

Advanced Techniques in Medical Imaging: A Comparative Study of Dimensionality Reduction Methods

K. Vijay Anand, R. Srinivas Yadav, G. Rama Krishna

Department of Electrical and Computer Engineering, Indian Institute of Technology Madras, India; Department of Biomedical Engineering, Amrita Vishwa Vidyapeetham, India

Abstract

This Paper represents a method for the automatic segmentation of the brain in magnetic resonance (MR) images of the human head. The method identifies brain areas of interest, including the gyri and other subcortical structures that were manually delineated in a set of labelled training images. Principal components analysis (PCA) is applied to the training ensemble in order to learn a PCA atlas subspace, which is a dimensionality-reduced linear subspace of labelled brain images. We employ this subspace to segment and label previously unseen subject images. This is accomplished by finding the PCA atlas closest to an input subject image through orthogonal projection of the latter into the subspace. The PCA atlas is then non rigidly registered to the subject image and the non rigid transformation is used to transfer the labels from the former to the latter, thereby segmenting the subject image. Our method is compared with alternative methods and the results are validated using overlap and distance metrics.

Indexed keywords: MRI Images, Principal Component Analysis Atlas.

Article History: Received: 17 December 2024 | Accepted: 05 February 2025 | Published: 23 February 2025

INTRODUCTION

The segmentation of brain images involves localizing and delineating the brain's different anatomical structures within three-dimensional (3D) tomographic images of the head. Brain image segmentation is important for longitudinal studies investigating how medications and other interventions affect the shapes and volumes of brain tissues, for surgical planning and for diagnosing disease. Traditionally, the various parts of the brain are segmented by trained experts who manually identify and delineate regions of interest. This process can be very

Difficult and time consuming. It requires the design of complex protocols for identifying brain structures. The results from different manual segmentations using the same protocol may be inconsistent, which raises issues if a set of images in a study are delineated by different individuals. With the advent of non-invasive medical imaging devices that can easily acquire large quantities of tomographic data, manual segmentation becomes impractical for many studies. An alternative to manual segmentation is to develop computer algorithms that can automatically segment anatomical structures of interest in images of the brain. The automatic 3D brain image segmentation algorithm developed in this work involves computing a 3D atlas consisting of labelled voxels, from among a space of possible atlases learned from an ensemble of manually segmented training images, that is most suitable for segmenting an unlabeled subject image. This optimal atlas is then non rigidly registered to the subject image. Finally, the non rigid transformation is used to transfer the labels from the atlas to the subject image, thus segmenting the subject image. In greater detail, our procedure is as follows:



© 2025 The Authors. Open Access under CC BY 4.0.

How to cite: K. Vijay Anand, R. Srinivas Yadav, G. Rama Krishna (2025). Advanced Techniques in Medical Imaging: A Comparative Study of Dimensionality Reduction Methods. Journal of Computer Engineering, 14(2), 50–59. DOI:

<https://doi.org/10.5281/zenodo.19339510>

Registration is the process of establishing correspondences between two images by finding a transformation of a source image into the coordinate system of a target image. First a set of training data comprising brain intensity images and their corresponding labels are mutually registered to minimize the variance across the data. This ensures that the variation in the data is attributed mostly to the differences in the shapes and textures of the anatomical structures of the brain and not to their locations, orientations, and sizes. Principal components analysis (PCA) is then performed on the intensity and label training data to learn a dimensionally-reduced linear subspace that is capable of generating labelled images that interpolate the training data. We refer to this subspace as a PCA atlas subspace. Then, an orthogonal projection is performed to find the point in the PCA atlas subspace that is closest

to the subject image. The point is used to generate an optimal intensity image with its associated labels, which we call the PCA atlas. A non rigid registration algorithm is then used to deform the input subject image to match the intensity image component of the PCA atlas and the resulting non rigid transformation is finally used to map

the associated labels over to the subject image in order to complete the segmentation. We apply our algorithm to 3D magnetic resonance (MR) images of the brain.

