

HARDI Based Pattern Classifiers for the Identification of White Matter Pathologies

Luke Bloy¹, Madhura Ingalhalikar¹, Harini Eavani^{1,2}, Timothy P.L. Roberts³,
Robert T. Schultz², and Ragini Verma¹

¹ Section of Biomedical Image Analysis, Department of Radiology,
University of Pennsylvania

² Center for Autism Research, Children's Hospital of Philadelphia

³ Lurie Family Foundation's MEG Imaging Center, Department of Radiology,
Children's Hospital of Philadelphia
{Luke.Bloy,Ragini.Verma}@uphs.upenn.edu

Abstract. The paper presents a method for creating abnormality classifiers from high angular resolution diffusion imaging (HARDI) data. We utilized the fiber orientation distribution (FOD) diffusion model to represent the local WM architecture of each subject. The FOD images are then spatially normalized to a common template using a non-linear registration technique. Regions of homogeneous white matter architecture (ROIs) are determined by applying a parcellation algorithm to the population average FOD image. Orientation invariant features of each ROI's mean FOD are determined and concatenated into a feature vector to represent each subject. Principal component analysis (PCA) was used for dimensionality reduction and a linear support vector machine (SVM) classifier is trained on the PCA coefficients. The classifier assigns each test subject a probabilistic score indicating the likelihood of belonging to the patient group. The method was validated using a 5 fold validation scheme on a population containing autism spectrum disorder (ASD) patients and typically developing (TD) controls. A clear distinction between ASD patients and controls was obtained with a 77% accuracy.

Keywords: Diffusion Imaging, HARDI, FOD, Classification, SVM.

1 Introduction

High dimensional pattern classification methods like support vector machines (SVM) identify brain abnormality patterns that enhance group separability while quantifying the degree of pathological abnormality associated with each individual. This paper proposes a HARDI based pattern classification framework that creates abnormality classifiers using information concerning white matter (WM) architecture derived from homogeneous WM regions. This classification framework not only elucidates regions that are affected by pathology but also assigns each individual with a score indicating the degree of abnormality. Such a score may prove useful in conjunction with other clinical measures as a diagnostic tool

to predict the extent of disease as well as act as a biomarker to assess disease progression or treatment efficacy.

While the bulk of classification work concerning medical imaging analysis has centered on structural imaging[4,5], the application of classifiers to diffusion imaging is relatively recent [9,7,14,3], while there is no literature using HARDI data. Caan et al [3] performed principal component analysis (PCA) and linear discriminant analysis (LDA) on linear and fractional anisotropy (FA) images, derived from diffusion tensor images, to classify schizophrenia patients. Wang et al. [14] used a k nearest neighbor (kNN) classifier trained on full brain FA images, while Ingalhalikar et al [7] used non-linear SVM trained on regional FA and diffusivity for classification in autism spectrum disorder (ASD) and schizophrenia. Finally, Lange et al [9] used quadratic discriminant analysis and SVM to perform hypothesis driven classification in an ASD population based on a-priori regions.

While these methods have had success, they are limited by the use of the diffusion tensor (DTI) data model which is known to be ineffective in modeling regions of complex white matter, i.e. multiple fibers with different orientations, different partial volume fractions. High dimensional diffusion data models, such as the fiber orientation distribution function (FOD) [12], have been developed to make use of new acquisition protocols which acquire HARDI data. These new data models are better able to model complex WM regions and thus should prove useful for studying WM pathology. While Schnell et al [11], proposed a tissue segmentation method based on HARDI classification, we believe this is the first work to utilize HARDI data models to perform subject classification between healthy controls and a disease population

We propose a classification framework that makes use of the FOD HARDI data model to extract orientation invariant features from regions of homogeneous WM architecture. These regions are determined from a WM parcellation algorithm applied to the average FOD image of the control population. Principal component analysis (PCA) is used for dimensionality reduction and a linear support vector machine (SVM) is used to perform the classification. The linearity of the framework allows for the examination of the SVM decision weights in the original feature space to aid in interpretability of the classification results. The framework is applied and cross-validated using a 5-fold cross-validation paradigm using a dataset comprised of children diagnosed with Autism Spectrum Disorder (N=23) and typically developing controls (N=22). The high accuracy, specificity and sensitivity (all $\sim 77\%$) establish the applicability of classifier scores for aiding diagnosis and prognosis.

