

EDIL-SegRayDP: Training-Free Iris Segmentation via Segmentation-First Ray-Wise Dynamic Programming

Trong-Thua Huynh^{1,*}, De-Thu Huynh², Cong-Sang Duong¹

¹Posts and Telecommunications Institute of Technology, Ho Chi Minh, Viet Nam

²The Saigon International University, Ho Chi Minh, Viet Nam

Abstract

Iris segmentation remains a critical yet challenging stage in biometric recognition, especially under off-axis capture, eyelid and eyelash occlusion, specular reflections, and illumination variations that violate the circular and unobstructed assumptions of classical pipelines. We present EDIL-SegRayDP, a training-free and explainable iris segmentation framework that departs from the conventional localization-first paradigm by treating annulus recovery as the primary optimization objective. Rather than committing early to a global center/radius hypothesis and refining it afterward, the proposed method performs segmentation-first boundary recovery with segmentation-aware center rescue and fail-safe outer-boundary control. Occlusion is handled explicitly through geometry-normalized masking and validity-aware annulus construction, while all key parameters are defined in scale-normalized form for cross-dataset portability. Experiments under a fixed-configuration protocol on IITD and CASIA-IrisV4-Interval show strong non-CNN performance with CPU-only inference, achieving an iris-mask mean Dice of 0.9106 on IITD and 0.9377 on CASIA-IrisV4-Interval, with corresponding pupil Dice of 0.9763 and 0.9755. Additional full-benchmark evaluations on CASIA-IrisV4-Lamp and CASIA-IrisV4-Thousand further confirm the portability of the proposed framework across more challenging and larger-scale subsets. Under the evaluation protocol adopted in this study, these results compare favorably with a recent training-free reference, supporting EDIL-SegRayDP as a competitive and interpretable training-free alternative for iris segmentation under non-ideal imaging conditions.

Received on 18 April 2026; accepted on 19 May 2026; published on 19 May 2026

Keywords: Iris segmentation, Training-free, Dynamic programming, Occlusion modeling, Explainable biometrics

Copyright © 2026 Trong-Thua Huynh *et al.*, licensed to EAI. This is an open access article distributed under the terms of the [CC BY-NC-SA 4.0](#), which permits copying, redistributing, remixing, transformation, and building upon the material in any medium so long as the original work is properly cited.

doi:10.4108/eetinis.132.12656

1. Introduction

Iris recognition is widely used because iris texture is highly distinctive and largely stable over time. Seminal work by Daugman and Wildes established the foundations of modern iris biometrics, where reliable recognition critically depends on accurate iris segmentation [1, 2]. Segmentation aims to isolate the iris annulus from surrounding structures such as the pupil, sclera, eyelids/eyelashes, and specular reflections, since errors at this stage propagate directly to normalization and matching. In practice, however, iris segmentation remains challenging under non-ideal acquisition conditions. Off-axis capture and near-field imaging can

produce visibly non-circular pupil/limbal boundaries, while eyelid/eyelash occlusion and specular highlights may remove or corrupt large portions of boundary evidence. These factors frequently degrade gradient-driven pipelines and cause unstable masks under illumination variation and partial occlusion [3].

Existing approaches broadly fall into two categories. Classical pipelines rely on geometric assumptions (typically circular or elliptical contours) combined with edge/gradient search and voting strategies, offering interpretability and low computational cost, but they can be brittle when boundary evidence is incomplete or when the limbal boundary deviates from a global shape model [4]. In contrast, deep learning methods perform end-to-end segmentation and often improve robustness by learning priors from

*Corresponding author: Email: thuaht@ptit.edu.vn

annotated masks; representative examples include iris-specific architectures and cross-modality designs for NIR/VIS robustness [5, 6]. However, learning-based solutions introduce practical costs in data annotation and maintenance, and may require retraining or careful re-validation under domain shifts.

Beyond the data requirement, many pipelines are also *localization-first*: they estimate a global center/radius model and then refine or post-process the segmentation. Under heavy occlusion or weak limbic gradients, early localization failures can irreversibly bias downstream steps. This observation motivates a *segmentation-first* design in which the boundary is recovered as the primary optimization objective, and center/radius hypotheses are validated (and rescued) using segmentation evidence.

Despite decades of research, existing pipelines still struggle to simultaneously achieve robustness, interpretability, and training-free deployment under non-ideal imaging conditions. To address this gap, we propose *EDIL-SegRayDP*, a segmentation-first framework that formulates iris segmentation as guarded geometry-constrained optimization with ray-wise dynamic programming, scale-normalized evidence, and explicit occlusion handling.

Unlike localization-first pipelines that commit early to a global geometric hypothesis, the proposed method treats annulus plausibility as the primary objective and uses geometric rescue only as a supporting mechanism. This distinction matters because many catastrophic failures arise from early commitment to an implausible center or outer search region. Accordingly, the contribution of *EDIL-SegRayDP* is not merely a more robust classical pipeline, but a reformulation of training-free iris segmentation as guarded segmentation-first optimization, in which geometric localization serves annulus plausibility rather than acting as the final objective. This reformulation is central to reducing catastrophic failure under weak or corrupted boundary evidence.

1.1. Motivation and Gap

Real-world eye images frequently violate the idealized assumptions behind classical iris segmentation. In off-axis or near-field capture, pupil and limbal boundaries can become noticeably non-circular, and eyelids, eyelashes, and specular highlights may remove or corrupt large portions of boundary evidence. Consequently, pipelines that commit early to a single global circle/ellipse estimated from incomplete gradients often produce biased fits and unstable masks under non-ideal conditions [3].

Learning-based segmenters alleviate many of these issues by learning strong shape and appearance priors from annotated data [5]. However, they introduce practical costs (pixel-level labeling and training/maintenance overhead) and may degrade under domain shifts across sensors, lighting, and capture protocols.

Moreover, their failure modes are often less transparent, complicating debugging and validation when deterministic behavior and explainability are required. As a result, there remains a practical need for a method that is not only robust under non-ideal conditions, but also deterministic, interpretable, and deployable without a training stage.

We therefore target the following gap: a *deterministic*, *parameter-light*, and *training-free* iris segmentation method that:

1. supports non-circular and off-axis boundaries without relying on a single global circle/ellipse,
2. explicitly models eyelid/eyelash/specular occlusions in an angularly varying manner,
3. performs *segmentation-first* optimization to avoid early-commitment failures, and
4. uses dimensionless, geometry-normalized parameters so that behavior is consistent across image scales and sensors.

EDIL-SegRayDP is designed to address this gap through a single deterministic optimization framework in which annulus plausibility, angularly varying boundary recovery, and guarded geometric rescue are coupled rather than handled as loosely connected post-processing stages.

1.2. Contributions

Our main contributions are:

1. A segmentation-first, training-free iris segmentation framework that reformulates annulus recovery as guarded geometry-constrained optimization, rather than a localization-first pipeline followed by boundary refinement.
2. A segmentation-aware geometric rescue mechanism in which center hypotheses are validated by early annulus plausibility and replaced when they would otherwise induce catastrophic downstream failure.
3. A fail-safe outer-boundary design that couples adaptive search-band widening with an outer-floor constraint to reduce pathological inward collapse under weak limbic evidence.

4. Geometry-normalized occlusion handling and validity-aware annulus construction, enabling angularly varying recovery under eyelid, eyelash, and specular corruption while preserving cross-dataset portability.
5. A fixed-configuration evaluation on IITD, CASIA-IrisV4-Interval, CASIA-IrisV4-Lamp, and CASIA-IrisV4-Thousand, including ablation, failure analysis, and CPU-only runtime characterization, to assess robustness beyond protocol-matched core benchmarks.

The remainder of this paper is organized as follows: Section 2 reviews related iris segmentation literature. Section 3 presents EDIL-SegRayDP in detail. Section 4 describes the evaluation protocol. Section 5 reports quantitative results and analysis. Section 6 concludes the paper.

2. Related Work

2.1. Geometry-Based Iris Localization and Segmentation

Classical iris segmentation is predominantly geometry-driven: it first localizes the pupil and limbus, then constructs an iris annulus for downstream normalization and matching. Early seminal pipelines assume near-circular boundaries and rely on strong radial intensity changes, such as Daugman's integro-differential operator and Wildes' circular Hough-transform-based localization [1, 2, 7]. Practical implementations further combine edge cues, voting schemes, and post-processing heuristics to obtain a usable annulus efficiently, as exemplified by widely used open toolkits [8].

