망막 혈관 분할을 위한 소수-샷 도메인 적응에 관한 연구

Muhammad Omer[1], Sardar Jaffar Ali[2], 추현승[3]
[1]소프트웨어학과, 성균관대학교
[2]AI시스템공학과, 성균관대학교
[3]전자전기컴퓨터공학부, 성균관대학교
email : {omer389, jaffar.ali, choo}@skku.edu

Few-shot Domain Adaptation for Retinal Vessel Segmentation

Muhammad Omer[1], Sardar Jaffar Ali[2], Choo Hyunseung[3]
[1]Dept. of Computer Science and Engineering, Sungkyunkwan University
[2]Dept. of AI Systems Engineering, Sungkyunkwan University
[3]Dept. of Electrical and Computer Engineering, Sungkyunkwan University

요 약

We propose a few-shot domain adaptation scheme that fine-tunes a DRIVE-trained U-Net (ResNet-34) on CHASEDB1 using only k labeled target images. On zero-shot transfer the model achieves Dice 0.647 / IoU 0.479 on CHASEDB1; with just k=5 labels it reaches Dice 0.732 / IoU 0.578, demonstrating strong label-efficiency. The lightweight model enables near real-time inference, making the approach practical for point-of-care screening.

1. 서론

Retinal vessel segmentation underpins automated screening and quantitative assessment for diabetic retinopathy, hypertensive retinopathy, and glaucoma, among others. Despite the progress of deep learning, models trained on one site or device often deteriorate on another due to domain shift—differences in camera optics, illumination, resolution, and patient cohorts. In medical imaging, obtaining sufficiently many target-domain annotations is costly and slow, making rapid, safe adaptation a practical bottleneck.

Prior adaptation strategies include unsupervised domain adaptation with adversarial alignment or style transfer, self-training with pseudo-labels, and test-time adaptation. While effective in curated settings, these pipelines can be data-hungry, training-unstable, or operationally complex, which hinders deployment in clinics. Few-shot approaches alleviate labeling cost but often rely on meta-learning or multi-stage procedures that complicate maintenance. The contributions are as follows:

(1) A innovative and reproducible few-shot adaptation protocol for retinal vessel segmentation that requires only k labeled target images and no specialized losses.

(2) A label-efficiency characterization ($k \in \{1,3,5\}$) with consistent improvements over zero-shot transfer, plus qualitative overlays highlighting reductions in thin-vessel misses and disc-adjacent false positives.

(3) A compact, end-to-end baseline with public scripts and preprocessing steps, aimed at lowering the barrier to clinical adaptation across devices and sites.

2. 에이전트 개발도구의 요구사항

We study few-shot domain adaptation for retinal vessel segmentation: a U-Net[1] with a ResNet-34 encoder is first trained on DRIVE[2] (source), then lightly fine-tuned on k labeled images from CHASEDB1[3] (target), with k ∈ {1, 3, 5}. The focus is on a simple, deployable protocol that minimizes extra components (no adversarial losses, no meta-learning) while maximizing label-efficiency and reproducibility.

For source training we use the DRIVE training set and reserve 10% of its training images for validation. Because the common Kaggle mirror does not include the official DRIVE test vessel annotations, our primary evaluation is conducted on CHASEDB1: in each run we fix a random seed (42), sample k images for adaptation, keep 4 images for validation, and use the remainder for testing (the test size decreases slightly as k grows). All images are read in RGB, resized to 512×512, scaled to [0,1], and masks are read in grayscale and binarized at 0.5 with nearest-neighbor resizing.

The model is U-Net (ResNet-34, ImageNet initialization)[5] with a single-channel sigmoid output for vessels—chosen to balance accuracy and near real-time throughput at 512×512 on a commodity GPU. We optimize with AdamW[4] and BCEWithLogits loss. For source training, we use lr = 3×10^{-4} , weight decay = 1×10^{-4} , early stopping on DRIVE-val Dice. For few-shot adaptation, we fine-tune all weights on the k labeled CHASE images with lr = 1×10^{-5}, weight decay = 1×10^{-4} , 20 epochs and patience 5 on the fixed 4-image CHASE validation set (batch size 2; mixed precision optional).

Evaluation reports Dice and Intersection-over-Union (IoU) on the CHASEDB1 test split using a 0.5 probability threshold, averaged over images. We present both zero-shot transfer (DRIVE→CHASEDB1 without adaptation) and k-shot results after fine-tuning[6], and we log the exact split sizes for transparency.

3. 실험 및 결과

We evaluate zero-shot transfer (DRIVE→CHASEDB1 without adaptation) and few-shot adaptation with k ∈ {1, 3, 5} labeled CHASEDB1 images. On zero-shot, the model attains Dice 0.647 / IoU 0.479. With k=1/3/5, performance improves to Dice 0.617/0.723/0.732 and IoU 0.447/0.567/0.578, showing a clear, monotonic label-efficiency trend. Figure 1 visualizes the Dice vs. k curve.