Applications of Brain MR Image Segmentation :

The results from brain MR image segmentation can provide a wealth of information that is important in understanding and dealing with the brain. The insight that segmentation affords about the shape and volume of different brain structures helps in understanding how the brain changes over time. It enables us to gain an understanding of the effects of disease and medication on the brain and can enable us to detect the onset of disease through the classification of these changes. Segmentation also allows us to reconstruct three-dimensional models of the structures in the brain that can provide assistance to surgical planning and guidance. The following is a survey of some of the literature that makes use of the results from brain MR image segmentation, emphasizing its importance in understanding the brain. The segmentation of a 3D MR image of the head provides a delineation of all the voxels that make up a particular anatomical region. From this, a model of the shape and the volume of the region can be computed. A great deal of research has focused on the shape and volume changes that occur in the brain over periods of time. Giedd et al. (1999) examined a range of individuals from ages four to twenty and scanned them every two years with a maximum of five scans. The white and gray matter in these images was segmented using artificial neural networks and automatic registration to a template image in which these areas have been manually identified. The authors of this study were able to calculate the volumes of the white and gray matter to understand how they change over time. The white matter grew linearly over all regions of the brain, while the gray matter had a nonlinear growth that was specific to different areas of the brain. This discovery gave insight into how the different parts of the brain develop over time. Resnick et al. (2003) examined a group of healthy adults aged fifty-nine to eighty-five years old and acquired MR head scans every two years for a total of three scans. The extra cranial tissue, cerebellum, and brainstem regions were manually removed from the images by a trained expert. Then the locations of the white matter, gray matter, and cerebrospinal fluid (CSF) were found using an adaptive Bayesian segmentation algorithm. From this, the volumes of the white matter, gray matter, and CSF in different areas of the brain were calculated. The results showed that there was a significant decline in the volume of white and gray matter and an increase in the CSF filled ventricle volumes for healthy elderly adults. A related research thread examines how the brain is affected by disease over time. Thompson et al. (2003) compared sets of elderly individuals with and without Alzheimer's disease. Two MR scans of the head were acquired two years apart for each subject, and a series of image processing steps were applied to measure how the brain changes across the set of images over time. The images were registered to a standard brain imaging template and were each mapped into the same space. A Gaussian mixture model was then used to classify each image into white matter, gray matter and CSF. A three-dimensional cortical surface was extracted from this segmentation and further analysis was performed to study the changes in shape and tissue distribution over time. The analysis showed that cortical atrophy was significantly greater in brains affected by Alzheimer's disease compared to those of control subjects. This study also allows an understanding of where and when

certain parts of the brain undergo atrophy, which can be used as a biological marker for the disease. The work presented in (Grimson et al., 1997) reports on an algorithm for segmenting structures in whole head MR images and how that information can assist in surgical planning. Their "augmented reality" method enables images of the segmentation of a patient's MR image to be projected onto the patient's head so that surgeons can visualize and localize the blood vessels of critical functional brain regions. The authors also describe methods for tracking the surgeons' instruments in relation to the segmented MR images. This exposes the surgeon to a three-dimensional reconstruction of the different segmented regions in relation to the location of their instruments within the brain. It shows the surgeons structures near the area in which they are working, which would otherwise be very difficult to see or locate. Gering et al. (2001) built on this surgical planning framework by creating a language to easily specify different anatomical structures and their locations in the scenes created from the imaging data. That work also contributed a procedure for fusing multi-modal imaging data to visualize what occurs during surgery and to plan an optimal trajectory through the brain. Additional information about the application of brain MR image segmentation to surgical planning is available in the literature.

BRAIN IMAGE SEGMENTATION

There has been a great deal of work on the segmentation of medical images of the brain. Most approaches rely on training data to learn the parameters of a model of the different structures in the brain. This is necessary because of the complex protocols used to specify where anatomical structures of interest appear in the brain images and because in many instances different areas do not have a natural, easily recognizable boundary. The methods used to segment normal structures in the brain typically differ from those used to segment tumors and

