
Evidential Deep Learning for Missing Boundary Detection in Topologically Constrained OCT Layer Segmentation

Botond Fazekas, Hrvoje Bogunović

Christian Doppler Laboratory for Artificial Intelligence in Retina
Institute of Artificial Intelligence, Center for Medical Data Science
Medical University of Vienna
Vienna, Austria
botond.fazekas@meduniwien.ac.at

Abstract

Optical coherence tomography (OCT) layer boundary regression methods provide sub-pixel precision and topological guarantees but fundamentally assume that every layer exists across all A-scans. This mathematical constraint fails in severe pathologies such as Geographic Atrophy (GA), where specific retinal layers disappear. We extend the topologically constrained SD-RetinaNet framework to jointly perform boundary regression and explicitly detect missing layers using uncertainty quantification. We introduce a Gaussian Negative Log-Likelihood (NLL) formulation to calibrate aleatoric uncertainty, capturing spatial boundary errors. Concurrently, we employ an Evidential Deep Learning (EDL) module to model epistemic uncertainty directly from the network outputs, allowing the network to detect regions with zero structural evidence for a layer. Our framework addresses the largely overlooked challenge of anatomical absence in boundary regression, combining sub-pixel localization with direct atrophy segmentation.

1 Introduction

Optical Coherence Tomography (OCT) is the gold standard for managing sight-threatening diseases like age-related macular degeneration (AMD) [1, 2]. Extracting fine retinal layer thicknesses is crucial for tracking disease progression. State-of-the-art layer segmentation increasingly utilizes boundary regression [3, 4], which achieves sub-pixel accuracy and implicitly prevents multi-surface topological violations by computing the expected vertical position from a column-wise probability mass function.

However, this formulation presents a fundamental limitation: it assumes every layer is present across the entire scan. Because the column-wise softmax must sum to one, the network falsely predicts boundary coordinates even when structural evidence is entirely absent. Detecting missing layers is highly relevant clinically, most notably in Geographic Atrophy (GA), where the localized loss of the photoreceptors and retinal pigment epithelium (RPE) serves as the primary endpoint for novel therapeutics [5].

Historically, automated atrophy detection follows three paradigms. *Direct lesion segmentation* identifies GA footprints via classification on *en face* projections [6, 7] or cross-sectional B-scans [8], but ignores the 3D layer geometry needed to monitor early structural thinning. *Pixel-wise layer segmentation* can theoretically capture atrophy by omitting the RPE prediction locally, yet lacks anatomical coherence, risking topological violations such as reversed layers or implausible holes.

Conversely, *boundary regression* guarantees strict topology but its continuous-surface constraint prevents it from accurately omitting missing layers.

To address this gap within the boundary regression paradigm, we introduce an uncertainty-aware extension to the topologically constrained SD-RetinaNet. We explicitly quantify two sources of uncertainty directly from the network’s spatial probability maps. First, we model *aleatoric uncertainty* using a Gaussian Negative Log-Likelihood (NLL) loss to calibrate the spatial variance of the boundary predictions, capturing ambiguity from speckle noise, low contrast, and shadowing from overlying hyperreflective areas. Second, we integrate epistemic uncertainty via Evidential Deep Learning (EDL) [9]. By placing a Dirichlet prior over the 1D spatial column, the network estimates the total structural evidence and identifies out-of-distribution scenarios where layers are missing. To our knowledge, this is the first evidential boundary regression framework to detect anatomical absence, bridging high-precision topology with robust atrophy detection.

2 Related Work

Retinal Biomarker Segmentation and Atrophy: Recent deep learning methods for OCT predominantly address layer and fluid lesion segmentation [10, 11, 4]. To overcome topological violations common in pixel-wise U-Nets, He et al. [3] introduced layer boundary regression (LBRM). While SD-RetinaNet [4] successfully integrated lesion constraints into this framework, neither method accounts for layers that vanish due to GA. Detection of GA has mostly been handled via direct regional segmentation [7], disconnected from sub-pixel layer morphology. Adapting deep regression models to correctly omit layers remains an unsolved challenge.

