

Minuscule Cell Detection in AS-OCT Images with Progressive Field-of-View Focusing

Boyu Chen¹, Ameenat Solebo², Daqian Shi^{1✉}, Jinge Wu¹, and Paul Taylor^{1✉}

¹ Institute of Health Informatics, University College London, London, UK
{daqian.shi, p.taylor}@ucl.ac.uk

² Great Ormond Street Institute of Child Health, University College London, London, UK

Abstract. Anterior Segment Optical Coherence Tomography (AS-OCT) is an emerging imaging technique with great potential for diagnosing anterior uveitis, a vision-threatening condition. This condition is characterized by the presence of inflammatory cells in the eye’s anterior chamber (AC). Automatic detection of these cells on AS-OCT images has attracted great attention. However, this task is challenging since each cell is minuscule (extremely small), representing less than 0.005% of the high-resolution image. Moreover, pixel-level noise introduced by OCT can be misclassified as cells, leading to false positive detections. These challenges make both traditional image processing algorithms and state-of-the-art (SOTA) deep learning object detection methods ineffective for this task. To that end, we propose a minuscule cell detection framework that progressively refines the field-of-view from the whole image to the AC region, and further to minuscule regions potentially containing individual cells. Our framework consists of: (1) a Field-of-Focus module that uses a vision foundation model to zero-shot segment the AC region, and (2) a Fine-grained Object Detection module that introduces Minuscule Region Proposal followed by our Cell Mamba to distinguish individual cells from noise. Experimental results demonstrate that our framework outperforms SOTA methods, improving F1 by around 7% over the best baseline and offering a more reliable alternative for cell detection. Our code is available at: <https://github.com/joeybyc/MCD>.

Keywords: AS-OCT · Anterior Uveitis · Cell Detection.

1 Introduction

Anterior Segment Optical Coherence Tomography (AS-OCT) is an imaging technique for visualizing the eye’s anterior chamber (AC), with particular potential in diagnosing a vision-threatening condition, anterior uveitis. This condition is characterized by inflammatory cells in the AC [27]. Automated detection of these cells using AS-OCT images has attracted substantial research interest [10].

However, this task is challenging due to the minuscule size of the cells. As shown in Fig. 1a (1), on x160 magnification of two regions of an AS-OCT image (blue boxes), the hyperreflective particles (green boxes in Fig. 1a (2)) represent

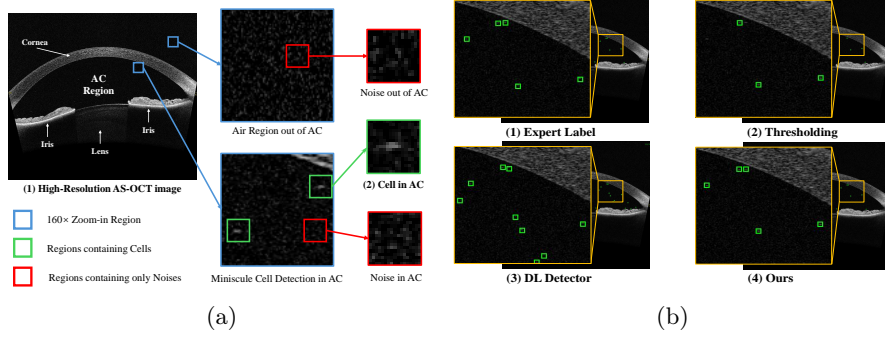


Fig. 1: (a) Cells and noise in the AS-OCT image. (b) Qualitative comparisons of cell detection results using different methods, with green boxes indicating the detected cells.

target cells. Each occupies less than 0.005% of the image. Moreover, pixel-level noise (highlighted in red boxes) in OCT images closely resembles the cells in texture, leading to false positive detections. Traditional denoising methods are ineffective as they indiscriminately remove both cells and noise.

Previous automated detection methods primarily rely on thresholding, which determines a decision boundary (threshold) for classifying pixels based on their intensity values (ranging from 0 for black to 255 for white). Pixels exceeding this threshold are classified as objects, while others are considered background. In particular, the Isodata algorithm [22], which is the default thresholding method in ImageJ¹, has been widely used [1, 3, 9, 19, 28].

However, these methods still require manual delineation of the AC region to exclude non-cell particles detected outside this area, making the process time-consuming and error-prone. Moreover, these methods could miss the cells with intensity values below their calculated decision boundary (see Fig. 1b (2)).

