

Drusen segmentation with sparse volumetric SD-OCT sampling

Abum Okemgbo^a, I. David Rein^b, Yiyang Wang^c, Amani Fawzi^d, Jacob Furst^c, and Daniela Raicu^c

^aUniversity of Pennsylvania, Philadelphia, Pennsylvania

^bDuke University, Durham, North Carolina

^cDePaul University, Chicago, Illinois

^dFeinberg School of Medicine, Northwestern University, Chicago, Illinois

ABSTRACT

Age-Related Macular Degeneration (AMD) is a common eye disease characterized by the build-up of drusen, small deposits of extracellular materials in the macula. Early detection of drusen is key to understanding the progression of AMD. Therefore, accurate and robust segmentation of drusen during AMD progression is important for automated detection, classification, diagnosis, and prognosis tasks. Spectral-domain optical coherence tomography (SD-OCT) is a popular macular imaging modality used for these tasks. However, because of the trade-off between resolution and speed, often clinical OCT scans will contain far fewer images per volume than the 100-200 images the drusen segmentation literature generally utilizes. To address this disparity, we develop a novel drusen segmentation algorithm for SD-OCT volumes with low volumetric resolution. We achieve comparable results to similar work, while using on average 16% of the volumetric information. We evaluate our segmentation approach on manually segmented images by two graders, and achieve median Dice coefficient scores of 0.75 and 0.66, respectively, which are close to our median inter-reader variability score of 0.75.

Keywords: Drusen, segmentation, retinal imaging, Age-Related Macular Degeneration, spectral-domain optical coherence tomography

1. INTRODUCTION

Over 1.75 million US citizens have Age-Related Macular Degeneration (AMD), and it is the leading cause of blindness in individuals over the age of 65. Additionally, almost 3 million people are projected to be diagnosed with AMD in the US by 2020.¹ The progression of AMD can be characterized by the formation and growth of drusen, small deposits of extracellular materials that form in the macula, just beneath the retina.² Accurate and robust segmentation of drusen during AMD progression is therefore important for automated detection, classification, diagnosis, and prognosis tasks.²⁻⁵

Research has been done for segmentation of drusen given color fundus images,^{6,7} however drusen segmentation using spectral-domain optical coherence tomography (SD-OCT)⁸ is less prevalent. SD-OCT is an in-vivo tomographic imaging technique that is widely available, very fast, and highly sensitive.⁹ For these reasons, it has become popular in ophthalmology settings.⁹ Therefore, analysis using SD-OCT images has the potential for broad application in many areas and domains.

There has been previous work exploring techniques for automated drusen segmentation from SD-OCT.^{5,10} Generally these methods create a difference map that highlights the areas in which layers in the macula deviate from their “healthy” positions. This work is promising, and we build upon it. Specifically, most drusen segmentation approaches rely heavily on volumetric information.^{5,11,12} However, some OCT devices output fewer images (B-scans, shown in Fig. 1) in order to increase the resolution of individual B-scans. For example, different scanning protocols from the Spectralis OCT device trade between producing higher resolution B-scans with

Further author information:

I.D.R.: E-mail: irving.rein@duke.edu

A.F.O.: E-mail: okemgbo@seas.upenn.edu




Figure 1: Relationship between an A-scan, a single depth scan, a B-scan, a set of A-scans aligned in an image, and a volume, the total collection of B-scans representing the macula of an eye.

more A-scans, and scanning quickly to produce a higher number of lower resolution B-scans.¹³ Our contribution has several components. One main goal is the development of an automated retinal pigment epithelium (RPE) and drusen segmentation algorithm that performs well despite a low SD-OCT vertical resolution, by carefully utilizing volumetric information to perform segmentation. The RPE and other relevant layers are displaying in Fig. 2. Additionally, our method has the advantage of not being learned,¹⁴ so we did not require a large amount of manually segmented volumes. This data needed for common supervised methods is very time-consuming to label by hand. Finally, our method is faster than deep learning approaches, and does not require a large amount of compute.^{15,16}




Figure 2: The inner limiting membrane (ILM) is shown in green, the retinal nerve fiber layer (RNFL) is shown in orange, the inner segment/outer segment (IS/OS) junction is in purple, the RPE is shown in red, and the Bruch's membrane (BM) is shown in blue.