However, these assumptions are frequently violated in non-ideal acquisition. Off-axis capture induces apparent non-circularity, while eyelid/eyelash occlusion and specular highlights corrupt the boundary evidence needed for robust circle/ellipse fitting. To address such conditions, prior geometry-based work explored more robust strategies for noisy and non-cooperative images, including deformable models and stronger low-level constraints [3, 4]. More recent geometry-oriented formulations incorporate polar-space modeling and robust fitting to better tolerate irregular boundaries and partial occlusion [9].

Related classical and hybrid training-free studies also explored robust localization and segmentation under low contrast, occlusion, and non-cooperative imaging conditions, further highlighting the continued relevance of interpretable geometry-based pipelines [10–12].

Recent benchmarking and robustness-oriented studies further emphasize that iris segmentation performance can vary substantially across datasets and acquisition conditions, particularly under occlusion, illumination instability, and boundary deformation [13, 14]. In particular, Sumi et al. [13] conducted a comprehensive evaluation of iris segmentation methods across multiple benchmark datasets and highlighted the persistent difficulty of achieving robust cross-dataset segmentation under heterogeneous imaging conditions. More recent robustness-oriented segmentation frameworks have also continued to emphasize stable localization and segmentation under challenging imaging conditions and heterogeneous acquisition environments [14]. These findings further reinforce the importance of robustness-aware evaluation and failure-aware segmentation design beyond average segmentation accuracy alone.

2.2. Learning-Based Iris Segmentation and Domain Robustness

Deep learning methods reframe iris segmentation as dense prediction and often provide strong performance by learning shape and appearance priors from annotated masks. Encoder-decoder architectures such as U-Net became a widely adopted baseline, including early end-to-end iris segmentation variants [15, 16]. Subsequent works propose iris-specific designs to improve boundary detail and robustness under non-cooperative capture, including multi-task attention and lightweight segmentation networks [5, 17, 18]. Cross-domain generalization remains a practical challenge, motivating approaches that explicitly target robustness across sensing modalities (e.g., NIR vs. VIS) [6].

In parallel, transformer-based segmentation has gained momentum due to improved global context modeling, including hierarchical attention backbones adapted to U-Net-like decoding [19, 20]. Foundation-model segmentation further broadens generality, and recent work adapts such models to iris segmentation with task-specific tuning and post-processing [21, 22]. Despite their accuracy, learning-based solutions typically require labeled masks and training infrastructure, and their behavior can be harder to audit or maintain when operational conditions shift.

2.3. Energy Minimization and Global Optimization for Non-Ideal Boundaries

Optimization-based segmentation provides a complementary axis between hand-crafted geometry and learned priors. By explicitly defining an energy that balances boundary evidence and regularization, these methods can enforce global consistency while remaining deterministic and interpretable. Classic examples

include active-contour style formulations for non-ideal irises [4], while polar-space energy minimization and robust fitting improve tolerance to low contrast and partial occlusion [9]. A key advantage is transparency: each objective term corresponds to a measurable signal (edge strength, smoothness/curvature, occlusion penalty), which enables systematic debugging and controlled deployment on resource-limited hardware.

This transparency is particularly attractive in biometric systems, where deterministic behavior, reproducibility, and interpretable failure analysis are often essential deployment requirements.

Nevertheless, many optimization-based pipelines still rely on early global hypotheses (e.g., a single circle/ellipse or a narrow search band) and may fail catastrophically when the initial estimate is biased. This motivates segmentation-first optimization designs, where intermediate segmentation outcomes can be used to rescue the geometry and prevent pathological boundary collapse.

2.4. Positioning of EDIL-SegRayDP

EDIL-SegRayDP follows the optimization-based philosophy but emphasizes three practical requirements that remain under-explored in prior pipelines: (i) *segmentation-first* ray-wise dynamic programming that allows angularly varying inner/outer boundaries rather than committing to a single global circle/ellipse, (ii) explicit and *geometry-normalized* occlusion handling with validity-aware annulus construction (eyelid/eyelash/specular treated as first-class factors), and (iii) *fail-safe boundary controls* that improve stability under weak or corrupted evidence, including segmentation-aware center rescue, adaptive search-band widening, and an outer-floor guard that prevents pathological outer-boundary collapse. All key parameters are defined in scale-normalized form, promoting portability across datasets and sensors without retraining.

As a result, EDIL-SegRayDP bridges classical geometric reasoning and globally consistent optimization while remaining fully training-free and CPU-deployable. Unlike recent learning-based robustness-oriented networks, the proposed framework remains fully training-free and does not rely on annotated masks or dataset-specific retraining for deployment across heterogeneous acquisition conditions. At the same time, unlike many earlier geometry-driven pipelines that depend heavily on a single global geometric hypothesis, EDIL-SegRayDP emphasizes segmentation plausibility, guarded optimization, and explicit failure prevention under unreliable boundary observations. Compared with recent active-contour-based training-free methods, the proposed framework does

not rely on a single global contour family for final annulus recovery, but instead performs angularly varying boundary optimization with explicit fail-safe controls.

Importantly, these fail-safe mechanisms are not introduced as isolated heuristics after boundary estimation, but are integrated into the optimization logic that determines whether and how annulus recovery proceeds under weak evidence.

Accordingly, EDIL-SegRayDP is positioned not as a minor refinement of localization-first geometry fitting, but as a segmentation-first optimization framework specifically designed for robust and interpretable iris segmentation under non-ideal imaging conditions.

3. Methodology

3.1. Notation and Overview

Let $I \in \mathbb{R}^{H \times W}$ denote a grayscale eye image. The objective is to estimate the pupil center $c = (c_x, c_y)$, together with the inner boundary $r_{\text{in}}(\theta)$ corresponding to the pupil-iris interface and the outer boundary $r_{\text{out}}(\theta)$ corresponding to the iris-limbus interface, for angles $\theta \in [0, 2\pi)$. The final iris region is represented by a valid annulus mask derived from the area between the two boundaries after excluding detected occlusions such as eyelids, eyelashes, and specular highlights. The proposed framework is fully training-free and uses dimensionless, geometry-normalized parameters to promote stable behavior across image scales and datasets.

At a high level, EDIL-SegRayDP is organized into two coupled modules, as illustrated in Fig. 1. The first module, *Preprocessing and Coarse Geometry*, normalizes the image, forms a plausible center prior, rescues suspicious center hypotheses when necessary, and produces a coarse annulus prior together with an adaptive outer search band. The second module, *Segmentation and Fail-safe Optimization*, validates the coarse annulus, selects a guarded or standard outer search mode, and recovers the final iris region through outer-first and inner constrained dynamic programming under geometric safeguards.

More specifically, the proposed pipeline proceeds through the following stages:

1. image preprocessing and evidence normalization;
2. multi-cue center initialization with segmentation-aware center rescue;
3. coarse inner/outer radius estimation with adaptive outer-band widening;
4. guarded outer-boundary optimization in polar coordinates;
5. constrained inner-boundary recovery;

