Prototype-Driven Robust Domain Generalization in Corneal Endothelial Cell Segmentation

Hongshuo Li[1,2†], Wenlong Li[3†], Shijia Zhou[1,2], Yalin Zheng[4], Hong Qi[3*], and Yitian Zhao[1,5*]

[1] Ningbo Institute of Materials Technology and Engineering, Chinese Academy of Sciences, Ningbo, China
yitian.zhao@nimte.ac.cn, doctorqihong@163.com
[2] University of Chinese Academy of Sciences, Beijing, China
[3] Department of Ophthalmology, Peking University Third Hospital, Beijing, China
[4] Department of Eye and Vision Science, University of Liverpool, Liverpool, UK
[5] School of Biomedical Engineering, ShanghaiTech University, Shanghai, China

Abstract. Accurate corneal endothelial cell (CEC) quantification is essential for the diagnosis and management of ocular diseases; however, current learning-based CECs segmentation is often hindered by a significant domain gap between homogeneous training sets and heterogeneous clinical data. Existing models frequently fail to generalize to images from unseen devices, centers, pathologies or species. We present ProtoCEC, a prototype-driven framework designed for robust domain generalization in confocal microscopy image segmentation. The architecture comprises a prototype collector for robust feature extraction and a prototype enhancer to improve class discriminability. At the inference stage, a prototype allocator facilitates dynamic feature-to-prototype matching to guide the segmentation mask generation. Our results confirm that ProtoCEC provides a significant performance boost over existing state-of-the-art architectures. We demonstrate the model's versatility through rigorous testing in four out-of-distribution scenarios: cross-center, cross-disease, cross-device, and cross-species, proving its efficacy as a generalized solution for confocal microscopy image analysis.

Keywords: Corneal endothelial cell · Segmentation · Prototype learning.

1 Introduction

The corneal endothelium is a single layer of cells on the innermost surface of the cornea. Because corneal endothelial cells (CECs) generally do not regenerate, diseases that damage them can lead to permanent swelling, cloudiness, and vision loss. In clinical settings, corneal confocal microscopy (CCM) provides the high-resolution *in vivo* imagery required for this assessment, the morphological quantification of CECs serves as a critical diagnostic biomarker. Metrics such as cell density, pleomorphism, and polymegathism are essential for managing corneal diseases like Fuchs' endothelial dystrophy and corneal endotheliitis, as

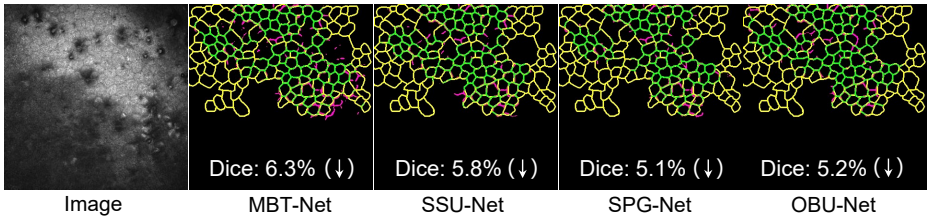

Fig. 1. Qualitative comparison showing the poor generalization of state-of-the-art methods [15, 18, 20, 16] when applied to unseen target domains. Qualitative comparison showing under- (yellow), over- (pink), and accurate- (green) segmentation. Labels indicate the decline in Dice score from the source to the target domain.

well as for monitoring corneal graft health [3, 7, 2, 8]. However, current quantification workflows rely heavily on manual or semi-automated analysis, which is labor-intensive and prone to subjectivity.

While deep learning-based segmentation offers a promising alternative, its clinical translation is obstructed by a critical technical deficiency: poor generalization capability [15, 18, 20, 16]. Current models suffer from significant performance attenuation when subjected to distributional shifts, failing to adapt to heterogeneous data sources defined by varying acquisition protocols, device instrumentation, or unseen pathologies [6]. This fragility stems from a fundamental reliance on small-scale, homogeneous training distributions. Existing CEC segmentation frameworks [15, 18, 20, 16] are typically optimized on narrow datasets (often restricted to single-center, single-disease cohorts). Consequently, these models overfit to domain-specific statistical biases rather than learning robust, invariant anatomical features, rendering them unreliable when deployed in diverse real-world environments.