Uncertainty Quantification in Deep Learning: Uncertainty in deep learning is categorized into aleatoric and epistemic uncertainty [12]. While Monte Carlo Dropout (MCD) captures epistemic uncertainty, it requires expensive multiple forward passes. Evidential Deep Learning (EDL) instead formulates learning as an evidence acquisition process [9, 13], quantifying epistemic uncertainty in a single forward pass. By adapting EDL to place Dirichlet priors over spatial column-wise probability maps, we leverage it to act as an anomaly detector for anatomical absence.

3 Methodology

Base Architecture: Our framework builds upon the SD-RetinaNet backbone [14], which employs a U-Net architecture equipped with an EfficientNet [15] encoder. To extract anatomical boundaries, the network utilizes an anatomy module that projects high-dimensional feature maps into 1D column-wise boundary representations over the image height H . Instead of enforcing a strict softmax normalization that forces probability maps to sum to one, we reframe the spatial localization task as a hybrid discrete-to-continuous evidential learning problem.

Discrete-to-Continuous Evidential Distribution: While standard boundary regression directly predicts a continuous value, our model treats the H vertical depth bins of each A-scan as a discrete probability space. To model epistemic uncertainty, we place a Dirichlet prior over these spatial bins. The network outputs non-negative evidence $e_h \geq 0$ for each bin $h \in H$, forming the Dirichlet parameters $\alpha_h = e_h + 1$. The total structural evidence is $S = \sum_{h=1}^H \alpha_h$, and the expected spatial probability mass function (PMF) is $p_h = \alpha_h / S$.

Uncertainty Calibration and Missing Layer Detection: We bridge this discrete evidential distribution to continuous boundary regression via spatial expectation. The sub-pixel boundary position is computed as $\mu = \sum_{h=1}^H p_h \cdot h$, and its aleatoric spatial variance is derived as $\sigma^2 = \sum_{h=1}^H p_h \cdot (h - \mu)^2$. This variance naturally expands in regions obscured by speckle noise or shadowing from hyperreflective areas. *Epistemic uncertainty* is dictated by the total evidence S . We leverage this to formulate our missing layer decision rule: if the predicted evidence S falls below a predefined threshold τ (empirically determined on a held-out validation set of 110 volumes) the layer boundary is classified as structurally missing. Otherwise, the network applies strict topological corrections to ensure non-crossing constraints among the valid layers.

Joint Optimization: The network is optimized using a composite loss function conditioned by a binary ground-truth mask indicating layer presence. In regions where the layer exists, we minimize the Gaussian Negative Log-Likelihood (NLL) utilizing the derived continuous spatial moments μ

and σ^2 against the target coordinate. This calibrates the spatial variance directly to the prediction error. We augment this with an evidence-maximizing objective that explicitly encourages the network to predict high total structural evidence when the layer is present. Conversely, in regions where the layer is anatomically missing, we apply an annealed Kullback-Leibler (KL) divergence penalty. This forces the predicted Dirichlet parameters toward a uniform, zero-evidence distribution ($KL(\alpha||\mathbf{1})$), effectively pushing the total evidence S to zero and training the network to confidently predict pathological absence.

4 Experimental Setup and Results




Figure 1: Comparative performance analysis. (a) Segmentation performance of the 2D projected atrophy area (GA footprint) against specialized baselines. (b) Maintenance of sub-pixel boundary localization accuracy compared to standard topological regression. (c) Qualitative B-scan highlighting the proposed method’s ability to identify missing layers (indicated by the transparent colored overlay). (d) En-face visualization of GA segmentation compared to ground truth (Green: correct, Red: over-segmentation, Blue: under-segmentation).

Dataset and Training Setup: We trained the proposed evidential model, the baseline SD-RetinaNet [4], and a state-of-the-art specialized 3D-to-2D GA segmentation model (Morano et al. [8]) on a shared private cohort consisting of 1100 OCT volumes with manual layer and lesion annotations. Of these, 55 volumes exhibited GA. The scans were acquired using Heidelberg Spectralis, Zeiss Cirrus, and Heidelberg HighRes devices, ensuring multi-vendor robustness.