2 Methods

The process of training and validating a classification framework like the one proposed here consists of a number of steps. Namely feature extraction and dimensionality reduction/feature selection, followed by training and cross-validation.

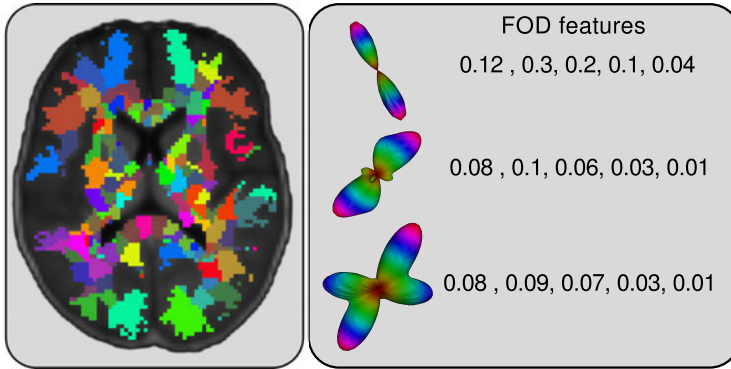


Fig. 1. Feature Extraction: Spatial regions of homogeneous WM are determined by parcelating the population average FOD image into regions of low variance, shown on the left. For each region the mean FOD is computed from which 5 orientation invariant FOD features are computed. The mean FOD and corresponding feature vectors are shown, on the right, for regions in the corpus callosum, as well as in anterior and posterior complex WM.

2.1 Feature Extraction

For each subject, feature extraction must be performed to extract a salient representation of the subject that will serve as a means of comparison. We are interested in identifying pathologies that manifest as abnormalities in the WM. We therefore concentrate on extracting features from spatially localized regions of homogeneous WM. This process entails modeling the diffusion process in each voxel using the fiber orientation distribution function (FOD) followed by spatially normalizing all subjects into a common template reference frame. Once all subjects are spatially normalized an average FOD image of the control population is computed and used to determine homogeneous WM regions of interest (ROIs). Finally orientation invariant features from each ROI are extracted and concatenated to build a feature vector representation of each subject.

Our process begins by using constrained spherical deconvolution[12] to compute an FOD image for each subject. The FOD diffusion model represents the voxel's DW-MRI signal as the spherical convolution of the FOD and the DW signal that would be measured for a single fiber bundle aligned along the z-axis. Information concerning both the orientation and partial volume of any constituent fiber bundles present in a voxel is represented via the FOD [12]. In this work we use the order 8 real spherical harmonic (RSH) representation of the FOD, thus each voxel is represented by the 45 RSH coefficients of its FOD.

Each FOD image is then spatially normalized to a template FOD image, an individual normal subject in this case, using a diffeomorphic demons-based FOD registration algorithm [1]. The process of spatial normalization defines a spatial and anatomical correspondence between subjects that allows a set of ROIs to be determined and to confidently represent corresponding areas of anatomy in each subject.

A population average FOD image, in the template space, is then computed from the registered FOD images of the normal subjects of the population. From this average FOD image we determine a set of WM ROIs with a homogeneous WM architecture using the methods described in [2]. This method utilizes normalized cuts spectral clustering to partition the WM into spatially connected regions which have a small FOD variance. Using this method we determined 883 WM ROIs, with an FOD variance below 0.08 (Shown in Figure 1).

