Analysis of macular OCT images using deformable registration

Min Chen,[1,2,*] Andrew Lang,[1] Howard S. Ying,[3] Peter A. Calabresi, [4] Jerry L. Prince,[1] and Aaron Carass[1]

[1]*Department of Electrical and Computer Engineering, The Johns Hopkins University, Baltimore, MD, 21218, USA*
[2]*Translational Neuroradiology Unit, National Institute of Neurological Disorders and Stroke, Bethesda, MD 20892, USA*
[3]*Wilmer Eye Institute, The Johns Hopkins School of Medicine, Baltimore, MD 21287, USA*
[4]*Department of Neurology, The Johns Hopkins School of Medicine, Baltimore, MD 21287, USA*
[*]mchen55@jhu.edu

Abstract: Optical coherence tomography (OCT) of the macula has become increasingly important in the investigation of retinal pathology. However, deformable image registration, which is used for aligning subjects for pairwise comparisons, population averaging, and atlas label transfer, has not been well–developed and demonstrated on OCT images. In this paper, we present a deformable image registration approach designed specifically for macular OCT images. The approach begins with an initial translation to align the fovea of each subject, followed by a linear rescaling to align the top and bottom retinal boundaries. Finally, the layers within the retina are aligned by a deformable registration using one-dimensional radial basis functions. The algorithm was validated using manual delineations of retinal layers in OCT images from a cohort consisting of healthy controls and patients diagnosed with multiple sclerosis (MS). We show that the algorithm overcomes the shortcomings of existing generic registration methods, which cannot be readily applied to OCT images. A successful deformable image registration algorithm for macular OCT opens up a variety of population based analysis techniques that are regularly used in other imaging modalities, such as spatial normalization, statistical atlas creation, and voxel based morphometry. Examples of these applications are provided to demonstrate the potential benefits such techniques can have on our understanding of retinal disease. In particular, included is a pilot study of localized volumetric changes between healthy controls and MS patients using the proposed registration algorithm.

© 2014 Optical Society of America

OCIS codes: (100.0100) Image processing; (170.4470) Ophthalmology; (170.4500) Optical coherence tomography.

References and links

1. J. G. Fujimoto, W. Drexler, J. S. Schuman, and C. K. Hitzenberger, "Optical coherence tomography (OCT) in ophthalmology: Introduction," Opt. Express **17**, 3978–3979 (2009).
2. E. M. Frohman, J. G. Fujimoto, T. C. Frohman, P. A. Calabresi, G. Cutter, and L. J. Balcer, "Optical coherence tomography: A window into the mechanisms of multiple sclerosis," Nat. Clin. Pract. Neuro. **4**, 664–675 (2008).
3. J. N. Ratchford, S. Saidha, E. S. Sotirchos, J. A. Oh, M. A. Seigo, C. Eckstein, M. K. Durbin, J. D. Oakley, S. A. Meyer, A. Conger, T. C. Frohman, S. D. Newsome, L. J. Balcer, E. M. Frohman, and P. A. Calabresi, "Active MS is associated with accelerated retinal ganglion cell/inner plexiform layer thinning," Neurology **80**, 47–54 (2013).

4. S. Saidha, E. S. Sotirchos, J. Oh, S. B. Syc, M. A. Seigo, N. Shiee, C. Eckstein, M. K. Durbin, J. D. Oakley, S. A. Meyer, T. C. Frohman, S. Newsome, J. N. Ratchford, L. J. Balcer, D. L. Pham, C. M. Crainiceanu, E. M. Frohman, D. S. Reich, and P. A. Calabresi, "Relationships between retinal axonal and neuronal measures and global central nervous system pathology in Multiple Sclerosis," JAMA Neurology **70**, 34–43 (2013).

5. P. A. Keane, P. J. Patel, S. Liakopoulos, F. M. Heussen, S. R. Sadda, and A. Tufail, "Evaluation of age-related macular degeneration with optical coherence tomography," Surv. Ophthalmol. **57**, 389–414 (2012).

6. G. Querques, R. Lattanzio, L. Querques, C. Del Turco, R. Forte, L. Pierro, E. H. Souied, and F. Bandello, "Enhanced depth imaging optical coherence tomography in Type 2 diabetes," Invest. Ophthalmol. Vis. Sci. **53**, 6017–6024 (2012).

7. Y. Lu, Z. Li, X. Zhang, B. Ming, J. Jia, R. Wang, and D. Ma, "Retinal nerve fiber layer structure abnormalities in early Alzheimer's disease: Evidence in optical coherence tomography," Neurosci. Lett. **480**, 69–72 (2010).

8. M. E. Hajee, W. F. March, D. R. Lazzaro, A. H. Wolintz, E. M. Shrier, S. Glazman, and I. G. Bodis-Wollner, "Inner retinal layer thinning in Parkinson disease," Arch. Ophthalmol. **127**, 737–741 (2009).

9. V. Guedes, J. S. Schuman, E. Hertzmark, G. Wollstein, A. Correnti, R. Mancini, D. Lederer, S. Voskanian, L. Velazquez, H. M. Pakter, T. Pedut-Kloizman, J. G. Fujimoto, and C. Mattox, "Optical coherence tomography measurement of macular and nerve fiber layer thickness in normal and glaucomatous human eyes," Ophthalmology **110**, 177–189 (2003).

10. D. Koozekanani, K. Boyer, and C. Roberts, "Retinal thickness measurements from optical coherence tomography using a Markov boundary model," IEEE Trans. Med. Imag. **20**, 900–916 (2001).

11. H. Ishikawa, D. M. Stein, G. Wollstein, S. Beaton, J. G. Fujimoto, and J. S. Schuman, "Macular segmentation with optical coherence tomography," Invest. Ophthalmol. Vis. Sci. **46**, 2012–2017 (2005).

12. M. K. Garvin, M. D. Abràmoff, X. Wu, S. R. Russell, T. L. Burns, and M. Sonka, "Automated 3-D intraretinal layer segmentation of macular spectral-domain optical coherence tomography images," IEEE Trans. Med. Imag. **28**, 1436–1447 (2009).

13. S. J. Chiu, X. T. Li, P. Nicholas, C. A. Toth, J. A. Izatt, and S. Farsiu, "Automatic segmentation of seven retinal layers in SDOCT images congruent with expert manual segmentation," Opt. Express **18**, 19413–19428 (2010).

14. A. Lang, A. Carass, E. Sotirchos, and J. L. Prince, "Segmentation of retinal OCT images using a random forest classifier," in "Proc. SPIE-MI 2013," (Lake Buena Vista, FL, 2013).

15. A. Lang, A. Carass, M. Hauser, E. S. Sotirchos, P. A. Calabresi, H. S. Ying, and J. L. Prince, "Retinal layer segmentation of macular OCT images using boundary classification," Biomed. Opt. Express **4**, 1133–1152 (2013).

16. A. Sotiras, C. Davatzikos, and N. Paragios, "Deformable medical image registration: A survey." IEEE Trans. Med. Imag. **32**, 1153–1190 (2013).

17. M. I. Miller, G. E. Christensen, Y. Amit, and U. Grenander, "Mathematical textbook of deformable neuroanatomies," Proc. Natl. Acad. Sci. **90**, 11944–11948 (1993).

18. J. Ashburner and K. J. Friston, "Voxel-based morphometry—the methods," NeuroImage **11**, 805821 (2000).

19. B. B. Avants, C. L. Epstein, M. Grossman, and J. C. Gee, "Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain," Med. Image Anal. **12**, 26–41 (2008).

20. M. Auer, P. Regitnig, and G. A. Holzapfel, "An automatic nonrigid registration for stained histological sections," IEEE Trans. Imag. Proc. **14**, 475–486 (2005).

21. K. K. Brock, M. B. Sharpe, L. A. Dawson, S. M. Kim, and D. A. Jaffray, "Accuracy of finite element model-based multi-organ deformable image registration," Med. Phys. **32**, 1647–1659 (2005).

22. W. Bai and M. Brady, "Motion correction and attenuation correction for respiratory gated PET images," IEEE Trans. Med. Imag. **30**, 351–365 (2011).

23. T. M. Jørgensen, J. Thomadsen, U. Christensen, W. Soliman, and B. Sander, "Enhancing the signal-to-noise ratio in ophthalmic optical coherence tomography by image registration—method and clinical examples," J. Biomed. Opt. **12**, 041208–041208 (2007).

24. J. Xu, H. Ishikawa, G. Wollstein, L. Kagemann, and J. S. Schuman, "Alignment of 3-d optical coherence tomography scans to correct eye movement using a particle filtering," IEEE Trans. Med. Imag. **31**, 1337–1345 (2012).

25. Y. M. Liew, R. A. McLaughlin, F. M. Wood, and D. D. Sampson, "Motion correction of in vivo three-dimensional optical coherence tomography of human skin using a fiducial marker," Biomed. Opt. Express **3**, 1774 (2012).

