Supplementary Material

S1. Dataset Overview

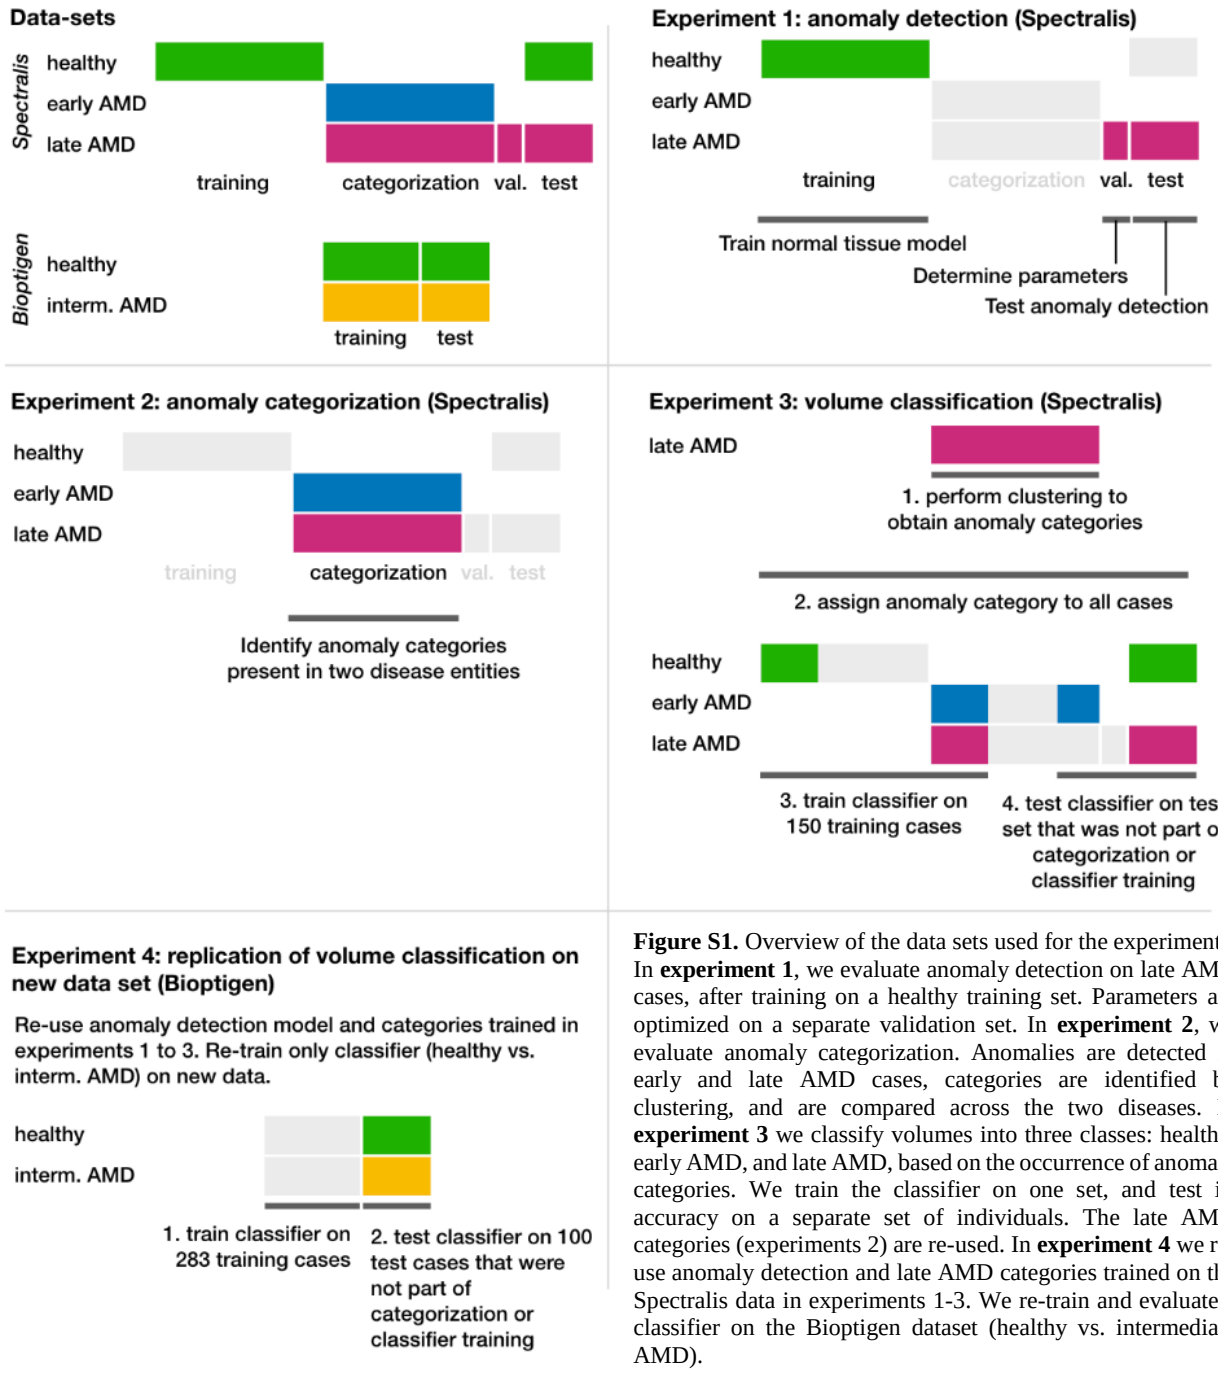

Figure S1. Overview of the data sets used for the experiments. In **experiment 1**, we evaluate anomaly detection on late AMD cases, after training on a healthy training set. Parameters are optimized on a separate validation set. In **experiment 2**, we evaluate anomaly categorization. Anomalies are detected in early and late AMD cases, categories are identified by clustering, and are compared across the two diseases. In **experiment 3** we classify volumes into three classes: healthy, early AMD, and late AMD, based on the occurrence of anomaly categories. We train the classifier on one set, and test its accuracy on a separate set of individuals. The late AMD categories (experiments 2) are re-used. In **experiment 4** we re-use anomaly detection and late AMD categories trained on the Spectralis data in experiments 1-3. We re-train and evaluate a classifier on the Bioptigen dataset (healthy vs. intermediate AMD).