abnormalities in the brain because the former assume that the model created from the training data provides a good representation of the brain's anatomy in general and they do not rely on features that are easily differentiated from normal brain appearance and structure. The methods used for segmenting normal brain images can roughly be categorized into voting methods, atlas based methods, and machine learning classification methods. Arno et al. (2005) registered every image in a set of manually labelled training image data to a subject image. The labels are transferred to the subject image from each of the images in the labelled training set and the most frequent label at each voxel is used as the final labelling. This process can be computationally expensive if the set of labelled images is very large because it needs to be performed for every subject image. Aljabar et al. (2007) use a similar approach, but they use a similarity criterion to select a subset of images from the training set that are similar to the subject image. This method relies heavily on the quality of the algorithm used to register the training data to the subject image. Compared with the other approaches, it retains the training data set in its entirety instead of creating a compact representation. The work in (Fischl et al., 2002) uses a probabilistic atlas to segment a novel subject image. It poses the problem in a Bayesian setting where the objective is to maximize the probability of a labelling and transformation of a subject image into the atlas space. This method allows the incorporation of prior information such as the spatial structure of the different anatomical areas and restrictions on the possible transformations of the subject image into the atlas space. It also incorporates Markov Random Fields (MRFs) to model the spatial relationship between the voxels of different anatomical parts. Babalola et al. (2008) used 3D volumetric Active Appearance Models (AAM) to segment images of the brain. They start with a global AAM to obtain the location of the different structures in the brain and then refine these locations with AAMs trained specifically for each area. They then use linear regression functions at each voxel in the image to compute the probability that it lies within the domain of a certain anatomical area. The AAMs that they use rely on a grid of corresponding points in the data that are transformed to produce different shapes and a set of shape-free volumetric intensities that are transformed to the grid shape for optimizing the locations of the AAM. This differs from (Fischl et al., 2002) in that the structural information is learned for each specific structure or label as a whole instead of only the relationship between a voxel and its neighbourhood found in MRFs. The method described by Tuetal. (2008) uses a statistical hybrid discriminative/ generative method to segment images. The discriminative component learns from thousands of features in the images to create an appearance model that can classify a voxel by assigning a particular label using local

information from the image in a probabilistic boosting tree (PBT) framework. The generative component uses PCA to create a global shape model of the different shapes. The components are used together to create an energy function that is minimized to find a final segmentation. The method presented in this thesis relates to (Fischl et al., 2002) in that the goal is to create an atlas that is tuned to the subject image, but we assume that once such an atlas is found, an additional high performance nonrigid registration will be performed to transform the subject image into the atlas space. PCA is used to create a linear subspace representing the training data. A subject image is registered to the mean training image and then orthogonally projected into the PCA atlas subspace to find the atlas image.

PRINCIPAL COMPONENT ANALYSIS

Principal Components Analysis (PCA) is a method that computes the mutually orthogonal directions of maximum variance in a collection of d -dimensional data. These f directions form an orthogonal basis that spans the data space. PCA is usually used for dimensionality reduction, by projecting the data on a subset $d < f$ of the basis vectors. The subset of vectors are termed the principal components, and they account for a large percentage of the variance. They thus represent each d -dimensional data point as a linear combination of the f principal components and, hence, by the d -dimensional vector of coefficients of the linear combination. Among its many uses, PCA can be used to compress data, to classify data, to visualize data trends in a lower-dimensional space, etc. Its application in the context of the image segmentation method proposed in this work is to capture within a compact representation the intensity and shape variation evident in the training data.