The vulnerability of existing frameworks to domain shifts is starkly apparent in cross-center evaluations. When subjected to distribution changes, standard models experience a catastrophic reduction in segmentation accuracy, evidenced by drastically lower Dice similarity coefficients compared to in-domain performance, as shown in Fig. 1 (a). Instead of maintaining robust feature extraction, these models frequently generate erroneous segmentation masks characterized by distinct failure modes: under-segmentation in low-contrast regions, spurious over-segmentation triggered by pathological noise, and a general inability to delineate precise cellular boundaries in unseen data distributions, as shown in Fig. 1 (b). Such instability reveals that existing methods lack the generalization capability required for reliable clinical deployment.

The fundamental challenge in this domain lies in the difficulty of learning robust feature representations from limited and homogeneous datasets. Rather than relying on implicit learning from scarce data, we propose a paradigm shift: the explicit extraction of robust feature prototypes that encapsulate the quintessential morphology of CECs. Functioning as stable "feature anchors," these prototypes guide the segmentation of unseen data via similarity matching,

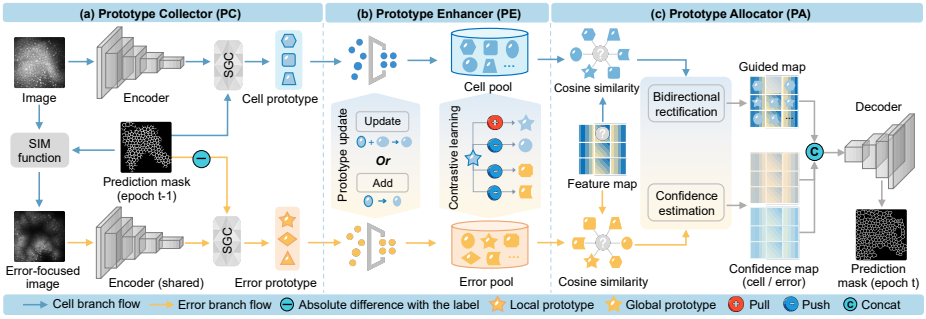

Fig. 2. Overall architecture of the proposed ProtoCEC. (a) Prototype collector decouples reliable cellular features from artifacts via difficulty-aware mining. (b) Prototype enhancer refines the feature space using contrastive learning with dual-negative filtering. (c) Prototype allocator performs bidirectional feature rectification (standardization) and suppression (attenuation) for robust segmentation.

effectively decoupling feature recognition from domain-specific noise. Driven by this intuition, we introduce ProtoCEC, a framework consisting of three synergistic components: a prototype collector, which mines the feature space for a diverse range of examples including hard-to-segment cases; a prototype enhancer, which enforces prototype independence and discriminative power; and a prototype allocator, which executes the final segmentation by aligning input features with the most appropriate prototypes. The main contribution of this work is the introduction of ProtoCEC, the first prototype-driven framework tailored for corneal endothelial cell segmentation. Designed to overcome the generalization bottleneck, it demonstrates exceptional robustness, consistently outperforming state-of-the-art methods across four challenging external scenarios: cross-center, cross-disease, cross-device, and cross-species settings.

2 Methodology

We present ProtoCEC, a framework tailored to overcome the limitations of domain shift by decoupling anatomical features from imaging artifacts. The architecture (Fig. 2) is composed of three distinct modules: a prototype collector for difficulty-aware feature mining, a prototype enhancer for contrastive feature calibration, and a prototype allocator for bidirectional feature rectification. The following subsections detail the design and interaction of these components.

2.1 Prototype Collector

The reliance on masked average pooling in standard prototypical networks [14, 11] leads to a loss of discriminability. While superpixel-guided clustering (SGC) [4] improves representation by creating multiple prototypes, it is flawed because it constrains feature mining to the ground-truth mask alone. In CCM imagery, this

limitation prevents the model from learning the distinct features of "hard" cells (under-segmentation) or the deceptive patterns of noise (over-segmentation). To resolve this, we propose a difficulty-aware mining strategy to improve SGC by leveraging information from all relevant regions, e.g., correct, under-segmented, and over-segmented, to build a more robust feature space.