S2. Volume Level Classification Examples

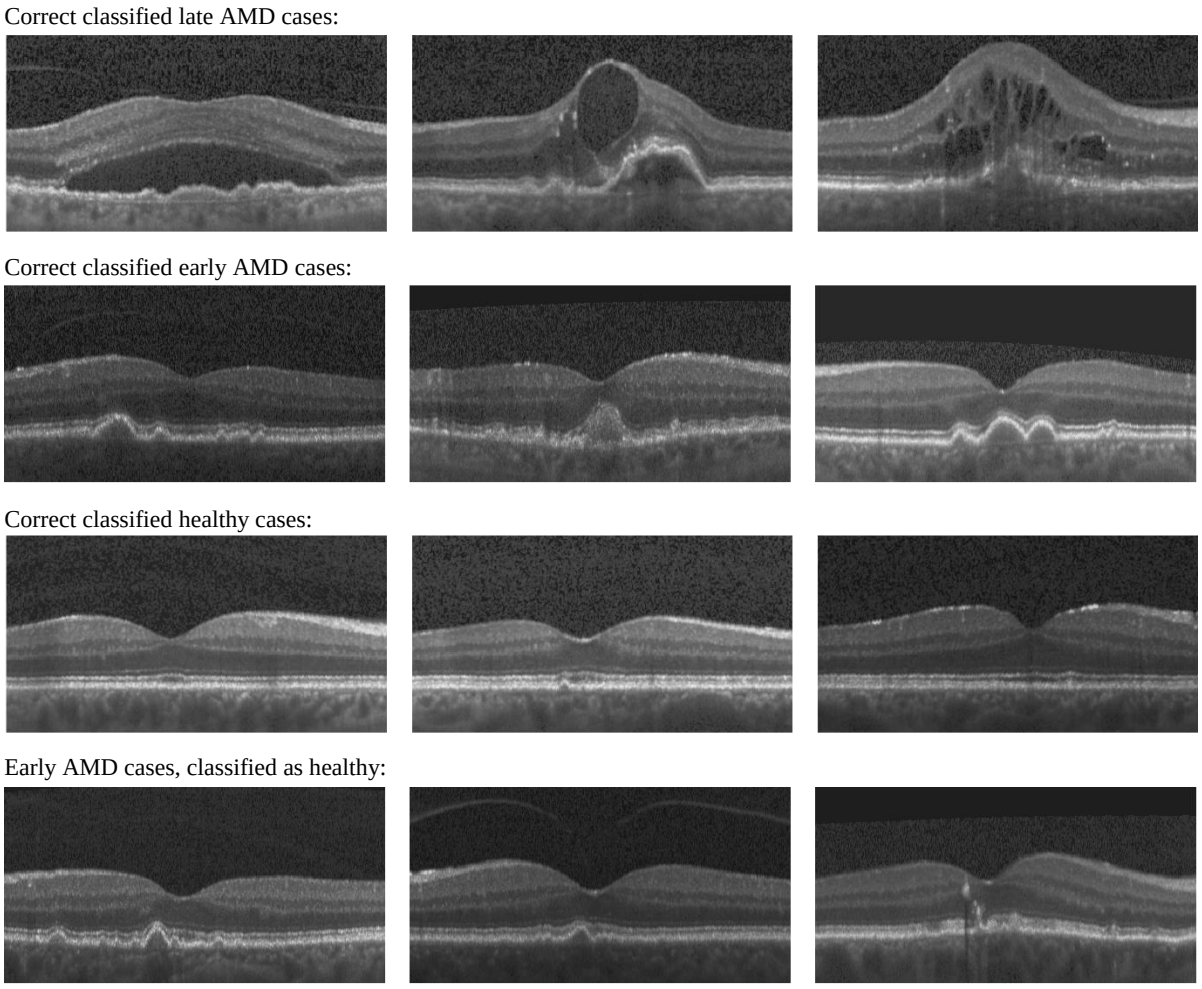

Figure S2. Exemplary cases of the volume level classification test set. While the first three rows show cases which were correctly classified, the bottom row illustrates incorrect classified early AMD volumes.

S3. Over-Segmentation

Generally, the goal of over-segmenting an image is to merge pixels into homogeneous groups of superpixels, while preserving boundaries of objects in an image. This allows to perform image analysis tasks on the greatly reduced number of superpixels as opposed to every pixel in the volume.

In our work, the runtime is reduced by a factor of 16 using superpixels, since the superpixels have an average size of 4-by-4. More precisely, we performed over-segmentation on each B-Scan separately, using the monoSLIC method. This approach transforms the image content to its monogenic signal, which enables to generate superpixels with high fidelity to local edge information while being of regular size and shape.

S4. Anomalous Regions - Manual Annotations

Regarding the manual annotations on the *late AMD* dataset, the whole volume was binary annotated B-Scan wise into normal and anomalous regions. A retina specialist manually marked all areas of the image that contained pathologic features. These may include for instance, but not limited to, drusen >63 microns, pigmentary changes, accumulations of any type of fluid, alterations of the retinal pigment epithelium including pigment epithelial detachment or atrophy, hyperreflective subretinal tissue and lipid exudates.

S5. Runtime

The inference times of all methods are provided in Table S1. Experiments were performed on a computer with an i7-3770K CPU and a TitanX GPU. Reported runtimes are valid for a single OCT volume with a pixel dimensionality of 512 x 496 x 49. Please note that the code is not optimized for runtime.

The runtime of the One-Class SVM depends on the feature-dimensionality as well as on the number of support vectors which are needed to describe the hyperplane in the feature space. The former results in a shorter runtime of PCA0.95, while the latter in a longer runtime of PCA256.

