

3D OCT imaging in clinical settings: toward quantitative measurements of retinal structures

Robert J. Zawadzki^{*a}, Alfred R. Fuller^b, Mingtao Zhao^c, David F. Wiley^b, Stacey S. Choi^a,
Bradley A. Bower^c, Bernd Hamann^b, Joseph A. Izatt^c and John S. Werner^a

^a Vision Science and Advanced Retinal Imaging Laboratory, Department of Ophthalmology & Vision Science,
UC Davis, 4860 Y Street, Suite 2400, Sacramento, CA, USA 95817;

^b Visualization and Computer Graphics Research Group, Institute for Data Analysis and Visualization (IDAV),
UC Davis, One Shields Avenue, Davis, CA 95616;

^c Department of Biomedical Engineering, Duke University, 101 Science Drive,
Durham, North Carolina 27708

ABSTRACT

The acquisition speed of current FD-OCT (Fourier Domain - Optical Coherence Tomography) instruments allows rapid screening of three-dimensional (3D) volumes of human retinas in clinical settings. To take advantage of this ability requires software used by physicians to be capable of displaying and accessing volumetric data as well as supporting post processing in order to access important quantitative information such as thickness maps and segmented volumes.

We describe our clinical FD-OCT system used to acquire 3D data from the human retina over the macula and optic nerve head. B-scans are registered to remove motion artifacts and post-processed with customized 3D visualization and analysis software. Our analysis software includes standard 3D visualization techniques along with a machine learning *support vector machine* (SVM) algorithm that allows a user to semi-automatically segment different retinal structures and layers. Our program makes possible measurements of the retinal layer thickness as well as volumes of structures of interest, despite the presence of noise and structural deformations associated with retinal pathology. 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.

Keywords: Optical Coherence Tomography, Image Processing, Ophthalmology, Optical Diagnostic for Medicine, Segmentation, Feature Finding

1. INTRODUCTION

Over the last decade, OCT [1-2] has become an essential tool in ophthalmology. New generations of OCT systems acquiring data in the Fourier domain [3-5] offer sensitivity and speed advantages over other standard OCT techniques working in the time domain, thereby making possible *in vivo* real-time and ultra-high resolution imaging. The speed advantage of FD-OCT systems for retinal volumetric measurements as well as first attempts for structure segmentation have been recently described by Wojtkowski et al. [6] and Mujat et al. [7]. Volumetric representations of the retina, as in other medical areas, have become an increasingly useful tool for ophthalmic clinical diagnosis and medical research. Standard OCT (logarithmic scale intensity-based) data sets are represented by a 3D rectilinear grid of scalar values (amount of back-scattered light stored for each voxel). Automated approaches for the classification and segmentation of these data sets aid a practitioner in isolating and examining areas of interest within these volumes. There are two main stumbling blocks with automated approaches. The first is that patient retinal structure can vary substantially among patients. Despite the challenges, interesting information can be gathered from the examination of diseased subjects. Thus, building a system that only handles the ideal case is of limited use. The other main stumbling block stems from inaccuracies in the data acquisition systems. OCT scans are composed of a “noisy” speckle field superimposed on imaged structures. This makes it exceedingly difficult to isolate regions solely based on scalar intensities. In addition, the term “area of interest” describes a vaguely defined concept. What the clinician is interested in isolating can greatly depend on the preference of the clinician and the disease in question.

*email: rjzawadzki@ucdavis.edu; phone: 1-916-734-5839; fax: 1-916-734-4543; web: <http://vsri.ucdavis.edu>

We introduce a new analysis tool for a volumetric data set that accounts for these complications. Instead of being a rectilinear grid of scalar values, we think of a volumetric data set as the superposition of (localized) normal distributions. This concept then models noise as a characteristic of a given distribution and allows us to use this characteristic in our classification algorithm. With this new understanding, we describe a machine learning approach for the classification and segmentation of an arbitrary *volumetric object*, defined as any coherent structure described by the superposition of a subset of localized distributions. We also describe how a clinician can easily use our tools to isolate and examine a volumetric object of interest.