To localize segmentation failures, we first compute an error mask $\mathbf{M}_{ER} = |\mathbf{P} - \mathbf{G}|$, where $\mathbf{P}$ is the prediction and $\mathbf{G}$ is groundtruth. We then apply Structural Importance Mapping (SIM) [5] to accentuate boundary details within these error-prone areas, generating a error-focused image:

$$\mathbf{I}_{ER} = \mathbf{I} \odot \hat{\mathbf{M}}_{\mathbf{ER}}, \quad \text{where } \hat{\mathbf{M}}_{\mathbf{ER}} = SIM(M_{ER}) = e^{-\beta \cdot \mathrm{SDF}(\mathbf{M}_{\mathbf{ER}})}, \tag{1}$$

where SDF$(\cdot)$ denotes the signed distance function. Based on this, prototypes are aggregated into two complementary pools. 1) Cell prototype pool ($\mathcal{P}_{cell}$): Derived from the predicted mask $\mathbf{P}$, this pool accumulates features from correctly segmented regions, capturing both standard and difficult cellular patterns that were successfully recovered. 2) Error prototype pool ($\mathcal{P}_{error}$): Extracted from the error mask $\mathbf{M}_{ER}$, this pool encodes the specific feature signatures responsible for failures, including patterns associated with over-segmentation (false positives) and under-segmentation (false negatives).

2.2 Prototype Enhancer

To prevent feature collapse and ensure prototypes are discriminative, we introduce a Prototype Enhancer to refine the collected pools via contrastive learning. For each query prototype $\mathbf{q}$ generated from the current batch, we select its most similar prototype in the corresponding pool as the positive key $\mathbf{k}^+$. Constructing an effective negative set $\mathcal{N}$ is critical. We propose a dual-threshold filtering mechanism to exclude both trivial negatives (similarity < 0.3) and potential false negatives (similarity > 0.8). The remaining samples force the model to learn fine-grained boundaries. The feature space is optimized using the prototype contrastive loss, formulated as:

$$\mathcal{L}_{con} = -\sum_{\mathbf{q}\in\mathcal{Q}} \log \frac{\exp(\mathrm{sim}(\mathbf{q}, \mathbf{k}^+)/\tau)}{\exp(\mathrm{sim}(\mathbf{q}, \mathbf{k}^+)/\tau) + \sum_{\mathbf{n}\in\mathcal{N}} \exp(\mathrm{sim}(\mathbf{q}, \mathbf{n})/\tau)}, \tag{2}$$

where $\tau = 0.1$ is the temperature scaling factor. To ensure the prototype library maintains temporal consistency and high representational quality, we adopt a momentum-based update strategy [1]. Instead of abrupt replacement, the most similar existing prototype $\mathbf{p}$ is updated by the incoming feature $\mathbf{p}_{\text{new}}$ according to: $\mathbf{p} \leftarrow \eta \cdot \mathbf{p} + (1 - \eta) \cdot \mathbf{p}_{\text{new}}$, with the momentum coefficient set to $\eta = 0.9$. Conversely, features with a similarity score below 0.3 are directly initialized as new prototypes. Furthermore, to maintain a compact and discriminative pool (capacity $N = 128$), we implement a frequency-based pruning policy: prototypes with the lowest matching frequency are prioritized for removal, ensuring the dictionary remains populated by the most relevant features.

2.3 Prototype Allocator

To translate the learned prototypes into spatial predictions, the prototype allocator is introduced to construct three spatially dense maps. Specifically, it produces a guided map to standardize semantic features across domains, alongside two confidence maps estimating pixel-level uncertainty to flag potential errors.

Bidirectional feature rectification We generate the guided map via a bidirectional allocation mechanism governed by the competitive similarity between encoder features and the prototype pools. For each spatial feature $\mathbf{f}$, we compute the maximum cosine similarity scores relative to the cell (s_{cell}) and artifact (s_{error}) pools. The dominant score dictates the rectification strategy. *Cell feature standardization:* If $s_{cell} > s_{error}$, the feature is deemed anatomically valid and standardized towards the canonical representation. We replace $\mathbf{f}$ with its nearest matching cell prototype: $\mathbf{f}' = \arg\max_{\mathbf{p} \in \mathcal{P}_{cell}} \text{sim}(\mathbf{f}, \mathbf{p})$. *Error feature suppression:* Conversely, if $s_{error} > s_{cell}$, the feature is identified as interference. We mitigate this noise through adaptive soft-gating: $\mathbf{f}' = \mathbf{f} \cdot (1 - s_{error})$.

