MorphoBoost: Morphology-Driven Boundary Enhancement Model for Accurate Segmentation of Langerhans Cells in Corneal Confocal Microscopy Images

Hongshuo Li[1], Ankai Dong[2], Tiande Zhang[1], Shijia Zhou[1], Yalin Zheng[3], Lei Mou[1*], and Yitian Zhao[1,4*]

[1] Ningbo Institute of Materials Technology and Engineering, Chinese Academy of Sciences, Ningbo, China
yitian.zhao@nimte.ac.cn, moulei@nimte.ac.cn
[2] School of Control and Computer Engineering, North China Electric Power University, Baoding, China
[3] Department of Eye and Vision Science, University of Liverpool, Liverpool, UK
[4] School of Biomedical Engineering, ShanghaiTech University, Shanghai, China

Abstract. Accurate segmentation of Langerhans cells (LCs) in corneal confocal microscopy (CCM) images is crucial for diagnosing and monitoring various ocular and systemic diseases. However, existing segmentation methods often struggle with the misidentification of activated LCs and inaccurate boundary delineation due to their complex morphological features and background noise. In this paper, we propose a novel segmentation framework, MorphoBoost, which integrates morphology-driven data augmentation and boundary optimization loss to address these challenges. MorphoBoost employs a "localization before segmentation" strategy, enhancing the diversity of activated LCs via spatial and appearance transformations, and refining segmentation boundaries through pixel-level and image-level optimizations. Our methods achieve state-of-the-art performance in segmenting both LCs types, especially activated ones. It establishes a new benchmark with a 17.10% increase in the Dice coefficient and a 5.71 decrease in modified Hausdorff distance over previous methods. This is bolstered by validation on clinical data.

Keywords: Morphology-driven augmentation · Langerhans cell segmentation · Dilation mask.

1 Introduction

Langerhans cells (LCs) are antigen-presenting cells primarily responsible for immune surveillance and regulation. They exhibit two states: an oval-shaped non-activated state and a curved-shaped activated state [11]. The morphological characteristics of the LCs such as maturity, density, size [5, 12, 13, 15, 6, 22], are closely related to many ocular or systemic diseases. In clinical practice, LCs images acquired by corneal confocal microscopy (CCM), are commonly used to

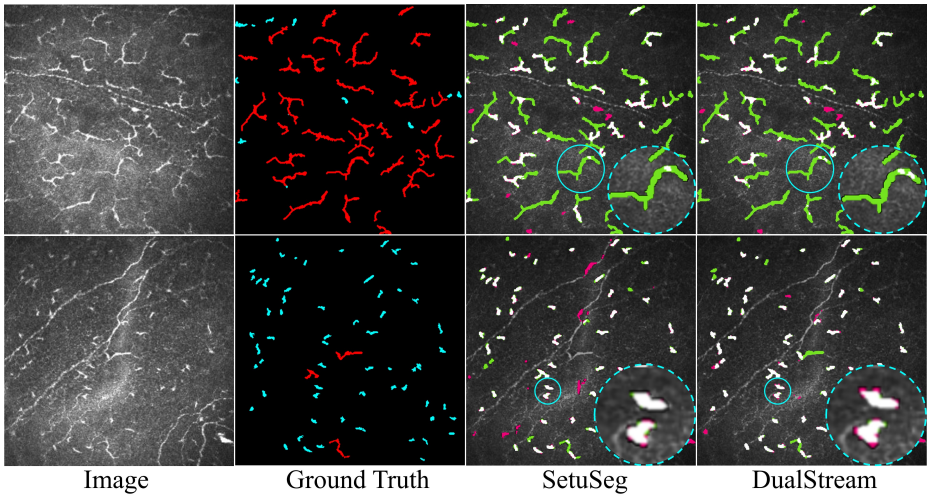

Fig. 1. Typical LCs segmentation issues in CCM images. Top row: misidentification of activated LCs; Bottom row: inaccurate boundary delineation. The automated results were obtained by two recent methods (SetuSeg [17] and DualStream [20]). Blue: non-activated LCs; Red: activated LCs; Green: undersegmentation; Pink: oversegmentation.

