Adaptation of a support vector machine algorithm for segmentation and visualization of retinal structures in volumetric optical coherence tomography data sets

Robert J. Zawadzki
University of California, Davis
Vision Science and Advanced Retinal Imaging
    Laboratory (VSRI) and Dept. of Ophthalmology
    and Vision Science
4860 Y Street, Suite 2400
Sacramento, California 95817

Alfred R. Fuller
David F. Wiley
Bernd Hamann
University of California, Davis
Institute for Data Analysis and Visualization (IDAV)
    and Dept. of Computer Science
One Shields Avenue
Davis, California, 95616

Stacey S. Choi
John S. Werner
University of California, Davis
Vision Science and Advanced Retinal Imaging
    Laboratory (VSRI) and Dept. of Ophthalmology
    and Vision Science
4860 Y Street, Suite 2400
Sacramento, California 95817

Abstract. Recent developments in Fourier domain—optical coherence tomography (Fd-OCT) have increased the acquisition speed of current ophthalmic Fd-OCT instruments sufficiently to allow the acquisition of volumetric data sets of human retinas in a clinical setting. The large size and three-dimensional (3D) nature of these data sets require that intelligent data processing, visualization, and analysis tools are used to take full advantage of the available information. Therefore, we have combined methods from volume visualization, and data analysis in support of better visualization and diagnosis of Fd-OCT retinal volumes. Custom-designed 3D visualization and analysis software is used to view retinal volumes reconstructed from registered B-scans. We use a support vector machine (SVM) to perform semiautomatic segmentation of retinal layers and structures for subsequent analysis including a comparison of measured layer thicknesses. We have modified the SVM to gracefully handle OCT speckle noise by treating it as a characteristic of the volumetric data. Our software has been tested successfully in clinical settings for its efficacy in assessing 3D retinal structures in healthy as well as diseased cases. Our tool facilitates diagnosis and treatment monitoring of retinal diseases. © *2007 Society of Photo-Optical Instrumentation Engineers.*
 [DOI: 10.1117/1.2772658]

Keywords: optical coherence tomography; ophthalmology; optical diagnostics for medicine; pattern recognition and feature extraction; support vector machine; volume visualization.

1 Introduction

Recent advances in Fd-OCT[1–3] make possible *in vivo* acquisition of ultrahigh-resolution volumetric retinal OCT data in clinical settings. This technology has led to new and powerful tools that have the potential to revolutionize the monitoring and treatment of many retinal and optic nerve diseases similar to the advancement achieved in other medical areas due to the application of clinical volumetric imaging. However, in order to fully realize this potential, new tools allowing the visualization and measurement of retinal features are required. Attempts at visualization of OCT volumetric retina data have recently been presented by several groups[4–6] and have included visualizations of highly magnified retinal structures imaged with adaptive optics (AO-OCT) systems.[7,8] Possible approaches to retinal layer segmentation have also been presented.[9–12]

In this paper we present a clinical Fd-OCT system that produces retinal volumes that are visualized and analyzed in real time with custom software. A fully automated approach that segments, classifies, and analyzes retinal layers would be ideal. However, the morphology of retinal layers varies dramatically from patient to patient and depends on the particular pathogenic changes of the disease in question. This causes problems for existing automatic retinal layer extraction methods used in clinical instruments.[13] Therefore, to simplify this task, we have developed a system based on a semiautomatic segmentation method. A clinician interactively specifies the location of a retinal layer on a few select slices in order to create a segmentation of the entire volume. Our method uses an SVM-based classification system,[14,15] which is a machine learning method used to predict results based on a given set of inputs. An SVM is tolerant of misclassification by the user and of physiological variation between patients and diseases and easily adapts to the varying data properties that constitute

Address all correspondence to Robert J. Zawadzki. Tel: 916–734–5839; Fax: 916–734–4543; E-mail: rjzawadzki@ucdavis.edu

a retinal layer. Also, our SVM approach gracefully handles the speckle noise that disturbs all data acquired through OCT by modeling it as a normal distribution characterized by a voxel's local data mean and variance. Once the layers are segmented, we can extract a thickness map, layer profile, and z-axis intensity projection or measure a volume of this structure. We compare our semiautomatic segmentations to manually segmented layers and test its performance for different retinal and optic nerve diseases.

2 Materials and Methods

The results presented in this paper have been acquired during the routine use of our clinical Fd-OCT system. Over the past two years, we have used this system to acquire volumetric scans from more than 200 subjects, with healthy and diseased retinas or optic nerve heads. More than half of these data sets have been used to create volumetric representations of the imaged retinal structures. SVM segmentation has been implemented with 30 normal and diseased retinas to evaluate classification of different retinal structures. Due to involuntary eye motion and reduced (or distorted) intensity of some OCT images, not all of the 3D scans are appropriate for volumetric reconstruction; the inappropriate scans are related to (especially among elderly patients) advanced cataract, significant eye aberrations, long eyelashes, and ptosis. As a standard method of qualifying retinal scans for volumetric reconstruction and subsequent segmentation, a movie showing all consecutive B-scans acquired in the volume is generated and viewed by the operator. The C-scan reconstructed from OCT images is also used to rate distortions caused by eye motion. Due to these problems, which are routine in any clinic, we have concentrated our efforts on a segmentation method that tolerates noisy data.