26. A. Giani, M. Pellegrini, A. Invernizzi, M. Cigada, and G. Staurenghi, "Aligning scan locations from consecutive spectral-domain optical coherence tomography examinations: A comparison among different strategies," Invest. Ophthalmol. Vis. Sci. **53**, 7637–7643 (2012).

27. M. Niemeijer, M. K. Garvin, K. Lee, B. van Ginneken, M. D. Abràmoff, and M. Sonka, "Registration of 3D spectral OCT volumes using 3D SIFT feature point matching," in "Proc. SPIE-MI 2009," (Lake Buena Vista, FL, 2009).

28. M. Niemeijer, K. Lee, M. K. Garvin, M. D. Abràmoff, and M. Sonka, "Registration of 3D spectral OCT volumes combining ICP with a graph-based approach," in "Proc. SPIE-MI 2012," (San Diego, CA, 2012).

29. A. A. Goshtasby, *2-D and 3-D Image Registration: For Medical, Remote Sensing, and Industrial Applications* (Wiley, 2005).

30. E. A. Maguire, D. G. Gadian, I. S. Johnsrude, C. D. Good, J. Ashburner, R. S. J. Frackowiak, and C. D. Frith, "Navigation-related structural change in the hippocampi of taxi drivers," Proc. Nat. Acad. Sci. **97**, 4398–4403 (2000).

31. C. D. Good, I. S. Johnsrude, J. Ashburner, R. N. A. Henson, K. J. Friston, and R. S. J. Frackowiak, "A voxel-based morphometric study of ageing in 465 normal adult human brains," NeuroImage **14**, 21–36 (2001).

32. A. F. Goldszal, C. Davatzikos, D. Pham, M. X. H. Yan, R. N. Bryan, and S. M. Resnick, "An image-processing system for qualitative and quantitative volumetric analysis of brain images," J. Computer Assisted Tomography **22**, 827–837 (1998).

33. C. Davatzikos, A. Genc, D. Xu, and S. M. Resnick, "Voxel-based morphometry using the RAVENS maps: Methods and validation using simulated longitudinal atrophy," NeuroImage **14**, 1361–1369 (2001).

34. S. M. Resnick, D. L. Pham, M. A. Kraut, A. Zonderman, and C. Davatzikos, "Longitudinal magnetic resonance imaging studies of older adults: A shrinking brain," J. Neurosci. **23**, 3295–3301 (2003).

35. M. Chen, A. Lang, E. Sotirchos, H. S. Ying, P. A. Calabresi, J. L. Prince, and A. Carass, "Deformable registration of macular oct using a-mode scan similarity," in "Biomedical Imaging (ISBI), 2013 IEEE 10th International Symposium on," (IEEE, 2013), pp. 476–479.

36. E. Gibson, M. Young, M. V. Sarunic, , and M. F. Beg, "Optic nerve head registration via hemispherical surface and volume registration," IEEE Trans. Biomed. Eng. **57**, 2592–2595 (2010).

37. B. Antony, M. D. Abràmoff, L. Tang, W. D. Ramdas, J. R. Vingerling, N. M. Jansonius, K. Lee, Y. H. Kwon, M. Sonka, and M. K. Garvin, "Automated 3-D method for the correction of axial artifacts in spectral-domain optical coherence tomography images," Biomed. Opt. Express **2**, 2403–2416 (2011).

38. Y. Zheng, R. Xiao, Y. Wang, and J. C. Gee, "A generative model for oct retinal layer segmentation by integrating graph-based multi-surface searching and image registration," in "16th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2013)," (Springer, 2013), pp. 428–435.

39. Y. Ou, A. Sotiras, N. Paragios, and C. Davatzikos, "Dramms: Deformable registration via attribute matching and mutual-saliency weighting," Med. Image Anal. **15**, 622–639 (2011).

40. A. N. Kuo, R. P. McNabb, S. J. Chiu, M. A. El-Dairi, S. Farsiu, C. A. Toth, and J. A. Izatt, "Correction of ocular shape in retinal optical coherence tomography and effect on current clinical measures," Am. J. Ophthalmol.**156**, 304–311 (2013).

41. G. K. Rohde, A. Aldroubi and B. M. Dawant, "The adaptive bases algorithm for intensity based nonrigid image registration," IEEE Trans. Med. Imag. **22**, 1470–1479 (2003).

42. A. Guimond, J. Meunier, and J.-P. Thirion, "Average brain models: A convergence study," Computer vision and image understanding **77**, 192–210 (2000).

43. B. Avants and J. C. Gee, "Geodesic estimation for large deformation anatomical shape averaging and interpolation," NeuroImage **23**, S139–S150 (2004).

44. S. Khullar, A. M. Michael, N. D. Cahill, K. A. Kiehl, G. Pearlson, S. A. Baum, and V. D. Calhoun, "ICA-fNORM: Spatial normalization of fMRI data using intrinsic group-ICA networks," Front. Syst. Neurosci **5** (2011).

45. D. Rueckert, A. F. Frangi, and J. A. Schnabel, "Automatic construction of 3-d statistical deformation models of the brain using nonrigid registration," IEEE Trans. Med. Imag. **22**, 1014–1025 (2003).

46. M. Chen, A. Carass, D. Reich, P. Calabresi, D. Pham, and J. Prince, "Voxel-wise displacement as independent features in classification of multiple sclerosis." in "Proc. SPIE-MI 2013," (Lake Buena Vista, FL, 2013).

47. S. Gerber, T. Tasdizen, P. T. Fletcher, S. Joshi, and R. Whitaker, "Manifold modeling for brain population analysis," Med. Image Anal. **14**, 643–653 (2010).

48. Y. Fan, S. M. Resnick, X. Wu, and C. Davatzikos, "Structural and functional biomarkers of prodromal Alzheimer's disease: A high-dimensional pattern classification study," NeuroImage **41**, 277–285 (2008).

49. S. Saidha, E. S. Sotirchos, M. A. Ibrahim, C. M. Crainiceanu, J. M. Gelfand, Y. J. Sepah, J. N. Ratchford, J. Oh, M. A. Seigo, S. D. Newsome, L. J. Balcer, E. M. Frohman, A. J. Green, Q. D. Nguyen, and P. A. Calabresi, "Microcystic macular oedema, thickness of the inner nuclear layer of the retina, and disease characteristics in multiple sclerosis: A retrospective study," The Lancet Neurology **11**, 963–972 (2012).

50. A. Klein, S. S. Ghosh, B. Avants, B. T. T. Yeo, B. Fischl, B. Ardekani, J. C. Gee, J. J. Mann, and R. V. Parsey, "Evaluation of volume-based and surface-based brain image registration methods," NeuroImage **51**, 214–220 (2010).

51. B. B. Avants, N. J. Tustison, G. Song, P. A. Cook, A. Klein, and J. C. Gee, "A reproducible evaluation of ANTs similarity metric performance in brain image registration," NeuroImage **54**, 2033–2044 (2011).

52. L. R. Dice, "Measures of the amount of ecologic association between species," Ecology **26**, 297–302 (1945).

53. K. J. Friston, J. Ashburner, S. Kiebel, T. Nichols, and W. Penny, eds., *Statistical Parametric Mapping: The Analysis of Functional Brain Images* (Academic Press, 2007).

54. J. B. Kerrison, T. Flynn, and W. R. Green, "Retinal pathologic changes in multiple sclerosis," Retina **14**, 445–451 (1994).

55. A. J. Green, S. McQuaid, S. L. Hauser, I. V. Allen, and R. Lyness, "Ocular pathology in multiple sclerosis: retinal atrophy and inflammation irrespective of disease duration," Brain **133**, 1591–1601 (2010).

56. B. J. Lujan, A. Roorda, R. W. Knighton, and J. Carroll, "Revealing henle's fiber layer using spectral domain optical coherence tomography," Invest. Ophthalmol. Visual Sci. **52**, 1486–1492 (2011).

57. B. C. Lucas, J. A. Bogovic, A. Carass, P.-L. Bazin, J. L. Prince, D. L. Pham, and B. A. Landman, "The Java Image Science Toolkit (JIST) for rapid prototyping and publishing of neuroimaging software," Neuroinformatics **8**, 5–17 (2010).

1. Introduction

Optical coherence tomography (OCT) is an imaging technology that enables objective analysis of the mechanisms of neurodegeneration in an important, yet isolated, portion of the central nervous system—the retina. It presents cellular level descriptions by providing micrometer (μm) resolution imaging of the retina based on the optical scattering properties of biological tissues. OCT offers several other distinguishing advantages such as patient comfort, ease of use, no ionizing radiation, quick acquisition, and low cost. These properties have allowed the ophthalmic community to better explore the retinal cell layers within the macula [1]. OCT has been particularly useful in the assessment and characterization of multiple sclerosis (MS) through the analysis of the thickness of various retinal layers [2–4]. It has also seen use for exploring numerous other diseases, including age-related macular degeneration [5], diabetes [6], Alzheimer's disease [7], Parkinson's disease [8], and glaucoma [9].