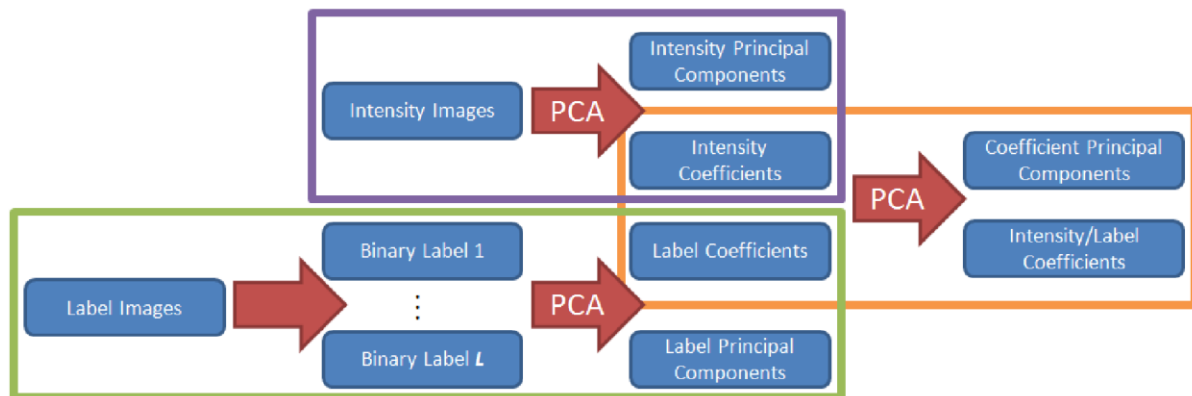
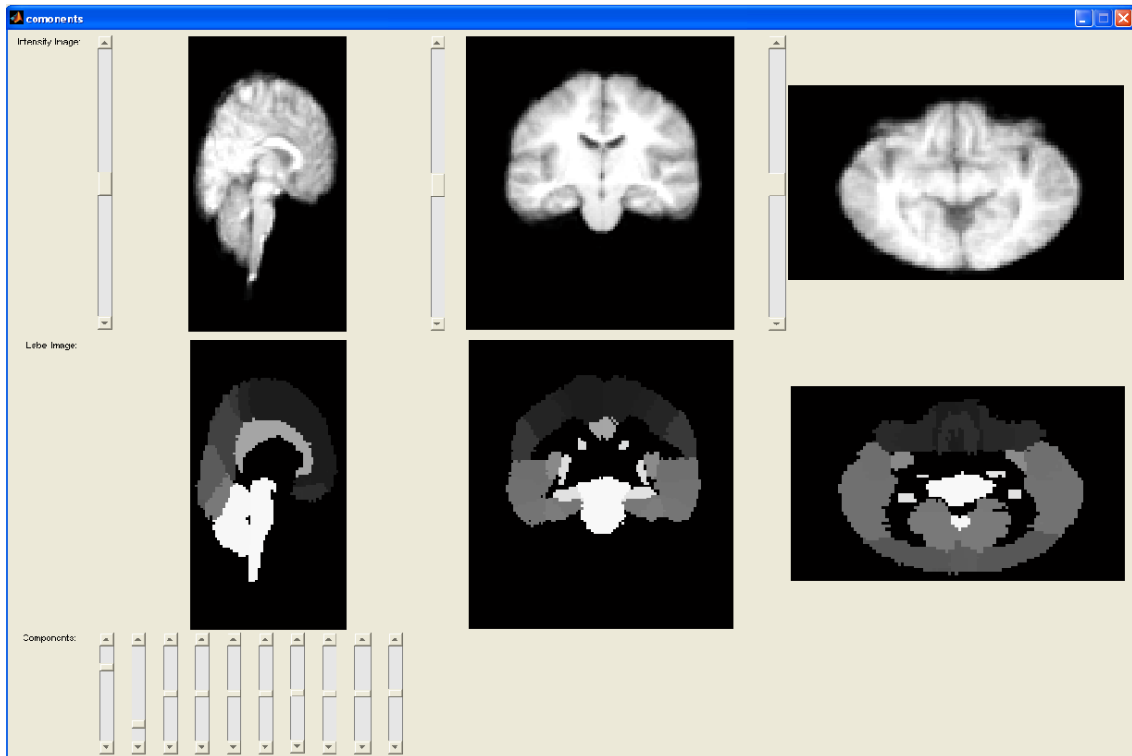