Method	**Time for feature-calculation**	**Time of one-class SVM**	**Number of Support Vectors**	**Number of Features**
PCA_{256}	0.4 sec.	450 sec.	28 271	256
$PCA_{0.95}$	0.25 sec.	75 sec.	16 650	63
DCAE	250 sec.	217 sec.	10 007	256
$DDAE_{ent}$	41 sec.	212 sec.	5 009	256

Table S1. The first two columns provide the runtime for feature computation as well as the time for calculating the prediction with the One-Class SVM. The last two column show the number of support vectors needed to describe the decision hyperplane and the feature dimension of the input.

S6. Publicly Available Dataset[1] – Preprocessing

The dataset consists of 384 Bioptigen SD-OCT volumes (269 AMD patients and 115 normal subjects) with a pixel dimensionality of 1000 x 512 x 100.

For one AMD case ("Farsiu_Ophthalmology_2013_AMD_Subject_1024.mat") our layer-segmentation algorithm failed due to bad scan quality, which is why we excluded it from our experiments. For all 383 remaining Bioptigen SD-OCT volumes, we conducted the following additional preprocessing steps for each B-Scan, preceding the preprocessing steps described in Section II-A:

- We performed non-local means filtering (NLMF) with a radius of the local patch of 3, a radius of the neighborhood search window of 3, and a strength of the NLMF filtering of 0.15.
- Resizing the B-scan from 512 to 496 pixels in the vertical dimension, and from 1000 to 512 in the horizontal dimension, using bilinear interpolation.
- Subsequently image adjustment is conducted. More specifically, the intensity values of the input B-scan are mapped to the range [0 1], where the bottom/lowest 10% pixel values are ignored: [10% 100%] → [0 1].
- Finally, scaling with a factor of 1.16 is performed in z-dimension, according to the definition provided online (*http://people.duke.edu/~sf59/RPEDC_Ophth_2013_dataset.htm*).

In Figure S4 these preprocessing steps are visualized.

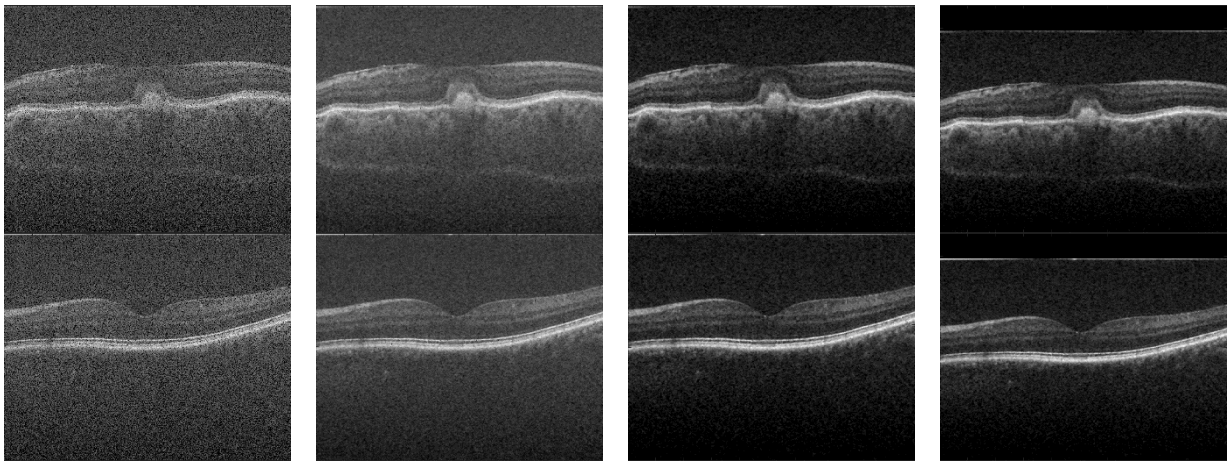

(a) (b) (c) (d)

Figure S4. The Preprocessing is shown for two B-scans, where the top and bottom row show an AMD and a healthy B-Scan, respectively. The original Bioptigen Scan (a), the result of applying NLMF (b), image adjustment (c) and scaling (d) are visualized.

[1] http://people.duke.edu/~sf59/RPEDC_Ophth_2013_dataset.htm