Deep learning (DL) object detection methods present another potential solution [21]. However, they struggle to capture the fine-grained details of cells during feature extraction and down-scaling operations, as each cell occupies an extremely small portion of the image. Consequently, DL detectors often misclassify pixel-level noise as cells due to their textural similarity (see Fig. 1b (3)).

To tackle these challenges, we introduce a progressive field-of-view focusing strategy. The core idea is to gradually refine the focus of analysis: first, by concentrating on the AC region where such cells are typically found; then, by zooming in on extremely small, fine-grained subregions that may contain individual cells; and finally, by examining the spatial patterns embedded within the latent representations derived from these minuscule subregions to accurately distinguish cells from background noise. To implement this approach, we present a Minuscule Cell Detection (MCD) framework, consisting of two modules:

¹ <https://imagej.net/ij/>

The Field-of-Focus (FoF) module leverages a foundation vision model to zero-shot segment the AC region, eliminating the need for manual intervention. Within this region, our Fine-grained Object Detection (FOD) module detects individual cells through two key components. We design the Minuscule Region Proposal to identify extremely small regions where cells might be present, capturing potential cells that thresholding methods would miss. Subsequently, our Cell Mamba learns to process fine-grained details within these minuscule regions to distinguish cells from noise, overcoming the limitation of conventional DL detectors in capturing features of such tiny objects. As a result, our framework not only detects more cells than thresholding methods but also achieves higher detection accuracy compared to existing DL detectors (see Fig. 1b (4)).

Our contributions are:

- 1) We propose MCD to detect minuscule inflammatory cells in AS-OCT images, where each cell represents less than 0.005% of the image. MCD achieves SOTA cell detection performance, improving F1 by around 7% over the best baseline and offering a more reliable alternative for cell detection.
- 2) We reveal a critical limitation in anterior uveitis studies using thresholding methods: these studies directly use the results detected by these methods for subsequent analysis without verification, assuming detected objects are cells. Our experiments demonstrate these methods are actually inaccurate, which has crucial clinical implications for future research.

2 Method

Fig. 2 illustrates our MCD framework, which comprises a Field-of-Focus (FoF) module and a Fine-grained Object Detection (FOD) module.

The FoF (bottom-left of Fig. 2) first segments the anterior chamber (AC) region. This resolves the limitation of threshold-based methods that cannot eliminate false positives outside the AC region without manual effort. Unlike conventional DL segmentation models (e.g., U-Net [23]) that require extensive annotated data, FoF utilizes a vision foundation model for zero-shot segmentation. These types of models can segment regions given an image and points (prompts) inside the object of interest. We design an Image-to-AC-Prompt (I2ACP) algorithm (section 2.1) to automatically generate these prompt points. In this work, we adopt Segment Anything Model (SAM) [16] as our foundation model to produce the AC segmentation mask.

Within the AC region, the Fine-grained Object Detection (FOD) module identifies cells through two components. First, a Minuscule Region Proposal (MiRP) algorithm (section 2.2) locates candidate boxes that potentially contain cells (middle of Fig. 2), aiming to identify cells that thresholding methods may miss. These regions preserve fine-grained details, which are processed by our Cell Mamba (section 2.3) to distinguish between actual cells and background noise (bottom-right of Fig. 2). This learns the fine-grained details that traditional DL detectors lose.

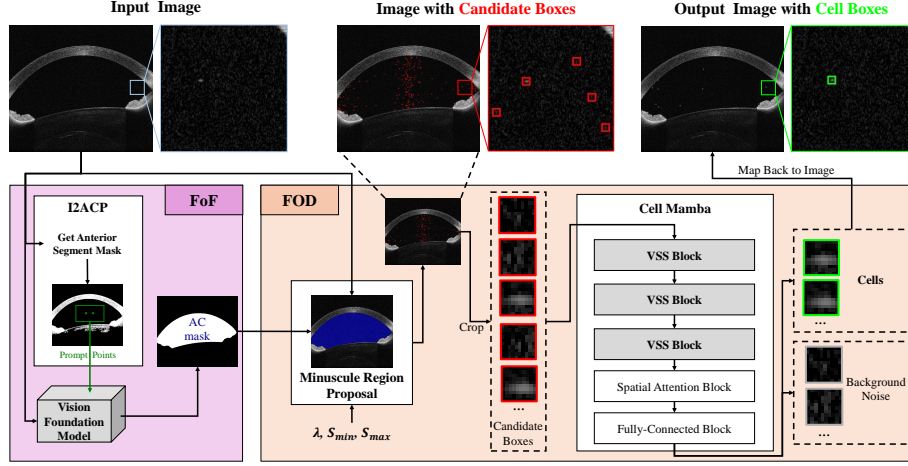


Fig. 2: The overall architecture of our minuscule cell detection framework.

