Accurate Iris Segmentation on Near Infrared Images

A novel, robust, human-level approach using template matching and ellipse-constrained active contours

Student

Andrin Rütsche

Introduction: Accurate anatomical iris segmentation is essential for gaze tracking. Like in many practical applications, we opt for NIR images of the eye instead of RGB images, because they show effectively reduced glare and exhibit better iris to pupil contrast in general. However specular reflections, eye lashes and eye lids cause deviations from a circular or elliptic contour assumption and thus impose a challenge for classic contour search algorithms like RANSAC, Hough-Transform or traditional active contour algorithms. Thus, they typically are not suited to generate subpixel precise iris and pupil contours. In contrast, modern ML-based models output pixelwise masks of the visible iris, which are useful for biometric segmentation; however, they cannot directly provide anatomical boundaries with subpixel precision without expensive training and/or dedicated post processing.

Approach / Technology: We propose an active contour framework that exploits eye anatomy and the radial structure of the pupil and iris boundary. A direction-aware radial edge map guides the contour, which is constrained to elliptical shapes. The contour initially samples two narrow angular sectors to lock onto a reliable boundary segment, then progressively expands its sampling range while suppressing low-confidence or inconsistent regions. This staged expansion steers the contour away from distractors such as eyelids, eyelashes, and iris texture. We further introduce a robust pupil-candidate proposal based on soft template matching and intensity statistics, yielding reliable initialization even in the presence of strong reflections. Image 1 displays the main structure of this pipeline.

Result: Evaluation on manually annotated NIR iris images from the CASIA V4 Interval dataset using IoU and F1 scores for both iris and pupil shows that the proposed method achieves human-level geometric accuracy. The reference annotations are provided by the IRISSEG-EP ground-truth ellipses introduced by Hofbauer et al., which include independent annotations from two human annotators. As shown in Table 1, the algorithm matches each annotator with comparable accuracy to the agreement observed between the two annotators themselves. This inter-annotator agreement defines a practical upper bound on achievable performance, reflecting the inherent uncertainty of manual iris annotation. Image 2 shows an OST internal image with a segmentation result in the lower tail of the observed performance distribution (IoU pupil: 0.937, IoU iris: 0.931), chosen to allow clear visual distinction between ground truth and prediction. The proposed pipeline shows reduced robustness for cases of strong pupillary constriction and for subjects with poor pupil–iris contrast, as commonly observed in elderly populations. Performance further degrades with increasing occlusion of the pupil and iris, as well as in the presence of glares located directly on either the pupil or iris boundary. Since pupil segmentation is used to initialize the subsequent iris fitting stage, failures in pupil localization increase the likelihood of downstream iris segmentation errors.
In conclusion, these results demonstrate that the proposed pipeline is a viable drop-in replacement for established iris segmentation techniques, including RANSAC, Hough-based methods, and generic active contours, with particular relevance to gaze analysis systems.

Image 1: Segmentation Pipeline
Own presentment

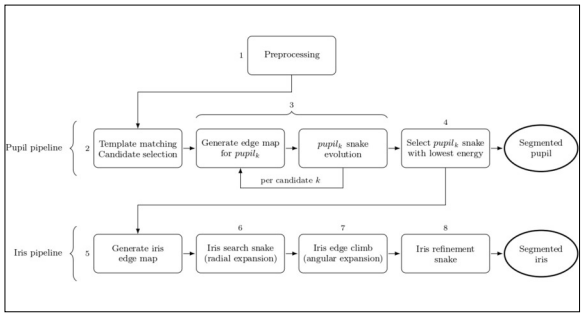

Image 2: Prediction output (green), groundtruth (red). The image has been blurred intentionaly to protect privacy.
Own presentment

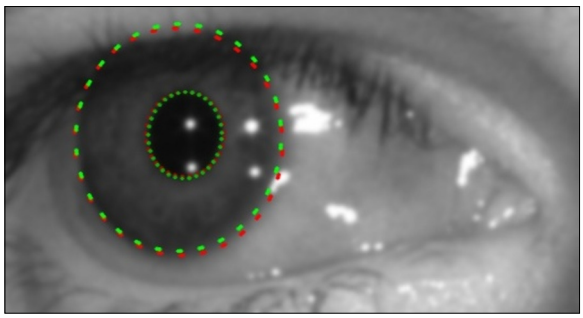

Table 1: Quantitative results, IoU and F1 scores on 2639 NIR eye images from the CASIA V4i dataset.
Own presentment

Comparison	F1 Pupil	F1 Iris	mIoU Pupil	mIoU Iris
Inter-Annotator (A → B)	0.984 ± 0.01	0.969 ± 0.01	0.968 ± 0.01	0.940 ± 0.03
Inter-Annotator (B → A)	0.984 ± 0.01	0.969 ± 0.01	0.968 ± 0.02	0.940 ± 0.03
Window Snake (WS → A)	0.984 ± 0.01	0.968 ± 0.02	0.969 ± 0.02	0.938 ± 0.03
Window Snake (WS → B)	0.984 ± 0.01	0.970 ± 0.01	0.965 ± 0.02	0.942 ± 0.03

Advisors
Prof. Dr. Martin Weisenhorn, Marc Benz

Subject Area
Image Processing and Computer Vision