.

assist ophthalmologists with rich pathological information, for studying disease-related alternations. Manual annotation of LCs is able to support the quantitative analysis, however, it is time-consuming and subjective. Moreover, the inter- and intra-observer variations introduced by manual labeling greatly diminish the accuracy of quantitative assessment of LCs [24]. Therefore, a fully automated and accurate LCs segmentation method is essentially needed.

Many studies [17, 20, 19, 2, 1] have been conducted to automatically segment the LCs. Early works mainly focus on designing manual feature descriptors, such as Gaussian filtering [19] and morphological operations [2, 1] to obtain geometric features for better LCs extraction. However, these methods rely heavily on low-level features and require parameter tuning by hand, and thus segmentation performance is limited. Deep learning-based models have recently been established to improve segmentation accuracy, by exploring high-level features rather than relying on low-level ones. For example, Setu et al. [17] used Mask R-CNN [10] to achieve joint optimization of detection and segmentation tasks. Wu et al. [20] introduced a dual-stream network with target-guiding and semantic-guiding modules to collaboratively optimize segmentation and detection tasks.

Although the above-mentioned deep learning methods have achieved promising segmentation performances, they still struggle with the misidentification of activated LCs and inaccurate boundary delineation. For example, as shown in the top row of Fig. 1, activated LCs are often falsely identified as corneal nerves [4].

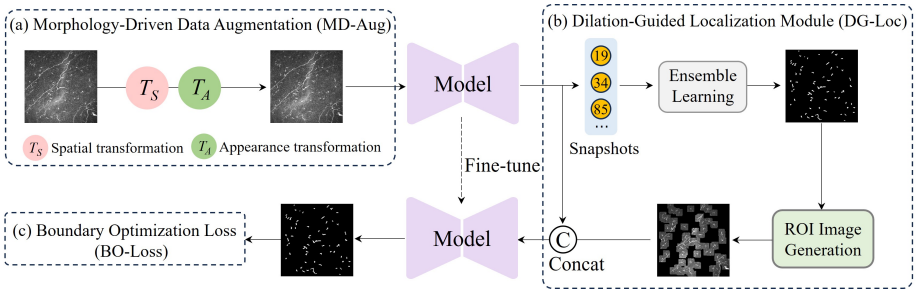

Fig. 2. The proposed method comprises three core components: MD-Aug, DG-Loc, and BO-Loss. Initially, data is augmented using MD-Aug to enhance diversity. The augmented data is then processed through DG-Loc to generate input data for the segmentation model. This input data is fed back into the model for fine-tuning, with a BO-Loss incorporated to enhance segmentation accuracy.

Additionally, many methods [17, 20] fail to preserve accurate boundaries, especially in low-quality CCM images, as shown in the bottom row of Fig. 1.

Inspired by [17, 20], we propose a morphology-driven boundary enhancement model named MorphoBoost that follows the "localization before segmentation" approach. In this framework, we design a morphology-driven data augmentation module and a boundary optimization loss to address the challenges of misidentification of activated LCs and inaccurate boundary delineation. The main contributions are summarized as follows:

(1) We propose a novel LCs segmentation method, MorphoBoost, which follows a "localization before segmentation" approach and integrates morphology-driven data augmentation and boundary optimization loss to address challenges related to misidentification and inaccurate boundary delineation.

(2) We design a Morphology-Driven Data Augmentation Module to enhance the diversity of activated LCs through spatial and appearance transformations. We also incorporate a Boundary Optimization Loss that integrates pixel-level and image-level components to refine segmentation boundaries.

(3) Our method sets a new benchmark in segmenting LCs, particularly activated ones, showcasing a 17.10% boost in the Dice coefficient and a 5.71 reduction in modified Hausdorff distance over prior techniques. Its efficacy is substantiated with clinical data validation.

2 Methods

2.1 Morphology-Driven Data Augmentation Module (MD-Aug)

The highly irregular morphologies of activated cells and their structural similarity to corneal nerves pose significant challenges for model learning. These challenges make it more difficult for models to accurately learn the features of activated LCs compared to non-activated ones. To enhance the model's ability

to identify activated LCs, we propose a Morphology-Driven Data Augmentation Module to augment the morphological diversity of activated cells. This module comprises two key components: spatial transformations and appearance transformations, both designed to comprehensively augment the feature diversity of activated LCs.