2.1 Experimental System

The base configuration of the experimental Fd-OCT system used for acquiring 3D volumes has been described previously.[7] For data presented in this paper, a different light source has been used: a superluminescent diode with an 855-nm central wavelength and an FWHM of 75 nm. The spectral bandwidth of this light source allows 4.5-μm axial resolution in the retina (refractive index, $n=1.38$). The high power of this light source supports use of a 90/10 fiber coupler, directing 10% of the light toward the eye, resulting in the power at the subject's cornea equal to 700 μW and allowing 90% throughput of the light back-reflected from the eye to the detector. With our current spectrometer design (200-mm focal length of the air-spaced doubled), the maximum axial range (seen after the Fourier transform) is about 3.6 mm in free space, corresponding to approximately 2.7 mm in the eye. Figure 1 shows a schematic of our clinical Fd-OCT system.

Each subject was imaged with several OCT scanning procedures including 3D scanning patterns centered at the fovea and ONH. We used two different scanning arrangements for 3D scanning patterns including (1) equally spaced 200 B-scans, with each based on 500 A-scans, and (2) equally spaced 100 B-scans based on 1000 A-scans. In both cases, volumes consisted of the same number of A-scans (100,000), and the time required to acquire a volume was 5.5 s for 50-μs CCD exposure time and 11.1 s for 100-μs exposure time.
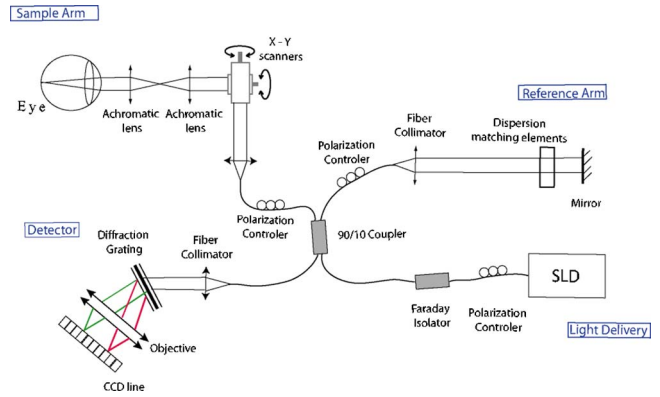

Fig. 1 Schematic of the clinical Fd-OCT system at UC Davis. SLD—superluminescent diode.

The longer exposure time was mainly used to increase image intensity and with 100 B-scans based on 1000 A-scans per frame acquisition mode, at the cost of more motion artifacts. Commercial Fd-OCT acquisition software from Bioptigen, Inc. permitted real-time display (at the acquisition speed) of the imaged retinal structures and saved the last acquired volume.

2.2 Image Processing

After each imaging session, saved spectral data are processed in LabVIEW using standard Fd-OCT procedures,[3,7,16] including spectral shaping, zero padding, and dispersion compensation. As a result, a set of high-quality OCT B-scans is saved in tiff format. It is possible to then analyze data and choose volumes with no eye motion or ones in which eye movement occurred only at the beginning or toward the end of the 3D acquisition, provided there was good fixation during the rest of the time (especially if the structure of clinical relevance was scanned without distortions). In these cases, the part of the image that was affected by eye motion is removed, resulting in a reduction in the number of B-scans and overall size of the reconstructed 3D volume. In addition, as a preprocessing step before volume visualization, OCT B-scans are registered using standard registration techniques. Figure 2 shows a single B-scan and also a cross section of several B-scans showing the alignment in different cross sections.

To correct for axial and lateral distortions, we used ImageJ,[17] with the Turboreg plug-in developed in the laboratory of P. Thévenaz at the Biomedical Imaging Group, Swiss Federal Institute of Technology Lausanne,[18] which computes rigid body transformations (translation and rotation) to minimize the difference between neighboring frames, as can be seen in the lower panel of Fig. 2.

2.3 Visualization and Analysis Software

To render the volumetric data, we used a technique similar to that described by Carbal et al.[19] Every time the volume is rendered, it is sliced into evenly spaced view-aligned planes. Each slice is rendered through a pixel shader using the voxel's data value, a color map, and alpha blending to create a volume rendering that is equivalent to a ray tracer that uses evenly spaced samples. This is done efficiently by implementing a highly optimized cube slicing algorithm and applying it
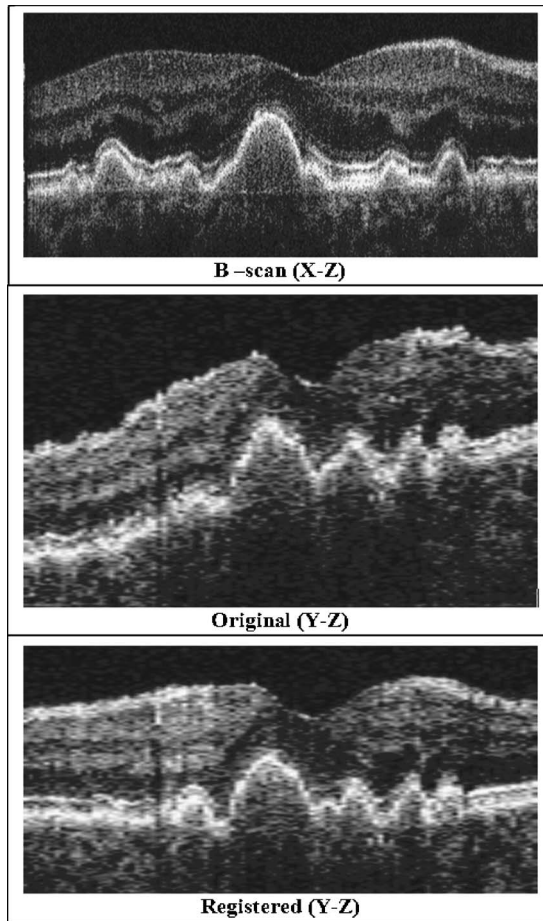
Fig. 2 Registration of the OCT frames. Top image shows one OCT B-scan. Middle image shows cross section through 98 unaligned B-scans. Bottom image shows the scans after registration with ImageJ.

to the bounding box of the volume. We also use adaptive slice spacing to balance performance/functionality and visual quality.

