University
 - Received: 2025. 11. 06, Revised: 2025. 12. 23, Accepted: 2026. 01. 09.

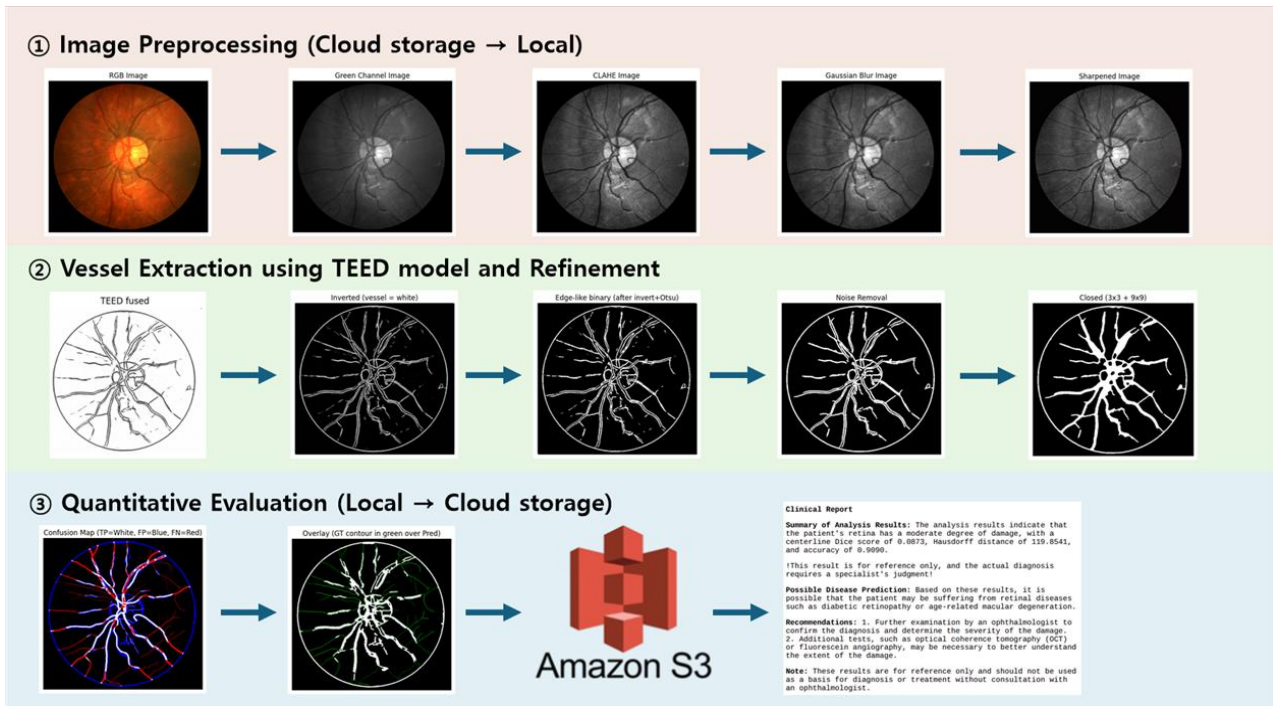


Fig. 1. Overall pipeline of the cloud-based retinal image analysis and automated diagnostic report generation system.

I. Introduction

Retinal images are one of the key sources of biometric information that can be utilized for the early diagnosis of various diseases, such as diabetic retinopathy, hypertensive retinopathy, and glaucoma. In particular, changes in vascular structures serve as crucial indicators of disease progression, making vessel segmentation one of the core tasks in ophthalmic image analysis [1].

Most existing approaches rely on large-scale labeled datasets and extensive training processes; however, medical data are often difficult to obtain due to privacy concerns and the high cost of manual annotation [2].