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 last year, more than 80 subjects, with healthy and diseased eyes, have been imaged with this system. This paper focuses on the results obtained with our 3D scanning procedures. Due to involuntary eye motion and reduced (or distorted) intensity of the some OCT images, not all of the 3D scans are appropriate for volumetric reconstruction; the main factors being (especially among elderly patients) advanced cataract, significant eye aberrations, long eye lashes and ptosis. As a standard method of qualifying retinal volumes 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 estimate distortions caused by eye motion. Due to these problems, we have concentrated our efforts on a segmentation method that tolerates noisy data.

2.1 Experimental system

The experimental FD-OCT system used for acquiring 3D volumes was previously described [8]. For our experiments described here, a new Superluminescent diode (855@75nm) was used. The spectral bandwidth of this light source allows a 4.5 μm axial resolution of the retina ($n=1.38$). High power of the new light source allows use of a 90/10 fiber coupler, directing 10% of the light toward the eye resulting in the power at the subject's eye equal to 700 μW and allowing 90% throughput of the light back-reflected from the eye to the detector. With our current spectrometer design, the maximum axial range (seen after the Fourier transform) is 2.7 mm in free space and around 2 mm in the eye. Figure 1 shows a schematic of our Clinical FD-OCT system.

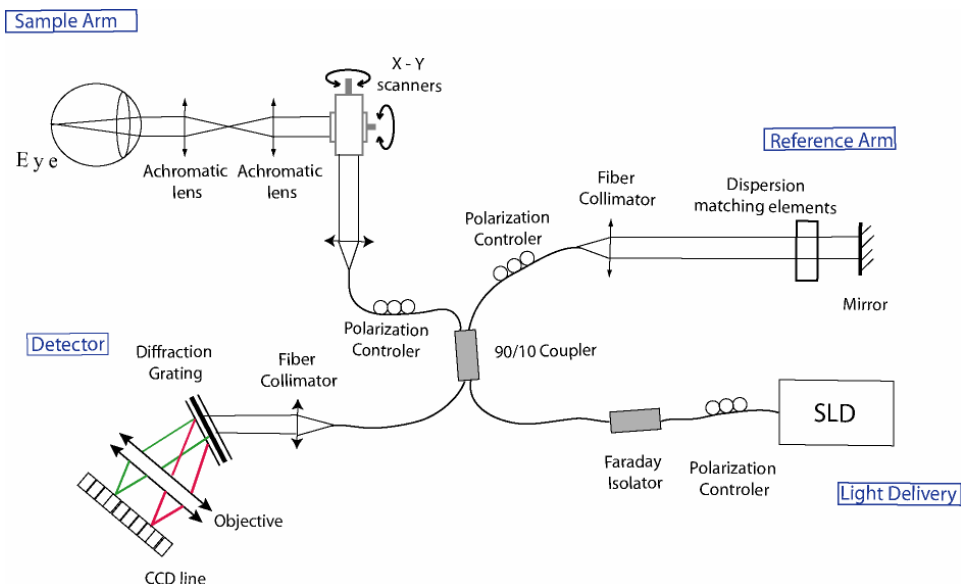


Figure 1: Clinical FD-OCT system at UC Davis. The reference arm matches the path length of the sample arm.

Each subject was imaged with several OCT scanning procedures including 5 mm vertical B-scans (1000 A-scans/B-scan), 8 mm vertical B-scans (1000 A-scans/B-scan), tiled by 5 degrees to scan through fovea and ONH, two 6 mm (1000 A-scans/B-scan) radial scanning patterns centered at the fovea and ONH and 4-6 mm 3D scanning patterns also

centered at the fovea and ONH. Figure 2 shows the relative locations of these patterns superimposed on a fundus image. We used two different scanning arrangements for 3D scanning patterns including 200 B-scans, where each was based on 500 A-scans, and 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 μ s CCD exposure time and 11.1s for 100 μ s exposure time. The longer exposure time was mainly used to increase image intensity, at the cost of more motion artifacts.