There has been substantial research on classification and segmentation of volumetric data. Most of this work focuses on visualization instead of explicit segmentation. The most common approach used in volume visualization is based on a *one-dimensional transfer function* to map scalar intensity to color and opacity. This effectively allows a region to be classified based solely on scalar intensity, hiding regions occluding areas of interest while highlighting the remaining areas with a chosen color scheme. This approach does not work for noisy data if it results in a region that cannot be classified solely by its scalar intensity (since these values occur somewhat randomly throughout the volume). This approach is also not applicable to volumes that contain volumetric objects of similar intensity values that occlude each other. A great deal of research has attempted to solve these inadequacies. Levoy et al.[20] used gradient magnitude to highlight material boundaries in volumetric data to better express structure. Kindlmann and Durkin[21] proposed the use of a *two-dimensional transfer function*, represented by a scatterplot, of scalar values versus gradient magnitudes. Others have described methods that help a user better utilize this 2D transfer function.[20,22,23] Iserhardt-

Bauer et al.[24] combined a 2D transfer with a region growing method that expands to fill regions of interest. Region growing methods are typically ideal for large meandering regions. Multidimensional transfer functions tend to operate better than lower-dimensional counterparts since the method can use more information to segment the data. The method of Bordignon et al.[25] demonstrates a multidimensional transfer function with a user interface based on star coordinates. This is a technique that maps multidimensional data to a 2D plane so that a user can interact with it. Pfister et al.[26] compared a variety of transfer function methods including some of those listed above. Despite these advances, 2D transfer functions remain inaccessible to the common user, and their effectiveness is still predicated on noiseless data. When it comes to discrete and explicit segmentation, medical imaging research has used artificial neural networks as a means to assist in these tasks.[27,28] However, support vector machines[14,15,29] have demonstrated better results to date than neural networks (at a larger computational cost) when applied to pattern recognition including object identification,[30] text categorization,[31] and face detection.[32] Tzeng et al.[33] compared the use of neural networks and support vector machines when trying to construct an *N-dimensional transfer function* that uses additional variables such as variance and color (if present) in addition to scalar values and gradient information. The ability of neural networks and SVMs to handle error (in the form of noise) makes them useful for noisy volumetric data sets. However, these methods still seek to classify a volumetric object solely based on discrete values, not on distributions. Our approach converts noise into a classifiable characteristic by using a specialized SVM that operates on distributions rather than scalar values.

2.3.1 SVM-based segmentation

Two main design decisions customize an SVM to a particular application. The first is the kernel used to map input (or world) space to feature space. As a consequence, the SVM can process problems that are not linearly separable in input space. In our case we chose a radial basis function kernel:

$$\text{Kernel}(\mathbf{u}, \mathbf{v}) = e^{-\gamma |\mathbf{u} - \mathbf{v}|^2}, \qquad (1)$$

where γ is a constant scaling exponential decay and $\mathbf{u}$ and $\mathbf{v}$ are two vectors in the feature space.

The second design decision to be made is the dimensionality (data characteristics) of our input space. The first (and most obvious) value we add to the input (or feature) vector is scalar intensity. It allows for efficient segmentation of regions having a low standard deviation (noiseless data). We also include the spatial location $\mathbf{r} = \{x, y, z\}$ of a voxel since we want to track spatial locality of data distributions, which allows the SVM to distinguish between features having similar data distribution characteristics but residing in different locations.

The method described by Tzeng et al.,[33] from which ours is derived, suggests that six neighbor intensity values $I(x \pm 1, y \pm 1, z \pm 1)$ should be considered at $\mathbf{r} = \{x, y, z\}$ to counter noise. They suggest that if a voxel value has been perturbed by noise, inclusion of its neighbors will help determine the "actual value." We found this approach to dramatically decrease the effectiveness of the SVM and lead to unintuitive results. We modified this approach to instead use the

mean intensity of the six neighbors. We also included the local variance (instead of the standard deviation) both to let the SVM gauge the accuracy of the mean and to characterize the local distribution. Our variance estimate is accurate as long as the set of neighboring voxels used to compute the variance belong to the same feature (i.e., are not on the border of two or more features, which would give an inaccurate representation of the variance for a particular layer). We detect border regions through a locally approximated gradient magnitude calculation. In summary, for a voxel $\mathbf{r}=\{x,y,z\}$ with scalar intensity value $I(\mathbf{r})$, the data characteristics we use for our input vector are

$$I(\mathbf{r}), \tag{2}$$

$$x, \tag{3}$$

$$y, \tag{4}$$

$$z, \tag{5}$$

$$\langle I(\mathbf{r})\rangle = \frac{1}{|N|}\sum_{n\in N} I(\mathbf{r}_n), \tag{6}$$

$$\sigma I(\mathbf{r}) = \frac{1}{|N|}\sum_{n\in N}(I(\mathbf{r}_n)-\langle I(\mathbf{r})\rangle)^2, \tag{7}$$

$$\langle \mathrm{grad}\,I(\mathbf{r})\rangle = \langle \mathrm{grad}_x I(\mathbf{r}), \mathrm{grad}_y I(\mathbf{r}), \mathrm{grad}_z I(\mathbf{r})\rangle$$
$$= \left|\frac{1}{|N|+1}\{\mathrm{grad}\,I(\mathbf{r})+\sum_{n\in N}\mathrm{grad}\,I(\mathbf{r}_n)\}\right|, \tag{8}$$

where N is a set of neighboring voxels around $\mathbf{r}$ and $\mathbf{r}_n : n \in N$ is a voxel in this set. Thus, for each $\mathbf{r}$ we store its intensity, location of a local approximation of the mean value $\langle I\rangle$, a local approximation of the variance σI, and a local approximation of the gradient $\mathrm{grad}\,I$ using a local difference operator.