While automated and semi-automated methods for analyzing and segmenting retinal layers in OCT have been introduced regularly over the past decade [10–15], there has been very little development of image registration methods for OCT. *Image registration* is the task of transforming different images into the same geometric space or coordinate system. This allows for both *intersubject registration*, the alignment of images from different subjects, and *longitudinal registration*, the alignment of multiple images from the same subject at different time points. There has been extensive literature surrounding image registration and its application in medical imaging [16]. It has been widely used in a variety of neuroanatomy settings [17–19] and for various vital organs [20–22]. Despite this rich history, the development of registration techniques in OCT has been relatively limited. Existing works are restricted almost exclusively to *rigid* registrations [23–28], where only rigid body transformations are allowed in the alignment. Such transformation models are generally inadequate for the retina, since the differences between retinal layers from different subjects (or even the same subject over time) cannot be fully represented by such a low-dimensional model. To capture these differences, an *affine* or *deformable* registration is necessary, where affine or nonlinear freeform transformations are used in the model, respectively. Deformable registration is also divided into subgroups, including parametric models which represent transformations using basis functions, and physical models, which use transformations allowed in elastic, viscous, or viscoelastic materials. Each of these subgroups have their own strengths and weaknesses (see [29] for a detailed discussion).

One potential application of retinal OCT registration is to combine macular scans from a population of subjects to create a macular stereotaxic space, often referred to as a *normalized space*, which can serve as a standard reference space for comparing between different subjects. Such a space directly enables a number of advanced analytic techniques such as voxel based morphometry (VBM) [18], which allows the exploration of local tissue composition differences between two populations, and has seen extensive use in analyzing functional and structural differences in the brain [30, 31]. A normalized space can also be used to study localized volumetric changes between populations, for example, by using the Regional Analysis of Volumes Examined in Normalized Space (RAVENS) [32, 33] method. RAVENS has been used to accurately quantify localized volume differences in the brain [34], and readily generates visualizations, called RAVENS maps, of the difference between two populations. The technique uses the deformations learned from each subject registration to compute the local expansions and contractions of tissue volume relative to the normalized space. In retinal OCT, these types

of analyses can allow local investigation of layer volume changes, which is not possible from solely examining layer thickness calculations from segmentation results.

To our knowledge, excluding our own preliminary work [35], only three other algorithms have been reported that apply a nonlinear deformable model to register retinal OCT data. Gibson et al. [36] presented a method that segments and extracts the optic nerve head surface from the OCT image, and then registers the surfaces using a combined surface and volumetric registration. The method was validated by comparing the overlap of the optical nerve head after the registration. Antony et al. [37] introduced a method where the surface between the inner and outer segments of the photoreceptor cells are used with a thin-plate spline to correct for axial artifacts. Lastly, Zheng et al. [38] presented a segmentation approach that uses the SyN [19] registration algorithm to aid in segmentation. The method uses approximate segmentations of the retinal layers to divide the image into regions, and then registers each region individually with training data to estimate a more accurate segmentation. Of these three methods, only Gibson et al. provided an inter-subject registration that allows for the construction of a normalized space and population analysis. The other two methods primarily used the registration internally to aid their distortion correction and segmentation algorithm. All three of these deformable registration methods required a segmentation of the OCT data to guide the registration algorithm. This requirement restricts the registration accuracy to that of the segmentation. Since OCT segmentations only provide information about the positions of layer boundaries, these registration algorithms are potentially inaccurate within the segmented regions.

Intensity-based volumetric registration, i.e., where the alignment is performed by matching the voxel intensities between the two images, is an alternative to existing segmentation-based methods. To our knowledge, the only intensity-based volumetric deformable registration that has been developed and validated specifically on macular OCT images is our preliminary work [35]. There are a number of generic volumetric registration algorithms (most of which were developed for brain image registration) that can be applied to OCT data. However, our investigations have found that such algorithms tend to be unreliable when applied to the whole OCT image, without separating and registering individual layers as in Zheng et al [38]. Figure 1 shows two OCT images and the results of registering a subject image to a target image using both SyN [19] and DRAMMS [39], two highly-rated deformable registration algorithms developed for brain image registration, with default settings.

The lack of deformable registration tools for retinal OCT can be attributed to several factors. First, OCT (particularly spectral-domain OCT) is a relatively new modality in comparison to other medical imaging modalities such as magnetic resonance imaging (MRI), computed tomography, and ultrasound. Hence, the demand and necessity for algorithms and techniques designed for OCT images is fairly recent. Second, there are a number of characteristics inherent to OCT images which make adaptation and development of existing registration algorithms challenging. These characteristics include poor signal-to-noise ratio (SNR) in comparison to other imaging modalities, which makes it difficult for intensity based methods to find correct correspondences. The data is often extremely anisotropic, where the gap between the cross-sectional slices (B-scans) can be upwards of 30 times larger than the resolution within each B-scan. This is generally not compatible with the regularization inherent in most transformation models in existing algorithms, and also creates significant challenges when performing interpolation. Third, the geometry of OCT images are not well defined between different subjects. The scan lines (A-scans) in the images are generally represented parallel to each other, while physically they should be fanning outwards from the center [40]. Since this geometry is specific to the optics of each individual's eye, it is unclear how transformations between retinal OCT images should be performed, particularly when deforming across different A-scans. Finally, the high resolution of OCT images create a large data size that poses a computational challenge. In the examples shown

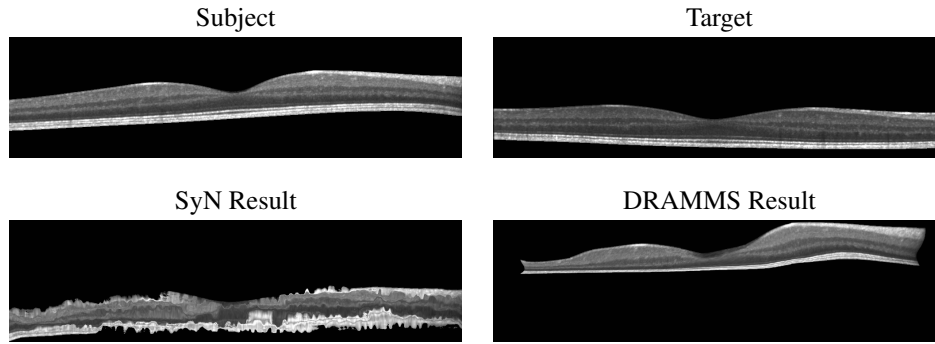| Subject	Target

SyN Result	DRAMMS Result

Fig. 1. The top row shows B-scans from two OCT images used as the subject and target. The second row shows example results when using the default settings for two generic deformable registration algorithms, SyN [19] and DRAMMS [39], to register the subject image to the target.

in Fig. 1, each algorithm took on average 8–10 hours to run; this is significantly longer than for a standard MR brain registration, which takes just over an hour using the same algorithms. Together, these traits lead to relatively poor performance when applying existing approaches directly to retinal OCT images. The successful development of deformable image registration techniques for OCT can offer the ability to greatly enhance our understanding of the complex interplay of macular degeneration with disease and aging from a population perspective.

2. Methods

2.1. Initial preprocessing

Prior to performing the OCT registration, our method requires two initial preprocessing steps that are applied to the OCT images. First, in order to reduce intensity outliers and develop a compatible scale for intensity comparison, we apply an intensity normalization process to each OCT volume. Second, we estimate and mask the images to the retinal boundaries by locating the inner limiting membrane (ILM) and Bruch's membrane (BrM), which also serve as landmarks for our affine registration step. The following subsections will briefly summarize these two steps. Both techniques have been used with success in previous work for layer segmentation [15], where more detailed explanations can be found. The figures shown in this paper have all been normalized and masked using these two techniques.

2.1.1. Intensity normalization

Our deformable registration algorithm is based on intensity differences between the subject and target images. Hence, we require that the images have similar intensity ranges, and that the values observed in a particular tissue type should be consistent across populations of images. However, OCT data often shows considerable intensity variability. Even B-scans within the same volume often have very different intensity ranges which cannot be attributed to tissue differences. Possible causes of this are (1) the automatic intensity rescaling performed by the scanner, which can be skewed by high intensity reflection artifacts, and (2) the automatic real-time averaging performed by the scanner, which can result in a specific B-scan undergoing more averaging, creating differences in its dynamic range.

To address these issues, we perform intensity normalization using a straightforward linear contrast stretching on each B-scan independently. Intensity values in the range $[0, I_{max}]$ are

linearly rescaled to $[0, 1]$, with values larger than I_{max} being capped at the maximum. The cutoff, I_{max}, is determined robustly by first median-filtering each individual A-scan within the same B-scan with a kernel size of 15 pixels (58 μm), and then setting I_{max} as 5% larger than the maximum intensity in the median filtered image.