Figure 2: The OCT scanning patterns and scanning areas for 3D volumetric measurements used during our standard data acquisition procedure superimposed on fundus images. Left: 5 mm and 8 mm line scanning patterns; center: 6 mm radial scanning patterns; right: 6 mm 3D scanning areas.

Our acquisition software permitted real-time display (at the acquisition speed) of the imaged retinal structures and saved the last 100 B-scans in the line and radial acquisition modes. In volumetric acquisition mode, the program was set to save the last acquired volume.

To visualize and analyze our 3D OCT data, we created a set of consecutive OCT B-scans saved as 2D images in tiff file format. This was later used as the input to our semi-automated registration tool that creates one file containing the whole volume, viewed subsequently in our visualization and analysis tool. It was possible to analyze cases 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 was removed resulting in a reduction in the number of B-scans and overall size of the reconstructed 3D volume.

2.2 Analysis software

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 since in this case a region cannot be solely classified 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. [9] used gradient magnitude to highlight material boundaries in volumetric data in order to better express structure. Kindlmann and Durkin [10] proposed the use of a *two-dimensional transfer function*, represented by a scatter plot, of scalar values versus gradient magnitudes. Other scientists have described methods that help a user better utilize this 2D transfer function [9,11]. Despite these advancements, 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 [12,13]. However, support vector machines (SVM) [14-16] have shown far better results than neural networks (at a larger computational cost) when applied to pattern recognition including object identification [17], text categorization [18], and face detection [19]. Tzeng et al [20] 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 distributions. Our approach converts noise into a classifiable characteristic by using a specialized SVM that operates on distributions rather than scalar values.

2.2.1 SVM-based segmentation

Machine learning methods seek to bridge the gap between a human's and a machine's abilities by learning and adapting to given stimuli. Concerning volumetric segmentation, it is a relatively easy task for a human to identify patterns and structures regardless of noise, deformation, or structural "oddities." In order to design a machine to be similarly effective, we must first understand how the human brain handles this class of stimuli. The human brain has the ability to see an object macroscopically. Given noise, our brain can adaptively classify a volumetric region by its general distribution instead of by its individual pixel intensities and, when confronted with abrupt changes in these distributions, a human still perceives the associated region borders. Classic SVM implementations do not have the ability to classify based on these characteristics. An SVM operates on one individual pixel at a time, with each operation completely independent of every other. While this paper does not propose a method to enable an SVM to see and extend global patterns, the described approach does enable an SVM to classify a region based on its distribution and the discrete distribution changes on the region's boundaries. By doing this, we enable an SVM to characterize a region in a volume following a human's approach.

The main design constraint needed to specialize an SVM to a given task involves the dimension of the input vector (or stimuli). This requires a concrete model of the input space. We view a volume as the superposition of localized normal distributions. Each object in this volume is then a subset of localized normal distributions, and it is the SVM's responsibility to approximate a function that can isolate these specific sets from the overall superposition. The first (and most obvious) value to insert into the input vector is intensity $I(\mathbf{r})$, $\mathbf{r} = (x, y, z)$. This allows for the quick and easy segmentation of regions with very low standard deviation (noiseless data) and is the most prominent characteristic of any volumetric data set. The input vector also must contain the 3D spatial location of a given point \mathbf{r} since we claim that these distributions are localized. This allows the SVM to distinguish between objects with similar distribution characteristic by their spatial location.