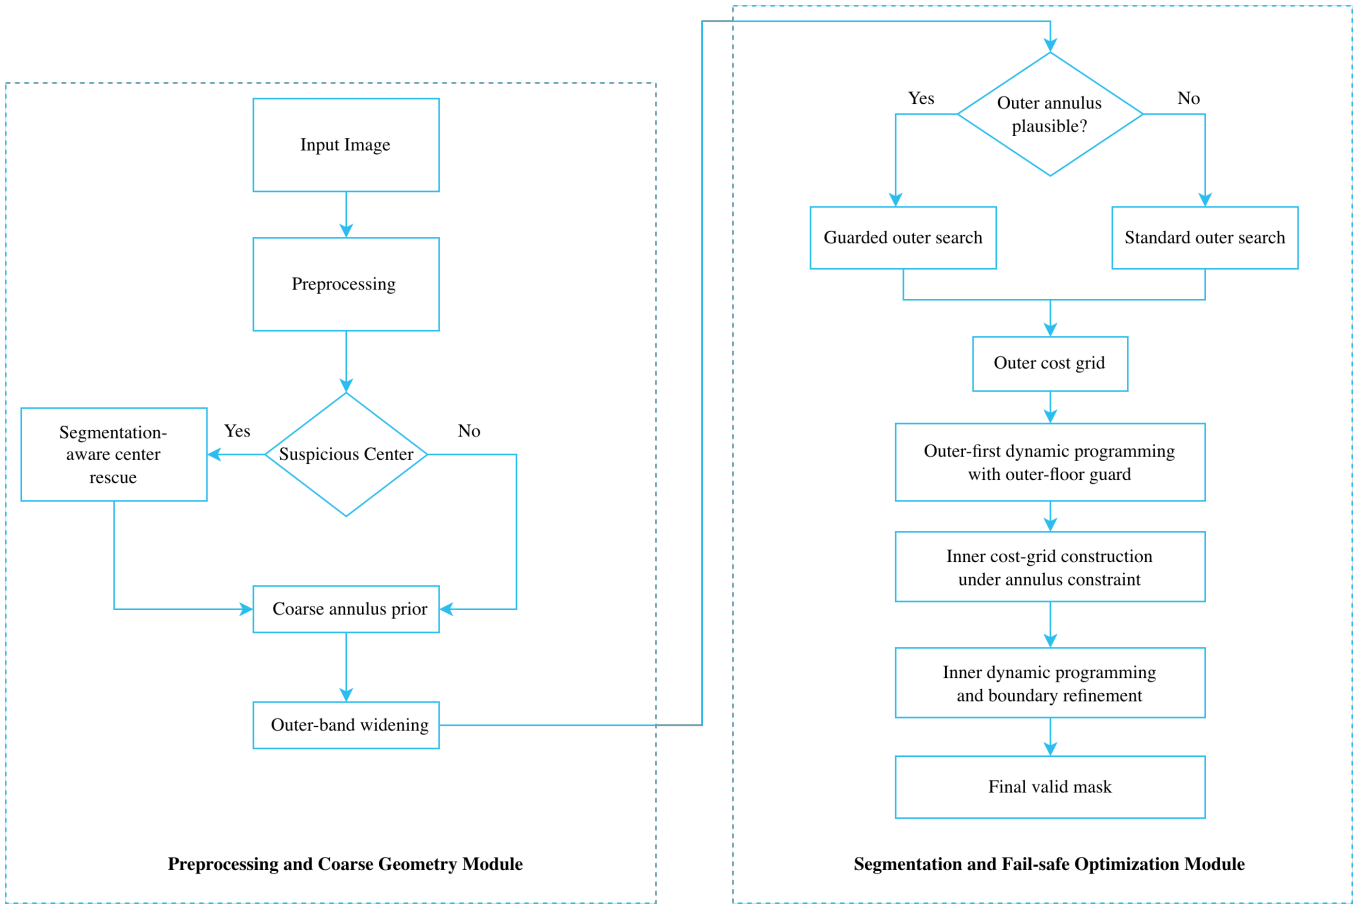


Figure 1. Architecture overview of EDIL-SegRayDP. The left module forms a plausible annulus prior through preprocessing, suspicious-center detection, segmentation-aware center rescue, and adaptive outer-band widening. The right module performs guarded boundary optimization by validating the coarse annulus, selecting an appropriate outer search strategy, and recovering the final iris region through outer-first and inner constrained dynamic programming

6. occlusion-aware annulus composition and final valid-mask construction.

3.2. Preprocessing and Evidence Normalization

The input image is first converted to grayscale when necessary and then normalized to reduce illumination variance. In the current implementation, local contrast is enhanced using contrast-limited adaptive histogram equalization (CLAHE), which improves the stability of subsequent radial evidence under non-uniform brightness. Specular highlights are detected by percentile-based intensity thresholding and refined by simple morphological cleanup. When enabled, these regions may be inpainted using a lightweight Telea-based scheme to reduce corrupted local gradients. Regardless of whether inpainting is applied, detected specular pixels are always recorded in the occlusion mask M_{occ} (Sec. 3.5).

To make later decisions less sensitive to absolute intensity scale, all gradient- or contrast-based evidence

maps are robustly normalized using percentile statistics. For an evidence map E (e.g., a difference-of-means or radial gradient response), the normalized evidence is defined as

$$\tilde{E} = \text{clip}\left(\frac{E - P_5(E)}{P_{95}(E) - P_5(E) + \epsilon}, 0, 1\right), \quad (1)$$

where $P_5(\cdot)$ and $P_{95}(\cdot)$ denote the 5th and 95th percentiles, respectively, and ϵ is a small constant introduced for numerical stability. This normalization suppresses outlier sensitivity and makes subsequent thresholds and weighting terms more consistent across datasets and acquisition conditions.

3.3. Segmentation-Aware Center Rescue and Coarse Radius Estimation

The proposed framework follows a *segmentation-first* principle: the pupil center is treated as a prior whose quality is judged by the plausibility of the resulting annulus segmentation, rather than being optimized as an isolated end objective. After obtaining an initial

fused center c , the pipeline performs a lightweight *center rescue* whenever the center hypothesis appears suspicious.

Suspicion test. Let $(\hat{r}_{in}, \hat{r}_{out})$ denote the coarse radii estimated at center c . The center is marked as suspicious if one of the following conditions holds: (i) a large fraction of rays intersects the image boundary before reaching \hat{r}_{out} ; (ii) the implied scale ratio $\hat{r}_{out}/\hat{r}_{in}$ falls outside a plausible range; or (iii) a low-resolution preview segmentation yields an abnormally small annulus area.

All suspicion criteria are implemented using dimensionless thresholds normalized by the coarse annulus scale, allowing the same configuration to be reused across datasets without dataset-specific retuning. These thresholds are fixed across datasets in all experiments reported in this paper.

Segmentation-aware rescue. Given a small set of alternative center candidates $\{c_i\}$, obtained from different geometric experts and/or local perturbations around c , we perform a lightweight preview outer-boundary search at reduced angular resolution and select the center that minimizes the preview objective:

$$c^* = \arg \min_{c_i} \mathcal{E}_{\text{prev}}(c_i), \quad (2)$$

where $\mathcal{E}_{\text{prev}}(c_i)$ denotes the minimal preview energy computed in a narrow band around the coarse outer-radius estimate of candidate c_i . This transforms center selection into an auxiliary optimization step driven directly by downstream annulus plausibility.

Coarse radii. For a selected center, a coarse pupil radius \hat{r}_{in} is estimated from strong inward edges around the dark pupil region. The coarse outer radius is then searched within a dimensionless range:

$$\hat{r}_{out} \in [\beta_{\min} \hat{r}_{in}, \beta_{\max} \hat{r}_{in}], \quad (3)$$

and is obtained from multi-scale gradient evidence (MSGE) aggregated by sector-wise radial confidence voting (SRCV), which stabilizes coarse outer-radius estimation under partial occlusion and weak limbic contrast. This stage provides a plausible coarse annulus prior before the full guarded optimization.

3.4. Segmentation-First Ray-Wise Dynamic Programming with Fail-Safe Outer Control

After selecting c^* and estimating coarse radii, EDIL-SegRayDP recovers *angularly varying* inner and outer boundaries, denoted by $r_{in}(\theta)$ and $r_{out}(\theta)$, on a discrete angular grid $\{\theta_j\}_{j=1}^{N_\theta}$. For each angle θ_j , radial evidence is sampled in a band centered around the corresponding coarse radius, producing a cost table $C_{out}(r, \theta_j)$ for the outer boundary and, analogously, $C_{in}(r, \theta_j)$ for the inner boundary. In practice, each

Algorithm 1: Segmentation-Aware Center Rescue

Input: Image I , center candidates $\{c_i\}$

Output: Selected center c^* and coarse radii $(\hat{r}_{in}, \hat{r}_{out})$

- 1 Compute an initial center prior c from multi-cue geometric fusion;
 - 2 Estimate coarse radii $(\hat{r}_{in}, \hat{r}_{out})$ at c ;
 - 3 **if** c is suspicious **then**
 - 4 **foreach** candidate center c_i **do**
 - 5 Estimate coarse radii at c_i ;
 - 6 Build a low-resolution preview outer band around $\hat{r}_{out}(c_i)$;
 - 7 Compute the preview annulus score $\mathcal{E}_{\text{prev}}(c_i)$;
 - 8 Select $c^* \leftarrow \arg \min_{c_i} \mathcal{E}_{\text{prev}}(c_i)$
 - 9 **else**
 - 10 $c^* \leftarrow c$;
 - 11 Re-estimate $(\hat{r}_{in}, \hat{r}_{out})$ at c^* ;
 - 12 **return** $c^*, (\hat{r}_{in}, \hat{r}_{out})$;
-

cost table combines normalized boundary evidence with geometric admissibility constraints and optional occlusion-aware penalties.