2.1.2. Retinal boundary detection and fovea localization

Our next preprocessing step is estimating and masking the images to the top and bottom boundaries of the retina, which are defined by the ILM and BrM, respectively. These retinal extents are found in the following manner. Each B-scan is independently Gaussian smoothed ($\sigma = 3$ pixels which is $\sigma_{(x,y)} = (17, 12)$ μm), followed by an image gradient computation along each A-scan using a Sobel kernel. In each A-scan, the two largest positive gradient values that are more than 25 pixels (97 μm) apart are taken as initial estimates of either the ILM or the inner segment (IS) outer segment (OS) boundary. We estimate the BrM, by searching up to 30 pixels (116 μm) below the IS-OS boundary for the largest negative gradient. These boundary estimates are refined by comparing to the median filtered collection of the gradients, and then Gaussian smoothing the result. Using this detection, we mask out non-retinal material and approximate the location of the fovea as the superior point of the thinnest portion of the retina within a search window at the center of the image.

2.2. Image registration method

The primary goal of image registration is to estimate a transformation that maps corresponding locations between a subject image $\mathcal{S}(\mathbf{x}')$ and a target image $\mathcal{T}(\mathbf{x})$. Here, $\mathbf{x}' = (x', y', z')$ and $\mathbf{x} = (x, y, z)$ describe 3D coordinates in the subject and target image domains, $\mathbb{D}_{\mathcal{S}}$ and $\mathbb{D}_{\mathcal{T}}$, respectively, and $\mathcal{S}(\mathbf{x}')$ and $\mathcal{T}(\mathbf{x})$ are the intensities of each image at those coordinates. For our OCT images, we assign the x, y, and z axes as the lateral, through-plane, and axial directions, respectively. This makes A-scan lines parallel to the z axis and B-scans as images parallel to the xz–plane.

We describe the transformation that the registration algorithm is attempting to solve as the mapping $\mathbf{v} \colon \mathbb{D}_{\mathcal{T}} \to \mathbb{D}_{\mathcal{S}}$. This is generally represented as a *pullback* vector field, $\mathbf{v}(\mathbf{x})$, where the vectors are rooted in the target domain and point to locations in the subject domain. The field is applied to $\mathcal{S}(\mathbf{x}')$ by pulling subject image intensities into the target domain. This produces the registration result, a transformed subject image, $\widetilde{\mathcal{S}}$, defined as

$$\widetilde{\mathcal{S}}(\mathbf{x}) = \mathcal{S}(\mathbf{v}(\mathbf{x})), \quad \forall \mathbf{x} \in \mathbb{D}_{\mathcal{T}}, \tag{1}$$

which has coordinates in the target domain. The goal of the registration algorithm is to find $\mathbf{v}$ such that the images $\widetilde{\mathcal{S}}$ and $\mathcal{T}$ are as similar as possible. This is performed by minimizing a cost function that evaluates how similar the intensities in $\mathcal{S}(\mathbf{v}(\mathbf{x}))$ are to $\mathcal{T}(\mathbf{x})$ at each $\mathbf{x}$, while constraining the transformation to be smooth and continuous (i.e., physically sensible).

As described in Section 1, the two major challenges of retinal OCT registration are (1) the image voxels tend to be highly anisotropic, and (2) the physical geometry of the image is not properly defined by how the A-scans are presented in the OCT. The former makes regularization and interpolation difficult, if not infeasible, for images with large B-scan separations. The latter obfuscates our ability to apply a sensible deformation that respects the physical space being imaged. This is particularly problematic when attempting to apply a deformation that crosses multiple A-scans. For example, a deformation parallel to the x or y axes in the OCT image is actually curved in physical space. The amount of curvature depends on the fanning of the A-scan during the acquisition, which depends on the optics of the eye being imaged. Such information is often not acquired with the OCT, which makes the actual deformation applied in the physical space ambiguous.

To address these concerns, we impose strong restrictions on the class of transformations our registration is allowed to estimate. In the x (lateral) and y (through-plane) directions, we permit only discrete translations such that A-scans from the subject always coincide with A-scans in the target (except for missing scans at the boundaries). This removes the need to interpolate intensities between two A-scans or two B-scans. Non-rigid transformations are only allowed in the z (axial) direction and will be constructed as a composition of individual (A-scan to A-scan) affine and deformable registration steps. This will permit accurate alignment of the features in each A-scan, including the retinal layers.

Our overall registration result is the composition of three steps: a 2D global translation of the whole volume (using discrete offsets only), a set of 1D affine transformations applied to each A-scan, and a set of 1D deformable transformations applied to each A-scan. These transformations are learned and applied in stages and then composed to construct the total 3D transformation $\mathbf{v}$. Let $\mathbf{r}$ represent the 2D global translation, and let $\mathbf{a}_{i,j}$ and $\mathbf{d}_{i,j}$ represent, respectively, the 1D affine and 1D deformable transformations applied to the A-scan indexed by the discrete x and y coordinates, $m = (i, j)$. We can write the total transformation as $\mathbf{v}(\mathbf{x}) = \mathbf{r} \circ \mathbf{a}_m \circ \mathbf{d}_m(\mathbf{x})$, where it is assumed that $\mathbf{x}$ belongs to the A-scan indexed by m. The deformed subject image in the template space is then given by

$$\widetilde{\mathcal{S}}(\mathbf{x}) = \mathcal{S}(\mathbf{r}(\mathbf{a}_m(\mathbf{d}_m(\mathbf{x})))), \quad \forall \mathbf{x} \in \mathbb{D}_{\mathcal{T}}. \tag{2}$$

In the following sections we present the method to carry out each of these steps.

2.2.1. 2D global translation

The goal of the 2D global translation is to align the subject such that its fovea is aligned with the same A-scan as the fovea in the target image. Let the positions of the foveae in the subject and target images be denoted by $\mathbf{f}_{\mathcal{S}}$ and $\mathbf{f}_{\mathcal{T}}$, respectively. These locations are found using the approach described in Section 2.1.2, which forces them to be located within A-scans in the two images. The rigid transformation is therefore described by

$$\mathbf{r}(\mathbf{x}) = \mathbf{x} - \mathbf{t}, \tag{3}$$

where

$$\mathbf{t} = \begin{bmatrix} (\mathbf{f}_{\mathcal{S}})_x - (\mathbf{f}_{\mathcal{T}})_x \\ (\mathbf{f}_{\mathcal{S}})_y - (\mathbf{f}_{\mathcal{T}})_y \\ 0 \end{bmatrix}. \tag{4}$$

The notation $(\mathbf{f}_{\mathcal{S}})_x$ and $(\mathbf{f}_{\mathcal{S}})_y$ denote the x and y-components of $\mathbf{f}_{\mathcal{S}}$, and likewise for $(\mathbf{f}_{\mathcal{T}})_x$ and $(\mathbf{f}_{\mathcal{T}})_y$. Figure 2(a) shows an example of a result after applying just this rigid 2D foveal translation to the subject image.

2.2.2. 1D affine transformation

The 1D affine transformation step assumes that the foveae are aligned and will now consider each A-scan separately. The goal is to find the 1D affine transformation that matches the ILM and BrM boundary locations (both found in Section 2.1.2) between each pair of foveal aligned A-scans in the subject and target images. Since there are two landmarks and two unknowns in a 1D affine transformation, the solution can be written in closed form. For each A-scan m the ILM and BrM boundaries are scalar locations along the A-scan. We refer to these as $\mathbf{i}_{\mathcal{S},m}$ and $\mathbf{b}_{\mathcal{S},m}$ for the subject image and $\mathbf{i}_{\mathcal{T},m}$ and $\mathbf{b}_{\mathcal{T},m}$ for the target image, respectively. The 1D affine transformation for that A-scan is then given by