Tzeng [13] suggested that a set of six-neighbor intensity values should be used to counter noise or intensity value anomaly. The idea is that if a function value is perturbed due to noise, the majority of the neighboring values should contain a more accurate value. Experimentally, this can lead to some unexpected results. By having these six-neighbor values as separate dimensions in our input vector, we encourage the creation of functions that have clauses similar to: "If the value to the right of me is X then I am in class A." This has the tendency to lead to small, scattered and seemingly inconsistent segmentations. Since our view of the volumetric data is mainly concerned with characterizing a distribution, we reinterpret the concept of "If my function value is off due to noise, most likely the majority of my neighbors have a more accurate value" to mean that "The mean of my neighbors approximates my actual value." With this assumption, we avoid undesirable clauses in our SVM's solution space by supplying the mean value of a given point's neighbors instead

of their individual values. We denote this mean value as $\langle I(\mathbf{r}) \rangle = \frac{1}{|N|} \sum_{n \in N} I(n)$, where the set N consists of a finite

number of voxels in the neighborhood of point \mathbf{r} .

A normal distribution is a function of two variables, the mean and the standard deviation. We are already including a localized estimate of the mean in our input vector with our new neighbor computation, but we still need to include the standard deviation. To reduce computational cost, we include a localized estimate of the variance instead of the standard deviation. If these values were the real mean and variance, this would be enough to classify an object based solely on its characteristic distribution. However, we can only perform a localized approximation based on the values of adjacent voxels. This approximation is accurate as long as we assume that objects are larger than a few voxels. Surprisingly, by making this local approximation, we found a better way for classifying volumetric objects by using the

value $\sigma I(\mathbf{r}) = \frac{1}{|N|} \sum_{n \in N} (I(n) - \langle I(\mathbf{r}) \rangle)^2$. The reason for this is two fold. First, data acquisition is imperfect, and the

measured intensity of a given object can change during the time when scanning the subject even if the object itself is uniform. This can occur due to spatial inaccuracies in the measuring device, changes in relative intensities in the data (due to normalization), or movement of the patient while scanning. Second, real objects cannot normally be characterized by a single distribution. Any drastic change in a distribution usually signifies a border between objects, but it is not uncommon for a single object's distribution to change slightly over the space it encompasses. By basing our input vector on localized distribution characteristics, we enable our SVM to classify a given object despite these inconsistencies. The SVM can now base its segmentation on ranges of distributions in given locations. It can even be based on how these ranges change with respect to space, for example, an object with a linear change in intensity. However, these localized approximations are not valid on the borders of these regions.

When creating an input vector for a single voxel in a volume, we can have two problem cases, one where every adjacent point is in the same (volumetric) object, and another where some neighboring points are in other (volumetric) objects. When the latter is the case, our local approximated mean and variance values are completely inaccurate. To enable the SVM to segment these regions intelligently, we need to provide it with an additional dimension that can distinguish these two cases by detecting object boundaries. Edge detection is usually accomplished by considering gradient magnitude at a voxel. Our method is similar, but we have to take into account the effects of noise in our gradient estimation scheme. We model gradient noise as a vector pointing in a random direction where the distributions of direction and magnitude have a zero-vector mean. By taking a larger voxel neighborhood into account than standard gradient calculations (usually only sampling six to 27 neighboring values), we find a localized approximation for the mean gradient at a given voxel, which is the *mean gradient + mean noise vector*. Since the mean noise vector is zero, this is equal to an approximation of the actual gradient. Our gradient estimate is defined as

$$\langle grad(r) \rangle = \langle gx(r), gy(r), gz(r) \rangle = \frac{1}{|N|+1} \left\{ grad(r) + \sum_{n \in N} grad(n) \right\}. \text{ Thus, the gradient (vector) estimate}$$

defines the last three dimensions to our input vector \mathbf{SV} , which becomes

$$SV = \{I(r), r, \langle I(r) \rangle, \sigma I(r), \langle grad(r) \rangle\}. \quad (1)$$