2.1 Image-to-AC-Prompt (I2ACP)

The I2ACP automatically generates prompts for vision foundation models to segment the AC region by exploiting two key anatomical features: (1) the anterior segment (AS) structure (cornea, iris, and lens, see Fig. 1a(1)) appears as high-intensity regions surrounding the AC, and (2) the AS centroid and its neighboring points naturally fall within the AC.

To generate the prompts, the I2ACP first creates a binary mask by thresholding the AS-OCT image at its mean intensity, capturing the bright AS structure (the mask inside the I2ACP block of Fig. 2). The largest connected component typically represents the AS. In cases where AS appears as two disconnected components, we compute the area ratio $r = A_2/A_1$ between the second-largest (A_2) and largest (A_1) components, merging them if $r > \mathcal{R}$. Based on the mask's centroid (x^{cent}, y^{cent}) , we generate a set of prompts p_i using controlled offsets:

$$p_i = (x^{cent} + \delta_i^w, y^{cent} + \delta_i^h)$$

These prompts (green points in the I2ACP block of Fig. 2) guide the vision foundation model to accurately segment the AC region.

2.2 Minuscule Region Proposal (MiRP) for Finding More Cells

MiRP identifies minuscule regions potentially containing cells within the segmented AC region. It exploits the observation that cells appear as bright, minuscule particles against darker backgrounds. Given a high-resolution AS-OCT image \mathcal{I} , an AC mask \mathcal{M}^{AC} from FoF, and cell size bounds $[\mathcal{S}_{min}, \mathcal{S}_{max}]$, MiRP first converts \mathcal{I} to grayscale \mathcal{G} .

We use Otsu’s method [20], a widely-used particle detection method in medical imaging, to calculate an initial thresholding decision boundary \mathcal{T}^{init} from \mathcal{G} . To capture lower-intensity cells that this boundary might miss, we introduce an adjustment factor λ :

$$\mathcal{T}^{opt} = \lambda \times \mathcal{T}^{init}$$

where $\lambda < 1.0$ lowers the decision boundary, enabling detection of dimmer particles overlooked by conventional thresholding methods (Fig. 1b(2)).

Using \mathcal{T}^{opt} , we generate a binary mask $\mathcal{M}^{bright} := \mathcal{G} > \mathcal{T}^{opt}$ and identify candidate cells by extracting connected components within the AC ($\mathcal{M}^{bright} \wedge \mathcal{M}^{AC}$). We then apply size constraints $[\mathcal{S}_{min}, \mathcal{S}_{max}]$ to filter out candidates too large or too small to be cells [9, 19]. For each remaining particle, we generate $w \times h$ bounding boxes centered at its centroid as candidate boxes (shown in the middle of Fig. 2).

While lowering the decision boundary helps detect more potential cells, it inevitably introduces noise particles that conventional thresholding would exclude. We propose Cell Mamba to tackle this by learning the fine-grained details that distinguish actual cells from background noise in these minuscule regions.

2.3 Cell Mamba for Fine-grained Feature Learning

Cell Mamba captures the fine-grained details essential for distinguishing minuscule cells from noise. Given a candidate box $\mathcal{X} \in \mathcal{R}^{w \times h \times 3}$ (where w , h , and 3 represent width, height, and number of channels), we extract features by using three stacked Visual State-Space blocks [18], denoted as $VSS_i(\cdot)$:

$$\mathcal{F}_{vss} = VSS_3(VSS_2(VSS_1(\mathcal{X})))$$

The resulting feature map $\mathcal{F}_{vss} \in \mathcal{R}^{W^v \times H^v \times C^v}$ (where W^v , H^v , and C^v denote the width, height, and channels of the feature map) captures essential fine-grained details. We then apply a spatial attention block inspired by [29] to model spatial relationships, generating an attention map $\mathcal{A} \in \mathcal{R}^{W^v \times H^v \times 1}$:

$$\mathcal{A} = \sigma(\text{Conv}([\text{AvgPool}(\mathcal{F}_{vss}); \text{MaxPool}(\mathcal{F}_{vss})]))$$

where $\sigma(\cdot)$ is the sigmoid activation, $\text{Conv}(\cdot)$ denotes a convolutional layer with kernel size 3 and padding 1, and $\text{AvgPool}(\cdot)$ and $\text{MaxPool}(\cdot)$ represent channel-wise average and maximum pooling operations. The attended feature map is obtained through element-wise multiplication:

$$\mathcal{F}_{sa} = \mathcal{A} \odot \mathcal{F}_{vss}$$

where \mathcal{A} is broadcast across all channels of \mathcal{F}_{vss} . Then, \mathcal{F}_{sa} is flattened and passed through fully connected layers $FC(\cdot)$ to predict if the box contains a cell:

$$\hat{Y} = FC(\text{Flatten}(\mathcal{F}_{sa}))$$

During training, we construct samples by cropping bounding boxes from annotated cell regions (positive samples) and non-cell areas (negative samples),

Table 1: Comparison of AC segmentation performance.

<i>Methods</i>	Swin-Unet [5]	UNet++ [32]	UNet [23]	nnUNet [11]	FoF
<i>IoU</i> (%)	95.88	95.63	95.98	96.23	96.53
<i>Dice</i> (%)	97.10	97.54	97.86	97.92	98.07

maintaining an imbalanced sampling ratio $\mathcal{N}^{pos} : \mathcal{N}^{neg}$. The network is optimized using cross-entropy loss:

$$\mathcal{L} = -\frac{1}{\mathcal{N}^b} \sum_{i=1}^{\mathcal{N}^b} \sum_{c=1}^2 Y_{i,c} \log(\hat{Y}_{i,c})$$

where \mathcal{N}^b is the batch size, i represents the index of each sample in a batch, $Y_{i,c}$ represents the ground truth label for class c , and $\hat{Y}_{i,c}$ is the predicted probability. Training terminates if validation loss shows no improvement for $\mathcal{N}^{patient}$ consecutive epochs.

3 Experiment

3.1 Datasets and Evaluation Metrics

We evaluated MCD using 1,376 high-resolution AS-OCT images (1598×1465 pixels) from the Imaging in Childhood Uveitis studies [2, 9, 26]. Six ophthalmic clinicians annotated the data via the Citizen Science project [13] on the Zooniverse², with final verification by a senior ophthalmologist. The dataset was split into 630 images for AC segmentation and 746 for cell detection evaluation.

We annotated 630 images for AC segmentation, using Intersection over Union (IoU) and the Dice coefficient as the metrics. We annotated 746 images for cell detection. Clinicians annotated by clicking cells, generating 10×10 pixel ground truth boxes centered at clicked points, each occupying < 0.005% of image area.

We use two matching criteria for a correct cell detection: (1) bounding box containing a ground truth point (denoted as "point") and (2) IoU>30% with ground truth box (denoted as "30"). The low IoU threshold accounts for significant IoU decrease in small object detection [8]. Each ground truth cell is matched only once to prevent double-counting.

To evaluate cell counting accuracy, we calculate two Mean Absolute Error (MAE) metrics across all images: MAE^{all} measures the average absolute difference between predicted and ground truth cell counts across all images, while MAE^c considers only correctly detected cells. To evaluate spatial localization accuracy of detection boxes, we calculate precision, recall, and F1-score based on the two matching criteria. All experiments used 40/10/50% training/validation/testing splits, repeated 5 times with different random splits.

² <https://www.zooniverse.org/>

Table 2: Comparison Between Predicted and Ground Truth Cell Counts. † indicates methods that process only 300×300 image patches instead of the whole image.