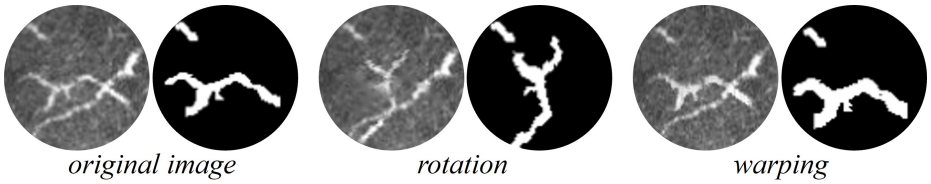

Fig. 3. Illustration of spatial transformations, including positional rotation and morphological warping.

For spatial transformations, we performed operations such as rotation and warping. Fig. 3 shows the results of the spatial transformations. The Shi-Tomasi corner detection algorithm [18] is utilized to identify critical locations, such as branch points and tips of the dendrites of LCs, which serve as control points for subsequent morphological transformations. Inspired by the Moving Least Squares (MLS) [14], we design a new approach to adjust the position and morphology of activated LCs. Specifically, we define control points $P = \{p_i\}_{i=1}^{n}$, where p_i represents the i-th detected corner. Target control points are then defined based on desired morphological changes, denoted as $\{q_i\}_{i=1}^{n}$. For any point u in the image, its transformed position $T(u)$ is calculated by minimizing the following weighted error:

$$E(u) = \sum_{i=1}^{n} \frac{1}{(1 + |u - p_i|)^2 \cdot |T(u) - q_i|^2} \tag{1}$$

The formula (2) enables local deformation of these cells by emphasizing contributions from points near control points, generating diverse synthetic images. This enhances the model's ability to capture detailed structures, increasing data diversity.

To preserve the overall morphology of activated LCs, regularization constraints are introduced to prevent excessive distortion during deformation. This ensures that the deformation remains as natural as possible within local regions and maintains the biological plausibility of the cells' morphology. The specific formula is as follows:

$$R = \sum_{u \in \Omega} \left(|T(u) - u| - \frac{1}{|\Omega|} \sum_{v \in \Omega} |T(v) - v| \right)^2 \tag{2}$$

Here, Ω represents the set of all points in the image, and $T(u)$ denotes the transformed position of point u. The term $\frac{1}{|\Omega|} \sum_{v \in \Omega} |T(v) - v|$ calculates the average

deformation magnitude across all points in the image. This regularization term works by comparing the local deformation magnitude $|T(u) - u|$ to the global average deformation magnitude. It ensures that local deformations are consistent with the overall deformation pattern, thereby maintaining the naturalness of the deformation while capturing the necessary structural details.

For appearance transformation, we apply a range of data augmentations, including color jittering, Gaussian noise, Gaussian blur, hue and saturation adjustments, random brightness and contrast modifications, and random flipping. These augmentations complement the spatial transformations by introducing variability in the visual appearance of the LCs, further enhancing the diversity of the activated LCs.

2.2 Dilation-Guided Localization Module (DG-Loc)

To enhance the efficiency and accuracy of segmenting LCs, we propose a Dilation-Guided Localization Module (DG-Loc) that dynamically generates a region of interest (ROI) image from initial segmentation results. This module reduces the image area processed by the segmentation network, mitigates background noise interference, and allows the model to focus on the complex morphological features of LCs.

ROI Image Generation. During model training, the DG-Loc retains N ($N = 4$) weight snapshots and employs ensemble learning to ensure the integrity of detected LCs within the image. This process generates a dilation mask M, which guides the creation of the ROI. A dynamic dilation kernel, whose size is adjusted through Gaussian sampling, is applied to the dilation mask M. This kernel is then multiplied with the original image I to obtain the ROI image I_{ROI}. This step effectively isolates the target regions within the image.

Concurrently, the output from the final epoch of the model is concatenated with I_{ROI} in the channel dimension to form the enhanced input image I' to obtain richer context. The enhanced ROI image I' is then fed back into the model for fine-tuning to achieve the final segmentation outcome. By dynamically generating the dilation mask and enhancing the ROI image with additional feature map information, the DG-Loc improves the model's ability to accurately segment LCs, even in the presence of complex morphological features and background noise.