$$\mathbf{a}_m(\mathbf{x}) = F_m \mathbf{x} + \mathbf{t}_m, \tag{5}$$

Subject Image
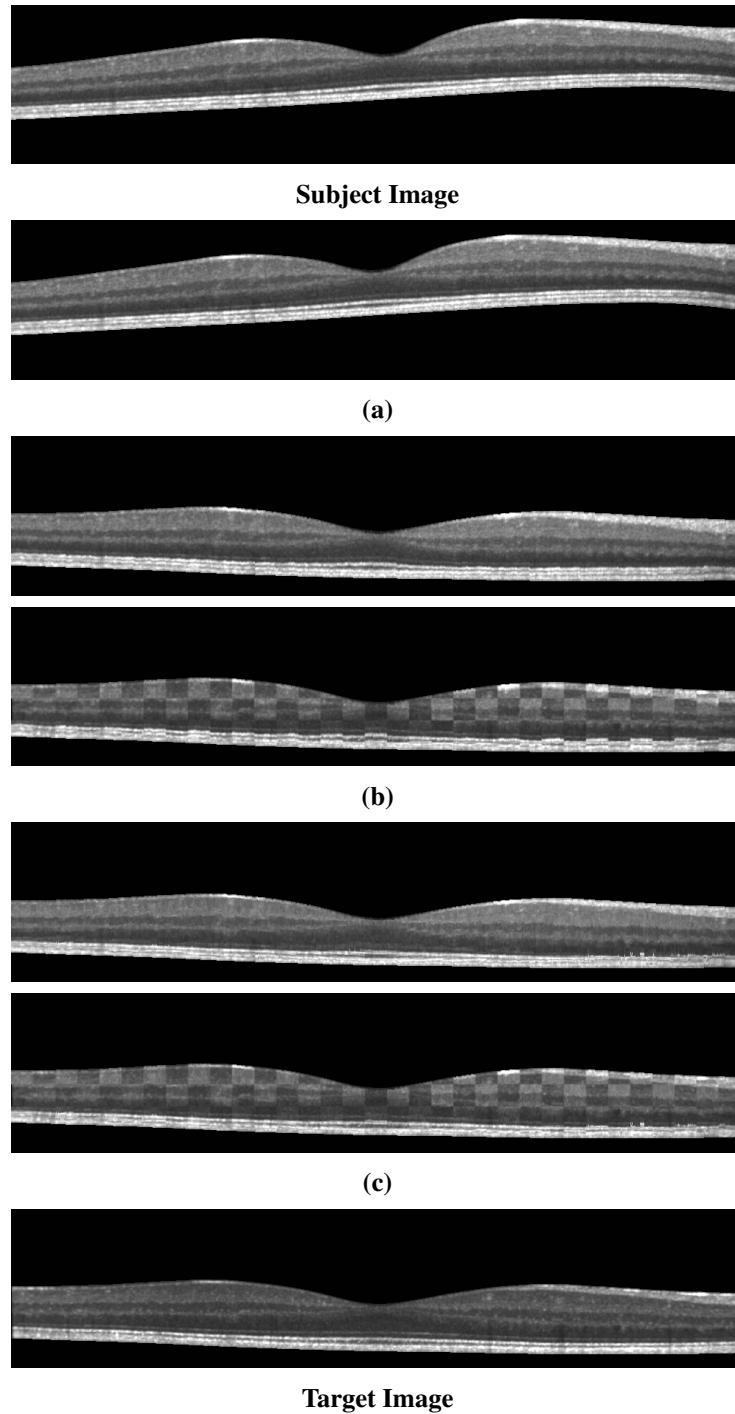
(a)

(b)

(c)

Target Image

Fig. 2. Shown are the outputs at each step of the registration algorithm, from the subject image $\mathcal{S}$ at the top, to the results of the registration after **(a)** the rigid alignment to the fovea, **(b)** the affine alignment of the A-scans and **(c)** the deformable registration. The target image $\mathcal{T}$ is shown at the bottom for reference. The bottom figures in **(b)** and **(c)** show checkerboard comparisons between each result and the target image.

where

$$F_m = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & r_m \end{bmatrix}, \tag{6}$$

and

$$\mathbf{t}_m = \begin{bmatrix} 0 \\ 0 \\ \dfrac{\mathbf{i}_{\mathcal{S},m} + \mathbf{b}_{\mathcal{S},m}}{2} - r_m \left(\dfrac{\mathbf{i}_{\mathcal{T},m} + \mathbf{b}_{\mathcal{T},m}}{2} \right) \end{bmatrix}. \tag{7}$$

The scale factor, r_m, is calculated from the ILM and BrM locations for each A-scan using

$$r_m = \frac{(\mathbf{b}_{\mathcal{S},m} - \mathbf{i}_{\mathcal{S},m})}{(\mathbf{b}_{\mathcal{T},m} - \mathbf{i}_{\mathcal{T},m})}. \tag{8}$$

We refer to the combination of the rigid translation step and this collection of 1D affine transformations at each A-scan as our *A-OCT* registration. Figure 2(b) shows the resulting image after applying this step to the foveal aligned result.

2.2.3. 1D deformable transformation

Following the A-OCT registration, the next step in our algorithm is to use a 1D deformable registration to further improve the alignment of the retinal layers. Unlike the 1D affine registration, where the transformation can be represented as a matrix, the deformable transformation $\mathbf{d}_m$ for each A-scan m is a free-form mapping at each $\mathbf{x}$ with the restriction that the deformation must be smooth and can only occur along the A-scan (z) direction. In this registration, $\mathbf{d}_m$ is modeled as a summation of radial basis functions (RBFs), $\Phi(\mathbf{x})$, using

$$\mathbf{d}_m(\mathbf{x}) = \mathbf{x} + \sum_i c_i \Phi(\mathbf{x} - \mathbf{x}_i), \tag{9}$$

where c_i and $\mathbf{x}_i$ determine the size and center of each RBF, respectively. For $\Phi(\mathbf{x})$, we choose the same RBF as presented in [41], except with the deformation restricted only to the z direction:

$$\Phi(\mathbf{x}) = \begin{bmatrix} 0 \\ 0 \\ \phi\left(\frac{\|\mathbf{x}\|}{s} \right) \end{bmatrix}, \quad \text{where} \quad \phi(r) = (1 - r)_+^4 (3r^3 + 12r^2 + 16r + 4), \tag{10}$$

for $(1 - r)_+ = \max(1 - r, 0)$, which has support s. This RBF has several beneficial properties such as smoothness, positive definiteness, and compact support. In addition, [41] showed that given the correct constraints on the size of each RBF, $\mathbf{d}_m$ is guaranteed to be homeomorphic. A homeomorphic deformation field allows the deformation to preserve the topology of the image and prevent *folding* and *tearing* of the underlying anatomy in the image.

To solve for $\mathbf{d}_m$, the algorithm uniformly places the RBFs along the A-scan and then iteratively minimizes the energy function

$$E_m^{\mathrm{SSD}} = \sum_{\{\mathbf{x} | (x,y) = m\}} \left(\mathcal{T}(\mathbf{x}) - \mathcal{S}_{\text{(A-OCT)}}(\mathbf{d}_m(\mathbf{x})) \right)^2, \tag{11}$$

which describes the intensity difference between the target image and the A-OCT result deformed by the current estimate of $\mathbf{d}_m$, across a particular A-scan. This is often referred to as the *sum of square differences* (SSD) similarity measure.

One disadvantage of optimizing the deformation field between only pairs of A-scans is that it can lead to discontinuities in the total deformation. It also ignores potentially useful neighboring information that can aid the optimization. To address this, we introduce a regularization term,

$$E_m^{\text{Reg}} = \sum_{\substack{r=-R \\ r\neq 0}}^{R} \sum_{\{\mathbf{x}|(x,y)=m+(r,0)\}} \frac{1}{|r|} \left(\mathcal{T}(\mathbf{x}) - \mathcal{S}_{\text{(A-OCT)}}(\mathbf{d}_m(\mathbf{x})) \right)^2, \quad (12)$$

where R is a parameter that determines how many adjacent A-scans, in the same B-scan, to use in the regularization. The current deformation $\mathbf{d}_m$ is applied at each adjacent A-scan $(m+(r,0))$ and then the function checks the SSD error produced by that deformation relative to the target image. Since we expect adjacent A-scans to be fairly continuous, if a deformation at m causes huge errors in adjacent neighbors, then that deformation is heavily penalized. A weight term is included to reduce the contribution of comparisons made further away from m, since we expect the deformation to be less applicable with distance.

We make two notes regarding this regularization term. First, the deformations of the adjacent A-scans with $\mathbf{d}_m$ are only used for the optimization of the deformation at the current A-scan. The deformations are not applied permanently to the adjacent A-scans. When the algorithm moves on to estimate a deformation at those adjacent A-scans, a new deformation is estimated. Second, the cost function does not regularize across B-scans, which follows our premise that the large separation between B-scans provide poor correspondences for the registration. Naturally, for data that do not suffer from this limitation, this regularization can be easily extended to use multiple B-scans as well.

The two energy terms are summed $(E^{\text{SSD}} + E^{\text{Reg}})$ to produce our total cost function, which we minimize at each A-scan to estimate $\mathbf{d}_m$. Brent's 1D line search optimization method is used to perform this minimization. After evaluating $\mathbf{d}_m$ at each A-scan, we have all the transformations required to create the final registration result, $\tilde{\mathcal{S}}$ from Eq. 2, which represent the subject image registered to the target domain. Figure 2(c) shows an example of this final result. We refer to the combination of all three registration steps as our *D-OCT* registration.