Global confidence estimation To capture global context, we synthesize two confidence maps that measure the total support for each class. Unlike the guided map, which relies on a single nearest prototype, these maps aggregate information from the entire library. The cell confidence map ($\mathbf{C}_{cell}$) is computed by summing the cosine similarities between the feature $\mathbf{f}$ and all entries in the cell pool: $\mathbf{C}_{cell} = \sum_{\mathbf{p} \in \mathcal{P}_{cell}} \text{sim}(\mathbf{f}, \mathbf{p})$. Symmetrically, we generate an error confidence map ($\mathbf{C}_{error}$) by aggregating similarities over $\mathcal{P}_{error}$, highlighting regions with high structural correspondence to known failure modes.

To form a rectified skip connection, the guided map and two global confidence maps are concatenated, replacing the noisy encoder features typical of U-Net architectures. By injecting standardized anatomical patterns and global context into the decoding pathway, we ensure reconstruction is driven by invariant features rather than domain artifacts. The decoder then performs standard upsampling to generate the final segmentation.

2.4 Objective Function and Inference

The total objective is defined as $\mathcal{L}_{total} = \mathcal{L}_{seg} + \lambda \mathcal{L}_{con}$. A key feature of our training strategy is the simultaneous optimization of "easy" and "hard" patterns. We calculate the segmentation loss $\mathcal{L}_{seg}$ on both the global image $\mathbf{I}$ (using ground truth $\mathbf{G}$) and the error-focused image $\mathbf{I}_{ER}$. For the latter, we utilize a specialized target $\mathbf{G}_{ER} = \text{Binarize}(\hat{\mathbf{M}}_{\mathbf{ER}}) \cdot \mathbf{G}$ to force the model to resolve ambiguous boundaries and artifacts.

Inference is streamlined for real-time application. The dynamic mining modules (PC and PE) are deactivated, and the prototype library is frozen. The input image $\mathbf{I}$ proceeds through a single efficient forward pass: from the encoder, through the PA (which references the stored, stable prototypes), to the decoder, yielding the final segmentation mask without iterative refinement.

3 Experiments

3.1 Datasets and implementation details

We evaluated transferability using one source training dataset and four external target datasets, covering diverse centers, pathologies, devices, and species.

• **CED** (source): Our in-house training set comprises 1,200 annotated images (384×384) from 467 patients acquired by Heidelberg HRT-III. It includes healthy controls, Fuchs' dystrophy, and corneal endotheliitis.

• **DATA-I** (cross-center): To assess robustness against site-specific variations, we collected 460 images (58 patients) from three unseen clinical centers, ensuring independence from the training distribution.

• **DATA-II** (cross-disease): To evaluate adaptability to unseen pathologies, We utilized 368 images (27 patients) featuring rare pathologies absent from training: posterior polymorphous corneal dystrophy, iridocorneal endothelial syndrome, congenital hereditary endothelial dystrophy, and post-penetrating keratoplasty eyes.

• **Rodrep** (cross-device): To test sensor invariance, we used 52 images acquired with a Nidek Confoscan 4 slit-scanning microscope [10], introducing a significant acquisition domain shift.

• **Alisarine** (cross-species): We evaluated biological transferability using 30 porcine corneal images from the Alisarine dataset [9].

We implemented the model in PyTorch using a ResNet-50 encoder and U-Net decoder on two NVIDIA RTX 3090 GPUs. Features from encoder blocks 2, 3, and 4 are used to build the prototype pool and apply the PA module as skip connections. We removed the final encoder ReLU to improve cosine similarity matching by extending the range to $[-1, 1]$. Training uses the AdamW optimizer (batch size 4, learning rate 10^{-4}) for 600 epochs with early stopping. Images are resized to 384×384. Key parameters are set to $N = 3$ prototypes, $\beta = 3$, and loss weight $\lambda = 0.2$. The CED dataset is randomly partitioned into training, validation, and test sets at an 8:1:1 ratio. During inference, the final predictions are generated using a sliding window strategy with a crop size of 384×384.