2.3 Boundary Optimization Loss (BO-Loss)

The complex morphology and ambiguous edges of LCs present significant challenges for accurately delineating segmentation boundaries. To enhance boundary precision and completeness, we propose a boundary optimization loss ($\mathcal{L}_{\mathrm{BO}}$) that integrates two components: a pixel-level loss and an image-level loss.

The pixel-level loss ($\mathcal{L}_{\mathrm{pixel\_level}}$) aligns segmentation boundaries with image edges by penalizing differences in segmentation predictions between adjacent pixels. It uses a pairwise potential function with a weighting factor ω_{ij} that depends on the intensity differences between neighboring pixels i and j:

$$\mathcal{L}_{\text{pixel\_level}} = \sum_{(i,j)\in\mathcal{N}} \omega_{ij} \cdot |y_i - y_j| \tag{3}$$

where $\omega_{ij} = \exp\left(-\frac{2\gamma^2}{|I_i - I_j|^2}\right)$, I_i and I_j are the pixel values of pixels i and j, $\gamma(\gamma = 1)$ controls the smoothness, and $\mathcal{N}$ denotes the 4-neighborhood set.

The image-level loss ($\mathcal{L}_{\text{image\_level}}$) refines the predicted boundaries by minimizing the difference between the predicted and true boundary maps, ensuring that the segmentation output closely matches the actual edges:

$$\mathcal{L}_{\text{image\_level}} = \sum_i \left(\frac{E_{\text{pred},i}}{\sum_j E_{\text{pred},j}} - \frac{E_{\text{true},i}}{\sum_j E_{\text{true},j}} \right)^2 \tag{4}$$

where $E_{\text{pred},i}$ is the pixel value in the predicted boundary map, and $E_{\text{true},i}$ is the pixel value in the true boundary map.

The total loss function is defined as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{Seg}} + \lambda(\mathcal{L}_{\text{pixel\_level}} + \mathcal{L}_{\text{image\_level}}) \tag{5}$$

where $\mathcal{L}_{\text{Seg}}$ is a linear combination of the cross-entropy loss and the Dice loss, with equal weighting, and λ is a hyperparameter that balances the impact of the segmentation loss and the boundary optimization loss.

3 Experiments

3.1 Datasets and Implementation Details

The **CORN-4S** dataset includes 560 corneal confocal microscopy images captured using a Heidelberg Retina Tomography III (HRT-III) system. Each image possesses a pixel resolution of 384×384, corresponding to a physical dimension of 400×400 μm^2 in tissue sampling. The **Clinical** dataset comprises 467 corneal confocal microscopy images captured using the same device from different hospitals. All Annotations were performed by a senior ophthalmologists using ITK-SNAP software to ensure accuracy. The experimental design incorporated five-fold cross-validation to ensure robustness of the evaluation. For model optimization, we employed the AdamW optimizer using a batch size of 4 and an initial learning rate of 1e-4. Training protocols included early stopping with patience of 10 epochs monitored on the validation loss to prevent model overfitting, with a total of 300 training epochs. The MorphoBoost was implemented using PyTorch 1.9.0, with all experiments conducted on an NVIDIA GeForce RTX 3090 GPU platform.

3.2 Comparison with the State-of-the-Art

We thoroughly assessed MorphoBoost alongside eight leading methods: four general medical segmentation approaches (CPFNet [8], CE-Net [9], TransUNet [3],

Table 1. Evaluation results of comparison methods in model performance.

Methods	non-activated		activated	
	Dice	MHD	Dice	MHD
TransUNet	81.46±4.14	5.42±2.12	36.41±5.33	21.42±7.19
CPFNet	84.32±2.34	4.01±2.53	45.87±3.24	18.78±8.23
CENet	82.25±3.56	4.78±1.87	43.36±4.79	20.67±7.89
MedNeXt	85.78±2.99	3.01±1.19	49.06±2.91	15.31±5.32
SetuSeg	84.81±4.65	3.89±1.45	46.88±4.62	18.02±5.34
DualStream	85.12±3.31	3.12±1.65	50.46±3.01	14.05±4.19
JFSA	86.06±5.02	2.97±0.98	51.15±3.44	13.69±6.36
SPG-Net	85.09±3.57	3.25±1.53	49.73±5.22	14.43±6.98
Proposed	**87.88±2.77**	**1.95±0.56**	**68.25±4.11**	**7.98±5.18**