2.3. Constructing a normalized space using deformable registration

An important application of deformable registration is the ability to construct a normalized space for analyzing a population of subjects. The normalized space is a common target space to which all images in a group of subjects are registered. This allows spatial correspondences between the subjects to be observed and analyzed together as a population. Often, an average atlas is constructed for this purpose, where the atlas represents the average anatomy of the population. Constructing such an atlas and using it as a normalized space via deformable registration is a well studied topic in brain MRI analysis. A number of methods have been proposed for creating such a space [42–44]. In generating our normalized space, we follow the method presented in [42] where the average atlas is found by iteratively registering each subject to an estimate of the average atlas, and then adjusting the atlas such that the average of all the deformations is closer to zero.

2.4. Regional analysis of volumes examined in normalized space

The ability to construct a normalized space opens up numerous existing techniques for analyzing images from a population perspective, such as voxel based morphometry [18], statistical deformation models [45], tissue density maps [32], disease classification [46], and manifold learning [47]. To demonstrate this capability, we apply our registration approach in conjunction with the tissue density based analysis known as RAVENS. This voxel-wise analysis was introduced for neuroimaging in [32], and has been used to show localized brain volume changes in

Alzheimer's disease [48]. It uses a tissue segmentation with the deformation learned from each registration to the normalized space to estimate the relative local volume changes between each subject and the atlas.

The relative local volume change is calculated by first taking each voxel in a segmentation and projecting it into the target space using the registration deformation fields. This process keeps track of how each voxel is distributed by the projection and allows the method to record the degree of compression and expansion at each voxel that the segmentations had to undergo in order to be registered into the normalized space. As a result, we obtain localized measures of volume changes for each subject relative to the average atlas. This provides the ability to locate differences in relative volume changes between control and disease populations. We use this RAVENS method in Section 5.3 to explore the volume differences in the macula between a healthy control cohort and a patient cohort diagnosed with MS.

3. Materials

3.1. Data

We use two pools of data for the various experiments and comparisons in the remainder of the paper. For clarity, we will state in each section which cohort was used. All of the data was acquired using a Spectralis OCT system (Heidelberg Engineering, Heidelberg, Germany). The automatic real-time function was enabled and set to 12 averages, with all scans having SNR of at least 20 dB. Macular raster scans ($20° \times 20°$) were acquired with 49 B-scans, each B-scan having 1024 A-scans with 496 pixels per A-scan. The B-scan resolution varied slightly between subjects and averaged 5.8 μm laterally and 3.9 μm axially. The through plane distance (slice separation) averaged 123.6 μm between images, resulting in an imaging area of approximately 6×6 mm. The entry position of the OCT scan was aimed towards the center of the pupil to reduce the rotation of the macula in the image. The research protocol was approved by the local Institutional Review Board, and written informed consent was obtained from all participants.

3.1.1. Validation cohort

We use a *validation cohort* consisting of OCT images of the right eyes from 45 subjects. The 45 subjects consisted of 26 patients diagnosed with MS while the remaining 19 subjects were healthy controls. The 26 MS patients were screened and found to be free of microcystic macular edema [49], which our registration does not account for. An internally developed protocol was used to manually label nine layer boundaries on all B-scans for all subjects. These nine boundaries partition an OCT data set into eight regions of interest: 1) retinal nerve fiber layer (RNFL), 2) ganglion cell layer and inner plexiform layer (GCIP), 3) inner nuclear layer (INL), 4) outer plexiform layer (OPL), 5) outer nuclear layer (ONL), 6) inner segment (IS), 7) outer segment (OS), and 8) retinal pigment epithelium (RPE) complex. These retina layers are used to demonstrate the accuracy of our registration method, by comparing the algorithm's ability to use the learned deformation to transfer known labels in the subject image to an unlabeled target image [17] (see Section 4 for details).

3.1.2. General cohort

Our *general cohort* consists of retinal images from 83 subjects, which consisted entirely of right eyes with 40 scans from healthy controls and 43 scans from MS patients. Layer segmentations for this cohort was generated automatically using a boundary classification and graph based method [15], which has been shown to be highly reliable when compared to manual segmentations. This data collection is used to demonstrate several applications of our deformable registration and normalized space (see Section 5.1 for details). Our use of automated segmentations was motivated by the desire to demonstrate a fully automated processing pipeline. Manual

segmentations, while more accurate, tend to present a bottleneck in large studies with hundreds of data sets.

4. Registration validation

The 45 OCT images and their associated boundary labels in the validation cohort (Section 3.1.1) were used to evaluate the accuracy of our registration. This involved performing 200 registrations with each algorithm by choosing five random subject images and registering each of them with the remaining 40 images as the target. The subject labels were then transformed onto the target domain using the learned deformation field, and compared against the manual segmentations for each target image. This gives an evaluation of each algorithm's ability to correctly align the retinal layers using registration.

We compared the performance of our A-OCT and D-OCT registration algorithms against SyN [19] a highly ranked [50] algorithm for general image registration, which is included in the ANTS [51] package. The Dice coefficient [52] was used to evaluate the accuracy of the results relative to the manual segmentations. This was calculated for each layer k using, $d_k = \frac{2|\mathcal{T}_k \cap \mathcal{S}_k|}{|\mathcal{T}_k| + |\mathcal{S}_k|}$, where $\mathcal{T}_k$ and $\mathcal{S}_k$ are the set of voxels labeled as layer k in the manual segmentation and the transferred labels, respectively. This metric is a measure of segmentation agreement and has a range of $[0, 1]$. A Dice coefficient of 1.0 corresponds to complete overlap between $\mathcal{T}_k$ and $\mathcal{S}_k$, while a score of 0.0 represents no overlap between the two. Table 1 shows the average Dice results over 200 registrations, for each layer, using each of the three algorithms. We also computed the average surface error for each layer boundary for the three algorithms, shown in Table 2. This measured the average absolute A-scan distance between each boundary surface in the transferred segmentation and the surface in the manual segmentation for the target image.

Standard two-tailed Student's t-tests (assuming unequal variances) at an α level of 0.01 were performed to check for significant improvements in the Dice and boundary error results between the algorithms. Significant improvements in both measures were found for all eight layers (and nine boundaries) and their mean when comparing SyN against either A-OCT or D-OCT. When comparing A-OCT against D-OCT, significant improvements in Dice were found for 5 of the 8 layers and decrease in boundary error for 7 of the 9 boundaries (see Table 1 and 2 for specifics). Overall trends showed that on average D-OCT performed better than or equivalent to A-OCT for all 8 layers.

Table 1. Dice overlap between segmentations transferred using a registration algorithm and the manual segmentation for eight retinal layers, averaged over 200 registrations (40 targets, 5 subjects). The deformable registrations were performed with SyN [19], A-OCT, and D-OCT. All eight layers and their mean were found to gained significant improvements (at an α level of 0.01) in Dice overlap when comparing SyN against either A-OCT or D-OCT. Asterisked (*) values on the D-OCT row indicate the layers that gained significant improvements in Dice when comparing A-OCT against D-OCT.

| | **Layers** | | | | | | | | |
	RNFL	GCIP	INL	OPL	ONL	IS	OS	RPE	Mean
SyN	0.51	0.56	0.35	0.37	0.49	0.25	0.27	0.42	0.40
A-OCT	0.83	0.80	0.55	0.63	0.83	0.63	0.71	0.85	0.73
D-OCT	0.84	0.82	0.61*	0.69*	0.85*	0.72*	0.76*	0.85	0.77*

Table 2. Average layer boundary surface errors (μm) between segmentations transferred using a registration algorithm and the manual segmentation for nine retinal layer boundaries, averaged over 200 registrations (40 targets, 5 subjects). The deformable registrations were performed with SyN [19], A-OCT, and D-OCT. All nine boundaries and their mean were found to have significantly less error (at an α level of 0.01) when comparing SyN against either A-OCT or D-OCT. Asterisked (*) values on the D-OCT rows indicate the boundaries that had significantly less error when comparing A-OCT against D-OCT.

	Layer Boundaries				
	ILM	**RNFL GCIP**	**GCIP INL**	**INL OPL**	**OPL ONL**
SyN	30.3	34.4	38.9	37.9	40.5
A-OCT	4.3	11.6	13.5	12.4	11.3
D-OCT	4.1*	10.3*	11.9*	10.6*	9.79*
	ONL IS	**IS OS**	**OS RPE**	**BrM**	**Mean**
SyN	38.0	37.1	36.2	31.3	36.8
A-OCT	6.9	6.4	7.2	5.1	9.3
D-OCT	5.8*	4.0*	6.9	4.9	8.0*

5. Applications

In this section we present three applications of our deformable registration algorithm for the purpose of studying population differences between patients with MS and healthy controls.

5.1. Average atlas and normalized space

One main application of deformable registration is the ability to create a normalized space that individual subjects can be moved into for comparison. A common approach to achieve this is to construct an average atlas that defines the normalized space. This atlas can then be used as the target image for registering future subject images into the normalized space.