3.2 Segmentation performance

We benchmark ProtoCEC against eight leading CEC segmentation methods: MBT-Net [15], SC-GAN [13], UG-Net [17], SSU-Net [18], TBNet [19], OBU-Net [16], SPG-Net [20], and FRB-FS [12].

Top row of Fig. 3 demonstrates the model's robustness across four distinct distribution shifts, where ProtoCEC consistently outperforms the second-best method: cross-center (+3.2%), cross-disease (+4.5%), cross-device (+7.7%), and cross-species (+3.6%). This superior generalizability stems from decoupling invariant anatomy from domain-specific interference: Unlike comparison method that prone to overfit domain textures, our *cell feature standardization* projects diverse inputs onto a stable manifold of canonical prototypes. This acts as an explicit domain adapter, ensuring the decoder receives consistent representations

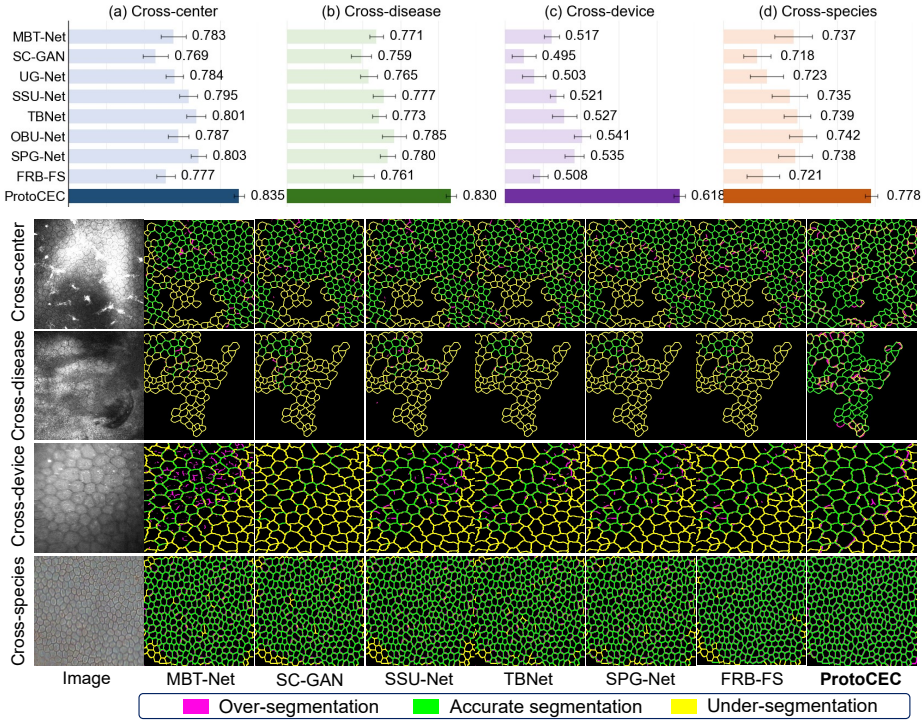

Fig. 3. Quantitative and qualitative evaluation of segmentation performance across diverse distribution shifts, comparing ProtoCEC against state-of-the-art models.

regardless of acquisition variations. The *error feature suppression* mechanism filters out unseen distractors in external distributions. Since domain-specific artifacts fail to match valid cellular prototypes, they are automatically attenuated, thereby immunizing the model against false predictions caused by distribution shifts.

Qualitative results (bottom row of Fig. 3) demonstrate robust generalization across four external scenarios. In cross-center transfers, the framework maintains consistent segmentation performance despite variations in acquisition protocols. For cross-disease validation, it exhibits high stability on unseen pathologies, preventing the structural hallucinations typical of baseline models. Under cross-device shifts, while competitors suffer from artifact-induced over-segmentation (e.g., keratic precipitates), our method effectively suppresses these false positives; although the rejection of unmapped out-of-distribution patterns may cause under-segmentation, this mechanism significantly reduces noise interference. Finally, in cross-species experiments, ProtoCEC eliminates the fragmentation and under-segmentation issues observed in dense cellular regions.