2.3.2 User interface and SVM training

The main advantage of using an SVM to perform segmentation is that it does not need to know what type of region it is segmenting a priori. Training data are provided at run time by a user through an intuitive interface, which is a part of our volumetric rendering software, to dynamically create a "segmentation function." Our system allows a user to quickly classify features by using a small number of specification points. We perform this specification on B-scans of the volume, similarly to Tzeng et al.[33] The user can interactively choose a B-scan from the volume and "paint" on that B-scan to mark points as "feature" or "background." The user is required only to draw points on the regions of interest and indicate regions not of interest. The difficulty of obtaining an accurate segmentation is not selecting training points for the SVM, but finding the smallest set of points that provide the best result. This turns into an iterative process of selecting points and testing the result of the SVM segmentation first on the single frame and then on the whole volume. The cost of choosing excessive numbers of points is an increase in processing time by the SVM.

2.3.3 Morphological analysis

The classification provided by the SVM is a partition of the entire volume so we can use it to extract relevant morphological data. For example, one can track the change in the volume of the given structure (e.g., cup of the ONH) that may be useful for retinal or optic nerve diagnosis. Another potentially useful metric is the density of a volumetric object and its standard deviation. Density in an OCT volume represents the intensity per voxel volume; thus, the amount of backscattered light that can be used as a metric characterizing the internal structure of a given retinal layer or structure.

2.3.4 Speed improvements

SVMs have a large computational cost that hinders a user's ability to iteratively create segmentation. We have developed techniques to reduce the computational cost for both SVM training and classification. Our techniques aim to reduce the time needed between iterations for creating and testing an SVM by a clinician. With the first method, the user tests the segmentation on individual slices (B-scans). This method usually requires only a few seconds (even for volumes of $1000 \times 500 \times 200$ resolution). The user can browse remaining slices and apply the trained SVM in order to preview the segmentation. A user can adjust the current SVM by marking misclassified voxels, retraining the SVM, and continuing. Once trained, an SVM classifies each voxel independently of other voxels. Thus, we take advantage of multiprocessor computers by using multithreading in order to significantly decrease computation time for the clinician.

Another effective method restricts the SVM to a subvolume of the data. Our application allows a user to specify axis-aligned clipping planes in order to specify a subregion. This approach can significantly reduce the number of voxels to be processed. Time is saved because a user does not have to train the SVM regarding "background" data existing outside this subvolume. Additionally, smaller training data reduce the mapping complexity of the SVM, resulting in both faster training and classification.

Assuming that objects are relatively large with respect to voxel size, we have implemented a checkerboard scheme. Thus, we only process every other voxel by the SVM. Each unclassified voxel is classified based on its neighbors. If no clear majority class is implied by the neighbors, we apply the SVM to that voxel to decide. This reduces classification time by about 40% and also leads to smoother object boundaries and fewer "orphaned" voxels, i.e., those that are disconnected from larger objects and appear to be noise.

3 Results

Our visualization software is regularly used to render retinal volumes and allows interactive viewing by our clinicians. Due to the use of hardware acceleration in our rendering software (we use common features of modern graphic cards that have been optimized for interactive media applications), this is done in real time at high resolution. We tested our SVM segmentation software on a variety of retinal volumes and retinal
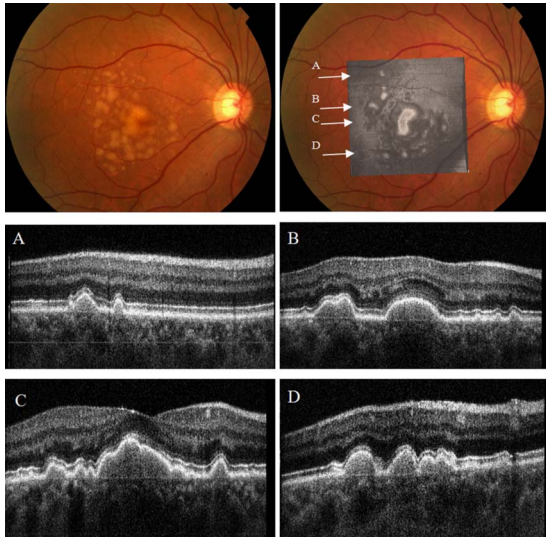
Fig. 3 Top left: fundus photo from a patient with non-exudative AMD; top right: same fundus photo with a superimposed C-scan reconstructed from the center depth on OCT volume. Arrows indicate the relative positions of the B-scans shown below.
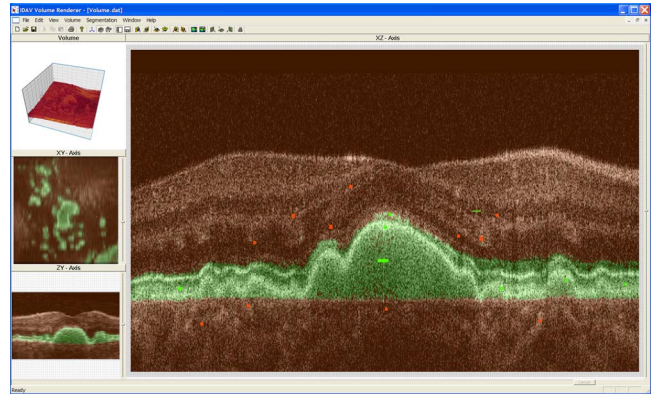

