



Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Coupled parametric model for estimation of visual field tests based on OCT macular thickness maps, and vice versa, in glaucoma care

Andy Tsai^{a,*}, Joseph Caprioli^b, Lucy Q. Shen^c

^a Department of Radiology, Children's Hospital Boston, 300 Longwood Avenue, Boston, MA 02115, United States

^b Glaucoma Division, Jules Stein Eye Institute, 100 Stein Plaza, Los Angeles, CA 90095, United States

^c Glaucoma Service, Massachusetts Eye and Ear Infirmary, 243 Charles Street, Boston, MA 02114, United States

ARTICLE INFO

Article history:

Received 29 July 2010

Received in revised form 16 May 2011

Accepted 19 May 2011

Available online xxxx

Keywords:

OCT

Visual field testing

Glaucoma

PCA

ABSTRACT

The current standard of care for glaucoma patients consists of functional assessment of vision via visual field (VF) testing, which is sensitive but subjective, time-consuming, and often unreliable. A new imaging technology, Fourier domain optical coherence tomography (OCT), is being introduced to assess the structural characteristics of the macula. This new complementary exam is efficient, objective, and reliable. A complex, but consistent, relationship exists between the structural information provided by macular OCT and the functional information gathered by VF maps. We propose a learning-based framework with the ability to estimate the VF map based on OCT macular thickness measurements as input (and vice versa). Central to this algorithmic framework is a coupled parametric model that captures not only the individual variabilities of the OCT macular thickness measurements and the VF maps, but also their co-dependencies. This model is derived by applying principal component analysis (PCA) to a library consisting of various pairs of OCT and VF maps. The parameters of this coupled model are obtained by solving a linear least squares problem. Next, these estimated parameters are used, in conjunction with the eigenvectors derived from PCA, to compute the estimate. The accuracy of this coupled parametric estimation model was evaluated by performing multiple leave-one-out cross validation experiments.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

Glaucoma is the leading cause of irreversible blindness, the fear of which ranks third among most people surveyed, followed only by cancer and heart disease (Prevent Blindness America Survey, 2002). Early detection and regular monitoring are essential to achieving optimal care for glaucoma patients. One important tool is the visual field (VF) exam, a functional test that is part of the standard of care for diagnosis and evaluation of glaucoma. It is performed regularly in the management of glaucoma. While the VF exam is sensitive to detect and monitor the disease; it is, unfortunately, subjective, time-consuming, and often unreliable. To offset these limitations, Fourier domain optical coherence tomography (OCT) has recently been introduced to provide important structural information about the posterior segment of the eye. This structural evaluation by OCT has been shown to be objective, efficient, and reliable (Schuman et al., 1996; Budenz et al., 2008; Leung et al., 2010; Wolf-Schnurrbusch et al., 2009). Several investigators have reported a significant but complex relationship between the structural damage as measured by OCT, and the functional deficits as

shown by VF testing in glaucoma patients (Ferreras et al., 2008; Harwerth et al., 2007; Hood and Kardon, 2007; Hood et al., 2007; Shen et al., 2010). However, this relationship has not been well defined to become clinically useful.

The ability to estimate the results of a functional test based on structural information, and vice versa, is the topic of this paper. Such predictive ability may improve the accuracy of glaucoma detection and monitoring, as the predicted VF map (based on OCT images) may be used to confirm real defects in the subjective VF test. Furthermore, this ability may reveal a functional-structural correspondence that effectively leads to a better understanding of glaucoma. In the following three subsections, we provide a short introduction to VF testing and OCT macular thickness maps; we briefly summarize our model and related research; and we argue that our novel methodology proposed here is superior to existing techniques.

1.1. Ophthalmologic evaluation in glaucoma

As mentioned earlier, two complementary exams for glaucoma patients include the functional assessment of vision via VF testing and OCT structural assessment of the posterior segment of the eye. These two exams are further discussed below.

* Corresponding author. Tel.: +1 617 817 5503.

E-mail address: andy.tsai@childrens.harvard.edu (A. Tsai).

1.1.1. Visual field function test

This exam maps the patient's central and peripheral visual function based on his or her response (e.g. pressing of a button) to seeing a flashing light of varying intensities presented at predetermined locations. The most common VF test has 54 test locations within the central 24 degrees of the field of view (standard automated perimetry 24-2, Humphrey Field Analyzer 750, Carl Zeiss Meditec, Inc., Dublin, CA). At the conclusion of the exam, the patient's responses are analyzed statistically and compared to a normative database to generate a detailed map (conventionally presented in units of decibels) noting the areas of VF defect. Locations labeled with negative values correspond to areas where the patient performed *below* the expected norm. The exponential magnitude of these negative values reflects how far below the norm the patient performed at the implicated location of the VF test. Similarly, VF locations labeled with positive values correspond to locations where the patient performed *above* the expected norm. Each eye is tested separately, yielding two VF maps for each patient. Only the central subset of the VF test locations which overlaps with the portion of the macula captured by the OCT is utilized for our analysis. This subset of VF points, in relationship to the entire 54 point spatial test locations, is encircled in green as shown in the second row of Fig. 1. After excluding the blind spot (blue¹ areas shown in Fig. 1), each of these VF maps contains a total of 34 data points. Fig. 1a is an example of a normal individual's 34-point VF map where all the numbers hover around 0, indicating that this particular patient performed at or close to the expected norm in these regions. The two dark blue pixels along the last column correspond to the patient's blind spots. In contrast, Fig. 1b is an example of a glaucoma patient's VF map demonstrating superior hemi-field defect. Notice that numerous points within the superior visual field take on values *well below* 0, indicating a significant VF defect.

1.1.2. Optical coherence tomography (OCT)

OCT is an emerging imaging modality based on the principles of low-coherence interferometry. With spectrometer detection method and a Fourier transform algorithm, Fourier domain OCT can perform 26,000 A-scans in the macula in one second. Each map generated consists of a 6 mm × 6 mm grid (36 mm²) centered on the fovea, with spacing of 0.25–0.5 mm within the grid, and a depth resolution of 5 μm. This fast and efficient data acquisition minimizes motion artifacts, while maximizing the signal-to-noise ratio (Gonzalez-Garcia et al., 2009; Wojtkowski et al., 2005). Images in the top row of Fig. 2 represent printouts from EMM5 macular scans (RTVue, Optovue Inc, Fremont, CA). The raw data were exported, re-centered to the fovea, normalized with the highest thickness values in each scan to account for inherent differences in the thicknesses of the maculas across different patients (such as may be due to age, ethnicity, and length of the eye). These thickness ratio maps are represented on the bottom row of Fig. 2 as colored contour maps. Fig. 2a shows the normalized macular OCT thickness map of an individual with a normal VF map as depicted in Fig. 1a. Notice the donut-shaped macula with a thickness close to a normalized value of 1, and the central area of decreased macular thickness corresponding to the fovea. In contrast, Fig. 2b illustrates the normalized macular OCT thickness map of a patient with glaucoma; the inferior macular thinning corresponds to the superior VF loss as shown in Fig. 1b.