3.3 Ablation studies

Impact of key components. As shown in Fig. 4(a), replacing SGC with the proposed prototype collector (PC) improves Dice from 0.741 to 0.765, validating the efficacy of hard example mining. Notably, integrating either the prototype enhancer (PE) or prototype allocator (PA) individually yields performance gains (0.735 and 0.748). Furthermore, their joint deployment outperforms these individual uses, peaking at 0.765 and demonstrating strong synergy for robust segmentation. This performance boost demonstrates a strong technical synergy between the components, confirming that their integration is essential for achieving robust segmentation across challenging domain shifts.

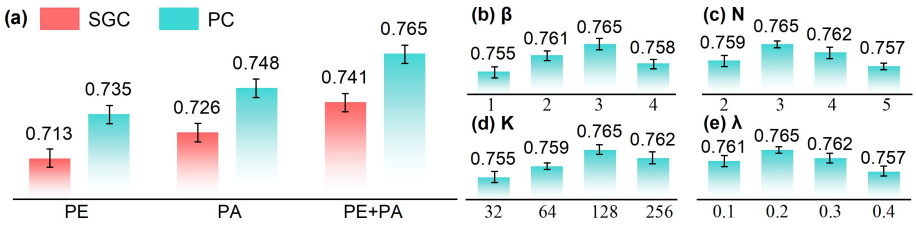

Fig. 4. Ablation study of the ProtoCEC framework. Mean Dice scores are averaged across four external datasets to evaluate: (a) the contribution of individual architectural components; and the model sensitivity to hyperparameters, including (b) scaling factor β, (c) prototype number N, (d) negative sample size K, and (e) loss weight λ.

Hyperparameter sensitivity. Fig. 4(b-e) identify the optimal settings ($\beta = 3, N = 3, K = 128, \lambda = 0.2$) for balancing discriminability and stability. Deviating from these peaks degrades performance; specifically, lower values restrict prototype capacity, while excessive values introduce over-regularization, disrupting feature alignment. These results highlight the sensitivity of the framework to its structural parameters and confirm that the selected configuration provides the most robust equilibrium for cross-domain generalization.

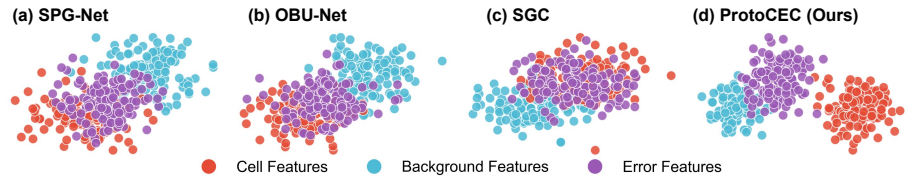

Fig. 5. t-SNE visualization of 128D feature representations across four external datasets.

Feature visualization. Fig. 5 illustrates the superior discriminative power of ProtoCEC, particularly regarding hard samples. While baseline methods ex-

hibit significant entanglement between cell features and deceptive artifacts (Error features), ProtoCEC successfully disentangles these representations by aligning Error features with the background cluster. This spatial separation in the feature space confirms our framework's ability to effectively identify and isolate deceptive patterns, ensuring they do not interfere with accurate cell segmentation. These visualizations provide clear evidence that our prototype-driven approach effectively purifies the feature space against domain-specific noise.

4 Conclusion

In this paper, we addressed the critical challenge of generalization in corneal endothelial cell segmentation by proposing ProtoCEC, a prototype-driven framework designed to learn robust feature representations from limited data. Through extensive validation across four distinct domain shift scenarios, including cross-center, cross-disease, cross-device, and cross-species settings, we demonstrated that ProtoCEC significantly outperforms existing state-of-the-art methods. Ultimately, this study provides a solid foundation for developing clinically reliable and generalizable ophthalmic AI systems.