The 40 healthy control images in the general cohort were used to construct an average atlas following the approach referenced in Section 2.3. Two iterations of the average atlas adjustment was applied in the atlas construction. Figure 3 shows two views of this average atlas. In the following sections, this atlas is used as the target image for moving each subject into the normalized space.

5.2. Statistical atlas

One useful application of a normalized macular OCT space is the ability to construct a statistical representation of spatial locations of retinal layers in the macula. By moving the retinal layer segmentations for each healthy control in the general cohort into this common space, we can empirically estimate the probability of a voxel belonging to a particular layer. Figure 4 shows each layer in a statistical atlas computed in this manner. Since everything is calculated in the average atlas space, the probabilities correspond spatially with the average atlas. Hence, if the average atlas is registered to a new image (or vice-verse), then the statistical atlas can be directly carried over using the same deformation. This provides a statistical estimate of the location of

each layer in the new image, which can be used with other segmentation and analysis methods by providing a prior probability of each layer in the new image.

5.3. RAVENS analysis of multiple sclerosis

For our final application, we evaluated RAVENS maps constructed by registering the entire general cohort (controls and MS patients) into our normalized atlas space. Figure 5 shows an example of a RAVENS map created between a subject image and the average atlas, using the approach described in Section 2.4. For our analysis, the registration was performed using the D-OCT registration and the segmentation of the retinal layers were automatically found using a boundary classification and graph based method [15].

The goal of this experiment is to perform a pilot study looking for significant differences in layer volume changes (relative to the normalize space) between the healthy control and MS cohorts. The RAVENS map for each subject was used with SPM [53] to perform a Student's T-test at each voxel in the image. False discovery rate (FDR) was used to correct for multiple hypothesis testing. This created a p-value map, where each voxel represents whether there is significant differences in volume changes between the control and MS cohorts at that location. Figure 6 shows a snapshot of this significance map (at an α level of 0.05) for all 8 layers combined. We see from the figure that this type of analysis shows better localization of differences between the two cohorts, in comparison to standard thickness measurements evaluated over a segmentation.

6. Discussion and conclusion

6.1. Evaluation against existing methods

To our knowledge, there has been no other openly available deformable registration algorithm designed and validated on macular OCT volumes. Hence, our method could only be compared against currently available generic registration algorithms that have been used with success in various other anatomical locations. Our results in Table 1 and 2 show that on average our D-OCT algorithm produced the most accurate and robust results for aligning retinal layers when comparing segmentations transferred by the registration against manual segmentations. Relative to SyN, both A-OCT and D-OCT performed significantly better when registering retinal layers, which is expected given that SyN was not designed for macular OCT.

Independent of the comparisons to SyN, our method was able to produce high levels of layer alignment relative to manual segmentations. The Dice coefficient tends to heavily penalize small errors in the segmentation when comparing thin structures such as retinal layers, because small shifts between thin objects can dramatically reduce their overlap. Hence, an overall mean Dice overlap of 0.77 for the layers is considered fairly high, given that each of the retinal layers is
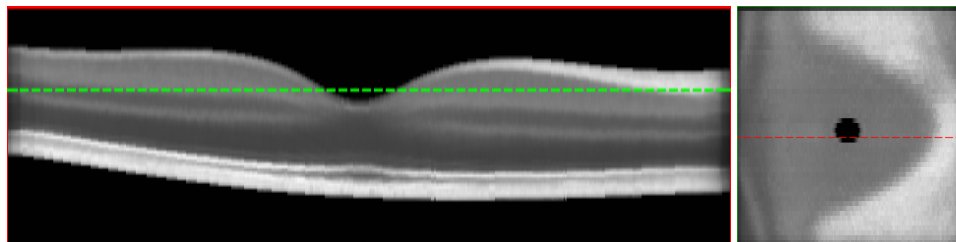

Fig. 3. B-scan (left) and en-face (right) views of an average atlas created from macular OCT images from 40 healthy control subjects. The green and red lines show the location of each view relative to the other. The vertical scale in the B-scan view is tripled to better show the details of the atlas.
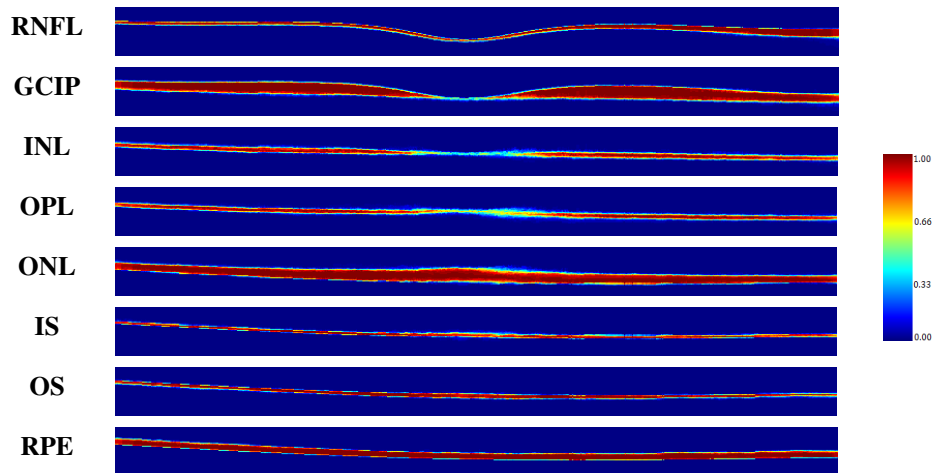
Fig. 4. Statistical atlases of the 8 retinal layers constructed from 40 healthy control subjects registered into a normalized space.

only around 10-20 voxels thick. The average surface errors for the layer boundaries were also fairly low for our algorithm, with an average of 8μm across all nine boundaries. We do note that this error is higher than state of the art segmentation algorithms [15], which is expected, since the registration algorithm must learn a deformation between the images, and is restricted by the regularization necessary to keep the deformation sensible.

Lastly, since A-OCT is solved directly and D-OCT is only optimized in a single direction at each A-scan, our registration approach is computationally very efficient. On a single 2.73 GHz processor with 8 GB of ram, running the A-OCT registration on $1024 \times 496 \times 49$ images took on average 3 minutes per registration (including pre-processing time). The D-OCT registration, which includes both pre-processing and A-OCT, had an average runtime of 106 minutes. This is a considerable improvement over SyN which had an average runtime of 9 hours on the same images, using the same machine.