Methods	MAE ^{all}	MAE ^c _{point}	MAE ^c ₃₀
† Faster-RCNN [21]	2.49	1.03	1.15
† Cascade-RCNN [4]	2.45	0.86	0.97
† RetinaNet [17]	1.74	1.12	1.24
† RepPoint [30]	2.88	1.18	1.29
Otsu [20]	1.59	1.59	1.69
Isodata [22]	1.57	1.56	1.66
MCD (Ours)	0.85	0.72	0.84

3.2 Baselines and Implementation Details

In the I2ACP, we generated two prompt points at offsets $(\delta_1^w, \delta_1^h) = (0, \delta' \times W)$ and $(\delta_2^w, \delta_2^h) = (0, -\delta' \times W)$ from the centroid of the AS. We used grid search over δ' [0–0.5] and \mathcal{R} [0.4–0.95] in 0.05 steps on the validation set, yielding optimal values of $\delta' = 0.1$ and $\mathcal{R} = 0.65$. In the MiRP, λ was searched from 0.70 to 1.00 in 0.01 steps on the validation set. As λ grows, MCD’s precision rises and recall declines, with F1 exhibiting a rise-peak-decline pattern. The optimal λ was selected based on the F1 peak and empirically set as 0.83. We set $\mathcal{S}_{max} = 25$ based on clinical observations that larger objects are unlikely to be cells [9, 19], and $\mathcal{S}_{min} = 1$ to include all potential cell candidates. Both w and h were set as 10. For our Cell Mamba, \mathcal{N}^{pos} , \mathcal{N}^{neg} , $\mathcal{N}^{patient}$, and \mathcal{N}^b were set to 1, 5, 30 and 128, respectively.

We compared FoF against SOTA segmentation baselines (UNet [23], UNet++ [32], nnUNet [11], Swin-UNet [5]) and evaluated cell detection against both threshold-based and DL-based approaches, considering only cells within segmented AC regions for a fair comparison.

Threshold-based methods include Otsu [20] and Isodata [22] (widely used in AS-OCT [1, 3, 9, 19, 28]). Both methods identify connected components (removing >25 and <2 pixels [9, 19]), with results converted to 10×10 pixel boxes at centroids. For DL-based methods, we evaluated Faster-RCNN [21], Cascade-RCNN [4], RetinaNet [17], RepPoint [30], DETR [6], Deformable-DETR [33], DINO [31], and YOLO11 [12]. Our initial attempts to train these models on full-resolution images failed to detect any cells as cells are miniscule, occupying less than 0.005% of the image area. Then, we divided images into 300×300 patches for training/inference, with results mapped back and standardized to 10×10 boxes. For brevity, we report only models achieving F1 scores $>55\%$.

4 Results and Discussion

Table 1 shows FoF, with zero-shot training, achieves a high IoU of 96.53% and a Dice coefficient of 98.07%, results comparable with the baseline models that require training.

Table 3: Cell detection results on Precision (%), Recall(%) and F1(%). † indicates methods that process only 300×300 image patches instead of the whole image.

Methods	Precision _{point}	Recall _{point}	F1 _{point}	Precision ₃₀	Recall ₃₀	F1 ₃₀
† Faster-RCNN [21]	52.27	78.56	62.77	50.56	75.99	60.72
† Cascade-RCNN [4]	54.93	82.11	65.82	53.33	79.71	63.9
† RetinaNet [17]	59.13	76.66	66.76	57.16	74.11	64.54
† RepPoint [30]	47.41	75.47	58.24	45.95	73.14	56.44
Otsu [20]	91.22	67.35	77.48	88.37	65.25	75.06
Isodata [22]	90.43	68.03	77.65	87.72	65.99	75.32
MCD	85.10	84.79	85.01	82.48	82.02	82.32

Table 2 shows that our MCD achieves the lowest cell counting error across all metrics, with MAE^{all} , MAE_{point}^c , MAE_{30}^c of 0.85, 0.72, 0.84, respectively. This means that, on average, our MCD’s cell count differs from the ground truth by less than one cell per image, indicating highly accurate quantification of inflammatory cells.

Table 3 shows that Our MCD achieves the highest F1 scores across all criteria: 85.01% for $F1_{point}$, and 82.32% for $F1_{30}$. Our MCD improves F1 by around 7% over the best baseline. This significant improvements demonstrates our MCD’s ability to effectively detect cells while maintaining a low false detection rate.

Previous thresholding detection methods used in anterior uveitis studies [1, 3, 9, 14, 15, 19, 24, 25] reported strong correlations between their detected cell counts and the clinical severity score of anterior uveitis. However, these studies proceeded without verifying the correctness of individual cell detections. Our evaluation reveals two major limitations in these methods: cell count inaccuracy (as shown in Table 2) and imprecise spatial localization of detected cells (as shown in Table 3). While these methods demonstrated correlations with clinical scores, these correlations were likely based on incomplete and spatially inaccurate cell detection. This imprecision could negatively impact future research that requires accurate cell counting or studies investigating the relationship between spatial cell distribution patterns and disease progression.