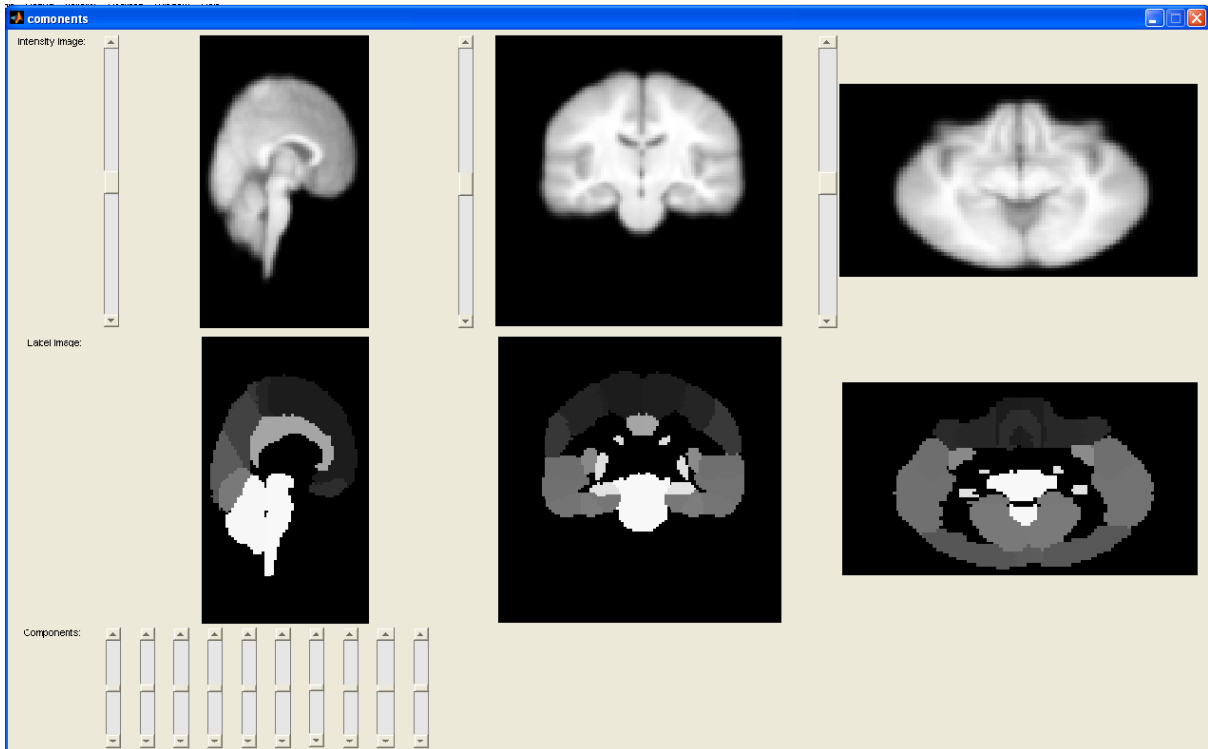