Another concern in the design of an SVM is the choice of the kernel that maps input space to feature space. An SVM finds a separating hyper-plane that maximizes the distance from the plane to the closest training data (known as *support vectors*). Since this hyper-plane is a shear plane, it can only classify linearly separable datasets. To overcome this deficiency, SVMs use an adjustable kernel that maps input vectors from possibly non-linearly separable input space to an adjustably larger-dimensional feature space. The more dimensions defining the feature space the slower the SVM. For our support vector machine, we chose to use a radial basis function. A radial basis function allows small circular sub-regions to be linearly separable in feature space. Since our vision of our volumetric object is a superposition of a localized distribution, this is an appropriate basis function. Additionally, we can extend the SVM model to handle the segmentation of multiple objects or single objects with discrete changes in their characteristic distributions. This is accomplished by using multiple SVMs. In the case of multiple objects, each object has its own SVM. In the case of an object with multiple characteristic distributions, we apply multiple SVMs to the single object.

2.2.2 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. The SVM uses training data to dynamically create a segmentation function. A user specifies the training data at runtime. We want to easily and quickly classify the fewest number of points possible to generate a “good” segmentation. To simplify the classification process and the task of specifying a 3D location with a 2D input device (a mouse), we perform this classification on 2D slices of the volume. The user can scroll a 2D plane through the volume and see a 2D image slice. The user can then “paint” on that plane to mark points as “object” or “background.” The user is required only to draw a green line through regions recognized as belonging to the desired object and “scribble out” regions to be excluded with a red line. Figure 3 shows our system’s interface having the nerve fiber layer segmented from the optic nerve head volume.

Through this user interface, the difficulty of obtaining a good segmentation is not based on selecting training voxels for the SVM, but rather in finding the smallest set of points that provides the best result. An iterative process of selecting points and evaluating the resulting segmentation accomplishes this. To make the search for good training data as effective as possible, it is critical to reduce the iteration time of the search. Unfortunately, SVMs are computationally expensive. The first and most obvious method to reduce iteration time is to allow the user to test the segmentation on a slice. Depending on resolution, this usually requires only a few seconds. This also allows the user to browse the volume and apply the trained SVM on various slices to get an idea of what the overall segmentation will look like. As the user marches through the volume in this manner, they can also make adjustments to the current segmentation by marking misclassified points. Another possible acceleration strategy is to make the segmentation process multi-threaded for computers supporting parallelization. Parallelization is possible since, after training the SVM, each classification of a point is independent.

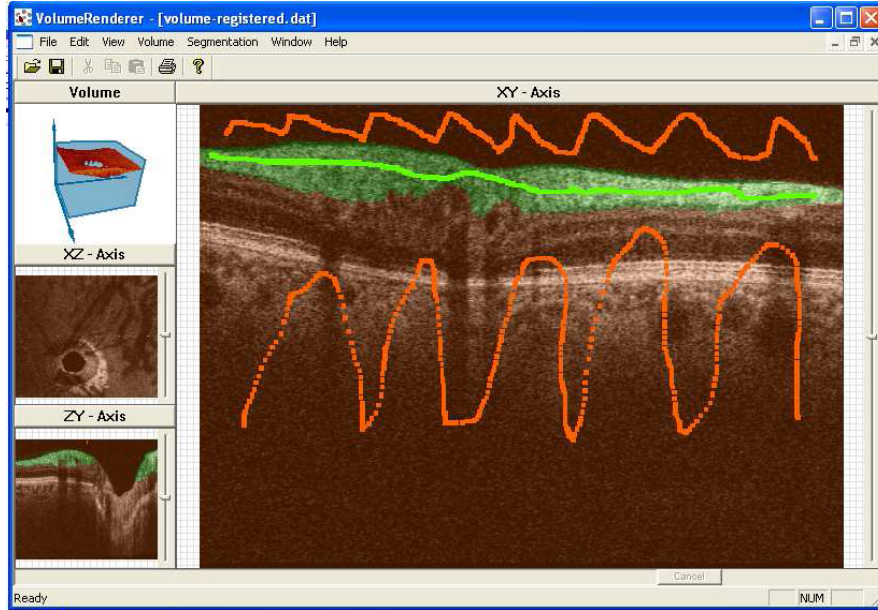


Figure 3: The user “paints” green lines with a mouse in a 2D slice to indicate points that belong to the desired object. Red lines are drawn over regions to indicate points that are not part of this object. These “annotated” data are then used to train an SVM, which classifies and segments a single volumetric object. Here, the nerve fiber layer of a retina as indicated by the green halo region around the green line. The segmentation is shown in the upper-left window pane.