In contrast, our MCD achieves good performance in detecting the cells. It opens up new possibilities for investigating the spatial distribution of inflammatory cells within the AC, which could enhance our understanding of disease pathogenesis and progression. Such insights could contribute to the development of more personalized treatment strategies, potentially improving patient outcomes. Furthermore, our framework could serve as a valuable reference for detecting extremely small objects in other medical imaging scenarios.

Despite our MCD’s superior performance in detecting minuscule cells in AS-OCT images, several limitations must be acknowledged. A challenge is the limited availability of public datasets for these tasks, making the development and release of comprehensive datasets essential for future research advances. Additionally, AS-OCT image quality can vary - artifacts and quality degradation may occur [7], potentially affecting MCD’s effectiveness.

5 Conclusion

We tackle the challenge of automatically detecting cells in AS-OCT images for anterior uveitis study by proposing a detection framework named MCD. We also reveals critical limitations in previous anterior uveitis studies that relied on threshold-based methods for cell detection, suggesting potential underestimation of inflammatory cell populations. MCD not only improves the technical aspects of cell detection but also holds promise for advancing clinical care and research in ocular inflammation.

Acknowledgments. This work was supported in part by the AWS Doctoral Scholarship in Digital Innovation (awarded through the UCL Centre for Digital Innovation) and an NIHR Clinician Scientist grant (CS-2018-18-ST2-005; ALS). We are grateful to the participating children and their families and the clinicians within the Great Ormond Street and Moorfields Eye Hospital Uveitis clinics.

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

References

1. Agarwal, A., Ashokkumar, D., Jacob, S., Agarwal, A., Saravanan, Y.: High-speed optical coherence tomography for imaging anterior chamber inflammatory reaction in uveitis: clinical correlation and grading. *American journal of ophthalmology* **147**(3), 413–416 (2009)
2. Akbarali, S., Rahi, J.S., Dick, A.D., Parkash, K., Etherton, K., Edelsten, C., Liu, X., Solebo, A.L.: Imaging-based uveitis surveillance in juvenile idiopathic arthritis: feasibility, acceptability, and diagnostic performance. *Arthritis & Rheumatology* **73**(2), 330–335 (2021)
3. Baghdasaryan, E., Tepelus, T.C., Marion, K.M., Huang, J., Huang, P., Sadda, S.R., Lee, O.L.: Analysis of ocular inflammation in anterior chamber—involving uveitis using swept-source anterior segment oct. *International Ophthalmology* **39**, 1793–1801 (2019)
4. Cai, Z., Vasconcelos, N.: Cascade r-cnn: Delving into high quality object detection. In: *CVPR’18*. pp. 6154–6162 (2018)
5. Cao, H., Wang, Y., Chen, J., Jiang, D., Zhang, X., Tian, Q., Wang, M.: Swin-unet: Unet-like pure transformer for medical image segmentation. In: *ECCV’22*. pp. 205–218. Springer (2022)
6. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: *ECCV’20*. pp. 213–229. Springer (2020)
7. Chen, B., Solebo, A.L., Taylor, P.: Automated image quality assessment for anterior segment optical coherence tomograph. In: *ISBI’23*. pp. 1–4. IEEE (2023)
8. Cheng, G., Yuan, X., Yao, X., Yan, K., Zeng, Q., Xie, X., Han, J.: Towards large-scale small object detection: Survey and benchmarks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023)