Experiment 1: Geographic Atrophy Footprint Detection: To evaluate the network’s ability to detect missing layers, we tested the atrophy en-face footprint extraction performance on the independent public OCTAVE dataset [16], comprising 198 Spectralis OCT volumes (3762 B-scans), 9 of which contain GA. Evaluated specifically on these 9 GA-positive volumes, our evidential boundary regression achieved highly competitive GA footprint segmentation performance with a Dice score of 0.6959 ± 0.1878 and a PR-AUC of 0.8746. This is directly comparable to the specialized direct-segmentation method by Morano et al. [8], which yielded a Dice of 0.7102 ± 0.2101 and a

PR-AUC of 0.9297. This demonstrates that our 1D epistemic uncertainty mechanism successfully bridges boundary regression and region-based GA footprint detection without requiring a separate 3D pixel-wise classifier or the multiple inference passes required by Monte Carlo Dropout (MCD) (see Figure 1a).

Experiment 2: Layer Segmentation Maintenance: To confirm that adding uncertainty constraints does not severely degrade standard topological layer regression, we evaluated boundary localization on the public AROI dataset [17], consisting of 1136 annotated B-scans from 24 AMD patients featuring three manually annotated retinal layers. The proposed method demonstrated a Mean Absolute Distance (MAD) of $5.85 \pm 2.82 \mu\text{m}$, compared to the original SD-RetinaNet baseline performance of $4.68 \pm 2.29 \mu\text{m}$. These results demonstrate that while the introduction of evidential constraints yields a slight increase in localization error (from $4.68 \mu\text{m}$ to $5.85 \mu\text{m}$), this represents an acceptable and necessary trade-off; sub-pixel precision remains highly robust while enabling the critical clinical capability of detecting missing pathological layers (see Figure 1b).

Experiment 3: Uncertainty Calibration: A core contribution of our Gaussian NLL formulation is superior spatial error calibration. We evaluated the aleatoric calibration on the AROI dataset by computing the Expected Calibration Error (ECE) and the Pearson correlation (r) between the predicted standard deviation (σ) and the absolute boundary error (MAD). Our evidential framework achieved an ECE of $4.1 \mu\text{m}$ and a strong correlation of $r = 0.4798$ ($p < 0.001$), vastly outperforming the uncalibrated softmax entropy of the baseline SD-RetinaNet (ECE: $6.2 \mu\text{m}$, $r = 0.3812$, $p < 0.001$). This indicates that our predicted aleatoric variance serves as a highly reliable proxy for boundary ambiguity caused by speckle noise and shadowing.

Qualitative Evaluation: Visually (Figure 1c), the baseline SD-RetinaNet incorrectly interpolates continuous boundaries across atrophic regions, maintaining false confidence despite the lack of structural evidence. In contrast, our evidential framework correctly drops boundary predictions within GA lesions. Furthermore, the predicted aleatoric variance (σ) visibly widens at lesion margins and areas of poor signal, faithfully representing structural ambiguity. In the en-face projections (Figure 1d), while our method accurately captures the core GA footprint, the observed over-segmentation primarily stems from detecting missing layers near the peripheral borders of the B-scans. These false positives correlate strongly with signal roll-off and scanning artifacts rather than true anatomical atrophy.

5 Conclusion

In this work, we introduced a novel uncertainty-aware framework that successfully bridges the gap between spatial evidential deep learning and 1D topological boundary regression. By explicitly modeling aleatoric spatial variance via a Gaussian NLL formulation and quantifying epistemic structural evidence through Dirichlet priors, our model overcomes the critical limitation of previous boundary regression techniques: the mathematical obligation to hallucinate layers in pathological regions. The proposed method reliably detects structural layer absence, preventing erroneous boundary predictions in areas of complete tissue loss like Geographic Atrophy, while providing highly calibrated error estimates in regions obscured by shadowing. Ultimately, this framework combines high-precision sub-pixel layer tracking with robust 2D atrophy footprint extraction in a single pass. While this extended abstract presents an ongoing work in progress, it establishes a reliable and mathematically sound paradigm for automated pathological OCT analysis. Current efforts are actively focused on further improving both the layer segmentation accuracy and the GA footprint detection performance. Future work will evaluate the framework on larger GA-positive cohorts to strengthen clinical claims, and develop spatial priors to counteract false positive absence detections near peripheral scan borders.