Fig. 4 User interface of volume visualization software with SVM training data (green for structure of interest and red for background) shown on the B-scan. The result of the frame segmentation based on these inputs can be seen on the B-scan. The result of the volume segmentation is marked green on the 2D cross sections; extracted volume is shown in the upper left window panel (color online only).

structures that have been segmented. SVM training and segmentation time varies dramatically depending on the complexity of the segmented layer or the feature, the number of voxels in the volume, and the accuracy of segmentation. Total segmentation time for a whole volume performed on a PC workstation with two Intel Xeon 3.6-GHz processors and 2 GB of main memory can be as short as a few minutes for small volumes and well-defined features to as long as 2 hours for large volumes and complex structures. Generally, the more training points needed to extract a feature, the more time-consuming is the segmentation. Thus, it is important for the operator to learn how to efficiently train the SVM to reduce time needed for segmentation. In order to quantify the accuracy of the segmentation, our visualization software has a built-in manual segmentation option, where the operator can draw the borders of the segmented feature on each B-scan. This process may be very time-consuming. However, it can be used to create a "reference" segmentation that can be directly compared to an SVM segmentation for verification purposes. Additionally, we can extract morphological data that have known normal values, such as retinal layer thickness maps or retinal layer profiles that can be directly used as a diagnostic tool in ophthalmology. Moreover, we can use our software to create z-axis projection intensity maps of extracted layers, allowing visualization of specific retinal features.

3.1 Visualization and Analysis of Volumetric Data

To illustrate the performance of our clinical Fd-OCT and our 3D data visualization and analysis software, we present four examples of visualization and analysis of the clinical data.

The first example is data acquired from a 55-year-old patient with non-exudative age-related macular degeneration (AMD). Figure 3 shows the patient's fundus photo used for regular clinical diagnosis. The top right image shows a reconstructed *en face* image (virtual C-scan) from the OCT volume (100 B-scans/1000 A-scans/B-scan acquired with 100-μs/line

exposure) superimposed on the fundus picture. This step allows precise registration of the acquired volume and estimation of the distortions caused by the subject's eye motions during the experiment.

To demonstrate the performance of our Fd-OCT system, four B-scans (labeled A–D) are shown in the lower part of Fig. 3. The location of the diseased structures can be easily seen in these images. As described above, this data set is used to create the interactive volumetric visualization using our software. Figure 4 shows a screen shot of our volume renderer and segmentation system. Note the points used to define the SVM segmentation.

Figure 5 shows an example of different 3D visualizations of this data set. We used a false color intensity based on a black-red-yellow color lookup table with black representing
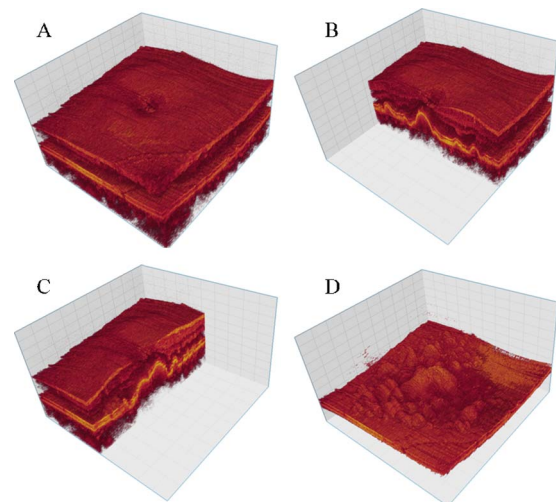

Fig. 5 3D visualization of the volumetric data shown in Fig. 4. (a) whole volume; (b) left part of the volume removed by x-z clipping plane; (c) left part of the volume removed by y-z clipping plane; (d) visualization of segmented part of the volume RPE and photoreceptor-outer segment complex.
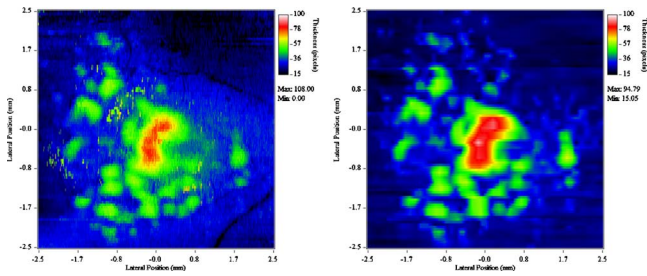
Fig. 6 Left: false color representation of the thickness map extracted from SVM-based segmentation of the RPE and photoreceptor-outer segment complex. Right: false color representation of the thickness map created from manual-based segmentation of the same structure (color online only).

Fig. 7 Evaluation of the volumetric data set acquired over foveal region of a volunteer with early stage dry AMD. Upper left: 3D rendering of the data; upper right: screen shot from the user interface of the SVM-based segmentation software with segmented structure (RPE + photoreceptors outer segments) highlighted in green; middle left projected on z-axis intensity of the original volume; bottom left projected on z-axis intensity of the SVM segmented structure (RPE + photoreceptors outer segments); bottom right: false color thickness map of the SVM segmented layer (color online only).

low-intensity voxels and yellow representing high-intensity voxels. Transparency of each voxel is based on its intensity and can be interactively set by the operator.