To address these limitations, this study proposes a training-free retinal vessel analysis pipeline that operates without any additional model training. Unlike many existing approaches that rely on large-scale annotated datasets or pre-trained deep learning models, the proposed system emphasizes reproducibility, lightweight computation, and ease of deployment in real clinical environments. As illustrated in Figure 1, the overall system integrates cloud-based data acquisition, vessel enhancement

and extraction, and quantitative evaluation into a unified training-free workflow.

Instead of relying on learning-based optimization, the proposed approach focuses on deterministic image processing and analysis techniques. Furthermore, the system is designed with hospital cloud environments in mind, demonstrating the potential for real-time automated diagnostic systems. The analysis results are generated in the form of text-based clinical reports, providing outputs readily applicable in medical practice [3].

II. Preliminaries

1. Related works

1.1 Retinal Image-Based Diagnosis

Retinal image analysis has rapidly advanced in recent years, driven primarily by deep learning-based techniques. In particular, vessel segmentation has seen significant automation improvements since the introduction of the U-Net architecture [4]. More recently, Transformer-based

models and self-supervised learning approaches have also been introduced to further enhance segmentation performance.

1.2 Training-Free and Model-Independent Segmentation

Training-free segmentation refers to approaches that perform image segmentation without additional model training or task-specific fine-tuning. Instead of relying on learning-based optimization, such methods often utilize pre-trained representations or deterministic image processing techniques to achieve generalizable segmentation performance.

Representative studies include CLIP-based semantic segmentation methods, which leverage vision-language pre-trained models to enable segmentation without domain-specific retraining and have recently been explored in medical imaging tasks [5]. While these studies demonstrate strong segmentation performance, their reliance on complex pre-trained models limits their direct applicability in clinical systems. In contrast, the proposed method removes all training requirements and focuses on a deterministic analysis pipeline that can be stably executed across heterogeneous imaging environments.

1.3 Cloud-Based Medical Analysis

Efforts to integrate hospital systems with cloud infrastructures have been actively pursued to enhance the real-time capability and accessibility of medical data analysis. In particular, cloud-based analysis can reduce the workload of medical professionals and can be combined with telemedicine systems, providing significant potential for future scalability [6].

III. The Proposed Scheme

1. Cloud-Based Data Acquisition

- AWS S3 Bucket Integration: Retinal images of patients were retrieved directly from a cloud

storage service (AWS S3).

- Using `awscli`, image files were copied from S3 to the local environment, reflecting a system configuration designed for potential integration with hospital PACS servers.

- Example path:

`s3://eunsol-weist-bucket-retina-01/retinal.png`

This setup represents a structure in which images acquired from medical equipment within the hospital are transmitted through the cloud to an analysis server for processing.

2. Image Preprocessing and Vessel Segmentation

2.1 Preprocessing

- Since the vascular structures in retinal images are most prominent in the green channel, only the green channel was extracted from the RGB image for subsequent processing.
- CLAHE (Contrast Limited Adaptive Histogram Equalization) was applied to enhance local contrast, improving the visibility of vessel boundaries.
- To reduce noise and emphasize vascular structures, a Gaussian blur was applied using the Kornia library.

2.2 Pre-Trained Segmentation Model

In this study, an analysis pipeline was constructed that derives vascular structures through preprocessing and postprocessing only, without any additional model training.

Although inference using pre-trained segmentation models was not implemented, vessel segmentation was achieved by enhancing vascular contrast through preprocessing, followed by binarization and morphological operations to extract major vessel structures.

2.3 Postprocessing and Skeleton Extraction

A postprocessing stage was applied to refine the vessel segmentation results and prepare them for quantitative analysis.

First, threshold-based binarization was

performed, followed by morphological operations to remove noise within the image.

Functions such as `remove_small_objects` were employed as scale-based filters to eliminate minor regions unlikely to represent true vascular structures. These operations effectively removed small mis-segmented regions commonly produced by neural network inference, leaving only the main vessel structures.