Outer-first optimization. The outer boundary is recovered first by minimizing a global energy over the discrete path $\{r_j\}$:

$$\min_{\{r_j\}} \sum_{j=1}^{N_\theta} \left(C_{out}(r_j, \theta_j) + \lambda_s |r_j - r_{j-1}| + \lambda_c |(r_j - r_{j-1}) - (r_{j-1} - r_{j-2})| \right), \quad (4)$$

subject to circular wrap-around consistency, i.e., $\theta_{N_\theta+1} = \theta_1$. This optimization is implemented by circular dynamic programming on the polar lattice.

Fail-safe outer-boundary control. When limbic evidence is weak or partially corrupted, a plain DP solver may prefer a smooth but anatomically implausible small-radius path. To reduce this failure mode, the proposed method introduces two safeguards.

(i) *Adaptive band widening.* If the coarse outer annulus is implausibly small relative to the inner radius or image scale, the method enlarges the outer search band once before constructing the full outer cost grid:

$$\Delta_{out} \leftarrow \kappa \Delta_{out}, \quad \kappa > 1, \quad (5)$$

thereby providing a more realistic search region for the subsequent outer-boundary optimization.

(ii) *Outer-floor guard.* A conservative lower bound is imposed on the admissible outer radius:

$$r_{\text{out}}(\theta_j) \geq r_{\text{floor}}(\theta_j) \triangleq \max(\rho \hat{r}_{\text{out}}, \hat{r}_{\text{in}} + m), \quad (6)$$

where $\rho \in (0, 1)$ is a dimensionless retention ratio applied to the coarse outer-radius estimate (set to $\rho = 0.85$ in all experiments), and m is a geometric margin that enforces a minimum radial separation between the inner and outer boundaries at every angle. In the current implementation, m is conservatively fixed to 8 pixels across all datasets and is not tuned per dataset. Since the evaluated datasets exhibit relatively comparable iris-image scales, this simple fixed margin was found sufficient to stabilize annulus geometry and reduce degenerate or excessively thin annuli under weak limbal evidence, while preserving anatomically plausible iris sectors for downstream normalization and validity-aware mask construction.

Inner boundary with geometric coupling. Once $r_{\text{out}}(\theta)$ has been recovered, the inner boundary is estimated using a similar DP objective on C_{in} , while enforcing the annulus constraint

$$r_{\text{in}}(\theta_j) \leq r_{\text{out}}(\theta_j) - m, \quad (7)$$

which preserves a valid annular geometry by enforcing a minimum separation margin between the inner and outer boundaries, thereby stabilizing pupil-boundary recovery under weak or partially corrupted boundary evidence.

The plausibility test for guarded recovery reuses the same scale-normalized geometric criteria introduced in the suspicion test, now applied to the selected coarse annulus prior.

The guarded recovery first checks whether the coarse annulus is plausible, then solves the outer and inner boundaries in an outer-first manner. Specifically, the outer-boundary path is recovered as

$$\{r_{\text{out}}^*(\theta_j)\}_{j=1}^{N_\theta} = \arg \min_{\{r_j\} \subset \mathcal{R}_{\text{out}}} \mathcal{E}_{\text{out}}(\{r_j\}), \quad (8)$$

where \mathcal{R}_{out} denotes the admissible set of outer-boundary paths satisfying

$$r_j \geq r_{\text{floor}}(\theta_j), \quad \forall j,$$

and \mathcal{E}_{out} is the dynamic-programming energy defined in Eq. (4).

After recovering the outer boundary, the inner-boundary path is obtained as

$$\{r_{\text{in}}^*(\theta_j)\}_{j=1}^{N_\theta} = \arg \min_{\{r_j\} \subset \mathcal{R}_{\text{in}}} \mathcal{E}_{\text{in}}(\{r_j\}), \quad (9)$$

where \mathcal{R}_{in} denotes the admissible inner-boundary paths satisfying

$$r_j \leq r_{\text{out}}^*(\theta_j) - m, \quad \forall j.$$

Algorithm 2: Guarded Outer-First Boundary Recovery

Input: $I, c^*, \hat{r}_{\text{in}}, \hat{r}_{\text{out}}$
Output: $r_{\text{in}}(\theta), r_{\text{out}}(\theta), M_{\text{valid}}$

- 1 **if** coarse annulus is implausible **then**
- 2 $\Delta_{\text{out}} \leftarrow \kappa \Delta_{\text{out}};$
- 3 mode \leftarrow guarded;
- 4 **else**
- 5 mode \leftarrow standard;
- 6 $r_{\text{floor}}(\theta) \leftarrow \max(\rho \hat{r}_{\text{out}}, \hat{r}_{\text{in}} + m);$
- 7 Build outer cost grid $C_{\text{out}}(r, \theta)$ for $r \geq r_{\text{floor}}(\theta);$
- 8 Recover $r_{\text{out}}(\theta)$ by circular DP;
- 9 Build inner cost grid $C_{\text{in}}(r, \theta)$ under
 $r_{\text{in}}(\theta) \leq r_{\text{out}}(\theta) - m;$
- 10 Recover $r_{\text{in}}(\theta)$ by constrained DP;
- 11 $M_{\text{ann}} \leftarrow M_{\text{iris}} \wedge \neg M_{\text{pupil}};$
- 12 $M_{\text{valid}} \leftarrow M_{\text{ann}} \wedge \neg M_{\text{occ}};$
- 13 **return** $r_{\text{in}}(\theta), r_{\text{out}}(\theta), M_{\text{valid}};$

For each image, the dominant optimization stage has time complexity $\mathcal{O}(N_\theta N_r K)$, where N_θ is the number of angular samples, N_r is the radial search range, and K is the transition bandwidth. In practice, this complexity remains manageable because the search bands are restricted by coarse annulus priors and fail-safe geometric constraints.

Proposition 1 (Global optimality of the DP solution). Let $C(r, \theta_j)$ be any per-angle cost table defined on a finite radius set \mathcal{R} . For the first-order smoothness energy in Eq. (4) with a bounded neighborhood transition $|r_j - r_{j-1}| \leq K$, the dynamic programming recursion returns a globally optimal discrete path $\{r_j^*\}_{j=1}^{N_\theta}$ minimizing the total energy under the same constraints.

Proof. The objective is an additive chain-structured energy over θ_j with local pairwise terms depending only on adjacent states (r_{j-1}, r_j) , with optional second-order regularization handled by state augmentation. Therefore, the Bellman optimality principle applies: the optimal prefix ending at radius r_j must contain an optimal prefix ending at a feasible predecessor radius r_{j-1} . By enumerating all feasible predecessors within the bounded transition window, the dynamic programming recursion yields the global optimum by induction over the angular index. \square

3.5. Geometry-Normalized Occlusion Modeling and Validity-Aware Mask Construction

EDIL-SegRayDP explicitly models major occlusion sources and constructs a validity-aware annulus mask. Let M_{iris} and M_{pupil} denote the rasterized masks induced by the recovered outer and inner boundaries

$r_{\text{out}}(\theta)$ and $r_{\text{in}}(\theta)$, respectively. The annulus mask is defined as

$$M_{\text{ann}} = M_{\text{iris}} \wedge \neg M_{\text{pupil}}. \quad (10)$$

Occlusion components. The occlusion mask M_{occ} is formed as the union of three components: (i) specular highlights, detected as percentile-based bright outliers and optionally inpainted for evidence stabilization; (ii) eyelid occlusion, modeled through geometry-normalized eyelid curves and validity bands; and (iii) eyelash occlusion, detected as dark outliers restricted to the annulus interior. All thresholds are expressed in dimensionless form as ratios of \hat{r}_{out} or \hat{r}_{in} , which improves consistency across scales and datasets.

Validity-aware final mask. The final valid iris mask used for evaluation is

$$M_{\text{valid}} = M_{\text{ann}} \wedge \neg M_{\text{occ}}. \quad (11)$$

A small number of deterministic clean-up steps, such as connected-component filtering or mild morphology, may be applied to remove tiny fragments while preserving annulus topology.

The following lemma (Lemma 2) formalizes the scale-invariance property of the proposed pipeline. Here, $s > 0$ denotes an arbitrary uniform isotropic image-rescaling factor applied to the input image; it is *not* a tunable hyperparameter of the method, but a theoretical analysis variable introduced to show that the normalized geometric decisions remain invariant under resolution changes.