To test the performance of our SVM segmentation algorithm, we used it to segment the retinal pigmented epithelium (RPE) and photoreceptor layers, structures that show the main changes associated with disease progression in AMD. The bumpy surface is due to large drusen, a defining feature of early stage AMD. A thickness map of these segmented structures is created from the SVM segmentation. Figure 6 shows the corresponding thickness maps from the SVM segmentation and from manual segmentation. As can be seen from the images in Fig. 6, the SVM method leads to clear separation of these layers, allowing its visualization as well as estimation of the thickness of the area of interest. The differences between the SVM and manual segmentations suggest that some refinement of the SVM may be necessary to better segment high density structures.

Figure 7 shows another SVM-based segmentation of the RPE-photoreceptor-outer segment complex from the retinal foveal region of a 56-year-old patient with early-stage dry AMD. In contrast to the first patient presented, the RPE-photoreceptor-complex disruptions are subtle. Figure 7 shows results generated with our custom visualization software (including z-axis intensity projection).

Our algorithm was able to segment these structures. However, since the size of these disrupted regions was small and there was intensity variation across the volume, mainly at the corners of the image, not all drusen were identified in the thickness map. When the size of disruption is bigger, as in the previous example, the SVM segmentation algorithm has been able to pick out the features of interest very accurately.

In the next two examples, we tested the SVM segmentation algorithm in segmenting the RNFL around the ONH. Thinning of RNFL is known to be a good indicator of possible onset of glaucoma; therefore, being able to accurately segment and measure its thickness and follow it over time would be a good diagnosis and monitoring tool in glaucoma management. Figure 8 shows the segmentation results of the ONH from a healthy 30-year-old volunteer with our SVM segmentation software.

We also segmented the same structure in a 48-year-old glaucoma patient with a moderate amount of visual field loss. The results are shown in Fig. 9.
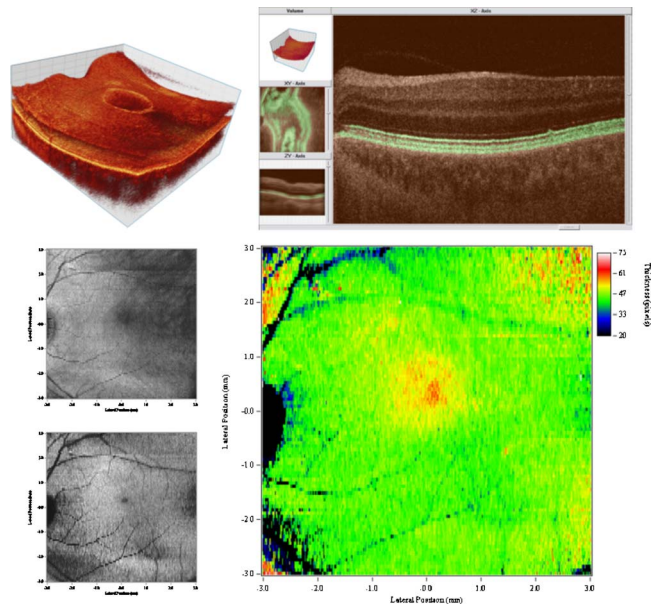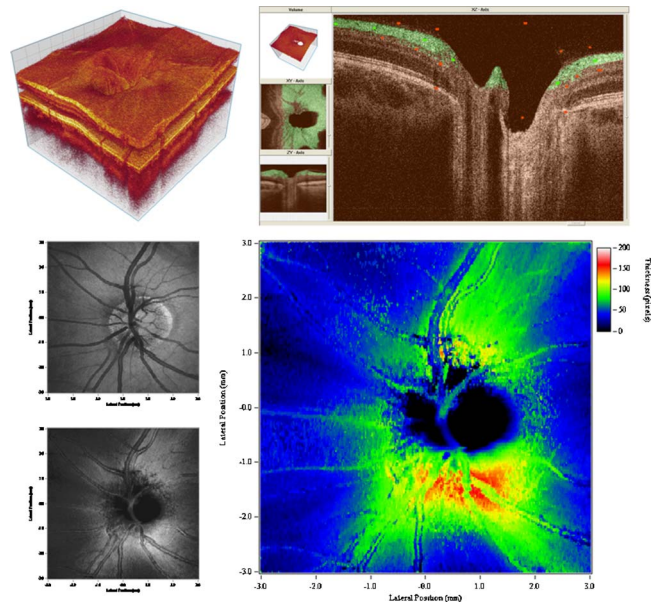

Fig. 8 Evaluation of the volumetric data set acquired over ONH region of the healthy volunteer. Upper left: 3D rendering of the data; upper right: screen shot from the user interface of the SVM-based segmentation software with segmented structure (RNFL) highlighted in green; middle left projected on z-axis intensity of the original volume; bottom left projected on z-axis intensity of the SVM-segmented structure (RNFL); bottom right: false color thickness map of the RNFL extracted from the SVM segmentation (color online only).
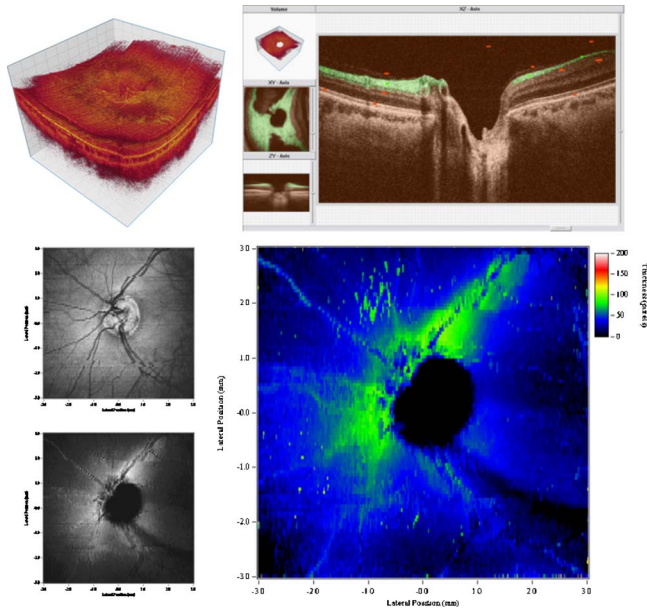
Fig. 9 Evaluation of a volumetric data set acquired over the ONH region of a volunteer with moderate glaucoma. Upper left: a 3D rendering of the data; upper right shows a screen shot from the user interface of the SVM-based segmentation software with segmented structure (RNFL) highlighted in green; middle left projected on *z*-axis intensity of the original volume; bottom left projected on *z*-axis intensity of the SVM-segmented structure (RNFL); bottom right; false color thickness map of the RNFL extracted from the SVM segmentation (color online only).