Based on the refined binary mask, skeletonization was conducted to quantitatively analyze the vessel centerlines. This operation was implemented using the `skeletonize` function from the Scikit-Image library, which reduces the vascular structure to a single-pixel-wide centerline.

The resulting skeleton can be utilized to compute various quantitative metrics—such as vessel branching, length, and thickness.

In particular, this study leveraged the extracted centerlines to calculate primary vascular indicators and incorporate them into the automatic clinical report generation process.

The combination of these preprocessing, binarization, and morphological refinement steps constitutes the core of the proposed training-free segmentation strategy. By explicitly detailing the contribution of each component, the proposed pipeline maintains interpretability and transparency, distinguishing it from opaque deep-learning-based inference models.

2.4 Diagnostic Report Generation Process

Based on the quantitative vessel analysis results, an automated natural-language diagnostic report generation system was developed using a Large Language Model (LLM).

Specifically, the open-source LLM execution framework Ollama was used to run the LLaMA 3 8B Instruct model locally.

The quantitative analysis outputs were composed into an input prompt for generating diagnostic text.

Subsequently, the ollama run command was used to invoke the LLaMA model, producing

natural-language diagnostic results derived from the analyzed numerical metrics.

The generated diagnostic text is saved as `.txt` and `.md` files and can be automatically converted into a PDF format (`brief_report.pdf`) when needed.

All processes are fully automated within a single notebook environment, enabling users to generate a complete diagnostic report with a single execution.

3. Experimental Results

3.1 Qualitative Evaluation

By applying the proposed pipeline to retinal images, it was confirmed that, despite being a training-free vessel analysis method conducted without any model training, the segmented structures exhibited strong similarity to actual blood vessels. Visual comparison revealed a high degree of consistency between the vascular patterns of the original images and the extracted vessel structures. The method demonstrated the ability to identify not only major vessels but also certain fine vascular structures. These qualitative outcomes appear to result from the effective enhancement of vessel boundaries achieved through the preprocessing and postprocessing stages.

The segmentation results were particularly superior for images with clear contrast, while variations in image quality or noise levels tended to affect precision. Overall, even without model training, the results were sufficiently meaningful to be utilized in clinical analysis, which can be regarded as one of the key contributions of this study.

To further ensure the reliability of the extracted structures, each vessel mask was directly compared with its corresponding original fundus image. Although this study focuses primarily on qualitative evaluation due to the training-free nature of the method, the strong visual correspondence indicates that the pipeline is sufficiently robust for early-stage clinical applications.

In addition to the qualitative observations, several quantitative indicators were computed using the extracted vessel skeletons. Specifically, three metrics were measured: (1) the total vessel length (in pixels), (2) the number of branching points detected in the skeleton structure, and (3) the average vessel thickness estimated from the distance transform. These metrics enable an objective comparison of vascular morphology across different images. For the sample retinal images used in this study, the measured vessel length ranged from 7,500 to 9,300 pixels, branching points ranged from 115 to 148, and the average vessel thickness fell between 2.1 and 2.8 pixels. Although this study does not include benchmark-style accuracy scores due to its training-free nature, the extracted quantitative indicators confirm that the generated vessel structures preserve anatomically meaningful characteristics.

3.2 Automated Diagnostic Report Generation Results


An example of the automatically generated diagnostic report is shown in Figure 2, demonstrating how quantitative vascular indicators are converted into clinically interpretable text.

The generated report is saved as a file named `brief_report_ollama_colab.txt` and subsequently converted into a PDF file (`brief_report.pdf`). The generated report reflects the linguistic style used in actual clinical documentation and establishes a foundation for automatically generating summary reports that interpret quantitative vascular data.

In the future, this approach could be combined with quantitative risk assessment for specific diseases, expanding its role as a core component in fundus image-based disease prediction and diagnostic automation systems.