2. METHODS

2.1 Data

All images were captured using the same SD-OCT device (Spectralis, Heidelberg Engineering, Heidelberg, Germany). It contains 52 de-identified individual subjects and 12,272 B-scans, which comprise 459 SD-OCT volumes. 75 images have a resolution of 384×496 , 8964 images have a resolution of 512×496 , 3238 images have a resolution of 768×496 , and 296 images have a resolution of 1024×496 . Additionally, we have a varying number of B-scan slices per volume. Two-thirds of images have 25 slices per volume, and the rest have various numbers of slices up to 121 slices.

2.2 Retinal pigment epithelium (RPE) segmentation

Our RPE segmentation algorithm follows previous work,¹⁷ with several modifications for our context. Because of the high reflectivity of the RPE, we begin by selecting the highest intensity pixel in each A-scan, shown as the output of step (i) in Fig. 3. Then, we remove the RNFL from our images as demonstrated in step (ii) of Fig. 3. Similarly to the RPE, the RNFL is often highly reflective, so its removal is important for reducing the number of outlier pixels from our initial RPE estimation as represented in Fig. 4. To remove the RNFL, when selecting high-intensity pixels we only consider pixels at least 50 pixels below the ILM as valid possibilities for the RPE pixels. The ILM segmentation comes from the proprietary software from the Spectralis SD-OCT device.¹³ Because the ILM is consistently only a few pixels above the RNFL, searching below 50 pixels below the ILM guarantees our exclusion of the RNFL, without affecting valid RPE pixels. Our current estimate of the RPE can be represented as a topographic map, shown in Fig. 3, step (iii). To refine our initial RPE segmentation, we use an iterative intensity-based approach, where each iteration is characterized by a line-search for the highest intensity pixels in each A-scan, followed by several median filters with windows of varying sizes across the RPE topographic map in step (iv). Specifically, after the RNFL removal, we apply a close filter to the RPE map with a 5x5 structuring element, followed by a median filter with dimension 2x31. The first dimension indicates the between-slice filter size, and the second dimension indicates the horizontal size across A-scans within each B-scan. We then perform a line-search from 40 pixels below the current RPE estimate in each A-scan to 30 pixels above the RPE, and return the position of the median of the 3 highest intensity pixels in the search range. Because of the thickness of the RPE, this helps remove noise from the highly reflective inner segment/outer segment (IS/OS) junction. We then apply a median filter of size 1x21.

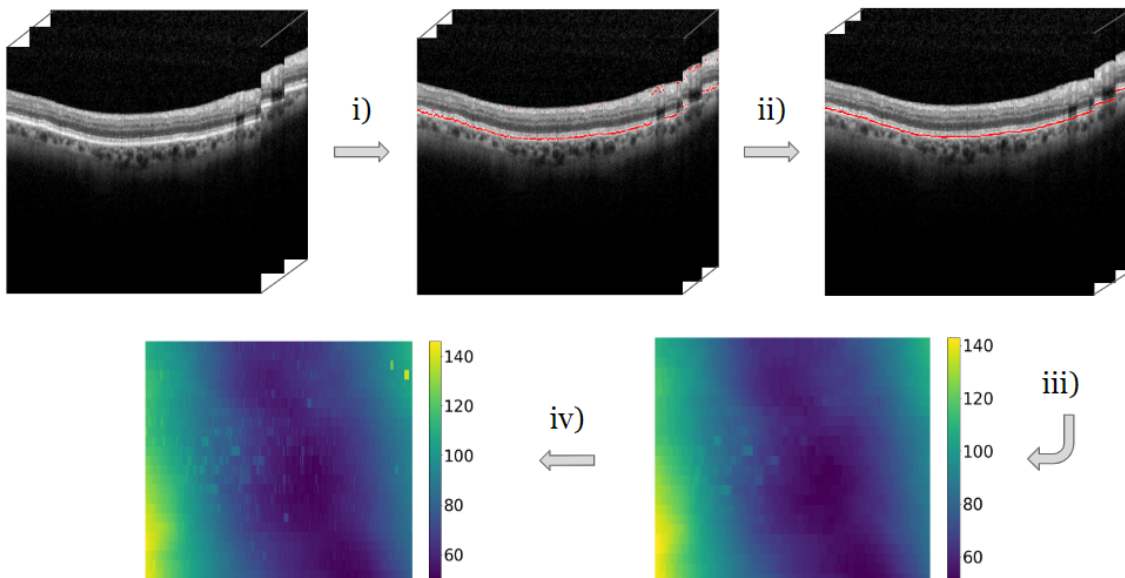


Figure 3: An overview of the RPE identification process. i) Initial high intensity pixel selection in each A-scan. Pixels in red represent the current RPE estimate. ii) Removal of high intensity RNFL pixels. iii) Topographic map representation of the RPE. The estimate of the RPE depth at each pixel is represented by the color in the map; deeper estimates are blue, while higher values are yellow. iv) Filtering and refinement steps.