1.2. Related research

One important class of unsupervised learning algorithms is principal component analysis (PCA). It is a classic statistical technique first described in 1901 (Pearson, 1901). PCA can

automatically learn and capture the complex patterns and relationships that characterize a training database and leverage this information for various applications (Jolliffe, 2002).

As reported by Tsai et al. (2004), this PCA-based approach can be extended to simultaneously handle multiple images of differing types within a single linear model. In this 2004 paper, PCA was applied to a sample covariance matrix consisting of example images drawn from different families of images. The derived eigenvectors not only capture the variations present within each family of images; but also the co-variations that exist among the various families. The approach we take in this paper bears some semblance to this particular aspect of their methodology.

When two images are strongly related to one another, and one of the images is unknown (as in our application), an estimation-theoretic approach can be taken to estimate that unknown image by incorporating this pair of images as an additional entry in a training database, and judiciously labeling the unknown image as the missing data. Viewed in this light, this type of problem is transformed into the estimation of this missing data within the training database. Various PCA-based methods have been reported in the literature that can provide an estimate of these missing data, each with their own way of handling missing values. For example, a probabilistic PCA, formulated in an EM-framework, was introduced in separate studies by Roweis (1998), and Tipping and Bishop (1999). Their algorithms simultaneously provide an estimation of the principal subspace and a reconstruction of the missing data. More recently, a PCA-based variational Bayesian subspace learning approach is proposed in the Raiko paper (Raiko et al., 2007) that can handle large scale problems with numerous missing values. Several other PCA-based approaches that can deal with missing data within the training database are presented in the book by Little and Rubin (2002). In general, however, the computational costs of all these methods are too expensive for this particular application, as this type of approach calls for re-calculation of the principal components with each query (rather than simply utilizing the same principal components calculated from the initial off-line training).

An algorithm developed by Essock et al. (2007) that is based on both the PCA method and wavelet-Fourier analysis is particularly pertinent to this paper as it is applied to a clinical problem that is similar to ours. Specifically, the goal of Essock's algorithm is to provide a binary assessment of whether a patient with ocular hypertension will later develop visual field loss or not. To this end, wavelet-Fourier analysis is applied to a library database consisting of retinal nerve fiber layer thickness estimates from both glaucomatous and healthy individuals. This wavelet-base approach extracts feature vectors that are then stacked row by row to form a feature matrix. PCA is then applied to transform this feature matrix into a smaller number of uncorrelated variables. Next, Fisher's linear discriminant analysis is employed to calculate the discriminant function based on these uncorrelated variables. For binary classification, this discriminant function is applied to another data set (outside of the training database), consisting of patients with ocular hypertension, to predict the subset of this patient population that will eventually convert to glaucoma and develop visual loss, and those who will not.

1.3. Contributions

The main contribution of the work presented here lies in the successful adaptation of the coupling technique developed by Tsai et al. (2004) for image estimation (as opposed to image segmentation for which it was originally formulated). In this paper, we propose a parametric linear model that utilizes this coupling technique to capture the complex structural-functional relationship that exists between the OCT thickness ratio map and its paired

¹ For interpretation of color in Figs. 1–18, the reader is referred to the web version of this article.

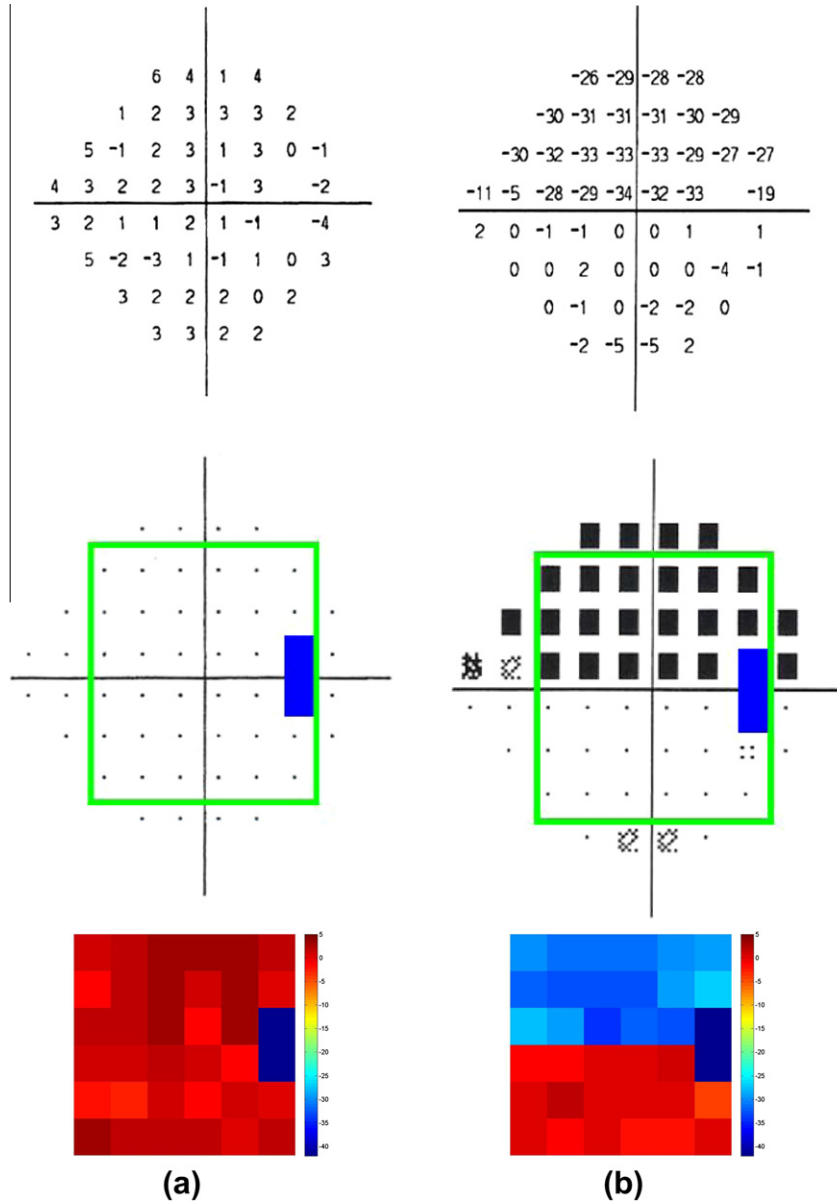


Fig. 1. Two examples of VF maps, one is from a normal individual (a), and one is from an individual with superior hemi-field deficit (b). First row are the total deviation plots. Second row illustrates the central subset of the VF test points used for this study, as encircled by the green square. Processed VF images are shown on the last row. The dark blue areas indicate the blind spots.