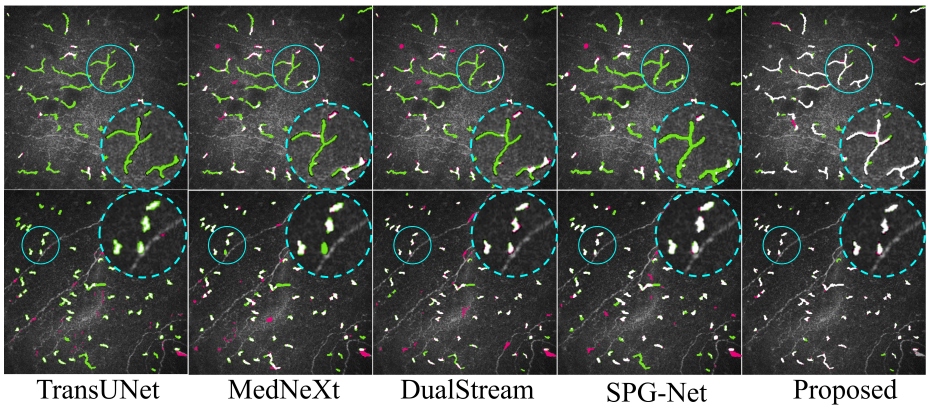

Fig. 4. Qualitative comparison between different methods. Under-segmentation is shown in green, and over-segmentation is shown in red.

MedNeXt [16]) and four cell-specific segmentation methods (SetuSeg [17], JFSA [25], DualStream [21], SPG-Net [23]). Evaluations were based on Dice coefficient (Dice) and Modified Hausdorff Distance (MHD) [7] metrics. As evidenced by Table 1, our method achieves state-of-the-art performance across all evaluation metrics in both LCs types, demonstrating absolute improvements of 17.10% and 5.71 respectively for activated LCs and 1.82% and 1.02 respectively for non-activated LCs over the suboptimal method. Compared methods overlook activated LCs, leading to lower recognition accuracy for activated LCs, whereas differences are subtle for non-activated LCs. This suggests that activated LCs features are harder to learn, underscoring the effectiveness of our MD-Aug. While methods like DualStream, JFSA, and SPG-Net use boundary regularization to enhance MHD for both LCs types, our MorphoBoost method shows a more significant improvement, indicating its superior boundary delineation capability.

Table 2. Ablation results of three key components on model performance.

DG-Loc	MD-Aug	BO-Loss	non-activated		activated	
			Dice	MHD	Dice	MHD
×	×	×	81.98±3.26	4.87±2.46	40.71±4.12	18.42±6.29
✓	×	×	84.87±5.34	3.56±3.53	51.93±6.14	15.78±7.23
✓	✓	×	85.45±3.26	2.64±2.37	62.18±4.39	12.66±5.89
✓	✓	✓	**87.88±2.77**	**1.95±0.56**	**68.25±4.11**	**7.98±5.18**

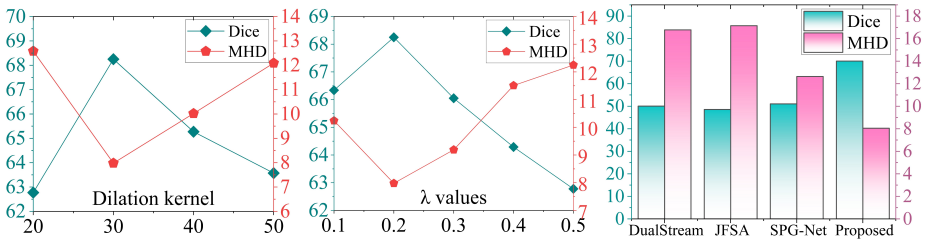

Fig. 5. Visualization of ablation results on activated LCs: (a) Impact of the dilation kernel size on model performance, (b) Selection of optimal hyperparameter λ, (c) Generalization validation on real clinical datasets.

To demonstrate MorphoBoost's superior performance, Fig. 4 compares visualizations from several baseline methods. Traditional approaches struggle to accurately identify activated LCs and delineate LCs boundaries. In contrast, MorphoBoost enhances recognition accuracy and boundary precision by diversifying activated LCs morphology and employing boundary optimization loss, thus establishing a more precise LCs segmentation model.