Figure 1: PCA Processing of Label and Intensity Images

The representation of the data by a subset of the basis will not reconstruct the data exactly, and there is a trade-off between the compactness of the representation and the error that it incurs. It displays the sum of squared differences error between a training image and its PCA representation as the number of eigenvectors (principal components) is increased. The eigenvectors are ordered in decreasing order of their corresponding Eigen values. As the number of eigenvectors increases, those that are added later account for less of the variance than the earlier ones.

PCA Atlas Subspace GUI

A graphical user interface (GUI) was developed using MATLAB to explore and visualize the PCA atlas subspace. It works with the data described in Section 3.1. It allows a user to interactively synthesize three-dimensional intensity images along and their associated labels. These PCA atlas images are generated using combined label-intensity coefficients and the PCA basis described earlier in this section. The GUI allows a user to vary the coefficients corresponding to the ten eigenvectors with the largest Eigen values using sliders. Figure 2 shows all the coefficients at zero, which yields the mean PCA atlas. Figure 2 shows four examples of different coefficients that were selected to synthesize different intensity and label images. Forty training images were used to create the PCA atlas subspace in this example, but only the coefficients associated with the 10 Eigen vectors with the largest Eigen values are retained in the GUI, while the 30 less significant ones are set to zero.



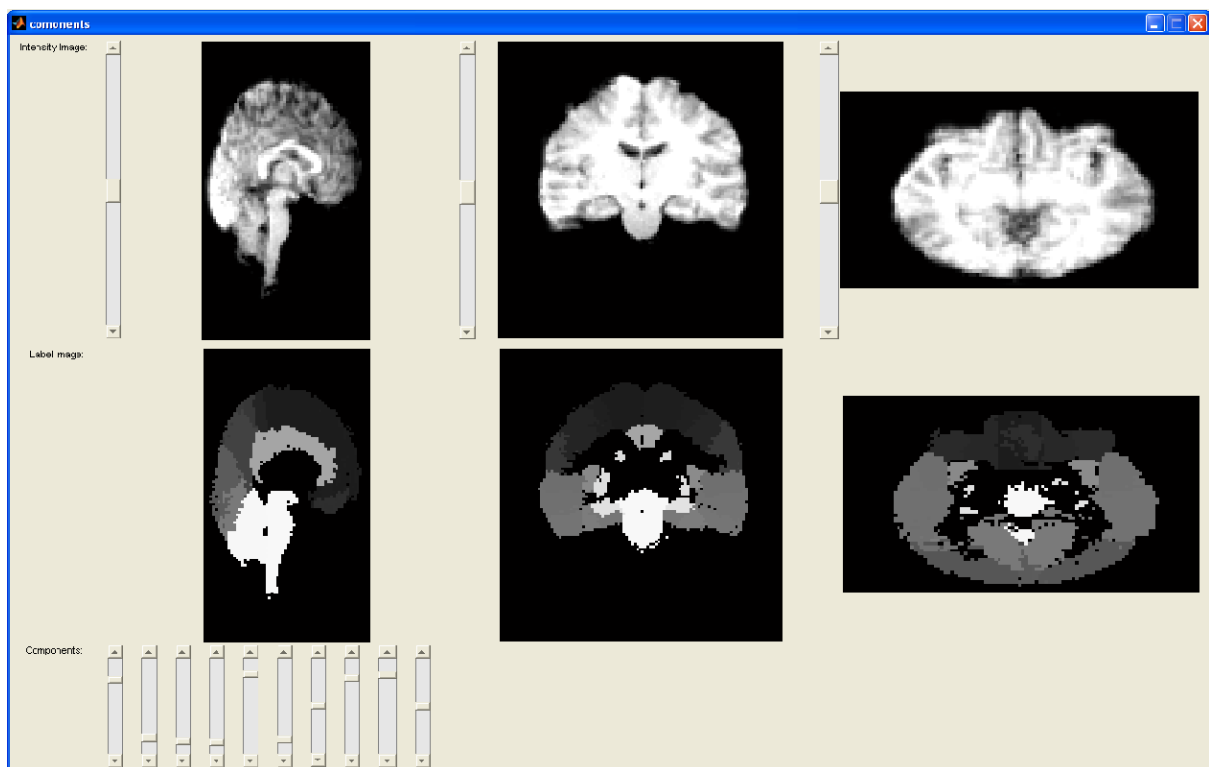
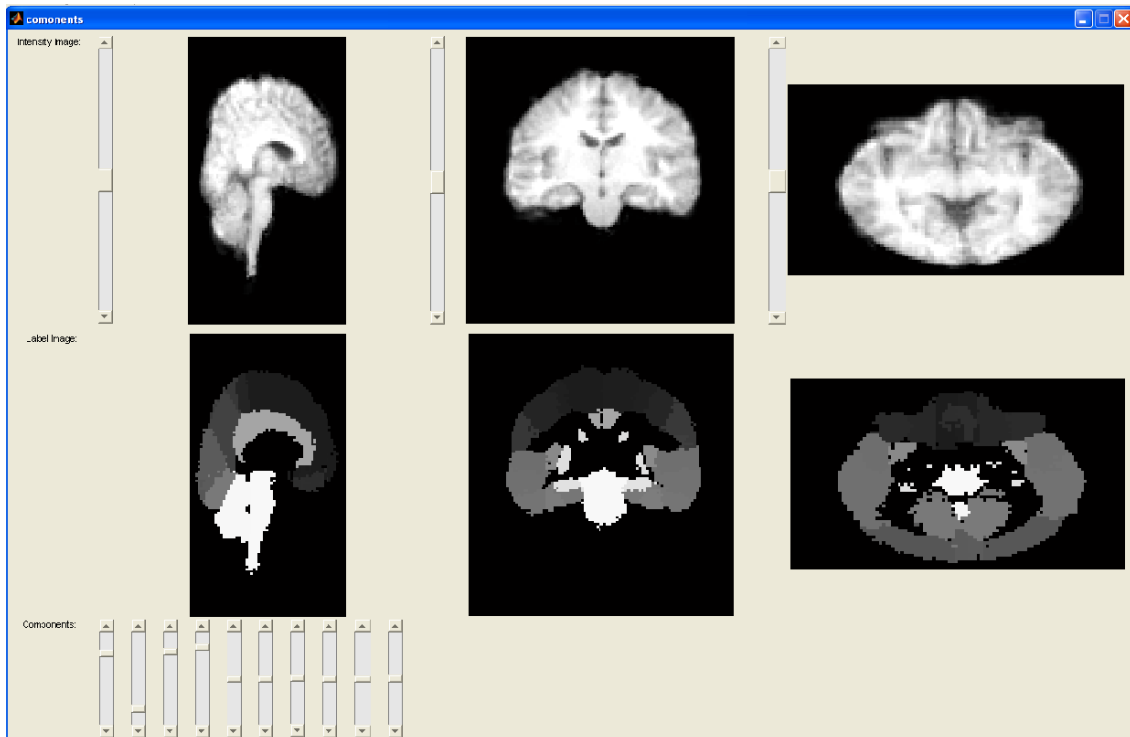