In practice, EDIL-SegRayDP operates on the native image resolution (i.e., effectively $s = 1$) without explicit image rescaling. The practical implication of the lemma is that the same fixed configuration can be reused across datasets acquired at different image resolutions without dataset-specific parameter readjustment, since the principal geometric thresholds are expressed as dimensionless ratios relative to the estimated iris radii.

Lemma 2 (Scale-normalized invariance). Assume that the input image is rescaled by a factor $s > 0$, and that all geometric quantities, including center coordinates and radii, scale accordingly. If the decision thresholds and hyperparameters in EDIL-SegRayDP are expressed as dimensionless ratios of r_{in} and r_{out} , then the accepted hypotheses and normalized boundary decisions remain invariant to the scale factor s , up to the same geometric rescaling.

Proof. Each threshold in the pipeline is applied to quantities of the form $\|c_i - c_{\text{ref}}\|/r_{\text{in}}$, $r_{\text{out}}/r_{\text{in}}$, $d_{\text{max}}/r_{\text{out}}$, or $\varepsilon r_{\text{out}}/r_{\text{out}}$. Under a uniform rescaling by s , both numerator and denominator scale by the same factor, so each dimensionless ratio remains unchanged. Consequently, all accept/reject decisions are preserved in normalized coordinates, and the recovered boundaries scale by the same factor. \square

4. Experimental Setup

4.1. Datasets and Comparison Basis

We evaluate EDIL-SegRayDP under a core comparative benchmark on two near-infrared iris datasets for which our current benchmark pipeline and mask handling are fully verified: IITD and CASIA-IrisV4-Interval. We further report expanded full-benchmark results on CASIA-IrisV4-Lamp and CASIA-IrisV4-Thousand under the same fixed-configuration protocol. Since the proposed method is entirely training-free, all available images with valid annotations are used directly for evaluation, without any train/validation/test split.

IITD is used as a challenging benchmark due to frequent weak outer-boundary evidence, irregular annulus geometry, and cases that are prone to catastrophic failures in localization-first classical pipelines. In our current full-set evaluation, the benchmark contains 2240 images.

CASIA-IrisV4-Interval represents a more cooperative near-infrared setting, but still includes realistic artifacts such as specular reflections, mild blur, and illumination variation. In our benchmark, the full-set evaluation contains 2639 images.

CASIA-IrisV4-Lamp is a more difficult near-infrared subset characterized by stronger illumination variation, weaker and less stable outer-boundary evidence, and more frequent non-ideal artifacts. In our benchmark, the full-set evaluation contains 16212 images.

CASIA-IrisV4-Thousand provides a larger-scale near-infrared benchmark with broader subject diversity and realistic capture variability. In our benchmark, the full-set evaluation contains 19977 images.

To position our method against prior classical work, we adopt as the main reference a recent training-free iris-segmentation paper based on active-contour models and occlusion removal, validated on IITD, CASIA-IrisV3, and CASIA-IrisV4. That work reports pupil/iris Dice values of 0.971/0.905 on IITD, 0.950/0.898 on CASIA-IrisV3, and 0.960/0.900 on CASIA-IrisV4. In the present paper, we compare directly on the overlapping benchmark family, with strict dataset overlap on IITD and an indicative comparison between our verified CASIA-IrisV4-Interval benchmark and the broader CASIA-IrisV4 results reported in the reference study.

4.2. Evaluation Protocol

EDIL-SegRayDP is evaluated under a *full-set, fixed-configuration* protocol. No learning stage, parameter fitting, or dataset-specific fine-tuning is performed. The same dimensionless parameter configuration is used across IITD, CASIA-IrisV4-Interval, CASIA-IrisV4-Lamp, and CASIA-IrisV4-Thousand, including the center-rescue trigger, annulus plausibility checks,

adaptive single-pass outer-band widening, and outer-floor guarding.

For each image, the pipeline predicts an inner boundary, an outer boundary, a pupil mask, and a valid iris mask. The valid iris mask is defined as the annulus after removing detected specular highlights, eyelid regions, and eyelash regions. This evaluation design is consistent with the segmentation-first philosophy of the proposed method: center hypotheses are not judged in isolation, but through the plausibility of the final annulus they support.

Following our benchmark scripts, we additionally report the number of *failed* cases, defined as images whose iris-mask IoU is below 0.70. This failure count is useful for classical methods because average scores alone can hide catastrophic collapses on difficult images. The fixed-configuration protocol is intentionally adopted to evaluate portability rather than dataset-specific optimization.

No dataset-specific threshold adjustment, post hoc parameter retuning, or subset-dependent search-range modification was introduced after inspecting the results on individual datasets.

4.3. Implementation Details

The proposed method is implemented as a fully deterministic CPU-only pipeline. All images are processed at their native resolutions without GPU acceleration or deep-learning inference. In the default implementation, the polar optimization uses $N_\theta = 360$ angular rays.

The inner and outer search bands are initialized from coarse radius estimates and subsequently refined through circular dynamic programming with smoothness constraints. When the coarse outer annulus is judged implausibly small, the method performs a single adaptive widening of the outer search band before constructing the full dynamic-programming cost grid. In addition, the admissible outer search range is constrained by a conservative geometric floor to reduce pathological inward collapse under weak limbal evidence. All related controls remain fixed across datasets in order to evaluate cross-dataset portability rather than dataset-specific tuning.

The average processing time is approximately 498 ms/image on IITD, 701 ms/image on CASIA-IrisV4-Interval (320×280), 1682 ms/image on CASIA-IrisV4-Lamp, and 1711 ms/image on CASIA-IrisV4-Thousand (640×480), measured on an Intel Core i7-13700 CPU in single-thread mode with batch size 1. The increased runtime on CASIA-IrisV4-Lamp and CASIA-IrisV4-Thousand is mainly attributed to the larger image resolution and the wider admissible radial search regions required under more challenging acquisition

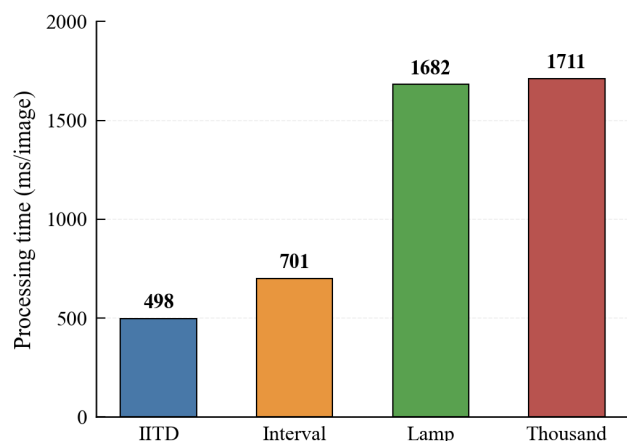


Figure 2. Average processing time of EDIL-SegRayDP under CPU-only single-thread inference across the evaluated datasets

conditions. The runtime characteristics across datasets are summarized in Fig. 2.

Although the current implementation is not optimized for real-time deployment, the framework remains fully deterministic, explainable, and deployable without GPU acceleration or a training stage.

4.4. Evaluation Metrics

We report standard pixel-wise segmentation metrics, including Dice coefficient, Intersection-over-Union (IoU), and pixel accuracy (Acc). Let Y and \hat{Y} denote the ground-truth and predicted binary masks, respectively. The metrics are defined as

$$\text{Dice} = \frac{2|Y \cap \hat{Y}|}{|Y| + |\hat{Y}|}, \quad \text{IoU} = \frac{|Y \cap \hat{Y}|}{|Y \cup \hat{Y}|}, \quad (12)$$

and

$$\text{Acc} = \frac{|Y \cap \hat{Y}| + |\bar{Y} \cap \bar{\hat{Y}}|}{|Y| + |\hat{Y}|}. \quad (13)$$

In addition to iris-mask metrics, we also report **pupil Dice** to verify that the inner boundary is recovered accurately rather than only serving as an auxiliary component of outer-annulus estimation. Pixel accuracy is reported for completeness, in line with the comparison style used by the reference paper.