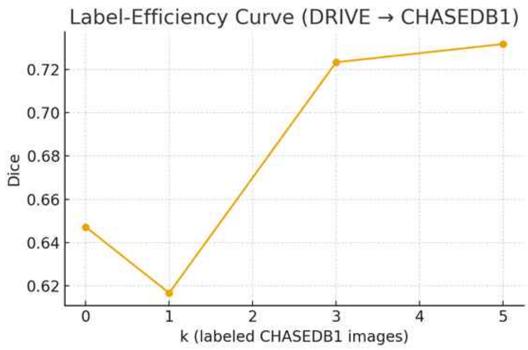

Figure 1. Shows the dice vs k curve, the best dice value is at k=5.

For completeness, we keep the protocol minimal (no adversarial losses, no heavy augmentation) and fix the random seed to stabilize the k-shot selection. Test-set size decreases slightly as k increases because k images are moved into the adaptation set; exact split sizes are logged. Inference at 512×512 remains near real-time on a commodity GPU, supporting point-of-care scenarios (screening/triage). Figures 2-3 provide compact bar charts for Dice and IoU. Qualitative overlays (baseline vs. k=5) show more continuous centerlines and recovery of thin peripheral vessels while reducing disc-adjacent false positives.

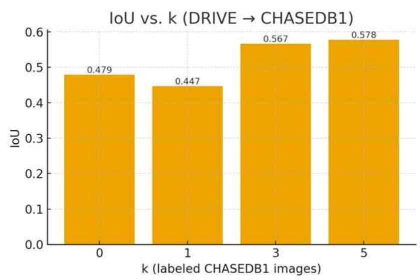

Figure 2. Shows IoU values with respect to k values.

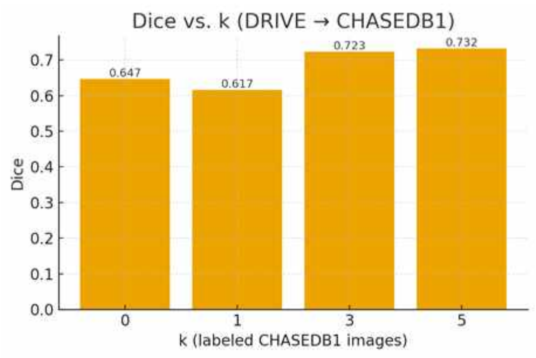

Figure 3. Shows Dice values with respect to k values.

4. 결론

We introduced a minimal, deployment-oriented few-shot domain adaptation pipeline for retinal vessel segmentation that fine-tunes a DRIVE-trained U-Net on k labeled CHASEDB1 images. Despite its simplicity (no adversarial losses or complex staging), the approach delivers consistent, label-efficient gains, improving from Dice 0.647 / IoU 0.479 in zero-shot transfer to Dice 0.732 / IoU 0.578 with just k = 5, while preserving near real-time inference at 512×512. These properties directly translate to healthcare impact: lower annotation cost, faster onboarding of new cameras and sites, and a straightforward maintenance path for clinical IT/MLOps, making the method a practical baseline for point-of-care screening and continuous service updates. Looking ahead, we see a clear path to responsible adoption: (i) multi-center validation across additional datasets and devices to stress-test robustness, (ii) uncertainty-aware outputs and calibration (e.g., ECE/Brier) to support human-in-the-loop review, (iii) lightweight enhancements such as batch-norm adaptation, test-time augmentation, or active selection of k images, and (iv) integration with PACS/DICOM and standard QA for clinical rollout. Taken together, our results suggest a high-utility, low-friction route for adapting vessel segmentation models in real healthcare settings.

Acknowledgement

This work was partly supported by the Korea government (MSIT), IITP, Korea, under the ICT Creative Consilience program (IITP-2025-RS-2020-II201821, 40%), (RS-2024-00459512, Development of Brain Disease (Stroke) Prediction Model based on Fundus Image Analysis, 20%), (RS-2021-II212068, Artificial Intelligence Innovation Hub, 20%), (RS-2019-II190421, Artificial Intelligence Graduate School Program(Sungkyunkwan University), 20%).

참고문헌

[1] O. Ronneberger, P. Fischer, T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” MICCAI, 9351, pp. 234-241. 2015

[2] J. Staal, M. D. Abramoff, M. Niemeijer, M. A. Viergever, B. van Ginneken, “Ridge-Based Vessel Segmentation in Color Images of the Retina,” IEEE Transactions on Medical Imaging, 23(4), pp. 501-509, 2004.

[3] M. M. Fraz, P. Remagnino, A. Hoppe, B. Uyyanonvara, A. R. Rudnicka, C. G. Owen, S. A. Barman, “An Ensemble Classification-Based Approach Applied to Retinal Blood Vessel Segmentation,” IEEE Transactions on Biomedical Engineering, 59(9), pp. 2538-2548, 2012

[4] I. Loshchilov, F. Hutter, “Decoupled Weight Decay Regularization,” ICLR, 2019. (AdamW)

[5] P. Yakubovskiy, “Segmentation Models PyTorch,” Software/Repository, 2020.

[6] D. Wang, E. Shelhamer, S. Liu, B. Raj, T. Darrell, “Tent: Fully Test-Time Adaptation by Entropy Minimization,” ICLR, 2021