Additionally, since the order in which these points are processed does not matter, optimizing memory access can further improve performance. It is also possible to reduce the number of points that need to be processed by specifying an area of interest. Often, a volumetric object exists only in a sub-region of the volume. Our application allows the user to specify an axis-aligned bounding box with clipping planes to isolate this sub-region. This localizes the SVM to a specific region of the volume and significantly reduces the number of points to be processed. Additionally, because we do not have to train the SVM to mask out extraneous structures that exist outside this bounding box, it reduces the amount of training data needed for a good segmentation. A smaller training data set tends to reduce the complexity of the mapping from input space to the higher-dimensional (but linearly separable) feature space used in the SVM, and results in both faster training of the SVM and faster classification. Assuming that objects are relatively large in relation to the size of a voxel, we have implemented three other optimizing techniques to increase the speed of the SVM.

The first technique is based on a checkerboard scheme using the principle “If all my neighbors are of class A, then I probably am of class A too.” We only process every other pixel with the SVM. Once this step is completed we consider the unprocessed pixels and set their classification to that which is most dominant in a set of neighbors. If there is not a clear “winning class” in this classification step, we apply the SVM to that point to decide. This reduces classification time of an entire volume by about 40%. This also leads to smoother boundaries between objects and less “orphaned” voxels in the middle of objects. Figure 4 shows an example of the 2D checkerboard algorithm.




Figure 4: Using a 2D checkerboard algorithm, we claim that the center pixel in the left image has a 75% probability of belonging to class A, while the center pixel in the right image has a 50% probability of belonging to class A or B. The center pixel in the left image is classified as A. The center pixel in the right image is classified by applying the SVM to it.

Another option for acceleration is to reduce the resolution of the input space and perform classification at a lower resolution. We then map that classification back to the original high-resolution data. Our application accomplishes this reduction in resolution through wavelet compression [21-23]. The compression step has the additional effect of reducing the amount of noise in the volume data and gives rise to both a smaller training data set and faster classification. The main drawback of this method is that it produces jagged segmentation regions, and an additional smoothing step may be required. Applying a smoothing filter to the data is another technique that produces similar results but with greater granularity. Our application currently implements a basic 3D low-pass filter. This has the effect of reducing the complexity of the volume, and in doing so reduces both the size of the training data set and complexity of the resulting SVM. The filter segmentation is then directly mapped back to the original data set.

2.2.3 Morphological analysis

Since the classification provided by the SVM is a partition of the entire volume we can use it to extract relevant morphological data. The volume of an object is extremely useful for medical analysis. One example is quantitatively tracking the growth of a brain tumor over time. An ophthalmologist can use the volume of retinal structures to monitor disease as well as treatment progression. Another potentially useful metric is the density of a volumetric object and its standard deviation. Density in an OCT scan represents the optical density captured by back scattering characterizing a given retinal layer.

3. RESULTS

We tested our method for a variety of OCT retina scans. Table 1 summarizes results including run times. These results are based on an Intel Xeon 3.6GZ workstation with 2GB of main memory.

	No optimization	Checkerboard
Unfiltered Data		
Slice	10.87 s	7.19 s
Volume	15.61 min	8.08 min
Gaussian Filter		
Slice	7.14 s	5.68 s
Volume	6.2 min	4.0 min
Low-pass Filter		
Slice	5.93 s	4.62 s
Volume	3.8 min	1.9 min

Table 1: Performance data. The training data set was created in approximately two minutes and contained 6,283 data points. The data set consists of 140 slices, each having a resolution of 485x350. Tests were performed on an Intel Xeon 3.6 GHz workstation having 2GB of main memory.