5. Results and Discussion

5.1. Main Quantitative Comparison

Table 1 reports the full-set results of EDIL-SegRayDP on IITD and CASIA-IrisV4-Interval. We compare our method against the adopted training-free reference, which serves as the main classical baseline in this study. On IITD, EDIL-SegRayDP achieves an iris Dice of 0.9106 and a pupil Dice of 0.9763, improving over the

Table 1. Quantitative comparison with prior classical techniques and the adopted training-free reference. The last row of each dataset block reports EDIL-SegRayDP

Database	Method	Images	Dice \uparrow		Accuracy \uparrow	
			Pupil	Iris	Pupil	Iris
IITD	Kumar et al. [23]	2092	0.9308	–	0.9306	–
	Sardar et al. [24]	2240	–	0.9850	–	–
	Gautam et al. [10]	2240	0.9431	–	0.9922	–
	Umer et al. [25]	2240	–	–	–	0.9848
	Jan et al. [11]	1120	–	–	–	0.9860
	Waisy et al. [26]	1120	–	–	–	0.9987
	Grissel et al. [12]	1120	0.9710	0.9050	0.9960	0.9470
	Our Proposal	2240	0.9763	0.9106	0.9958	0.9481
	CASIA-IrisV4-Interval	Gangwar et al. [27]	2639	–	–	–
Wu et al. [28]		2639	–	0.9707	–	0.9836
Sardar et al. [24]		2639	–	0.9870	–	–
Daugman [7]		2639	–	–	0.7840	0.4217
Masek [8]		1568	–	–	–	0.8300
Petrovska et al. [29]		2639	–	–	–	0.9340
Grissel et al. [12]		2639	0.9695	0.9000	0.9960	0.9470
Our Proposal		2639	0.9755	0.9377	0.9964	0.9693

The adopted reference results for Grissel et al. are reported on the broader CASIA-IrisV4 setting, whereas our current verified benchmark is conducted on CASIA-IrisV4-Interval.

reference values of 0.9050 and 0.9710, respectively. On CASIA-IrisV4-Interval, the proposed method achieves an iris Dice of 0.9377 and a pupil Dice of 0.9755, compared with the reference values of 0.9000 and 0.9695 reported for CASIA-IrisV4.

These results indicate that the proposed method is competitive with the adopted training-free reference on both benchmark families, with the strongest protocol-matched evidence obtained on IITD and a favorable indicative comparison on CASIA-IrisV4-Interval relative to the broader CASIA-IrisV4 results reported in the reference study. The improvement is moderate on IITD but consistent across both iris and pupil Dice, indicating that the gain is not limited to only one component of the segmentation task. On CASIA-IrisV4-Interval, the margin is more pronounced for

Table 2. Full benchmark summary of EDIL-SegRayDP across all evaluated datasets under the fixed-configuration protocol

Dataset	Images	Iris Dice \uparrow	Pupil Dice \uparrow	Iris Acc \uparrow	Fail (%) \downarrow
IITD	2240	0.9106	0.9763	0.9481	1.50
CASIA-IrisV4-Interval	2639	0.9377	0.9755	0.9693	2.50
CASIA-IrisV4-Lamp	16212	0.8950	0.8750	0.9870	8.00
CASIA-IrisV4-Thousand	19977	0.9180	0.8750	0.9910	5.50

iris segmentation, where EDIL-SegRayDP improves the reported Dice from 0.9000 to 0.9377. The pupil Dice on CASIA-IrisV4-Interval is also higher than the adopted reference value, suggesting that the proposed boundary recovery remains stable for both the inner and outer contours.

To complement the core comparison in Table 1, Table 2 summarizes the fixed-configuration performance of EDIL-SegRayDP across all four evaluated datasets. On CASIA-IrisV4-Lamp, the proposed method achieves an iris Dice of 0.8950 and a pupil Dice of 0.8750, confirming that Lamp is the most challenging subset among the evaluated benchmarks. On CASIA-IrisV4-Thousand, the method attains an iris Dice of 0.9180 and a pupil Dice of 0.8750, indicating stronger outer-annulus recovery than on Lamp while still showing reduced inner-boundary accuracy compared with IITD and CASIA-IrisV4-Interval.

Taken together, these expanded results suggest that the proposed fixed-configuration design is portable beyond the core comparison datasets and degrades in a structured manner rather than collapsing unpredictably under more challenging acquisition conditions.

Figure 3 provides representative qualitative examples on IITD and CASIA-IrisV4-Interval. The visual results are consistent with the quantitative findings: the proposed method recovers stable inner and outer boundaries across both cooperative and more challenging cases, while preserving anatomically plausible annulus geometry under partial occlusion and specular corruption.

5.2. Ablation of the Fail-Safe Design

To test whether the robustness gain comes from the proposed segmentation-first fail-safe design rather than from any single isolated heuristic, Table 3 reports a staged ablation study of the proposed fail-safe design. On IITD, each component contributes progressively to improved robustness, with the most pronounced gains obtained from adaptive widening and the outer-floor guard. In particular, the fail count decreases from 68

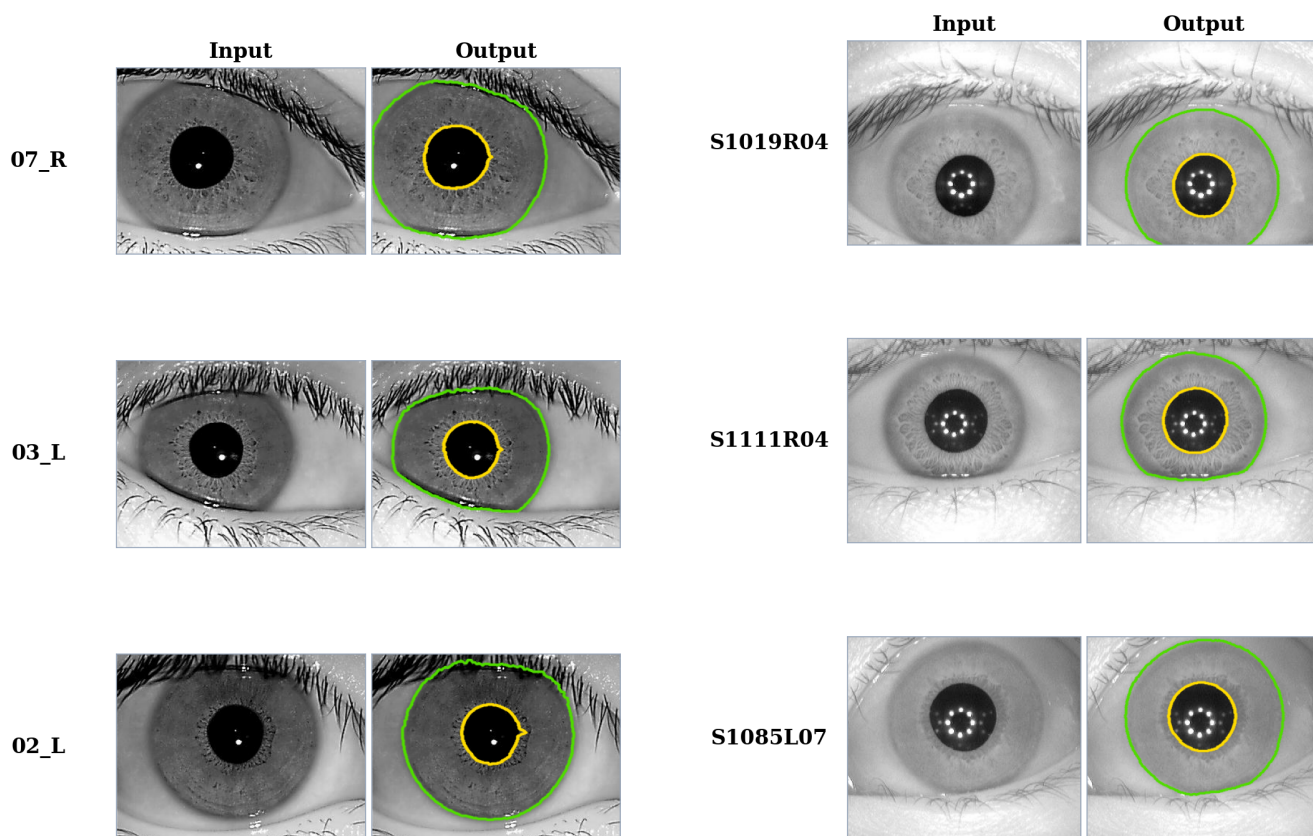


Figure 3. Qualitative segmentation examples of EDIL-SegRayDP on two datasets. (a) IITD examples. (b) CASIA-IrisV4-Interval examples. The yellow contour denotes the recovered inner boundary, and the green contour denotes the recovered outer boundary

in the base variant to 34 in the final guarded variant, while iris Dice improves from 0.8449 to 0.9106 and iris accuracy increases from 0.9243 to 0.9481. This trend supports our claim that the main failure mode on IITD is catastrophic outer-boundary collapse under weak or corrupted evidence.