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

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

References

1. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020)
2. Joseph, N., Benetz, B.A., Chirra, P., Menegay, H., Oellerich, S., Baydoun, L., Melles, G.R., Lass, J.H., Wilson, D.L.: Machine learning analysis of postkeratoplasty endothelial cell images for the prediction of future graft rejection. Translational Vision Science & Technology **12**(2), 22–22 (2023)
3. Li, D., Li, Z., Duan, H., Wang, X., Lin, Z., Dai, K., Qi, Q., Dong, Y., Lin, P., Su, W., et al.: A novel single-cell level evaluation method for corneal endothelial cell function. npj Digital Medicine (2025)
4. Li, G., Jampani, V., Sevilla-Lara, L., Sun, D., et al.: Adaptive prototype learning and allocation for few-shot segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8334–8343 (2021)
5. Li, H., Ma, B., Mou, L., Liu, Y., Zheng, Q., Qi, H., Zhao, Y.: Progressive distillation for incremental learning in corneal confocal microscopy segmentation. IEEE Transactions on Medical Imaging (2025)
6. Liu, S., Kandakji, L., Stupnicki, A., Sumodhee, D., Leucci, M.T., Hau, S., Balal, S., Okonkwo, A., Moghul, I., Kanda, S.P., et al.: Current applications of artificial intelligence for fuchs endothelial corneal dystrophy: a systematic review. Translational Vision Science & Technology **14**(6), 12–12 (2025)

7. Prada, A.M., Quintero, F., Mendoza, K., Galvis, V., Tello, A., Romero, L.A., Marrugo, A.G.: Assessing fuchs corneal endothelial dystrophy using artificial intelligence–derived morphometric parameters from specular microscopy images. Cornea **43**(9), 1080–1087 (2024)
8. Qu, J.H., Qin, X.R., Peng, R.M., Xiao, G.G., Cheng, J., Gu, S.F., Wang, H.K., Hong, J.: A fully automated segmentation and morphometric parameter estimation system for assessing corneal endothelial cell images. American Journal of Ophthalmology **239**, 142–153 (2022)
9. Ruggeri, A., Scarpa, F., De Luca, M., Meltendorf, C., et al.: A system for the automatic estimation of morphometric parameters of corneal endothelium in alizarine red-stained images. British Journal of Ophthalmology **94**(5), 643–647 (2010)
10. Selig, B., Vermeer, K.A., Rieger, B., Hillenaar, T., Luengo Hendriks, C.L.: Fully automatic evaluation of the corneal endothelium from in vivo confocal microscopy. BMC medical imaging **15**(1), 13 (2015)
11. Tian, Z., Zhao, H., Shu, M., Yang, Z., Li, R., Jia, J.: Prior guided feature enrichment network for few-shot segmentation. IEEE transactions on pattern analysis and machine intelligence **44**(2), 1050–1065 (2020)
12. Wang, T., Nan, X., Wang, Y., Yan, Y., Gao, Z., Liu, J.: Enhanced corneal endothelial cell segmentation via frequency-selected residual fourier diffusion models. In: ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 1–5. IEEE (2025)
13. Xi, R., Zhang, Y., Bai, R., Higashita, R., Liu, J.: Sc-gan: Structure consistent gan for modality transfer with fft and multi-scale perception. In: 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2023)
14. Zhang, C., Lin, G., Liu, F., Yao, R., Shen, C.: Canet: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5217–5226 (2019)
15. Zhang, Y., Higashita, R., Fu, H., Xu, Y., Zhang, Y., Liu, H., Zhang, J., Liu, J.: A multi-branch hybrid transformer network for corneal endothelial cell segmentation. In: International conference on medical image computing and computer-assisted intervention. pp. 99–108. Springer (2021)
16. Zhang, Y., Higashita, R., Zeng, L., Li, J., Xi, R., Liu, T., Fu, H., Towey, D., et al.: Online bayesian approximation based uncertainty aware model for ophthalmic image segmentation. IEEE Journal of Biomedical and Health Informatics (2025)
17. Zhang, Y., Wang, W., Wang, B., Cai, Z., Towey, D., Bai, R., Higashita, R., Liu, J.: Ug-net: Corneal endothelial cell segmentation based on uncertainty estimation and soft spatial attention. In: 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2023)
18. Zhang, Y., Xi, R., Fu, H., Towey, D., Bai, R., Higashita, R., Liu, J.: Elongated physiological structure segmentation via spatial and scale uncertainty-aware network. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 323–332. Springer (2023)
19. Zhang, Y., Xi, R., Wang, W., Li, H., Hu, L., Lin, H., Towey, D., Bai, R., Fu, H., Higashita, R., et al.: Low-contrast medical image segmentation via transformer and boundary perception. IEEE Transactions on Emerging Topics in Computational Intelligence **8**(3), 2297–2309 (2024)
20. 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)