VF map. Central to this model is the application of PCA to a sample covariance matrix composed of various pairs of OCT and VF images to arrive at a set of eigenvectors that can serve as the orthogonal basis function for image representation. By design, this orthogonal basis function efficiently encodes the variabilities within the library database, and naturally captures the co-dependencies between the OCT and the VF images. The parameters applied to our model correspond to the weighting scheme for this set of orthogonal basis functions, and are estimated by solving a linear least squares problem.

2. PCA-based coupled parametric model

We begin the description of our model with definition notations. Let Ψ^{oct} and Ψ^{vf} denote the macular OCT thickness ratio map and its paired VF map, respectively. Suppose there are n such example pair images in our training library database as denoted by

$\{\Psi_1^{oct}, \Psi_2^{oct}, \dots, \Psi_n^{oct}\}$ and $\{\Psi_1^{vf}, \Psi_2^{vf}, \dots, \Psi_n^{vf}\}$. The n OCT ratio maps are aligned with one another because of the manner by which they are obtained (all centered over the fovea). In a similar fashion, the n VF maps are aligned with one another. Hence, alignment is not an issue for either the OCT or the VF images, as they are all performed/obtained with reference to the same fixed structure.

We first compute the mean image functions $\bar{\Psi}^{oct}$ and $\bar{\Psi}^{vf}$, representing the mean OCT ratio map and mean VF map of our library database, respectively. To extract the variability of the OCT and VF images, the mean images $\bar{\Psi}^{oct}$ and $\bar{\Psi}^{vf}$ are subtracted from each of the n OCT and VF images, respectively. This operation yields n mean-offset OCT and n mean-offset VF images as denoted by $\tilde{\Psi}^{oct}$ and $\tilde{\Psi}^{vf}$, respectively, as shown below:

$$n \text{ mean-offset OCT images : } \{\tilde{\Psi}_1^{oct}, \tilde{\Psi}_2^{oct}, \dots, \tilde{\Psi}_n^{oct}\} \quad (1)$$

$$n \text{ mean-offset VF images : } \{\tilde{\Psi}_1^{vf}, \tilde{\Psi}_2^{vf}, \dots, \tilde{\Psi}_n^{vf}\} \quad (2)$$




Fig. 2. Two examples of OCT thickness maps, one is from an individual with a normal VF (a), and one is from a glaucoma patient with inferior macular thinning (b). Cross sections of full-thickness retina are shown in black boxes on the first row, with their locations indicated by red lines. Processed OCT images are shown in the second row.

where

$$\tilde{\Psi}_i^{oct} = \Psi_i^{oct} - \bar{\Phi}^{oct} \quad \forall i \in \{1, \dots, n\} \quad (3)$$

$$\tilde{\Psi}_i^{vf} = \Psi_i^{vf} - \bar{\Phi}^{vf} \quad \forall i \in \{1, \dots, n\}. \quad (4)$$

To prepare the training library, we form n column vectors $\{\tilde{\Psi}_1, \dots, \tilde{\Psi}_n\}$, where each column vector $\tilde{\Psi}_i$ consists of P samples of $\tilde{\Psi}_i^{oct}$ directly concatenated on top of Q samples of $\tilde{\Psi}_i^{vf}$. Therefore, each $\tilde{\Psi}_i$ consists of $N = P + Q$ elements. Specifically, in our setup, for OCT and VF imaging pair i , the 501×501 pixel size OCT image is used to generate a column of $P = 251,001$ lexicographically ordered pixels (where the columns of the image grid are sequentially concatenated on top of one another to form one large column). This same strategy is applied to transform the 6×6 VF image to generate a column of $Q = 34$ lexicographically ordered pixels (after excluding the blind spot which occupies two test locations in the VF). These two column vectors are concatenated on top of each other to form a large column vector $\tilde{\Psi}_i$ of dimension $N = 251,035$. This process is repeated n times, one for each OCT and VF image pair, to form a tall and skinny matrix of size $251,035 \times n$. We define this matrix S as

$$S = [\tilde{\Psi}_1 \ \tilde{\Psi}_2 \ \dots \ \tilde{\Psi}_n]. \quad (5)$$

A simple schematic diagram to illustrate the formation of S is shown in Fig. 3.

Eigenvalue decomposition is employed to factor $\frac{1}{n}SS^T$ as:

$$\frac{1}{n}SS^T = X\Lambda X^T \quad (6)$$

where X is a tall and skinny rectangular $N \times n$ matrix whose columns represent the n principal variational modes or eigenvectors of the paired OCT and VF images, and Λ is an $n \times n$ diagonal matrix whose diagonal elements, denoted by $\lambda_1, \dots, \lambda_n$, represent the corresponding non-zero eigenvalues. Because these eigenvectors are

derived by performing principal component analysis on the matrix S , which contains not only information about the variabilities of the OCT and VF images, but also about their co-dependencies, the eigenvectors derived will naturally demonstrate strong coupling between the OCT and the VF images. This idea is central to the success of our algorithm. Further, as a direct result of concatenating the vectorized OCT image representation directly on top of the vectorized VF image representation, the $N \times n$ -dimensional matrix X can be easily partitioned into an upper and a lower submatrix. Specifically,

$$X = \begin{bmatrix} X^{oct} \\ X^{vf} \end{bmatrix} \quad (7)$$

where X^{oct} is of dimension $P \times n$, and X^{vf} is of dimension $Q \times n$. As we will show later, these submatrices X^{oct} and X^{vf} are important for image reconstruction.

From a computational standpoint, calculating the eigenvectors and eigenvalues of the $N \times N$ matrix $\frac{1}{n}SS^T$ is simply not practical. To offset this difficulty, in standard practice, we typically compute the eigenvectors and eigenvalues from a much smaller $n \times n$ matrix W given by:

$$W = \frac{1}{n}S^TS. \quad (8)$$

Thus, if \mathbf{d} is an eigenvector of W with corresponding eigenvalue λ , then $S\mathbf{d}$ is an eigenvector of $\frac{1}{n}SS^T$ with eigenvalue λ (see proof on p. 59 in (Leventon, 2000)).