This two-step process of a line-search followed by a median filter is repeated two more times, with search ranges of 10 pixels above and below the RPE, and 5 pixels above and below the RPE, respectively, and with median filter sizes of 1x5 and 1x2, respectively. These filter sizes were selected heuristically to iteratively refine and narrow down the RPE pixel selection process. We initially have large filters, to broadly find the RPE position, and we slowly decrease the search ranges and median filter sizes to refine the estimation in greater detail. Finally, we apply a 1D averaging filter of size 10 to each RPE estimate in the volume. We use the non-averaged RPE estimate on both sides of the image to avoid boundary effects.

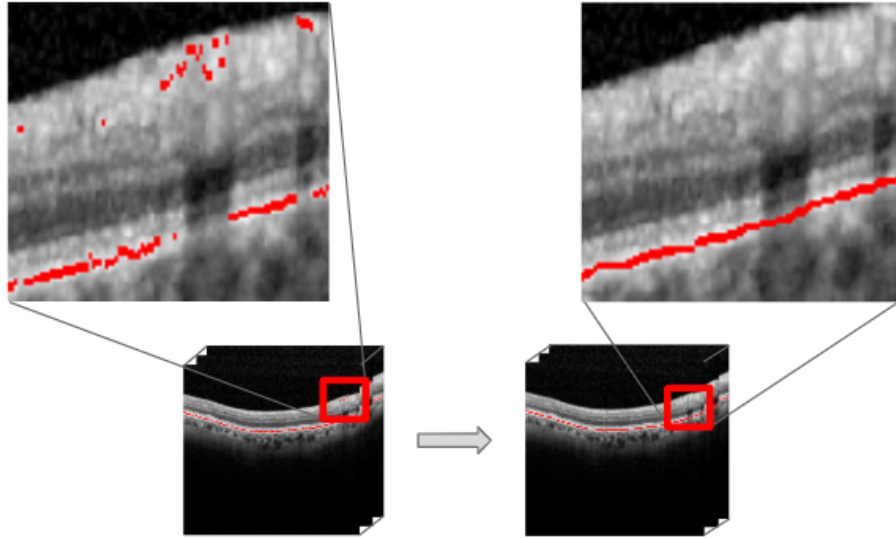


Figure 4: A close-up of step (ii) of Fig. 3, illustrating the removal of high-intensity pixels in the RNFL.

2.3 Drusen segmentation

Once we have an estimated RPE segmentation, we begin the drusen segmentation process. First, we take the difference between the RPE segmentation map (a) resulting from the proposed segmentation and BM segmentation map (b) resulting from using proprietary software from the Spectralis SD-OCT device¹³ to obtain a difference map, shown in the first image of Fig. 5 in step (i). This highlights the areas where the RPE deviates from its “healthy” shape. This healthy shape is understood as the natural concave shape of the RPE.

We apply the mean-shift algorithm¹⁸ to remove non-elevated sections of the RPE in step (ii). The mean-shift algorithm returns clusters representing different height levels in the difference map. For instance, the highest height level denotes a grouping of the most elevated points. The lowest level represents areas where the RPE is not elevated, which we remove from our drusen segmentation. Before continuing, we establish a minimum druse height to separate small changes in the RPE from actual drusen in step (iii). This helps remove noise from the highly reflective IS/OS junction, and helps individually segment connected drusen.

Each 4-connected component above the height threshold is indexed as an initial individual drusen segmentation. Now that we have markers for individual drusen, we refine the segmentation to more accurately represent the areas of each one. In step (iv), for each druse we explore the neighboring locations within each B-scan and iteratively add the nearest elevated regions to the druse segmentation. Once we reach a non-elevated location, we stop expanding the druse in that given direction. Finally, in step (v) we unflatten the obtained drusen segmentation by projecting the difference map of the segmentations back onto the original volumes. The final segmented drusen are the polygons formed between the RPE and the BM in the segmented areas, as shown after step (v).