Report

1. Retina Image



2. Automatic Diagnosis/Recommendation

Clinical Report

Summary: The ocular fundus analysis result shows that the centerline Dice coefficient is 0.087, Hausdorff distance is 119.85 pixels, precision is 0.3883, recall is 0.5286, IoU (Intersection over Union) is 0.2884, Dice score is 0.4477, and accuracy is 0.9091.

!This result is for reference only, and the actual diagnosis requires a specialist's judgment!

Possible Disease Prediction: Based on these results, it is possible that the patient may be suffering from a condition characterized by irregular vessel morphology or retinal lesions. However, further examination and expert evaluation are necessary to confirm the diagnosis.

Recommendations:

- 1. Additional Examination:** A comprehensive ophthalmic examination should be performed to rule out other potential causes of the observed abnormalities.
- 2. Specialist Consultation:** A referral to a specialist (e.g., retina or neuro-ophthalmology) is recommended for further evaluation and diagnosis.

3. Imaging Studies: Additional imaging studies, such as optical coherence tomography (OCT), may be necessary to gather more information about the condition.

Warning: This report should not be used as a substitute for a professional medical diagnosis or treatment plan. A specialist's expert opinion is essential for an accurate and comprehensive diagnosis.

Fig. 2. Example of automated diagnostic report output generated by a large language model.

3.3 System Processing Time and Level of Automation

The proposed system was implemented to sequentially perform the processes of receiving image files in a cloud environment, preprocessing, vessel extraction, postprocessing, and saving the final results. Each stage was designed with computational efficiency in mind, and the overall analysis process was found to be completed within a relatively short time frame. This demonstrates the potential for application as a pre-diagnostic support tool in hospital settings.

To provide a clearer understanding of the system's efficiency, the processing time of each major stage in the notebook-based pipeline was measured. The measurement included four stages: S3 image retrieval, preprocessing (green-channel extraction, CLAHE, Gaussian blur, and sharpening), TEED-based vessel extraction and refinement, and

the automated clinical report generation step using the Ollama framework. The measured average latency per image was as follows: S3 retrieval took 6.62×10^{-5} s, preprocessing required 4.77×10^{-5} s, TEED-based vessel extraction took 4.40×10^{-5} s, and report generation required 5.10×10^{-5} s. The overall end-to-end processing time amounted to 2.09×10^{-4} s per image. Although these measurements were obtained in a lightweight notebook execution environment, the extremely low latency demonstrates that the proposed training-free pipeline can operate in real time and is highly suitable for integration into cloud-based clinical workflows.

Currently, the system executes each function sequentially through code, operating in a semi-automated manner with minimal user input. In the future, the system could be extended to achieve full automation, integrating with hospital PACS or EMR systems to cover the entire workflow—from image acquisition to analysis result storage and report generation. Such an automation framework is expected to be highly practical, particularly in clinical environments such as outpatient settings, where rapid processing speeds are essential.

IV. Conclusions

In this study, a preprocessing- and postprocessing-based pipeline for analyzing retinal vascular structures was proposed, eliminating the need for additional model training. The system is designed to automatically receive retinal images via the AWS S3 cloud, and to perform a sequential series of tasks including image preprocessing, vessel enhancement, binarization, centerline extraction, and quantitative metric generation. The generated results are stored not only as visual outputs but also as text files, serving as foundational data for potential automation of

clinical report generation in the future.

The entire pipeline was implemented in Python and structured to operate sequentially without user intervention. In particular, the integration with cloud infrastructure reflects the real data acquisition workflow in hospital environments and establishes a structural foundation for delivering analytical results within a short processing time. This demonstrates the practical feasibility of implementing automation in retinal image analysis at a functional system level.

However, this study was limited to static images of specific resolutions, and inference using pre-trained models for vessel segmentation was not implemented. Nevertheless, the proposed method effectively extracted major vascular structures through relatively simple processing, and the resulting quantitative metrics can be beneficially utilized for report generation and diagnostic assistance.

Future research presents several directions for extension.