On CASIA-IrisV4-Interval, the gains are smaller but still consistent. Center rescue and adaptive widening improve both Dice and failure rate, whereas the outer-floor guard produces almost no additional change beyond the widened configuration. This behavior is reasonable because CASIA-IrisV4-Interval is a more cooperative dataset, where anatomically implausible inward collapse occurs less frequently than on IITD. Overall, the ablation results confirm that the proposed robustness gains arise from the interaction of the fail-safe components rather than from a single isolated module.

5.3. Dataset-Wise Analysis

A broader trend across the evaluated datasets is that EDIL-SegRayDP performs best on the more cooperative subsets and degrades gracefully on more challenging ones. In particular, CASIA-IrisV4-Interval

Table 3. Ablation study of the proposed fail-safe design

Variant	Dataset	Iris Dice	Pupil Dice	Iris Acc	Pupil Acc	Fail
Base model	IITD	0.8449	0.9633	0.9243	0.9950	68
+ Center rescue	IITD	0.8552	0.9758	0.9277	0.9958	64
+ Adaptive widening	IITD	0.8840	0.9759	0.9397	0.9958	46
+ Outer-floor guard	IITD	0.9106	0.9763	0.9481	0.9958	34
Base model	CASIA-IrisV4-Interval	0.9298	0.9644	0.9664	0.9961	109
+ Center rescue	CASIA-IrisV4-Interval	0.9346	0.9654	0.9682	0.9964	89
+ Adaptive widening	CASIA-IrisV4-Interval	0.9378	0.9655	0.9693	0.9964	75
+ Outer-floor guard	CASIA-IrisV4-Interval	0.9377	0.9655	0.9693	0.9964	66

Table 4. Failure statistics under the full-set evaluation protocol. A failure is defined as an iris-mask IoU below 0.70

Dataset	No. images	Failures	Rate (%)
IITD	2240	34	1.50
CASIA-IrisV4-Interval	2639	66	2.50
CASIA-IrisV4-Lamp	16212	1297	8.00
CASIA-IrisV4-Thousand	19977	1099	5.50

yields the strongest overall performance among the CASIA subsets, whereas CASIA-IrisV4-Lamp is the most difficult, with the lowest iris Dice and the highest failure rate. CASIA-IrisV4-Thousand occupies an intermediate position: its iris segmentation remains relatively strong, but the pupil Dice is still noticeably lower than on IITD and Interval. This pattern suggests that the proposed segmentation-first optimization is robust for outer-annulus recovery at scale, but that inner-boundary estimation remains more sensitive to difficult acquisition conditions.

Even so, IITD remains an important benchmark because it better exposes the catastrophic-failure modes of localization-first methods under a protocol-matched comparison with prior work. In such cases, a biased center or an underestimated coarse outer radius may force all subsequent radial searches toward an implausible annulus. The fact that EDIL-SegRayDP still improves upon the adopted reference Dice values on IITD therefore supports the practical value of the proposed fail-safe design, rather than merely reflecting easier image conditions.

5.4. Failure Analysis

To complement the average metrics, Table 4 reports the number of failed cases, defined as images whose iris-mask IoU is below 0.70. This metric is especially informative for deterministic geometric methods because average Dice alone may hide a small subset of catastrophic failures.

Fig. 4 presents representative failure cases from the evaluated datasets under challenging non-ideal imaging conditions. The examples illustrate several typical failure modes, including severe eyelid or eyelash occlusion, extremely weak limbic boundaries, strong illumination artifacts, and degraded annulus geometry. In these situations, the recovered annulus may partially collapse inward or fail to align with the true iris contour, leading to substantial IoU degradation despite the deterministic geometric constraints.

Although EDIL-SegRayDP achieves strong average results, the nonzero failure counts indicate that difficult

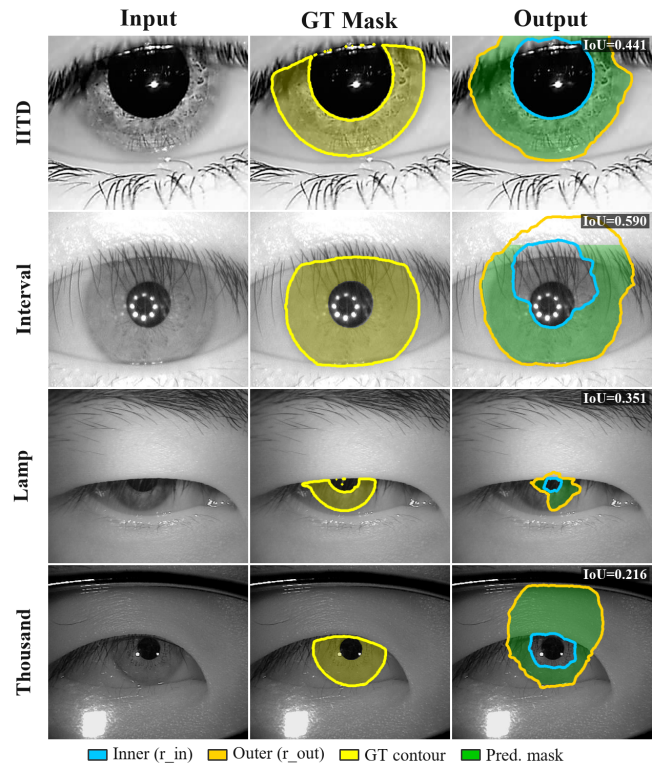


Figure 4. Representative failure cases from IITD, CASIA-IrisV4-Interval, CASIA-IrisV4-Lamp, and CASIA-IrisV4-Thousand under challenging non-ideal imaging conditions. Each row shows the input image, ground-truth mask, and predicted segmentation result. Typical failure modes include severe occlusion, weak outer-boundary evidence, illumination corruption, and annulus collapse under degraded geometric conditions

cases still remain. Typical failures arise under severe eyelid/eyelash occlusion, very weak outer-boundary evidence, or annulus geometry that is strongly degraded by illumination and off-axis effects. The expanded benchmark further shows that the failure rate rises substantially on CASIA-IrisV4-Lamp and remains non-negligible on CASIA-IrisV4-Thousand, confirming that weak boundary evidence and non-ideal annulus geometry remain the main stressors for deterministic geometric segmentation. Nevertheless, the failure rates on IITD and CASIA-IrisV4-Interval remain relatively low over the full-set protocol, which suggests that the proposed safeguards substantially reduce catastrophic collapse compared with a plain classical pipeline.

5.5. Effect of the Segmentation-First Fail-Safe Design

The quantitative gains are consistent with the actual design of EDIL-SegRayDP. First, the *segmentation-aware center rescue* reduces errors caused by implausible initial center hypotheses. Instead of accepting the first center estimate unconditionally, the method evaluates

whether that center supports a plausible annulus and switches to a better candidate when necessary. This is particularly important in images where the pupil is partially corrupted or where local dark regions may mislead a conventional center detector. Second, the *adaptive single-pass outer-band widening* improves robustness when the coarse outer estimate is biased inward. In a localization-first pipeline, such underestimation can force the subsequent optimization to search in the wrong radial range, often producing a collapsed outer boundary. By widening the search band once before full dynamic programming, the proposed method gives the outer-boundary solver a more plausible search region without introducing an expensive iterative loop. Third, the *outer-floor guard* prevents anatomically implausible inward collapse when limbic evidence is weak but locally smooth. This control is especially important on IITD, where weak or corrupted outer-boundary cues occur more frequently. Taken together, these mechanisms explain why the proposed method improves the adopted training-free reference while remaining fully deterministic, interpretable, and CPU-deployable.

5.6. Discussion and Limitations

The present results support the main claim of this work: a carefully designed training-free pipeline can remain competitive with recent non-learning methods when center priors, annulus plausibility, and fail-safe outer-boundary control are coupled in a segmentation-first manner. In particular, the results suggest that the performance gain does not come from a single isolated module, but from the robust interaction between center selection, coarse annulus estimation, guarded dynamic programming, and explicit occlusion handling.