Table 1 shows that our SVM optimizations improved the speed by up to a factor of two in one particular application. Unfortunately, the quality of the segmentation is more difficult to discern. By visible inspection, the difference in quality between using no optimizations and using the checkerboard optimization is negligible. (The checkerboard optimization mostly reduces the number of outliers). The low-pass filter produces segmentations that over-estimate region size but have clear boundaries. The unfiltered segmentation contains numerous misclassified points within the segmented region and requires a smoothing step to create concise regions. Using a Gaussian filter, on the other hand, creates an extremely concise and accurate segmentation. It seems that a reasonable compromise between speed and accuracy can be achieved when using a Gaussian filter. To quantify these inferences would require hand segmentation of the entire volume. A Hausdorff error could then be computed to assess the accuracy of a segmentation. Additionally, with this particular application, we can extract specialized morphological data that has known normal values. A retinal layer thickness map, for example, is commonly used as a diagnostic tool in ophthalmology. We can use our segmentation to extract this thickness map from a diseased eye and compare the results to the known normal ranges. Besides extracting raw data, our tool provides a useful means of viewing this thickness map for diagnosis. This thickness map can be rendered as a volume and shown using a traditional 2D color map and/or a height map for visual inspection.

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 an example analysis of the clinical data acquired for a patient with non-exudative age-related macular degeneration (AMD). Figure 5 shows a standard fundus photo of the patient's eye that is commonly used in clinical settings. The right image shows a reconstructed en-face image from 3D volume OCT (100 B-scans / 1000 a-scans scanning pattern with 100 μ s/line exposure) superimposed on the fundus, allowing registration of the acquired volume and estimation of the distortions caused by the subject's eye motions during the experiment.




Figure 5: Left: standard fundus photo from a patient with non-exudative AMD; right: same fundus photo with an OCT (C-scan) superimposed. Arrows indicate the relative positions of the B-scans shown in Figure 7.

To demonstrate the performance of our FD OCT system, four B-scans (labeled A, B, C and D) are shown in Figure 6. The location of the diseased structures can be easily seen in these images.




Figure 6: 4 B-scans chosen from a set of 100 B-scans. The retinal position of these B-scans is shown relative to the fundus view shown in Figure 6.

As described above, this set of images has been used to create the volumetric file that can be visualized using our volume rendering software. Figure 7 shows an example of a 3D visualization of this data set.




Figure 7: 3D visualization of the volumetric data shown in Figures 6. Left: whole volume; center: left part of the volume removed by X-Z clipping plane; right: upper part of the volume removed by X-Y clipping plane.

We also used our SVM segmentation algorithm to segment the RPE and photoreceptor layers, structures that show the main changes associated with the disease progression in AMD. As can be seen from the images in Figure 8, the SVM method leads to clear separation of these layers allowing better visualization as well as estimation of the thickness in the area of interest.




Figure 8: Left: 3D visualization of the SVM-segmented RPE and Photoreceptor layers; right: 3D representation of the thickness map of the structure shown in the left image.

4. CONCLUSIONS

A major limitation in segmentation of OCT scans is that blood vessels cause drastic changes in the characteristic distribution within a single layer by “dimming” scalar values that exist after the imaging beam has been attenuated. As seen in Figure 9, these blood vessels do not allow much light to pass through them, which means that data in regions around and below them are obscured, distorted, or occluded. Our eyes 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. The only way to cope with these anomalies 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 only useful 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, to come up with SVM-based segmentation methods that can handle this problem well.




Figure 9: Red circles indicate blood vessels that distort the intensity values around them and dim the area below them.

One must take care when performing analysis based upon morphological data obtained from SVM segmentation since the quality of a segmentation depends heavily upon the training data and SVM quality. Thus, a practitioner must take care to establish an error estimate for their proper morphological measurements and analysis.