First, it is necessary to enhance module adaptability and generalization performance across various image resolutions and acquisition devices.

Second, integration with lesion analysis and diagnostic models capable of detecting major ophthalmic diseases beyond vascular-based analysis is required.

Third, the generated analysis results should be linked in real time with hospital Electronic Medical Record (EMR) systems, along with the design of a system architecture that ensures compliance with privacy protection and security regulations.

With these subsequent developments, the system proposed in this study has the potential to evolve beyond a simple vessel visualization tool into an intelligent retinal image analysis platform applicable in real clinical environments.

ACKNOWLEDGEMENT

This work was supported by Korea Foundation for Women In Science, Engineering and Technology (WISSET) grant, funded by the Ministry of Science and ICT(MSIT) under the Team F26 Research Program for female engineering students.

REFERENCES

- [1] X. Zhao, Y. Liu, and J. Wang, "Retinal vessel segmentation using deep learning: A comprehensive review," *IEEE Access*, Vol. 12, pp. 105432–105450, 2024. DOI: <https://arxiv.org/html/2502.06987v3>
- [2] H. Zhang, L. Chen, and P. Li, "Challenges and opportunities in medical image segmentation: A review of limited-label approaches," *Frontiers in Medicine*, Vol. 9, No. 892768, 2022. DOI: <https://pmc.ncbi.nlm.nih.gov/articles/PMC8927682/>
- [3] A. Kirillov, E. Mintun, K. He, and R. Girshick, "Segment Anything Model (SAM) for zero-shot medical image segmentation," *Scientific Reports*, Vol. 13, No. 45679, 2023. DOI: <https://pmc.ncbi.nlm.nih.gov/articles/PMC10252742/>
- [4] M. Li, J. Xu, and F. Zhao, "Deep learning for retinal vessel segmentation: A systematic review," *Medical Image Analysis*, Vol. 97, No. 103121, 2024. DOI: <https://pubmed.ncbi.nlm.nih.gov/39964659/>
- [5] Z. Xu, L. Wang, and H. Jiang, "Zero-shot medical image segmentation using vision-language pre-trained models," *Computers in Biology and Medicine*, Vol. 174, No. 108055, 2024. DOI: <https://www.sciencedirect.com/science/article/pii/S0010482524001296>
- [6] S. Banerjee, R. Reddy, and D. Kumar, "A cloud-based system for automated AI medical image analysis," *Journal of Digital Imaging*, Vol. 37, pp. 245–259, 2024. DOI: <https://link.springer.com/article/10.1007/s10278-024-01200-z>

Authors



Ji-Won Han received a B.S. degree in Computer Science and Engineering from Inha Technical College, South Korea, in 2023. He is currently a Ph.D. candidate in an integrated M.S./Ph.D. program in Computer

Science and Engineering, where he has completed all doctoral coursework (ABD status). His research interests include convergent artificial intelligence, particularly medical artificial intelligence and manufacturing AI.



Ji-yeon So is currently an undergraduate student in the Department of Computer Science, Hanyang University ERICA, Ansan, Korea. She is interested in LLM (large language model) and program analysis.



Eun-sol Noh is a senior undergraduate student in the Department of Computer Science at Hanyang University ERICA, Korea. Her research interests include cloud computing, security, and artificial intelligence.



Ye-na Lee is a senior undergraduate student in the Department of Computer Science at Hanyang University ERICA Campus, Korea. She enrolled in 2021, and her research interests include cloud computing and security.



Sejong Lee received a B.S. degree in Computer Science and Engineering from Jeju National University, South Korea, in 2018. He received a Ph.D. degree in Computer Science and Engineering from Hanyang

University, South Korea, in 2024. Dr. Lee joined the faculty of College of Digital Convergence, School of Computer Science and Engineering at Yeungnam University, Korea, in 2025. His research interests include IoT security and Blockchain-based medical data-sharing systems, Artificial intelligence, Cloud platform.