At the same time, several limitations should be acknowledged. First, the protocol-matched comparison with prior work currently focuses on IITD and CASIA-IrisV4-Interval, where overlap with the adopted training-free reference is the clearest. We additionally report fixed-configuration benchmark results on CASIA-IrisV4-Lamp and CASIA-IrisV4-Thousand, but these expanded evaluations are not yet accompanied by fully protocol-matched comparisons against prior work. Second, while the average performance is strong, the remaining failure cases show that severe occlusion and extremely weak limbic contrast are still challenging for a deterministic geometry-based method. In addition, the current CPU-only implementation requires approximately 701 ms/image on CASIA-IrisV4-Interval and 1700 ms/image on CASIA-IrisV4-Lamp and CASIA-IrisV4-Thousand, indicating that the present system prioritizes deterministic robustness and interpretability over real-time execution.

Even with these limitations, the present study demonstrates that a carefully guarded deterministic formulation can remain competitive on difficult iris benchmarks without training data, retraining, or opaque learned priors, which is precisely the regime targeted by this work.

Future work will extend the same fixed-configuration evaluation to additional datasets and further analyze runtime and boundary-level localization quality.

6. Conclusion

We presented EDIL-SegRayDP, a training-free and explainable iris segmentation framework built around a segmentation-first optimization strategy. The method combines segmentation-aware center rescue, coarse annulus plausibility checking, adaptive single-pass outer-band widening, and an outer-floor guard to improve robustness under weak or corrupted boundary evidence. Full-set experiments on IITD and CASIA-IrisV4-Interval, together with expanded benchmark results on CASIA-IrisV4-Lamp and CASIA-IrisV4-Thousand, show that the proposed method achieves competitive classical performance. These results suggest that the main value of EDIL-SegRayDP lies not only in competitive classical accuracy, but in showing that segmentation-first, fail-safe geometric optimization is a viable alternative to localization-first and learning-based pipelines when interpretability, deterministic behavior, and training-free deployment are primary requirements. Future work will extend the same fixed-configuration evaluation to additional datasets and further analyze runtime and boundary-level localization quality.

Acknowledgements. This research is supported by the Posts and Telecommunications Institute of Technology (PTIT), Vietnam.

Data Availability

The data that support the findings of this study are available from the official data providers, subject to their terms and conditions. The CASIA dataset is available at <https://hycasia.github.io/dataset/casia-irisv4/>. The IITD dataset is available at <https://shorturl.at/dy4Az>.

References

- [1] Daugman J. How iris recognition works. *IEEE Transactions on Circuits and Systems for Video Technology*. 2004;14(1):21-30.
- [2] Wildes RP. Iris recognition: an emerging biometric technology. *Proceedings of the IEEE*. 1997;85(9):1348-63.
- [3] Tan T, He Z, Sun Z. Efficient and robust segmentation of noisy iris images for non-cooperative iris recognition. *Image and Vision Computing*. 2010;28(2):223-30.

- [4] Ross A, Shah S. Segmenting Non-Ideal Irises Using Geodesic Active Contours. In: Biometrics Symposium: Special Session on Research at the Biometric Consortium Conference; 2006. p. 1-6.
- [5] Wang C, Muhammad J, Wang Y, He Z, Sun Z. Towards Complete and Accurate Iris Segmentation Using Deep Multi-Task Attention Network for Non-Cooperative Iris Recognition. *IEEE Transactions on Information Forensics and Security*. 2020;15:2944-59.
- [6] Sharma G, Jaswal G, Nigam A, Ramachandra R. FISNET: A Learnable Fusion-Based Iris Segmentation Network Improving Robustness Across NIR and VIS Modalities. *IEEE Access*. 2025;13:101472-90.
- [7] Daugman JG. High confidence visual recognition of persons by a test of statistical independence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 1993;15(11):1148-61.
- [8] Masek L. Recognition of Human Iris Patterns for Biometric Identification [Master's Thesis]. The University of Western Australia; 2003.
- [9] Labati RD, Sansone C, Scotti F. Non-ideal iris segmentation using polar spline RANSAC and low-level energy minimization. *Computer Vision and Image Understanding*. 2020;201:102787.
- [10] Gautam G, Mukhopadhyay S. An adaptive localization of pupil degraded by eyelash occlusion and poor contrast. *Multimedia Tools and Applications*. 2019;78(6):6655-77.
- [11] Jan F, Min-Allah N, Agha S, Usman I, Khan I. A robust iris localization scheme for the iris recognition. *Multimedia Tools and Applications*. 2021;80(3):4579-605.
- [12] Mathias GP, Gagan JH, Mallya BV, Kumar JRH. A unified approach for automated segmentation of pupil and iris in on-axis images. *Computer Methods and Programs in Biomedicine Update*. 2022;2:100084.
- [13] Sumi MR, Das P, Hossain A, Dey S, Schuckers S. A Comprehensive Evaluation of Iris Segmentation on Benchmarking Datasets. *Sensors*. 2024;24(21):7079.
- [14] Xia C, Yan Y, Zhang R, Liu Y, Wang Q. A Light Spatial-Frequency Network for Robust Iris Segmentation and Localization. *Applied Soft Computing*. 2025;175:113009.
- [15] Ronneberger O, Fischer P, Brox T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In: *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*; 2015. p. 234-41.
- [16] Lozej J, Meden B, Štruc V. End-to-End Iris Segmentation Using U-Net. In: *IEEE International Work Conference on Bioinspired Intelligence (IWObI)*; 2018. p. 1-8.
- [17] Jha RR, Jaswal G, Gupta D, Saini S, Nigam A. PixISegNet: Pixel-level iris segmentation network using convolutional encoder-decoder with stacked hourglass bottleneck. *IET Biometrics*. 2020;9(1):11-24.
- [18] Arsalan M, Kim DS, Lee MB, Owais M, Park KR. IrisDenseNet: Robust Iris Segmentation Using Densely Connected Fully Convolutional Networks in the Images by Visible Light and Near-Infrared Light Camera Sensors. *Sensors*. 2018;18(5):1501.
- [19] Liu Z, Lin Y, Cao Y, Hu H, Wei Y, Zhang Z, et al. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*; 2021. p. 10012-22.
- [20] Cao H, Wang Y, Chen J, Jiang D, Zhang X, Tian Q, et al. Swin-Unet: Unet-like Pure Transformer for Medical Image Segmentation. In: *Computer Vision – ECCV 2022 Workshops*; 2022. p. 205-18.
- [21] Kirillov A, Mintun E, Ravi N, Mao H, Rolland C, Gustafson L, et al. Segment Anything. *arXiv preprint arXiv:230402643*. 2023. Available from: <https://arxiv.org/abs/2304.02643>.
- [22] Farmanifard P, Ross A. Iris-SAM: Iris Segmentation Using a Foundation Model. In: *International Conference on Pattern Recognition and Artificial Intelligence*. Springer Nature Singapore; 2024. p. 394-409.
- [23] Kumar JRH, Teotia K. Automatic Pupil Segmentation Based On Circular Active Discs. In: *TENCON 2019 - 2019 IEEE Region 10 Conference (TENCON)*; 2019. p. 1932-6.
- [24] Sardar M, Banerjee S, Mitra S. Iris Segmentation Using Interactive Deep Learning. *IEEE Access*. 2020;8:219322-30.
- [25] Umer S, Dhara BC, Chanda B. Iris recognition using multiscale morphologic features. *Pattern Recognition Letters*. 2015;65:67-74.
- [26] Al-Waisy AS, Qahwaji R, Ipson S, Al-Fahdawi S, Nagem TAM. A multi-biometric iris recognition system based on a deep learning approach. *Pattern Analysis and Applications*. 2018;21(3):783-802.
- [27] Gangwar AK, Joshi A, Singh A, Alonso-Fernandez F, Bigun J. IrisSeg: A Fast and Robust Iris Segmentation Framework for Non-Ideal Iris Images. In: *2016 International Conference on Biometrics (ICB)*. Halmstad, Sweden; 2016. p. 1-8.
- [28] Wu X, Zhao L. Study on Iris Segmentation Algorithm Based on Dense U-Net. *IEEE Access*. 2019;7:123959-68.
- [29] Petrovska D, Mayoue A. Description and Documentation of the BioSecure Software Library. *BioSecure*; 2007.