Figure 2: PCA Mean Image and Labels

RESULTS AND DISCUSSION

Figure 3.1 shows horizontal, coronal, and sagittal slices of a typical brain image with the manually segmented, gold standard labels overlaid on the image in yellow. Figure 3.2 shows the same slices

overlaid with the labels from the segmentation that was automatically computed by the algorithm that we developed in this work.




Figure 3.1: Manual Segmentation Results




Figure 3.2: PCA-segmentation Results

CONCLUSION

Our novel segmentation algorithm based on PCA atlases operates by learning a combined statistical model of the training intensity images and their corresponding label images-the linear, PCA atlas subspace-finding the point at which an unlabeled subject image projects into this subspace, and nonlinearly registering the intensity and label images associated with that projected point-the PCA atlas-to a subject image in order to segment and label the image. The nonrigid registration is needed in general because there will be a mismatch between the intensity image component of the PCA atlas and the input subject image. We tested our algorithm against the two competing methods. Its performance is similar to that of the voting strategy in all but the Sensitivity measure, meaning that it had the fewest false negatives or that it included more voxels in the boundary of the regions that were true positives than the other algorithms. It was able to describe the information from the training data in a very compact form, which can be advantageous when there is a huge amount of training data.

In such a case, the voting strategy would have an exorbitantly long running time and would be impractical for computing timely image segmentations because it would need to register all the images together for each subject image. Our method would be able to store this information in a basis comprising a significantly smaller number of vectors than the number of training examples while maintaining a high level of accuracy. Once the statistical model is created, it can be used to find the location of any subject image in its statistical space and not increase the segmentation running time as would the voting strategy. Another important aspect of these algorithms is the time that they take to perform a complete segmentation. Since the voting procedure requires as many registrations as there are images in the training set, it took more than two hours to complete the segmentation using 39 training images in order to generate the results. This is required for any segmentation using the voting algorithm, since any subject image is required to be in correspondence with all the training data and this cannot be computed beforehand. Our method took four hours to complete a segmentation using the full-resolution training data, including the offline training and online segmentation phases of the algorithm. The segmentation phase takes roughly 30 minutes to perform by itself. Also, once our algorithm has been trained, it can be applied to any subject image that is well represented by the training data.

FUTURE WORK

The performance of our method is dependent on how well it is able to capture information from the training data. In particular, when a model is built for the label images contained in the training set, the generation of labels from the basis is highly dependent on the Thresholding that takes place. In its current implementation, the threshold is selected as the midpoint of the range of possible values that a generated label can take. This parameter may need to be optimized for better performance. Also, when deciding the ownership of a particular voxel, random noise is added so as not to favour one over the other, but another strategy, such as splitting the voxels evenly between the number of competing regions, may be better. This is an issue that can be investigated in future work. The nonrigid registration algorithm has a significant effect on the segmentation results. Our framework can be kept as is and different registration algorithms can be employed. It would be informative to evaluate quantitatively the performance of different nonrigid registration algorithms in order to determine how dependent our method is on them. Another aspect of our method that requires further investigation is the way that the label data are analyzed. Currently, a binary image is created for each of the labels in the image space and a separate PCA space is created for each binary image space. One alternative to this would be to fit a mesh to a particular label and register the points on the mesh to each of the training label images. Then a PCA space can be created by processing the points in the mesh, instead of treating a binary image as a point in a high dimensional space. This would help to ease the computational burden of generating the labels and it would place more importance on the surface of the anatomical regions in question instead of the 3D volume structure. This notion can also be extended to the intensities. We could focus our attention only on the image intensity values that occur at the surface interface. PCA can be performed on these intensities to create a statistical model. Compared to using the entire voxel space, which contains useful intensity information, this alternative method would focus on relatively few intensities. The underlying idea is that if the model is positioned initially in approximately the correct location, it can use its learned knowledge to deform its surface parameters to better fit the intensities of a subject image. However, by considering intensity information only at the surface, the model may restrict its intensity information so much that it will be unable to converge onto the correct boundary. This could be improved by using the gradient of the intensity image instead of the actual intensities. A neighbourhood of intensities may be used to compute a smoothed gradient. This affords intensities sampled from these regions more information about the neighbourhood of the mesh and will hopefully improve its ability to capture how the model changes as it approaches the correct boundary. In the algorithm presented in this work, the intensity images and label images were affinely registered together as a group. Then a statistical model was generated for the corresponding intensities and labels. Another possible way to model the intensities would be to register them together nonlinearly to the same structural image. The structural images in our case are represented as label images. The mean label image can be used as a target to register the remaining training data. The transformation would thus