Because of the projection operation, the weighted sum of these eigenvectors can be used to generate a reconstructed image vector, which, in turn, can then be *unconcatenated* to yield two separate vectorized image representations – one to reconstruct the OCT image, and one to reconstruct the VF image. More precisely, the first P samples of the reconstructed image vector can be rearranged (by undoing the earlier lexicographical concatenation of the grid




Fig. 3. Diagram illustrating the formation of the matrix S and the column vectors $\tilde{\psi}_i$.

columns) to form a 501×501 rectangular OCT image. Likewise, the last Q samples of the reconstructed image vector can be rearranged in a similar fashion to form a 6×6 rectangular VF image (after adding back in the two blind spots).

3. Image estimation

We begin this section by introducing a framework for estimating VF maps using the OCT thickness ratio maps as input. Suppose we have noisy measurement of *both* the OCT and the accompanying VF images. Denote the lexicographically vectorized representation of these two noisy images as Y^{oct} and Y^{vf} , respectively. To be precise, the dimensions of Y^{oct} are $P \times 1$, and the dimensions of Y^{vf} are $Q \times 1$. Concatenating them on top of one another forms the $N \times 1$ dimensional measurement vector Y :

$$Y = \begin{bmatrix} Y^{oct} \\ Y^{vf} \end{bmatrix}. \quad (9)$$

As we mentioned earlier, both of these measurements are contaminated by noise. We model the OCT measurement noise as independent identically distributed (iid) zero-mean Gaussian process with standard deviation of σ_{oct} . In a similar fashion, we model the VF measurement noise as iid zero-mean Gaussian process with standard deviation of σ_{vf} . In our learning-based coupled parametric model, the explicit dependence of these measurements Y on the eigenvectors X is given by the following measurement model:

$$Y = X\alpha + v \quad (10)$$

where

$$v = \begin{bmatrix} v^{oct} \\ v^{vf} \end{bmatrix} \quad (11)$$

with

$$v_i^{oct} \sim G(0, \sigma_{oct}^2) \quad \forall i \in \{1, \dots, P\} \quad (12)$$

$$v_i^{vf} \sim G(0, \sigma_{vf}^2) \quad \forall i \in \{P+1, \dots, Q\}. \quad (13)$$

The $n \times 1$ parameter α determines the weighting scheme of the eigenvectors in X that best approximates the measurements Y .

Using Eqs. 7, 9, and 11 to expand Eq. (10), we obtain the following measurement model:

$$\begin{bmatrix} Y^{oct} \\ Y^{vf} \end{bmatrix} = \begin{bmatrix} X^{oct} \\ X^{vf} \end{bmatrix} \alpha + \begin{bmatrix} v^{oct} \\ v^{vf} \end{bmatrix}. \quad (14)$$

To calculate α for the above equation, we solve the following quadratic minimization problem:

$$\hat{\alpha} = \arg \min_{\alpha} \|Y - X\alpha\|^2 \quad (15)$$

$$= \arg \min_{\alpha} \left\| \begin{bmatrix} Y^{oct} \\ Y^{vf} \end{bmatrix} - \begin{bmatrix} X^{oct} \\ X^{vf} \end{bmatrix} \alpha \right\|^2. \quad (16)$$

Because the columns of the matrix X are linearly independent (which is, by design, as they are derived based on PCA), this minimization problem has a unique solution given by:

$$\hat{\alpha} = \left(\frac{1}{\sigma_{oct}^2} X^{octT} X^{oct} + \frac{1}{\sigma_{vf}^2} X^{vfT} X^{vf} \right)^{-1} \cdot \left(\frac{1}{\sigma_{oct}^2} X^{octT} Y^{oct} + \frac{1}{\sigma_{vf}^2} X^{vfT} Y^{vf} \right). \quad (17)$$

Recall that in formulating our problem earlier, we assumed that both the OCT and VF measurements are available. However, in reality, we only have the OCT measurements and not the VF measurements, as the goal is to estimate the VF measurements based on the provided OCT measurements. Hence, the standard deviation of the measurement noise σ_{vf} in Eq. (17) is set to infinity, which is one way to indicate complete lack of information regarding Y^{vf} . With $\sigma_{vf} = \infty$, Eq. (17) reduces to:

$$\hat{\alpha}^{oct} = (X^{octT} X^{oct})^{-1} \cdot (X^{octT} Y^{oct}) \quad (18)$$

The $\hat{\alpha}^{oct}$ derived based on Eq. (18) is used in conjunction with the eigenvectors X^{vf} to estimate Y^{vf} . Specifically, an estimation of Y^{vf} is given by:

$$\hat{Y}^{vf} = X^{vf} \hat{\alpha}^{oct}. \quad (19)$$

Thus far, we have shown the methodology by which we estimate the VF measurements based on OCT measurements. Importantly,




Fig. 4. VF database.

this method is flexible enough to accommodate the reverse scenario – that is, estimating OCT measurements based on VF measurements, by altering our algorithm in a simple and straightforward manner. Specifically, instead of setting σ_{vf} to infinity, we set σ_{oct} to infinity so that Eq. (17) becomes:

$$\hat{\alpha}^{vf} = (X^{vfT} X^{vf})^{-1} \cdot (X^{vfT} Y^{vf}). \quad (20)$$

Since we are now interested in estimating Y^{oct} , it is given by the following:

$$\hat{Y}^{oct} = X^{oct} \hat{\alpha}^{vf}. \quad (21)$$

4. Results

In this section, we demonstrate the performance of our coupled parametric estimation model. We begin in Section 4.1 by displaying the VF and OCT training databases that we employ for this paper as well as the calculated principal modes associated with these databases. In Section 4.2, we qualitatively and quantitatively assess the performance of our model by showing the results of our leave-one-out experiments in the estimation of VF images based on OCT data. Next, we switch the roles of the VF maps and the OCT thickness ratio maps by performing leave-one-out experiments in the estimation of OCT images based on VF data, the results of which are displayed in Section 4.3.

4.1. Training databases and their principal modes

A total of 14 patients with glaucomatous superior VF loss to a single eye underwent VF testing and OCT imaging of both eyes to yield 28 pairs of VF and OCT data. Patients with glaucomatous superior VF defect were chosen because this is, by far, the most

common VF defect in glaucoma. As a result, it is easier to collect a richer database for training and testing. Of these 28 pairs of data, 8 pairs of data were discarded by the glaucoma specialists because of one of the following 3 reasons: the VF data were unreliable, the OCT image had poor image quality, and the VF/OCT pair did not show the expected anatomic correspondence. The 20 remaining VF and OCT data sets were used in the algorithm and are displayed in Figs. 4 and 5, respectively.