Acknowledgements

The financial support by the Christian Doppler Research Association, Austrian Federal Ministry of Economy, Energy and Tourism, the National Foundation for Research, Technology and Development, and Heidelberg Engineering is gratefully acknowledged.

References

- [1] Neil M. Bressler. Age-Related Macular Degeneration Is the Leading Cause of Blindness... *Jama*, 291(15):1900–1901, 2004.
- [2] U. Schmidt-Erfurth, S. Klmscha, S. M. Waldstein, and H. Bogunovic. A view of the current and future role of optical coherence tomography in the management of age-related macular degeneration. *Eye*, 31(1):26–44, January 2017. ISSN 1476-5454. doi: 10.1038/eye.2016.227.
- [3] Yufan He, Aaron Carass, Yihao Liu, Bruno M. Jedynek, Sharon D. Solomon, Shiv Saidha, Peter A. Calabresi, and Jerry L. Prince. Structured layer surface segmentation for retina OCT using fully convolutional regression networks. *Medical Image Analysis*, 68:101856, February 2021. ISSN 1361-8415. doi: 10.1016/j.media.2020.101856.
- [4] Botond Fazekas, Guilherme Aresta, Dmitrii Lachinov, Sophie Riedl, Julia Mai, Ursula Schmidt-Erfurth, and Hrvoje Bogunovic. SD-LayerNet: Semi-supervised Retinal Layer Segmentation in OCT Using Disentangled Representation with Anatomical Priors. In Linwei Wang, Qi Dou, P. Thomas Fletcher, Stefanie Speidel, and Shuo Li, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, Lecture Notes in Computer Science, pages 320–329, Cham, 2022. Springer Nature Switzerland. ISBN 9783031164521. doi: 10.1007/978-3-031-16452-1_31.
- [5] Spencer C. Cleland, Sri Meghana Konda, Ronald P. Danis, Yijun Huang, Dawn J. Myers, Barbara A. Blodi, and Amitha Domalpally. Quantification of Geographic Atrophy Using Spectral Domain OCT in Age-Related Macular Degeneration. *Ophthalmology Retina*, 5(1): 41–48, January 2021. ISSN 2468-7219, 2468-6530. doi: 10.1016/j.oret.2020.07.006.
- [6] Dmitrii Lachinov, Philipp Seeböck, Julia Mai, Felix Goldbach, Ursula Schmidt-Erfurth, and Hrvoje Bogunovic. Projective Skip-Connections for Segmentation Along a Subset of Dimensions in Retinal OCT. In Marleen de Bruijne, Philippe C. Cattin, Stéphane Cotin, Nicolas Padoy, Stefanie Speidel, Yefeng Zheng, and Caroline Essert, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, pages 431–441, Cham, 2021. Springer International Publishing. ISBN 9783030871932. doi: 10.1007/978-3-030-87193-2_41.
- [7] Bart Liefers, Paul Taylor, Abdulrahman Alsaedi, Clare Bailey, Konstantinos Balaskas, Narendra Dhingra, Catherine A. Egan, Filipa Gomes Rodrigues, Cristina González Gonzalo, Tjebo F. C. Heeren, Andrew Lotery, Philipp L. Müller, Abraham Olvera-Barrios, Bobby Paul, Roy Schwartz, Darren S. Thomas, Alasdair N. Warwick, Adnan Tufail, and Clara I. Sánchez. Quantification of Key Retinal Features in Early and Late Age-Related Macular Degeneration Using Deep Learning. *American Journal of Ophthalmology*, 226:1–12, June 2021. ISSN 0002-9394. doi: 10.1016/j.ajo.2020.12.034.
- [8] José Morano, Guilherme Aresta, Dmitrii Lachinov, Julia Mai, Ursula Schmidt-Erfurth, and Hrvoje Bogunović. Self-supervised learning via inter-modal reconstruction and feature projection networks for label-efficient 3d-to-2d segmentation. In Hayit Greenspan, Anant Madabhushi, Parvin Mousavi, Septimiu Salcudean, James Duncan, Tanveer Syeda-Mahmood, and Russell Taylor, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*, pages 589–599, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-43901-8.
- [9] Murat Sensoy, Lance Kaplan, and Melih Kandemir. Evidential Deep Learning to Quantify Classification Uncertainty. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/a981f2b708044d6fb4a71a1463242520-Abstract.html>.
- [10] Abhijit Guha Roy, Sailesh Conjeti, Sri Phani Krishna Karri, Debdoot Sheet, Amin Katouzian, Christian Wachinger, and Nassir Navab. ReLayNet: retinal layer and fluid segmentation of macular optical coherence tomography using fully convolutional networks. *Biomedical Optics Express*, 8(8):3627–3642, August 2017. ISSN 2156-7085. doi: 10.1364/BOE.8.003627.
- [11] Hrvoje Bogunovic, Freerk Venhuizen, Sophie Klmscha, Stefanos Apostolopoulos, Alireza Bab-Hadiashar, Ulas Bagci, Mirza Faisal Beg, Loza Bekalo, Qiang Chen, Carlos Ciller, Karthik Gopinath, Amirali K. Gostar, Kiwan Jeon, Zexuan Ji, Sung Ho Kang, Dara D. Koozekanani,