put each training label image into close correspondence with the mean label image. These transformations can be applied to the intensity images. PCA of the intensity images can then be performed to account for the change texture. A separate PCA is performed on the unregistered label images so that its texture and structure are modeled separately. To generate a new image, a label image and texture image are synthesized. Then the texture image is transformed from the mean label shape to the synthesized label shape. This type of modeling of the data is similar to the Active Appearance Model (Cootes et al., 1998). Another method that might improve the segmentation would be to create a separate statistical model for each of the regions being segmented. Each specific anatomical region would then have more freedom to adapt to fit a subject image, because it would be independent of the other models. A bounding box could be created around each region to bound the area that the statistical model takes into account, or the masking method that we use can be adapted to help restrict the model to only the most essential voxels in the image. Finally, we would like to develop potentially more powerful nonlinear generalizations of the linear (matrix) algebraic PCA atlas subspace method developed in this work, through the use of multi linear (tensor) algebra and the multi linear PCA (Vasilescu and Terzopoulos, 2003, 2007). This should make it possible to learn multi linear PCA (MPCA) atlas subspaces from more extensive training image datasets associated with multiple causal factors, such as disease type, age, and gender.

REFERENCES

- [1] Aljabar, P., Heckemann, R., Hammers, A., Hajnal, J., and Rueckert, D. (2007). Classifier selection strategies for label fusion using large atlas databases. *Medical Image Computing and Computer Assisted Intervention-MICCAI 2007*, pages 523-531. 11
- [2] Arno, K., Brett, M., Satrajit, G., Jason, T., and Hirsch, J. (2005). Mindboggle: Automated brain labeling with multiple atlases. *BMC Medical Imaging*, 5, 11, 35.
- [3] Ardekani, B., Guckemus, S., Bachman, A., Hoptman, M., Wojtaszek, M., and J., N. (2005). Quantitative comparison of algorithms for inter-subject registration of 3d volumetric brain mri scans. *Journal of Neuroscience Methods*, 142(1):67-76. 16
- [4] Arthur, D. and Vassilvitskii, S. (2007). k-means++: The advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, page 1035. Society for Industrial and Applied Mathematics.7
- [5] Babalola, K., Cootes, T., Twining, C., Petrovic, V., and Taylor, C. (2008). 3D brain segmentation using active appearance models and local regressors. *Medical Image Computing and Computer-Assisted Intervention-MICCAI 2008*, pages 401-408. 12
- [6] Babalola, K., Patenaude, B., Aljabar, P., Schnabel, J., Kennedy, D., Crum, W., Smith, S., Cootes, T., Jenkinson, M., and Rueckert, D. (2009). An evaluation of four automatic methods of segmenting the subcortical structures in the brain. *Neuro Image*. 36
- [7] Bhatia, K., Hajnal, J., Puri, B., Edwards, A., and Rueckert, D. (2004). Consistent group wise non-rigid registration for atlas construction. In *IEEE International Symposium on Biomedical Imaging: Nano to Macro, 2004*, pages 908-911. 17
- [8] Comaniciu, D. and Meer, P. (2002). Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 24(5):603-619. 8