We next show the training phase of our coupled parametric model. Specifically, we applied the methodology as described in Section 2 to the paired VF and OCT databases illustrated in Figs. 4 and 5 to arrive at 20 eigenvectors that are used to effectively capture the variations and the co-dependencies between the VF maps and the macular OCT thickness ratio maps. The cumulative sum of the eigenvalues is plotted in Fig. 6, respectively. Notice that the first mode captures approximately 25% of the variation, and this amount tapers toward 0 with each additional mode.

To provide some intuition behind how the eigenvectors capture the variabilities and co-dependencies of the OCT and VF pair images, the first three eigenvectors are varied by different amounts and shown in Figs. 7–9. Each figure demonstrates the effect of varying a particular principal mode in changing the appearances of the macular OCT thickness ratio map and the VF map. For example, notice that by varying the first principal mode from $-0.6\sqrt{\lambda_1}$ to $0.6\sqrt{\lambda_1}$, a large superior hemi-field defect is created in VF, which is associated with thinning in the lower half of the macula (Fig. 7), an expected relationship based on the anatomy of the eye. Because PCA maximizes the largest singular value, the first principal mode is the pattern that accounts for the greatest amount of the variance–covariance structure. By varying the second principal mode from $-0.6\sqrt{\lambda_2}$ to $0.6\sqrt{\lambda_2}$, a supero-nasal VF defect is created, which in turn, is associated with predominant thinning in the peripheral inferior macula (Fig. 8). The third eigenvector is predominantly associated with changes in the width and radius of the donut

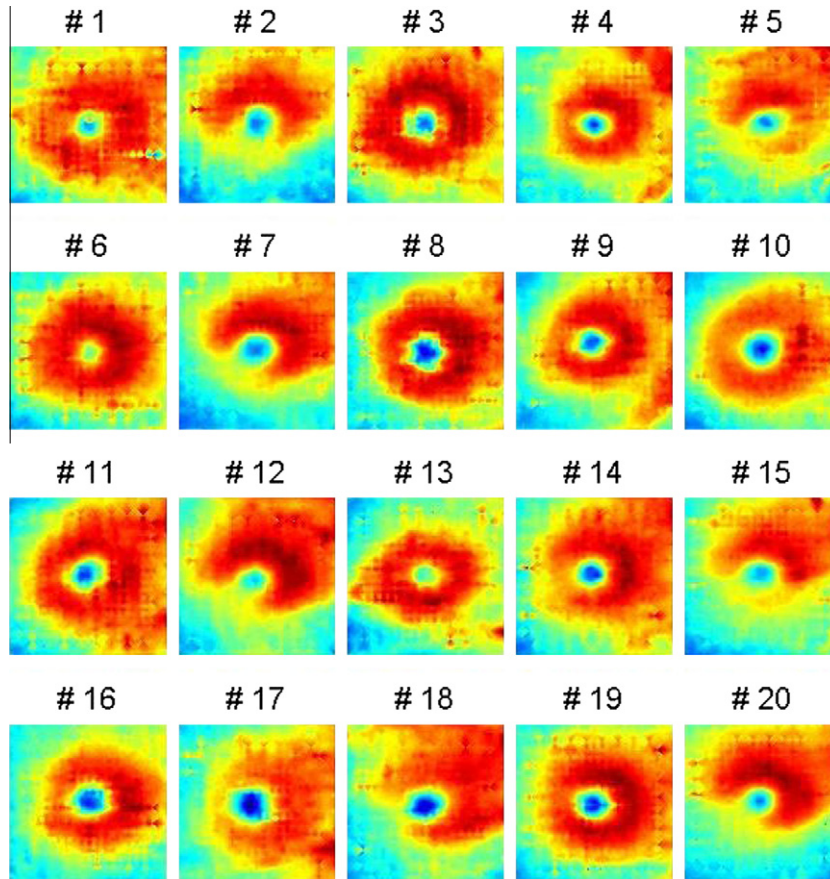


Fig. 5. OCT database.




Fig. 6. Cumulative distribution of the eigenvalues.

appearing macula (Fig. 9). As the structural–functional relationship is complex, with each successive eigenvector, the coupling between VF maps and OCT images are more subtle and harder to identify visually.

4.2. Estimation of VF maps based on OCT images of the macula

To evaluate the performance of our coupled parametric model in estimating VF maps based on OCT images as described in




Fig. 7. First eigenmode.




Fig. 8. Second eigenmode.




Fig. 9. Third eigenmode.

Eqs. (18) and (19), we performed 20 leave-one-out experiments. The estimates from these experiments were then compared to actual measurements.

One method to determine the optimal number of eigenvectors to use in our model is to calculate the root-mean-square error (RMSE) between the measured and the estimated VF maps, as defined by

$$RMSE^{vf} = \sqrt{\frac{\sum_{i=1}^Q (\hat{Y}_i^{vf} - Y_i^{vf})^2}{Q}}, \quad (22)$$

while the number of eigenvectors employed is varied from 1 to 19. The idea behind this approach is to find the ideal number of eigenvectors to employ so that the estimated VF image best mimics the measured VF image. The results of this exercise are shown in Fig. 10, which demonstrated that the lowest RMSE (5.37) is accomplished by utilizing just the largest eigenvector. We suspect that only the largest eigenvector is needed because the variability of our database is limited, as we restricted ourselves to only patients with glaucomatous superior VF defects. Expansion of the database with inclusion of other visual field deficits is needed to test this hypothesis.

A representative sample of the results from our 20 leave-one-out experiments, based on using the single largest eigenvector, is displayed in Fig. 11. Three images are shown in a row for each patient. The first column is the input OCT data, the second column is the output estimated VF image based on our algorithm, and the third column is the measured VF map. The estimated VF image is purposely positioned next to the actual VF image to facilitate a direct qualitative comparison. Qualitatively, the VF estimation results appear reasonably accurate overall.

To provide a quantitative assessment of the performance of our algorithm, we generated a scatter plot of estimated versus measured VF values for each of the 34 points in each of our 20 leave-one-out experiments (using just the single largest eigenvector), the results of which are shown in Fig. 12. A best fit line, shown in red, is plotted through these 680 points. The slope and y-intercept of this line are 0.743 and -1.077 respectively, with the ideal slope and y-intercept being 1 and 0 respectively. The R -value for this line fit is 0.844 with p -value of <0.0001 . This R -value is considered very good in this application, especially given the fact that our estimation results are based on relationships captured between an *objective* structural test and a *subjective* functional test.