2.4 Evaluation

We evaluated our results on a set of 89 manually annotated images. The images were each manually segmented by two experienced, non-medical graders. From each patient who was diagnosed with drusen, images with visible drusen biomarkers were selected. We removed images containing large pigment epithelium detachments (PEDs). The primary metric with which we evaluate our automated segmentation algorithm is the Sørensen-Dice coefficient,¹⁹ also known as the Dice similarity coefficient (DSC). In our context, for each B-scan we have two sets of pixels: one from our manual segmentation, and one from our automated segmentation. The DSC of these

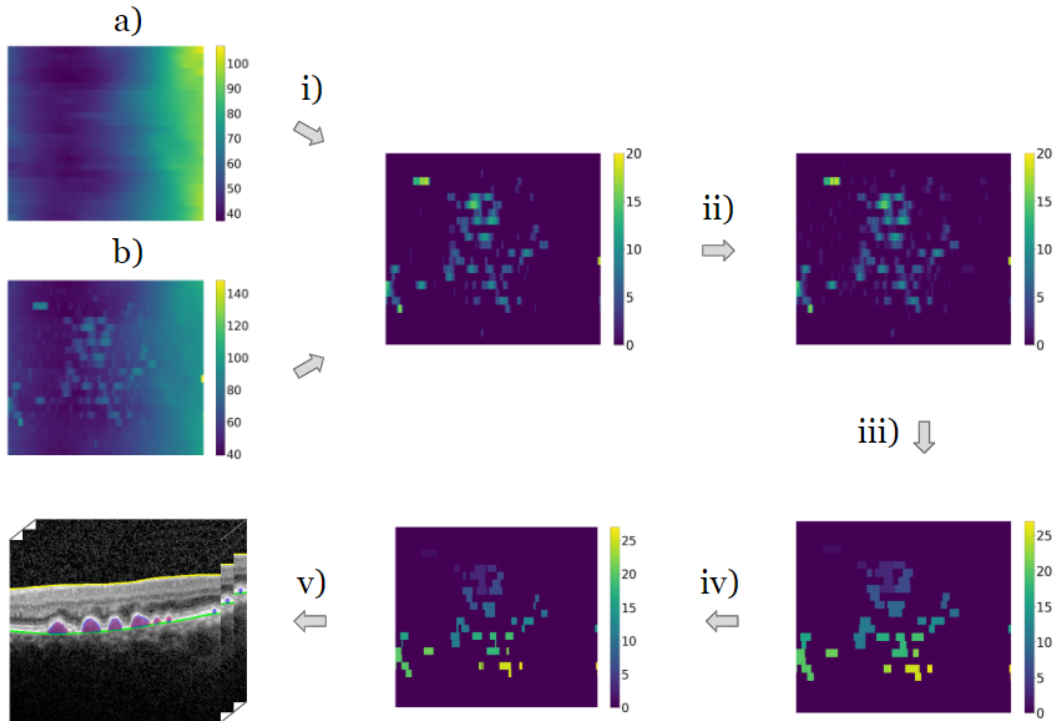


Figure 5: An overview of the drusen segmentation process. Initially, there are the RPE (a) and BM (b) maps. i) Subtraction map made from the BM and RPE. ii) Mean-shift threshold-based segmentation. iii) Height-based segmentation refinement. Each color represents a different druse. iv) Drusen region expansion. v) Remap values to original image.

two sets is equivalent to the common F1-score evaluation criteria, if we consider each pixel a sample. To provide a baseline for our evaluation, we compute the DSC between our manual annotations.

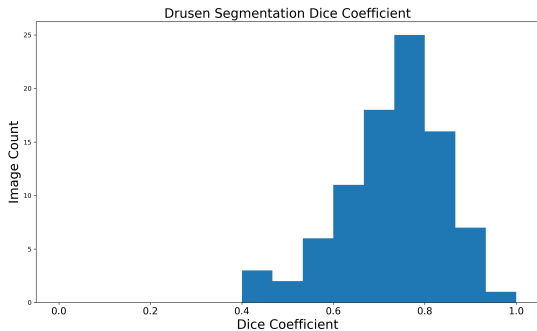
3. RESULTS

As shown in Table 1, our graders had high inter-reader DSC values, indicating significant overlap in their manual drusen segmentations. The DSC between our algorithm and each grader is very close to the inter-reader DSC, indicating that our algorithm is performing similarly to our graders. As expected, the distribution is slightly left-skewed, with median DSC values between our algorithm and each grader of 0.75 and 0.66 for grader 1 and grader 2, respectively. The inter-reader DSC is 0.75. We would not expect our algorithm to have a higher DSC with either grader compared to the inter-reader DSC. There is some work that reports 10 to 20 percent higher overlap coefficients,⁵ however, considering the significant reduction in volumetric resolution, these results represent a strong improvement over other methods for sparsely sampled SD-OCT volumes. Furthermore, the magnitude of the dice coefficients can be affected by the specific data set used. For example, if an image is difficult to manually segment, we would expect lower dice coefficients between the algorithms and the graders due to the added uncertainty in the segmentation. The full distribution of DSC values between our algorithm and each grader can be see in Fig. 6.