As can be seen from these last two examples, SVM-based segmentation was able to successfully differentiate RNFL from other retinal structures, allowing visualization of the RNFL thickness map that shows thinning of this layer for the subject with glaucoma. It confirms good performance of the SVM-based segmentation for thick and well-defined structures.

4 Conclusion and Discussion

A major limitation in segmentation of OCT volumes is that image intensity can vary depending on B-scan location and due to the shadowing effects caused by blood vessels. These effects are visualized in Fig. 10, where blood vessels do not allow much light to pass through them, which means that data
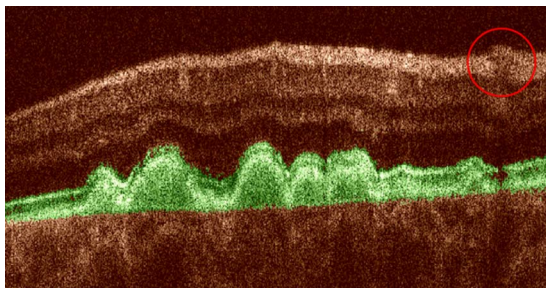

Fig. 10 OCT B-scan with segmented RPE-photoreceptor layers; red circle indicates blood vessels that distort the intensity values around and below them, disturbing the SVM segmentation (color online only).

in regions around and below them are obscured, distorted, or occluded. Our visual system can cope with this by recognizing global patterns and extrapolating them into these regions. But as stated earlier, SVM classification remains a local operation and is unable to "see" global structure. This makes segmentation in these areas difficult and inaccurate. Recently Mujat et al.[9] have presented a fully automated segmentation method that can overcome this problem—however, the computational cost is larger, resulting in *frame* segmentation time similar to our *volume* segmentation time. The only way to cope with these anomalies in the SVM is to specify a large amount of additional training data. When spending only a small amount of time generating training data, these inaccuracies are pronounced enough to make any morphological data extracted from these regions useful only as a first approximation. Filters partially alleviate this problem, but a better solution is desirable. We plan to direct our attention to this problem in the future, using SVM-based segmentation methods.

It should be noted that when one performs analyses using morphological data obtained from SVM segmentation, the quality of a segmentation depends heavily upon the training data and SVM quality (i.e., how well the SVM can predict based on the given input data).

We will refine our method to take into account certain global information such as, in the case of OCT retinal data, the number of layers that make up the retina and the fact that these layers normally do not intersect one another. Using such known constraints should further improve segmentation results.

References

1. M. Wojtkowski, T. Bajraszewski, P. Targowski, and A. Kowalczyk, "Real time *in vivo* imaging by high-speed spectral optical coherence tomography," *Opt. Lett.* **28**, 1745–1747 (2003).
2. R. Leitgeb, C. K. Hitzenberger, and A. F. Fercher, "Performance of Fourier domain vs. time domain optical coherence tomography," *Opt. Express* **11**, 889–894 (2003).
3. N. A. Nassif, B. Cense, B. H. Park, M. C. Pierce, S. H. Yun, B. E. Bouma, G. J. Tearney, T. C. Chen, and J. F. de Boer, "*In vivo* high-resolution video-rate spectral-domain optical coherence tomography of the human retina and optic nerve," *Opt. Express* **12**, 367–376 (2004).
4. M. Wojtkowski, V. Srinivasan, J. G. Fujimoto, T. Ko, J. S. Schuman, A. Kowalczyk, and J. S. Duker, "Three-dimensional retinal imaging with high-speed ultrahigh-resolution optical coherence tomography," *Ophthalmology* **112**, 1734–1746 (2005).
5. U. Schmidt-Erfurth, R. A. Leitgeb, S. Michels, B. Povazay, S. Sacu, B. Hermann, C. Ahlers, H. Sattmann, C. Scholda, A. F. Fercher, and W. Drexler, "Three-dimensional ultrahigh-resolution optical coherence tomography of macular diseases," *Invest. Ophthalmol. Visual Sci.* **46**, 3393–3402 (2005).
6. E. Götzinger, M. Pircher, and C. K. Hitzenberger, "High speed spectral domain polarization sensitive optical coherence tomography of the human retina," *Opt. Express* **13**, 10217–10229 (2005).