As mentioned previously, the VF measurements are subjective and often unreliable. However, these are the established and




Fig. 10. RMSE in VF prediction as a function of number of eigenvectors used.

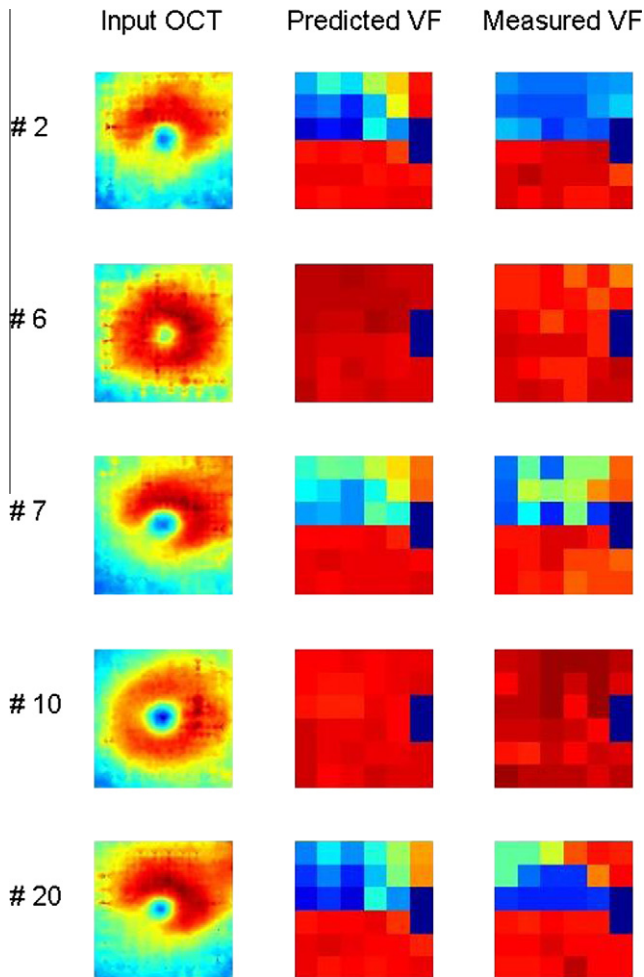


Fig. 11. Examples of VF prediction results.

clinically accepted measurements. One way to assess a new method of measurement with an established one is via the Bland–Altman plot. To visually judge the two methods via the Bland–Altman plot, the mean difference between these two methods (i.e. the bias), and the 95% limits of agreement are overlaid on the Bland–Altman plot as shown in Fig. 13. The bias here is -0.04 dB, and the standard deviation is 5.48 dB.² Visually, there is relatively good agreement between the measured and the estimated values.

4.3. Estimation of OCT images of the macula based on VF testing

In a similar fashion, we qualitatively and quantitatively evaluated the performance of our parametric model in estimating OCT thickness ratio maps based on VF maps. Specifically, we performed 20 leave-one-out experiments based on the model described by Eqs. (20) and (21), and then compared our estimates to actual measurements.

Again, to determine the optimal number of eigenvectors to use, we calculated the RMSE between the measured and the estimated OCT macular thickness maps, as defined by

$$RMSE^{oct} = \sqrt{\frac{\sum_{i=1}^P (\hat{Y}_i^{oct} - Y_i^{oct})^2}{P}}, \quad (23)$$

while the number of eigenvectors employed is varied from 1 to 19. The results of this exercise are shown in Fig. 14, which demonstrates that the lowest RMSE (0.081) is from using just the single largest eigenvector. As already mentioned, we suspect that only the largest eigenvector is needed because of the limited variability of our database.

Representative results from our 20 leave-one-out experiments, based on using the single largest eigenvector, is displayed in

² As a reminder, units of dB are used in the Bland–Altman plot because the inherent units of VF maps are in dB.




Fig. 12. Scatter plot of predicted versus measured VF.

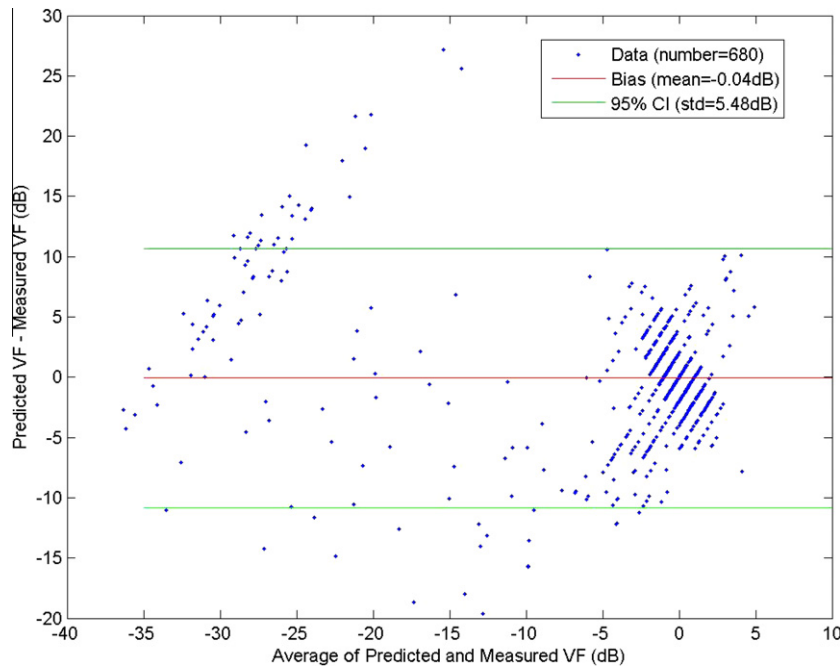


Fig. 13. Bland Altman plot of VF prediction.

Fig. 15. In a similar fashion as to the prior section to facilitate direct comparison, three images are shown in a row for each patient. The first column is the input VF data, the second column is the output estimated OCT image based on our algorithm, and the third column is the measured OCT map. These results appear reasonably accurate.

For a more quantitative assessment, we again generated a scatter plot of estimated versus actual OCT values for each of the 251,001 points in each of our 20 leave-one-out experiments (using only the single largest eigenvector), the result of which is shown in

Fig. 16. A line of best fit (shown in red) is plotted through the 5,020,020 points. The slope and y-intercept of this line is 0.801 and 0.127 respectively, with the ideal slope and y-intercept being 1 and 0 respectively. The R -value for this line fit is 0.893 with p -value of <0.0001 . Again, this R -value is considered very good in this application.