In Fig. 7, we show examples of our automated segmentation process. In image 7a, we can see an example of an image that was accurately segmented by our algorithm. In image image 7b, we display an image in which the drusen areas were segmented accurately, but four individual drusen were segmented as a single druse. This issue, where drusen close to each other are grouped together, is the most common failure mode for our algorithm.

Table 1: Distribution of DSC values between our algorithm and each grader.

Criteria	Grader 1	Grader 2	Inter-reader DSC
Minimum	0.42	0.25	0.47
Quartile Q1	0.67	0.59	0.70
Median	0.75	0.66	0.75
Quartile Q3	0.80	0.74	0.82
Maximum	0.96	0.91	0.92



(a) Grader 1



(b) Grader 2


Figure 6: The distribution of Dice coefficient values between our automated algorithm and each grader.

4. DISCUSSION


We made four specific changes to the RPE segmentation algorithm¹⁷ our approach is based on. First, we remove the RNFL. This layer is often highly reflective, causing the initial highest-intensity pixel selection process to locate many pixels in that layer. This helps improve our initial RPE estimate. After removing the RNFL and obtaining the RPE topographic map, we compute the morphological closing of the map instead of applying a top-hat transform (defined as the morphological opening subtracted from the map) to remove outlier pixels. The close filter gave us a more accurate initial segmentation of the RPE that was more robust to high intensity pixels in other layers of the macula. The third change is with regard to the median filter dimensions. As our volumes on average have only 25 B-scans, we cannot rely heavily on the assumption of smoothness of the RPE between slices. More specifically, since our B-scans can be up to 200 μm apart, many drusen are only captured in a single slice. Therefore, if we apply too much smoothing across slices, we would miss many smaller drusen. So, we reduced the y-axis (vertical) dimension of the various median filters. Finally, because of our reduced ability to use large smoothing filters across B-scans, we applied a windowed moving average filter at the end of the RPE segmentation process to better capture the smooth nature of the RPE. By qualitative assessment, these changes improve the RPE segmentation approach on volumes with low between-slice resolution.

5. CONCLUSION

We develop an algorithm that segments drusen in SD-OCT volumes with few B-scans. Our approach performs comparably to existing methods, but with far less information per volume. Our current evaluation criteria are based on the Dice coefficient, which suggests that the segmentations have a high overall overlap. However, we do not evaluate the quality of the individual segmentations. Specifically, we have not measured how well our algorithm matches automated individual segmentations to manual segmentations. This issue is most clearly represented by Fig. 7b. In this image, although the general overlap is relatively high, with a Dice coefficient of 0.67, several nearby drusen are incorrectly segmented as a single druse. Despite this, there are many key features



(a) Accurate segmentation



(b) Multiple drusen are incorrectly grouped together

Figure 7: Examples of the algorithm’s drusen segmentations (shown in red).

of AMD progression that can be computed without individual drusen counts, and the development of methods that segment individual drusen from sparsely sampled SD-OCT volumes remains a promising avenue of research.

ACKNOWLEDGMENTS

Our research is supported by NSF CISE REU Program Award Number 1659836.