- Donghuan Lu, Dustin Morley, Keshab K. Parhi, Hyoung Suk Park, Abdolreza Rashno, Marinko Sarunic, Saad Shaikh, Jayanthi Sivaswamy, Ruwan Tennakoon, Shivin Yadav, Sandro De Zanet, Sebastian M. Waldstein, Bianca S. Gerendas, Caroline Klaver, Clara I. Sánchez, and Ursula Schmidt-Erfurth. RETOUCH: The Retinal OCT Fluid Detection and Segmentation Benchmark and Challenge. *IEEE Transactions on Medical Imaging*, 38(8):1858–1874, August 2019. ISSN 1558-254X. doi: 10.1109/TMI.2019.2901398.
- [12] Alex Kendall and Yarin Gal. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/2650d6089a6d640c5e85b2b88265dc2b-Abstract.html>.
 - [13] Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus. Deep Evidential Regression. In *Advances in Neural Information Processing Systems*, volume 33, pages 14927–14937. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/aab085461de182608ee9f607f3f7d18f-Abstract.html>.
 - [14] Botond Fazekas, Guilherme Aresta, Philipp Seeböck, Julia Mai, Ursula Schmidt-Erfurth, and Hrvoje Bogunović. SD-RetinaNet: Topologically Constrained Semi-Supervised Retinal Lesion and Layer Segmentation in OCT. *IEEE Transactions on Medical Imaging*, pages 1–1, 2025. ISSN 1558-254X. doi: 10.1109/TMI.2025.3615240. URL <https://ieeexplore.ieee.org/document/11186218>.
 - [15] Mingxing Tan and Quoc Le. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In *Proceedings of the 36th International Conference on Machine Learning*, pages 6105–6114. PMLR, May 2019.
 - [16] Daniel S. Kermany, Wesley Poon, Anaya Bawiskar, Natasha Nehra, Orhun Davarci, Glori Das, Matthew Vasquez, Shlomit Schaal, Raksha Raghunathan, and Stephen T. C. Wong. Identifying Retinal Features Using a Self-Configuring CNN for Clinical Intervention. *Investigative Ophthalmology & Visual Science*, 66(6):55, June 2025. ISSN 1552-5783. doi: 10.1167/iovs.66.6.55.
 - [17] M. Melinščak, M. Radmilović, Z. Vatauvuk, and S. Lončarić. Aroi: Annotated retinal oct images database. In *2021 44th International Convention on Information, Communication and Electronic Technology (MIPRO)*, pages 371–376, 2021. doi: 10.23919/MIPRO52101.2021.9596934.