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 do not intersect one another. Using such known constraints should further improve segmentation results.

ACKNOWLEDGEMENTS

This research was supported by a grant from the National Eye Institute (EY 014743). We gratefully acknowledge the contributions of Don T. Miler and Yan Zhang from School of Optometry, Indiana University, Bloomington. We thank the members of the visualization and computer graphics research group, IDAV, University of California, Davis, for their help.

REFERENCES

1. D. Huang, E. A. Swanson, C.P. Lin, J.S. Schuman, W.g. Stinson, W. Chang, M.R. Flotte, K. Gregory, C.A. Puliafito, "Optical Coherence Tomography", *Science* **254**, 1178-1181 (1991)
2. A. F. Fercher, C.K. Hitzenberger, G. Kamp, Y. Elzaiat, "Measurement of intraocular distances by backscattering spectral interferometry", *Opt. Commun.* **117**, 43-48 (1995)
3. 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)
4. 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)
5. N.A. Nassif, B. Cense, B.H. Park, M.C. Pierce, S.H. Yun, B.E. Bouma, G.J. Tearney, t.C. Chen, 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)
6. 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)
7. 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)
8. 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)
9. M. Levoy, "Display of Surfaces from Volume Data" in *IEEE Computer Graphics and Applications* **8**, 29-37, (1988)
10. G. Kindlmann and J.W. Durkin, "Semi-Automatic Generation of Transfer Functions for Direct Volume Rendering" in *IEEE Symp. Volume Visualization*, 79-86, (1998)
11. 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.*, 255-262, (2001)

12. E. Gelenbe, Y. Feng, K. Ranga, and R. Krishnan, "Neural Networks for Volumetric MR Imaging of the Brain" in Proc. Int'l Workshop Neural Networks for Identification, Control, Robotics, and Signal/Image Processing, 194-202, (1996)
13. 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 Networks* **3**, 672-682, (1992)
14. B.E. Boser, I. Guyon, and V. Vapnik, "A Training Algorithm for Optimal Margin Classifiers" in Proc. Fifth Ann. Workshop Computational Learning Theory, 144-152, (1992)
15. C. Cortes and V. Vapnik, "Support Vector Network" *Machine Learning*, **20**, 273-297, (1995)
16. K.R. Muller, S. Mika, G. Ratsch, "An Introduction to Kernel-Based Learning Algorithms" *IEEE Transactions on Neural Networks*, **12**, (2001)
17. V. Blanz, B. Scholkopf, H. Buhmann, C. Burges, V. Vapnik, and T. Vetter, "Comparison of View-Based Object Recognition Algorithms Using Realistic 3D Models" in Proc. Int'l Conf. Artificial Neural Networks, 251-256, (1996)
18. J. Thorsten, "Text Categorization with Support Vector Machines: Learning with Many Relevant Features" in Proc. 10th European Conf. Machine Learning, 137-142, (1998)
19. E. Osuna, R. Freund, and F. Girosi, "Training Support Vector Machines: An Application to Face Detection" in Proc. '97 Conf. Computer Vision and Pattern Recognition, 130-137, (1997)
20. Fan-Yin Tzeng, Eric B. Lum, and Kwan-Liu Ma, "An Intelligent System Approach to Higher-Dimensional Classification of Volume Data" *IEEE Transactions on Visualization and Computer Graphics*, **11**, (2005)
21. E. Aboufadel and S. Schlicker, "Discovering Wavelets", John Wiley & Sons, Inc. New York, (1999)
22. C.S. Burrus and R. A. Gopinath, H. Guo, "Introduction to Wavelets and Wavelet Transforms", A Primer. Upper Saddle River, (1998)
23. A.K. Louis, P. Maas and A. Rieder, "Wavelets: Theory and Applications", John Wiley & Sons Ltd. Chichester, (1997)