7. R. Zawadzki, S. Jones, S. Olivier, M. Zhao, B. Bower, J. Izatt, S. Choi, S. Laut, and J. Werner, "Adaptive-optics optical coherence tomography for high-resolution and high-speed 3D retinal *in vivo* imaging," *Opt. Express* **13**, 8532–8546 (2005).
8. Y. Zhang, B. Cense, J. Rha, R. Jonnal, W. Gao, R. Zawadzki, J. Werner, S. Jones, S. Olivier, and D. Miller, "High-speed volumetric imaging of cone photoreceptors with adaptive optics spectral-domain optical coherence tomography," *Opt. Express* **14**, 4380–4394 (2006).
9. M. Mujat, R. Chan, B. Cense, B. Park, C. Joo, T. Akkin, T. Chen, and J. de Boer, "Retinal nerve fiber layer thickness map determined from optical coherence tomography images," *Opt. Express* **13**, 9480–9491 (2005).
10. D. Cabrera Fernández, H. M. Salinas, and C. A. Puliafito, "Automated detection of retinal layer structures on optical coherence tomography images," *Opt. Express* **13**, 10200–10216 (2005).
11. S. Makita, Y. Hong, M. Yamanari, T. Yatagai, and Y. Yasuno, "Optical coherence angiography," *Opt. Express* **14**, 7821–7840 (2006).
12. E. C. Lee, J. F. de Boer, M. Mujat, H. Lim, and S. H. Yun, "*In vivo* optical frequency domain imaging of human retina and choroid," *Opt. Express* **14**, 4403–4411 (2006).
13. S. R. Sadda, Z. Wu, A. C. Walsh, L. Richine, J. Dougall, R. Cortez, and L. D. LaBree, "Errors in retinal thickness measurements obtained by optical coherence tomography," *Ophthalmology* **113**, 285–293 (2006).
14. B. E. Boser, I. Guyon, and V. Vapnik, "A training algorithm for optimal margin classifiers," in *Proc. Fifth Ann. Workshop Computational Learning Theory*, pp. 144–152 (1992).
15. C. Cortes and V. Vapnik, "Support vector network," *Mach. Learn.* **20**, 273–297 (1995).
16. M. Wojtkowski, V. Srinivasan, T. Ko, J. Fujimoto, A. Kowalczyk, and J. Duker, "Ultra high-resolution high-speed Fourier domain optical coherence tomography and methods for dispersion compensation," *Opt. Express* **12**, 2404–2422 (2004).
17. W. S. Rasband, Image J, U.S. National Institutes of Health, Bethesda, Maryland, http://rsb.info.nih.gov/ij/, 1997–2006.
18. http://bigwww.epfl.ch/thevenaz/turboreg/.
19. B. Carbal, N. Cam, and J. Fornan, "Accelerated volume rendering and tomographic reconstruction using texture mapping hardware," in *VVS '94: Proc. 1994 Symp. Volume Visualization*, ACM Press, New York, pp. 91–98 (1994).
20. M. Levoy, "Display of surfaces from volume data," *IEEE Comput. Graphics Appl.* **8**, 29–37 (1988).
21. G. Kindlmann and J. W. Durkin, "Semi-automatic generation of transfer functions for direct volume rendering," in *IEEE Symp. Volume Visualization*, pp. 79–86 (1998).
22. J. Kniss, G. Kindlmann, and C. Hansen, "Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets," in *Proc. IEEE Visualization '01 Conf.*, pp. 255–262 (2001).
23. H. O. Shin, B. King, M. Galanski, and H. K. Matthies, "Use of 2D histograms for volume rendering of multidetector CT data: Development of a graphical user interface," *Acad. Radiol.* **11**(5), 544–550 (2004).
24. S. Iserhardt-Bauer, P. Hastreiter, B. F. Tomandl, and T. Ertl, "Evaluation of volume growing based segmentation of intracranial aneurysms combined with 2D transfer functions," in *Proc. SimVis 2006*, pp. 319–328 (2006).
25. A. L. Bordignon, R. Castro, H. Lopes, T. Lewiner, and G. Tavares, "Exploratory visualization based on multidimensional transfer functions and star coordinates," 19th Brazilian Symp. Comp. Graph. Image Process. (2006).
26. H. Pfister, B. Lorensen, C. Bajaj, G. Kindlmann, W. Schroeder, L. S. Avila, K. M. Raghu, R. Machiraju, and J. Lee, "The transfer function bake-off," *IEEE Comput. Graphics Appl.* **21**(3), 16–22 (2001).
27. E. Gelenbe, Y. Feng, K. Ranga, and R. Krishnan, "Neural networks for volumetric MR imaging of the brain," in *Proc. Int. Workshop Neural Networks for Identification, Control, Robotics, and Signal/Image Processing*, pp. 194–202 (1996).
28. L. O. Hall, A. M. Bensaid, L. P. Clarke, R. P. Velthuizen, M. S. Silbiger, and J. C. Bezdek, "A comparison of neural network and fuzzy clustering techniques in segmenting magnetic resonance images of the brain," *IEEE Trans. Neural Netw.*," **3**, 672–682 (1992).
29. K. R. Muller, S. Mika, and G. Ratsch, "An introduction to kernel-based learning algorithms," *IEEE Trans. Neural Netw.* **12** (2001).
30. V. Blanz, B. Scholkopf, H. Bulthoff, C. Burges, V. Vapnik, and T. Vetter, "Comparison of view-based object recognition algorithms using realistic 3D models," in *Proc. Int. Conf. Artificial Neural Networks*, pp. 251–256 (1996).
31. J. Thorsten, "Text categorization with support vector machines: Learning with many relevant features," in *Proc. 10th Euro. Conf. Machine Learning*, pp. 137–142 (1998).
32. E. Osuna, R. Freund, and F. Girosi, "Training support vector machines: An application to face detection," in *Proc. '97 Conf. Comp. Vis. Pat. Recog.*, pp. 130–137 (1997).
33. F.-Y. Tzeng, E. B. Lum, and K.-L. Ma, "An intelligent system approach to higher-dimensional classification of volume data," *IEEE Trans. Vis. Comput. Graph.* **11** (2005).