3.3 Ablation Study

We evaluated the impact of DG-Loc, MD-Aug, and BO-Loss on performance, as detailed in Table 2. Models without these components struggled with segmenting LCs. DG-Loc enhanced performance by 11.22% in Dice and 2.64 in MHD by generating ROI images that isolate LCs areas, improving the model's focus on LCs. MD-Aug increased accuracy by 10.25% in Dice and 3.12 in MHD by diversifying the morphology of activated LCs. The addition of BO-Loss led to peak performance, underscoring its critical role in precisely defining boundaries.

Additionally, we assessed the impact of the selection of dilation kernel on model performance. As shown in Fig. 5(a), A small dilation kernel size can deprive the model of sufficient context for precise segmentation, whereas a larger one might introduce excessive noise. A dilation size of 30 achieves the optimal balance for accurate segmentation. We evaluated the impact of hyper-parameter λ on model performance by identifying the optimal value of $\lambda = 0.2$ through grid search, as depicted in Fig. 5(b). Our method's robustness was further validated on real clinical datasets from additional centers (Fig. 5(c)). Our approach outper-

forms other state-of-the-art methods in identifying activated LCs, highlighting its effectiveness and generalizability for practical clinical use.

4 Conclusion

This study introduces a novel morphology-driven boundary enhancement model (MorphoBoost) for accurate LCs segmentation in CCM images. The proposed method effectively addresses the challenges of misidentification of activated LCs and inaccurate boundary delineation through a morphology-driven data augmentation module and a boundary optimization loss. It achieves state-of-the-art performance, marked by substantial improvements in segmentation accuracy, as confirmed by clinical data validation.

Acknowledgments. This work is supported by the National Natural Science Foundation of China (62401551, 62422122, 62272444).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Braiki, M., Benzinou, A., Nasreddine, K., Hymery, N.: Automatic human dendritic cells segmentation using k-means clustering and chan-vese active contour model. Computer methods and programs in biomedicine **195**, 105520 (2020)
2. Braiki, M., Benzinou, A., Nasreddine, K., Labidi, S., Hymery, N.: Segmentation of dendritic cells from microscopic images using mathematical morphology. In: 2016 2nd International Conference on Advanced Technologies for Signal and Image Processing (ATSIP). pp. 282–287. IEEE (2016)
3. Chen, J., Lu, Y., Yu, Q., Luo, X., Adeli, E., Wang, Y., Lu, L., Yuille, A.L., Zhou, Y.: Transunet: Transformers make strong encoders for medical image segmentation. arXiv preprint arXiv:2102.04306 (2021)
4. Chen, J., Mou, L., Ma, S., Fu, H., Guo, L., Zheng, Y., Zhang, J., Zhao, Y.: Nerveformer: A cross-sample aggregation network for corneal nerve segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 79–88. Springer (2022)
5. Chinnery, H.R., Zhang, X.Y., Wu, C.Y., Downie, L.E.: Corneal immune cell morphometry as an indicator of local and systemic pathology: a review. Clinical & experimental ophthalmology **49**(7), 729–740 (2021)
6. Downie, L.E., Zhang, X., Wu, M., Karunaratne, S., Loi, J.K., Senthil, K., Arshad, S., Bertram, K., Cunningham, A.L., Carnt, N., et al.: Redefining the human corneal immune compartment using dynamic intravital imaging. Proceedings of the National Academy of Sciences **120**(31), e2217795120 (2023)
7. Dubuisson, M.P., Jain, A.K.: A modified hausdorff distance for object matching. In: Proceedings of 12th international conference on pattern recognition. vol. 1, pp. 566–568. IEEE (1994)
8. Feng, S., Zhao, H., Shi, F., Cheng, X., Wang, M., Ma, Y., Xiang, D., Zhu, W., Chen, X.: Cpfnet: Context pyramid fusion network for medical image segmentation. IEEE transactions on medical imaging **39**(10), 3008–3018 (2020)