One way to assess the accuracy of our estimation method is via the Bland–Altman plot. To visually judge the two methods via the Bland–Altman plot, the bias and the 95% limits of agreement are overlaid on a Bland–Altman plot shown in Fig. 17. The bias here




Fig. 14. RMSE in OCT prediction as a function of number of eigenvectors used.




Fig. 15. Examples of OCT prediction results.

is 0.0006, and the standard deviation is 0.104. Again, there is relatively good agreement between the estimated and the measured OCT data.

5. Conclusion and future research directions

In this final section, we summarize the contributions of this paper, and provide suggestions for extensions and future research directions. In the preceding sections, we presented a mathematically principled framework that is capable of providing a reasonable estimate of the VF map when given the corresponding OCT image, and vice versa. One key attraction of this technique is its versatility as it can be used to estimate a functional test based on structural information, and vice versa. Moreover, the training phase of our formulation requires no supervision and is easy to implement. In this phase, we employ the method of PCA to decompose a covariance matrix into a series of orthogonal patterns, which embody (in a step down fashion) not only the variability in the OCT and VF images respectively, but also in their co-dependencies. The estimation phase of our algorithm is accomplished rapidly, as it simply involves solving a linear least squares equation for a small set of parameters. The results of the leave-one-out experiments clearly demonstrate the estimation capabilities of our model and show promise for new applications.

This coupled parametric model is simple to understand and implement; it is computationally efficient; and it offers both strong unsupervised learning and estimation capabilities. This versatile model has two distinct advantages over most of the techniques previously described: first, it does not require recalculation of the eigenvectors with each query; and second, it accommodates a wider subspace than the training database which admits query data that may be far away from the training set but still close to the principal subspace.

In terms of clinical utility, this is one of the first studies to use OCT macular thickness maps to predict VF defects. The preliminary data generated based on our approach can be used by clinicians for patients with superior VF defects, even at this immature stage. For example, our predicted visual field can be compared to the subjective VF test, and verify which defects are indeed real. It is known that when the VF defect is significant (i.e. <-40 dB), the defect can be quite variable (even if the patient is asked to retake the test). Hence, it is not unexpected to find differences between the predicted and measured VF values. In short, this algorithm may

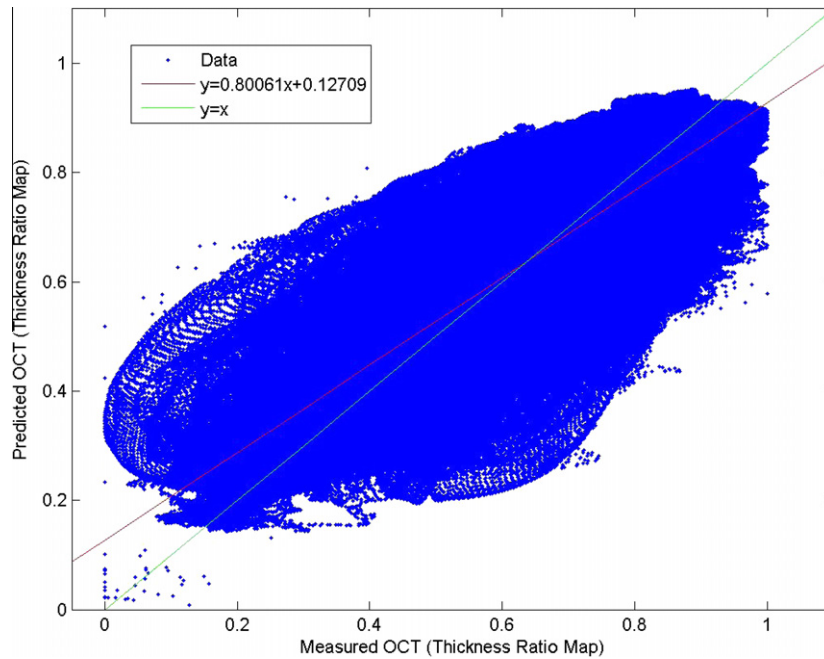


Fig. 16. Scatter plot of predicted versus measured OCT.

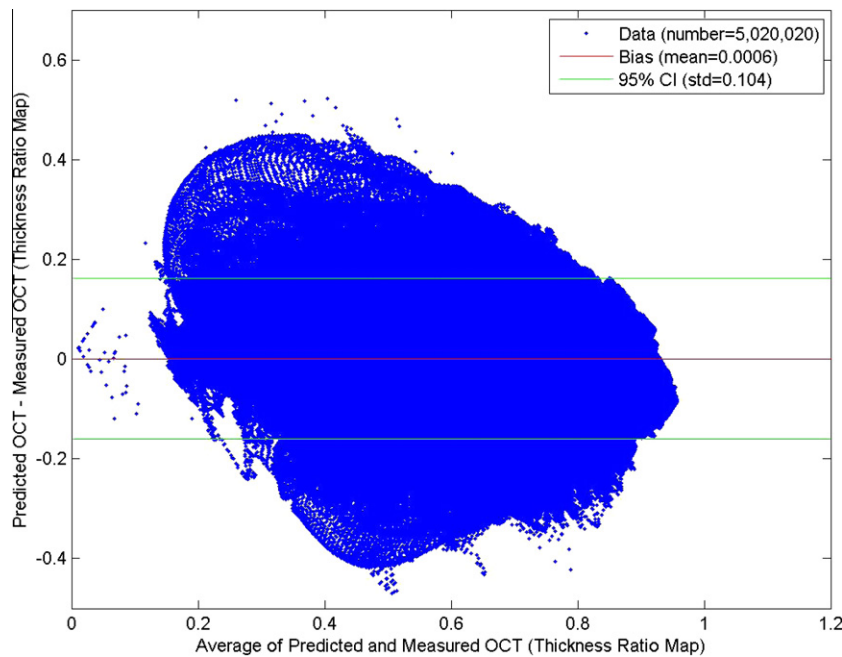


Fig. 17. Bland Altman plot of OCT prediction.

help to incorporate Fourier domain OCT of the macula into the clinical care for glaucoma patients.