6.2. Clinical relevance

Using our registration approach, we have constructed a normalized space and average atlas for macular OCT, which enables a number of population based analysis. In this work, we've shown two such applications. First, we have created a set of statistical atlases, which can be used as prior information for future segmentation and registration approaches. This allows such algorithms to improve its accuracy by comparing unknown images to these atlases in order to create robust initializations for the solution. These atlases can also potentially be used to develop more advanced atlases for comparison against normal anatomy, similar to standard growth charts. Second, we used the normalized space to perform population analysis by constructing RAVENS maps to analyze local volume differences between controls and patients with MS. Our results found significant clusters of volume change difference between healthy and MS patients in the RNFL and GCIP layer. This result corresponds with existing reports that suggest patients with MS undergo significant atrophy of the RNFL and GCIP [54, 55]. Our results also found much smaller clusters of differences in the INL, which is consistent with reports suggesting that INL atrophy is also present in MS, but only in about 40% of patients [49, 55]. However, one major advantage provided by our analysis is the ability to detect specifically in 3D where such changes are occurring, whereas previous methods relied on average OCT thickness analysis
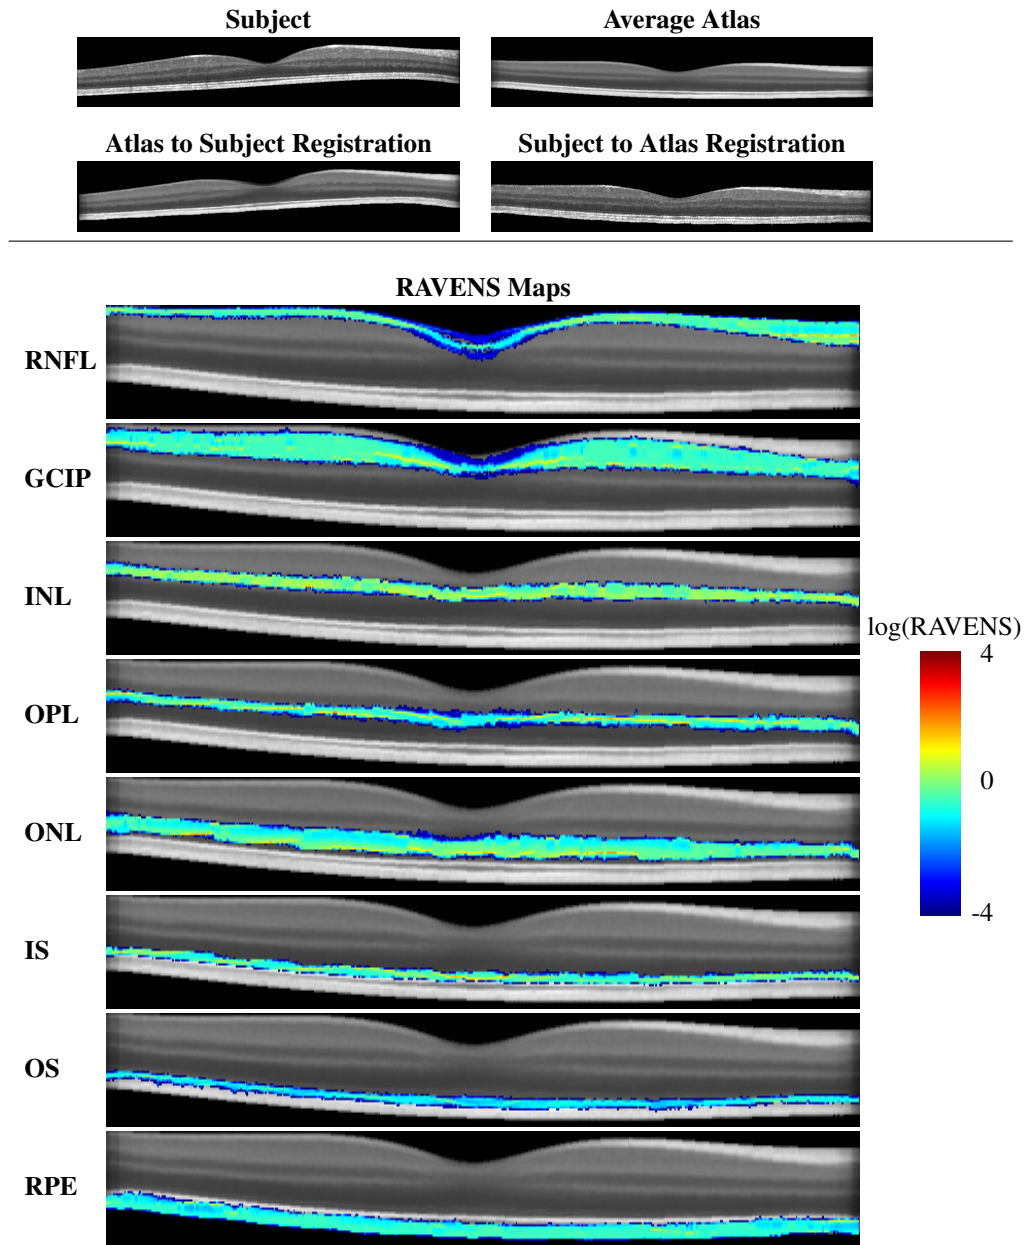
Fig. 5. Shown are a subject image and an average atlas, their registrations to each other, and the subject's RAVENS maps for each layer, overlaid on the average atlas. Each RAVENS map is shown in log scale to better illustrate where the tissue is expanding (blue) or compressing (red) relative to the average atlas. The vertical scale of the B-scan is doubled to better show the details of the maps.
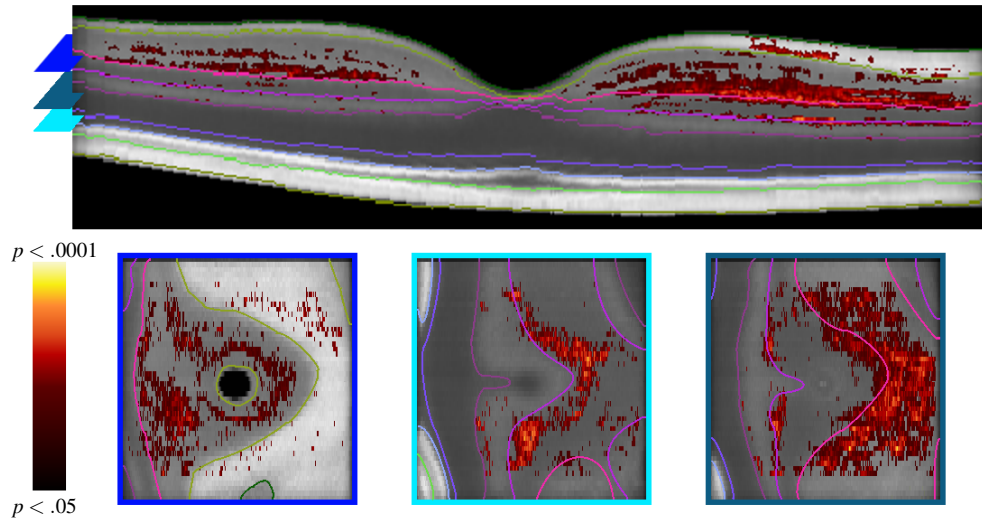
across 2D regions of the macula or post-mortem histological analysis. We see from Fig. 6 that the significant voxels are strongly localized towards the medial (nasal) side of the fovea. In addition we see that the highly significant differences are concentrated in the deeper areas of the GCIP layer. This type of analysis is unavailable when using standard approaches, such as thickness measurements, for evaluating macular OCT.

6.3. *Limitations and possible improvements*

There are several limitations to our registration algorithm that are worth discussing. First, a common limitation of most registration algorithms is when correspondences are inconsistent between the subject and target images. In retinal OCT, this can occur when particular layer structures are not visible in either the target or subject image. An example of this is shown in Fig. 2, where the hypo-intense region in the RPE (directly below the Verhoeff's membrane) is visible in the subject image, but not in the target image. Since there is no correspondence between the images, the algorithm relies on regularization for alignment, creating errors.

These errors can potentially be avoided by using better OCT acquisitions with higher resolution or SNR such that the same structures always appear in both images. Alternatively, if the problem is isolated to a few A-scans, such as in this case, then it might be possible to address it by increasing the weight of E_m^{Reg} relative to E_m^{SSD} in the final cost function. However, this is a trade-off that can force too much regularization across A-scans and result in lower overall accuracy.

Another example of inconsistent correspondences in retinal OCT is disrupted or missing layers due to pathology (detachment, edema, etc.). The locations of such pathology are often unpredictable, and may not be present in both the target and subject images. As a result, the algorithm may attempt to align layers with similar intensities to these pathologies, creating errors in the registration. To address these problems, further development of the algorithm will be necessary to detect and regularize the transformations in these situations.

Second, data limitations can directly affect the performance of the algorithm. Due to the sparse sampling of the B-scans in our Spectralis data, we are limited by the type of regularization and interpolation allowed in the algorithm. Given data with higher resolutions and denser sampling, we can potentially improve the accuracy of the registration and provide better acuity in the types of analysis we have demonstrated.

Third, errors during the data acquisition can affect the accuracy of the registration. One example of this is the appearance of a rotation of the retina, often caused by the OCT scan not being centered on the pupil during the acquisition [56]. When the retina is tilted in the OCT, our discrete translation of the A-scans becomes a less valid model for rigidly aligning the macula, which lowers our registration accuracy. In general, our algorithm cannot correct for this rotation, because rotating the OCT would cause the z-axis in the OCT to no longer represent the A-scans. This would make it difficult to justify any transformations along this axis, since the size of each voxel along the axis would now depend on the fan-beam separation, which depends on the cornea of the subject being imaged. Care must be taken to use data that is properly centered to minimize these errors.

Fourth, intensity inhomogeneity can still be a problem for the algorithm after intensity normalization, where shifts in intensity can be seen in the OCT even after preprocessing. One potential improvement to address this is to use a similarity metric that is more robust to intensity variabilities. Our preliminary results suggest that the cross-correlation measure may be a good replacement for SSD in future versions of the algorithm.

Lastly, since the detected ILM and BrM locations are used in the affine registration step, errors in their detection can cause poor initialization for the deformable registration step. While this would not be detrimental to the algorithm, it does mean the deformable registration would need to move the boundaries larger distances to find the correct alignment. This might cause the optimization to get caught in a local minima, resulting in errors in the final registration. Although the retinal mask is generally found very accurately and robustly in our experiments, certain pathology can make these boundaries difficult to detect. In such cases, the boundary detection algorithm may need to be modified.

6.4. Conclusion

We have presented both an affine and deformable registration method designed specifically for OCT images of the macula, which respect the acquisition physics of the imaging modality. Our validation using manual segmentations shows that our algorithm is considerably more accurate and robust for aligning retinal layers than existing generic registration algorithms. Our method opens up a number of applications for analysis of OCT that were previously not available. We have demonstrated three such examples through the creation of an average atlas, a statistical atlas, and a pilot study of local volume changes in the macula of healthy controls and MS patients. However, these are but a small example of the existing methods and techniques that utilize deformable registration and normalized spaces to perform population based analysis of health and disease. The software for our OCT registration algorithm and atlas construction approach can be freely downloaded as modules for the Java Image Science Toolkit (JIST) [57] software package at `https://www.nitrc.org/projects/toads-cruise/`.