9. Gu, Z., Cheng, J., Fu, H., Zhou, K., Hao, H., Zhao, Y., Zhang, T., Gao, S., Liu, J.: Ce-net: Context encoder network for 2d medical image segmentation. IEEE transactions on medical imaging **38**(10), 2281–2292 (2019)
10. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017)
11. Jamali, A., Kenyon, B., Ortiz, G., Abou-Slaybi, A., Sendra, V.G., Harris, D.L., Hamrah, P.: Plasmacytoid dendritic cells in the eye. Progress in retinal and eye research **80**, 100877 (2021)
12. Kheirkhah, A., Darabad, R.R., Cruzat, A., Hajrasouliha, A.R., Witkin, D., Wong, N., Dana, R., Hamrah, P.: Corneal epithelial immune dendritic cell alterations in subtypes of dry eye disease: a pilot in vivo confocal microscopic study. Investigative ophthalmology & visual science **56**(12), 7179–7185 (2015)
13. Lagali, N.S., Badian, R.A., Liu, X., Feldreich, T.R., Ärnlöv, J., Utheim, T.P., Dahlin, L.B., Rolandsson, O.: Dendritic cell maturation in the corneal epithelium with onset of type 2 diabetes is associated with tumor necrosis factor receptor superfamily member 9. Scientific reports **8**(1), 14248 (2018)
14. Lancaster, P.: Moving weighted least-squares methods. In: Polynomial and Spline Approximation. pp. 103–120. Springer, Dordrecht (1979). https://doi.org/10.1007/978-94-009-9443-0_7
15. Marsovszky, L., Resch, M.D., Németh, J., Toldi, G., Medgyesi, E., Kovács, L., Balog, A.: In vivo confocal microscopic evaluation of corneal langerhans cell density, and distribution and evaluation of dry eye in rheumatoid arthritis. Innate immunity **19**(4), 348–354 (2013)
16. Roy, S., Koehler, G., Ulrich, C., Baumgartner, M., Petersen, J., Isensee, F., Jaeger, P.F., Maier-Hein, K.H.: Mednext: transformer-driven scaling of convnets for medical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 405–415. Springer (2023)
17. Setu, M.A.K., Schmidt, S., Musial, G., Stern, M.E., Steven, P.: Segmentation and evaluation of corneal nerves and dendritic cells from in vivo confocal microscopy images using deep learning. Translational Vision Science & Technology **11**(6), 24–24 (2022)
18. Shi, J., Tomasi, C.: Good features to track. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 593–600. IEEE, Seattle, WA (1994)
19. Suberi, A.A.M., Zakaria, W.N.W., Tomari, R., Lim, K.P.: Optimization of overlapping dendritic cell segmentation in phase contrast microscopy images. In: 2016 IEEE EMBS Conference on Biomedical Engineering and Sciences (IECBES). pp. 246–250. IEEE (2016)
20. Wu, J., Gao, J., Lin, J., Huang, Z., Liu, Y., Chen, Z., Long, Q., Zhao, J., Ding, D.: A dual-stream network for langerhans' cells segmentation in ccm images. In: International Workshop on Ophthalmic Medical Image Analysis. pp. 126–135. Springer (2024)
21. Wu, J., Gao, J., Lin, J., Huang, Z., Liu, Y., Chen, Z., Long, Q., Zhao, J., Ding, D.: A dual-stream network for langerhans' cells segmentation in ccm images. In: International Workshop on Ophthalmic Medical Image Analysis. pp. 126–135. Springer (2024)
22. Wu, M., Zhang, X., Karunaratne, S., Lee, J.h., Lampugnani, E.R., Selva, K.J., Chung, A.W., Mueller, S.N., Chinnery, H.R., Downie, L.E.: Intravital imaging of the human cornea reveals the differential effects of season on innate and adaptive immune cell morphodynamics. Ophthalmology **131**(10), 1185–1195 (2024)

23. Zhang, Y., Xi, R., Zeng, L., Towey, D., Bai, R., Higashita, R., Liu, J.: Structural priors guided network for the corneal endothelial cell segmentation. IEEE Transactions on Medical Imaging **43**(1), 309–320 (2023)
24. Zhao, Y., Zheng, Y., Zhao, Y., Liu, Y., Chen, Z., Liu, P., Liu, J.: Uniqueness-driven saliency analysis for automated lesion detection with applications to retinal diseases. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part II 11. pp. 109–118. Springer (2018)
25. Zhu, Y., Yin, X., Meijering, E.: A compound loss function with shape aware weight map for microscopy cell segmentation. IEEE Transactions on Medical Imaging **42**(5), 1278–1288 (2022)