Unfortunately, our approach has a few shortcomings: first, the prediction capability is somewhat limited because the algorithm depends heavily upon a relatively small training library database for its implementation, and only superior VF defects were included in the training database. The larger and more accurate the database, the better our technique is expected to perform (as is true in most training-based algorithms). To offset this weakness, we are actively recruiting more patients to enrich our library database and expand the versatility of this estimation algorithm. Second, the

performance of our methodology is also limited by our use of the PCA which assumes our data set to be linear combinations of basis function. Perhaps non-linear methods such as kernel PCA (Scholkopf et al., 1998) may yield improved performances. Third, 2D correspondence maps between VF maps and OCT images are difficult to create because, when using PCA, we must first transform the 2D image matrices into 1D image vectors, and then compute the corresponding 1D eigenvectors from the sample matrix composed of a concatenation of 1D image vectors. The 1D eigenvectors are then transformed back to 2D matrix for image reconstruction. This matrix-to-vector conversion unfortunately causes

the loss of inherent structural information present within the original 2D images. In an effort to correct this problem, we are currently investigating the use of 2D PCA (Yang et al., 2004). Finally, we are exploring improved mechanisms of capturing co-dependencies among different families of images.

Acknowledgments

This work was supported in part by Department of Radiology, Children's Hospital Boston, American Glaucoma Society, unrestricted grants from Research to Prevent Blindness and Pfizer Inc., and Harvard Catalyst/the Harvard Clinical and Translational Science Center. The authors would like to thank Optovue, Inc. for providing technical assistance with RTVue, Dennis Mock for his role in data collection and processing, and Nancy Drinan for her help in the preparation of this manuscript. Finally, the authors would like to thank the anonymous reviewers for their valuable comments and thoughtful suggestions.

References

- Budenz, D.L., Chang, R.T., Fredette, M.J., Feuer, W.J., Anderson, D.R., 2008. Reproducibility of peripapillary retinal nerve fiber thickness measurements with Stratus OCT in glaucomatous eyes. *Ophthalmology* 115 (4), 661–666.
- Essock, E., Guvant, P., Zheng, Y., Garway-Heath, D., Kotecha, A., Spratt, A., 2007. Predicting visual field loss in ocular hypertensive patients using wavelet-fourier analysis of GDx scanning laser polarimetry. *Optometry and Vision Science* 84, E380–E387.
- Ferreras, A., Pablo, L.E., Garway-Heath, D.F., Paolo, F., Garvia-Feijoo, J., 2008. Mapping standard automated perimetry to the peripapillary retinal nerve fiber layer in glaucoma. *Investigative Ophthalmology & Visual Science* 49, 3018–3025.
- Gonzalez-Garcia, A.O., Vizzeri, G., Bowd, C., Medeiros, F.A., Zangwill, L.M., Weinreb, R.N., 2009. Reproducibility of RTVue retinal nerve fiber layer thickness and optic disc measurements and agreement with Stratus optical coherence tomography measurements. *American Journal of Ophthalmology* 147, 1067–1074, e1.
- Harwerth, R.S., Vilupuru, A.S., Rangaswamy, N.V., Smith, E.L., 2007. The relationship between nerve fiber layer and perimetry measurements. *Investigative Ophthalmology & Visual Science* 48, 763–773, 3rd.
- Hood, D.C., Kardon, R.H., 2007. A framework for comparing structural and functional measures of glaucomatous damage. *Progress in Retinal and Eye Research* 26, 688–710.
- Hood, D.C., Anderson, S.C., Wall, M., Kardon, R.H., 2007. Structure versus function in glaucoma: an application of a linear model. *Investigative Ophthalmology & Visual Science* 48, 3662–3668.
- Jolliffe, I., 2002. *Principal Component Analysis*. Springer-Verlag.
- Leung, C.K., Cheung, C.Y., Weinreb, R.N., Qiu, K., Liu, S., Li, H., Xu, G., Fan, N., Pang, C.P., Tse, K.K., Lam, D.S., 2010. Evaluation of retinal nerve fiber layer progression in glaucoma: a study on optical coherence tomography guided progression analysis. *Investigative Ophthalmology & Visual Science* 51, 217–222.
- Leventon, M., 2000. *Statistical models for medical image analysis*. MIT Ph.D. Thesis.
- Little, R., Rubin, D., 2002. *Statistical Analysis with Missing Data*. Wiley.
- Pearson, K., 1901. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine* 2, 559–572.
- Prevent Blindness America Survey, 2002. Prevent Blindness America, 211 West Wacker Drive, Suite 1700, Chicago, Illinois 60606.
- Raiko, T., Ilin, A., Karhunen, J., 2007. Principal component analysis for large scale problems with lots of missing values. *European Conference on Machine Learning*, 691–698.
- Roweis, S., 1998. EM algorithms for PCA and SPCA. *Advances in Neural Information Processing Systems*, 626–632.
- Scholkopf, B., Smola, A., Muller, K.R., 1998. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation* 10, 1299–1319.
- Schuman, J.S., Pedut-Kloizman, T., Hertzmark, E., Hee, M.R., Wilkins, J.R., Coker, J.G., Puliafito, C.A., Fujimoto, J.G., Swanson, E.A., 1996. Reproducibility of nerve fiber layer thickness measurements using optical coherence tomography. *Ophthalmology* 103, 1889–1898.
- Shen, L.Q., Mock, D., Sami, M., Caprioli, J., 2010. Regional macular thickness correlates with glaucomatous visual field loss. *American Glaucoma Society Annual Meeting*, March, 2010, Naples, FL.
- Tipping, M.E., Bishop, C.M., 1999. Probabilistic principal component analysis. *Journal of the Royal Statistical Society* 61, 611–622.
- Tsai, A., Wells, W., Tempny, C., Grimson, E., Willsky, A., 2004. Mutual information in coupled multi-shape model for medical image segmentation. *Medical Image Analysis* 8 (4), 429–445.
- Wojtkowski, M., Srinivasan, V., Fujimoto, J.G., Ko, T., Schuman, J.S., Kowalczyk, A., Duker, J.S., 2005. Three-dimensional retinal imaging with high-speed ultrahigh-resolution optical coherence tomography. *Ophthalmology* 112, 1734–1746.
- Wolf-Schnurrbusch, U.E., Ceklic, L., Brinkmann, C.K., Illiev, M.E., Frey, M., Rothenbuehler, S.P., Enzmann, V., Wolf, S., 2009. Macular thickness measurements in healthy eyes using six different optical coherence tomography instruments. *Investigative Ophthalmology* 50 (7), 3432–3437.
- Yang, J., Zhang, D., Frangi, A.F., Jing-yu, Y., 2004. Two dimensional PCA: a new approach to appearance-based face representation and recognition. *IEEE Transactions on PAMI* 26, 131–137.