REFERENCES

- [1] The Eye Diseases Prevalence Research Group, “Prevalence of Age-Related Macular Degeneration in the United States,” *JAMA Ophthalmology* **122**, 564–572 (04 2004).
- [2] Wang, J. J., Foran, S., Smith, W., and Mitchell, P., “Risk of Age-Related Macular Degeneration in Eyes With Macular Drusen or Hyperpigmentation: The Blue Mountains Eye Study Cohort,” *JAMA Ophthalmology* **121**, 658–663 (05 2003).
- [3] Age-Related Eye Disease Study Research Group, “The Age-Related Eye Disease Study Severity Scale for Age-Related Macular Degeneration: AREDS Report No. 17,” *JAMA Ophthalmology* **123**, 1484–1498 (11 2005).
- [4] Klein, R., Klein, B. E., Jensen, S. C., and Meuer, S. M., “The five-year incidence and progression of age-related maculopathy: The beaver dam eye study,” *Ophthalmology* **104**(1), 7 – 21 (1997).
- [5] de Sisternes, L., Jonna, G., Greven, M. A., Chen, Q., Leng, T., and Rubin, D. L., “Individual Drusen Segmentation and Repeatability and Reproducibility of Their Automated Quantification in Optical Coherence Tomography Images,” *Translational Vision Science & Technology* **6**, 12–12 (02 2017).
- [6] Mittal, D. and Kumari, K., “Automated detection and segmentation of drusen in retinal fundus images,” *Computers & Electrical Engineering* **47**, 82 – 95 (2015).
- [7] Khalid, S., Akram, M. U., Hassan, T., Jameel, A., and Khalil, T., “Automated segmentation and quantification of drusen in fundus and optical coherence tomography images for detection of armd,” *Journal of Digital Imaging* **31**, 464–476 (Aug 2018).
- [8] Huang, D., Swanson, E. A., Lin, C. P., Schuman, J. S., Stinson, W. G., Chang, W., Hee, M. R., Flotte, T., Gregory, K., and Puliafito, C. A., “Optical coherence tomography,” *Science (New York, N.Y.)* **254**, 1178–1181 (Nov 1991).
- [9] Yaqoob, Z., Wu, J., and Yang, C., “Spectral domain optical coherence tomography: a better oct imaging strategy,” *BioTechniques* **39**(6S), S6–S13 (2005).
- [10] Farsiu, S., J. Chiu, S., Izatt, J., and Toth, C., “Fast detection and segmentation of drusen in retinal optical coherence tomography images - art. no. 68440d,” *Progress in Biomedical Optics and Imaging - Proceedings of SPIE* **6844** (03 2008).
- [11] Chen, Q., Leng, T., Zheng, L., Kutzscher, L., Ma, J., de Sisternes, L., and Rubin, D. L., “Automated drusen segmentation and quantification in sd-oct images,” *Medical Image Analysis* **17**(8), 1058 – 1072 (2013).
- [12] Chiu, S. J., Izatt, J. A., O’Connell, R. V., Winter, K. P., Toth, C. A., and Farsiu, S., “Validated automatic segmentation of amd pathology including drusen and geographic atrophy in sd-oct images,” *Investigative ophthalmology & visual science* **53**(1), 53–61 (2012).
- [13] “Spectralis hardware operating manual,” (Aug 2017).

- [14] De Fauw, J., Ledsam, J. R., Romera-Paredes, B., Nikolov, S., Tomasev, N., Blackwell, S., Askham, H., Glorot, X., ODonoghue, B., Visentin, D., et al., “Clinically applicable deep learning for diagnosis and referral in retinal disease,” *Nature medicine* **24**(9), 1342 (2018).
- [15] Gorgi Zadeh, S., Wintergerst, M. W. M., Wiens, V., Thiele, S., Holz, F. G., Finger, R. P., and Schultz, T., “Cnns enable accurate and fast segmentation of drusen in optical coherence tomography,” in [*Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*], Cardoso, M. J., Arbel, T., Carneiro, G., Syeda-Mahmood, T., Tavares, J. M. R., Moradi, M., Bradley, A., Greenspan, H., Papa, J. P., Madabhushi, A., Nascimento, J. C., Cardoso, J. S., Belagiannis, V., and Lu, Z., eds., 65–73, Springer International Publishing, Cham (2017).
- [16] Fang, L., Cunefare, D., Wang, C., Guymer, R. H., Li, S., and Farsiu, S., “Automatic segmentation of nine retinal layer boundaries in oct images of non-exudative amd patients using deep learning and graph search,” *Biomedical optics express* **8**, 2732–2744 (Apr 2017). 28663902[pmid].
- [17] Fabritius, T., Makita, S., Miura, M., Myllylä, R., and Yasuno, Y., “Automated segmentation of the macula by optical coherence tomography,” *Opt. Express* **17**, 15659–15669 (08 2009).
- [18] Comaniciu, D. and Meer, P., “Mean shift: a robust approach toward feature space analysis,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **24**, 603–619 (May 2002).
- [19] Dice, L. R., “Measures of the amount of ecologic association between species,” *Ecology* **26**(3), 297–302 (1945).