9. Etherton, K., Rahi, J.S., Petrushkin, H., Dick, A.D., Akbarali, S., Pattani, R., Hau, S., Lacassagne, S., Liu, X., Denniston, A.K., et al.: Quantitative and qualitative assessment of anterior segment optical coherence tomography capture of disease state in childhood anterior uveitis. *British Journal of Ophthalmology* **107**(7), 966–972 (2023)
10. Invernizzi, A., Marchi, S., Aldigeri, R., Mastrofilippo, V., Viscogliosi, F., Soldani, A., Adani, C., Garoli, E., Viola, F., Fontana, L., et al.: Objective quantification of anterior chamber inflammation: measuring cells and flare by anterior segment optical coherence tomography. *Ophthalmology* **124**(11), 1670–1677 (2017)
11. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods* **18**(2), 203–211 (2021)
12. Jocher, G., Qiu, J., Chaurasia, A.: Ultralytics YOLO (Jan 2023), <https://github.com/ultralytics/ultralytics>, accessed: June 2025
13. Jones, M.L., Spiers, H.: The crowd storms the ivory tower. *Nature methods* **15**(8), 579–580 (2018)
14. Kang, T.S., Lee, Y., Lee, S., Kim, K., Lee, W.s., Lee, W., Kim, J.H., Han, Y.S.: Development of fully automated anterior chamber cell analysis based on image software. *Scientific Reports* **11**(1), 10670 (2021)
15. Keino, H., Aman, T., Furuya, R., Nakayama, M., Okada, A.A., Sunayama, W., Hatanaka, Y.: Automated quantitative analysis of anterior segment inflammation using swept-source anterior segment optical coherence tomography: a pilot study. *Diagnostics* **12**(11), 2703 (2022)
16. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: ICCV’23. pp. 4015–4026 (2023)
17. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV’17. pp. 2999–3007 (2017)
18. Liu, Y., Tian, Y., Zhao, Y., Yu, H., Xie, L., Wang, Y., Ye, Q., Jiao, J., Liu, Y.: Vmamba: Visual state space model. In: NeurIPS’24. pp. 103031–103063 (2024)
19. Lu, M., Wang, X., Lei, L., Deng, Y., Yang, T., Dai, Y., Li, Y., Gan, X., Hu, Y., Chen, H., et al.: Quantitative analysis of anterior chamber inflammation using the novel casia2 optical coherence tomography. *American Journal of Ophthalmology* **216**, 59–68 (2020)
20. Otsu, N., et al.: A threshold selection method from gray-level histograms. *Automatica* **11**(285-296), 23–27 (1975)
21. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence* **39**(6), 1137–1149 (2016)
22. Ridler, T., Calvard, S.: Picture thresholding using an iterative selection method. *IEEE Trans. Syst. Man Cybern* **8**(8), 630–632 (1978)
23. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI’15. pp. 234–241. Springer (2015)
24. Rose-Nussbaumer, J., Li, Y., Lin, P., Suhler, E., Asquith, M., Rosenbaum, J.T., Huang, D.: Aqueous cell differentiation in anterior uveitis using fourier-domain optical coherence tomography. *Investigative ophthalmology & visual science* **56**(3), 1430–1436 (2015)
25. Sharma, S., Lowder, C.Y., Vasanji, A., Baynes, K., Kaiser, P.K., Srivastava, S.K.: Automated analysis of anterior chamber inflammation by spectral-domain optical coherence tomography. *Ophthalmology* **122**(7), 1464–1470 (2015)

26. Solebo, A.L., Bellchambers, A., Kellett, S., Rahi, J.S., Dick, A.D.: Establishing the normative data set necessary for imaging-based childhood uveitis surveillance: A cross-sectional study. *Investigative Ophthalmology & Visual Science* **65**(1), 9–9 (2024)
27. Standardization of Uveitis Nomenclature (SUN) Working Group, et al.: Standardization of uveitis nomenclature for reporting clinical data. results of the first international workshop. *American journal of ophthalmology* **140**(3), 509–516 (2005)
28. Uthayananthan, P., Tanwar, N., Rahi, J.S., Dick, A.D., Solebo, A.L.: Imaging-based detection of anterior chamber inflammation: a comparative diagnostic accuracy study. *American Journal of Ophthalmology* **270**, 131–139 (2025)
29. Woo, S., Park, J., Lee, J.Y., Kweon, I.S.: Cbam: Convolutional block attention module. In: *ECCV'18*. pp. 3–19 (2018)
30. Yang, Z., Liu, S., Hu, H., Wang, L., Lin, S.: Reppoints: Point set representation for object detection. In: *ICCV'19*. pp. 9657–9666 (2019)
31. Zhang, H., Li, F., Liu, S., Zhang, L., Su, H., Zhu, J., Ni, L.M., Shum, H.Y.: Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605* (2022)
32. Zhou, Z., Rahman Siddiquee, M.M., Tajbakhsh, N., Liang, J.: Unet++: A nested u-net architecture for medical image segmentation. In: *DLMIA ML-CDS'18*. pp. 3–11. Springer (2018)
33. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159* (2020)