With spatial ROIs determined, we compute the feature vector representation for each subject by first determining the mean FOD in each ROI. From the mean FOD we compute the L2 norm of the RSH coefficients in each order (1 level). Because rotations of the FOD will transfer energy within an order but not from one order to another the collection of these L2 norms provides an orientation invariant representation [6] of the mean FOD. Thus for an order 8 FOD model, as the one used here, each ROI is represented by the p_l for l in 0, 2, 4, 6, 8. Where $p_l = \sqrt{\sum_{m=-l}^{m=l} \tilde{f}^2}$ is defined in terms of the RSH coefficients of the mean FOD (\tilde{f}) of the ROI. These 5 features are computed for each of the 883 WM ROIs and concatenating them yields a representation of each subject's WM by a 4415 element feature vector.

2.2 Dimensionality Reduction

When using high dimensional feature representations of subjects within a classification framework such as this, a critical task is that of feature selection or dimensionality reduction. Particularly when using small sample sizes, such as those commonly available to medical imaging studies, the reduction of the feature space dimensionality is essential to avoid over-fitting and for minimizing classification error. In this work, we use principal component analysis (PCA) to obtain a concise basis of the original feature space while still accounting for the majority of the population variance ($> 90\%$). The PCA features are linear combinations of the original FOD features, and the PCA basis describes an invertible linear operator which relates the two feature spaces. This allows for new subjects to be mapped into the PCA feature space and for vectors in the PCA feature space, such as the decision boundaries, to be mapped back into the original FOD feature space.

2.3 Support Vector Machine Training and Cross-Validation

Linear support vector machines (LSVMs) determine a hyperplane in the feature space, defined by the PCA features in this case, which optimally separates the dataset into patients and controls. Once the SVM has been trained, a new test subject (x) can be labeled, assigned to either the control or patient group, based on the distance between the subject and separating hyperplane. This distance is given by the decision function: $f(x) = w^t x + b$, where w is the SVM feature weights, describing the contribution of each feature, and b is the bias of the hyperplanes from the origin. The distance is used by the classifier to determine,

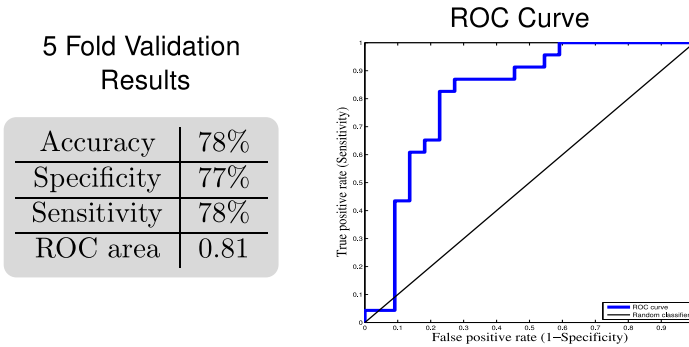


Fig. 2. The classification framework was applied to an ASD dataset consisting of 23 ASD patients and 22 controls. Results from a five-fold cross-validation paradigm are shown on the left, while the receiver operating characteristic (ROC) curve is shown on the right.

via Platt’s method [10], the probabilistic score for the subject and the subject is labeled, either patient or control, based on the sign of the score.

The goal of a classification method is to accurately and robustly classify unseen test subjects. In medical imaging studies the small sample size ($N=45$ in this case) often prohibits the division of the dataset into a single training and testing dataset, each of which accurately represent the entire dataset. For this reason we use a stratified 5-fold cross-validation method to validate our framework. This entails partitioning the original dataset into 5 pieces or folds each of which contain roughly the same proportion of patients and controls. A fold is chosen to serve as the test dataset while the classifier is trained on the remainder of the dataset. This process is repeated until each fold has acted as test data, yielding an abnormality score for each subject as well as a classification accuracy for each fold.

The overall classification method can then be evaluated based on the average accuracy achieved across the 5 folds, as well as by examining the receiver operating characteristic (ROC) curve. The ROC curve is a plot of the sensitivity vs (1-specificity) of the classifier as the discrimination threshold is varied.

3 Application of Classification Framework to a Clinical Population

In this work we apply our framework to the problem of classifying a population of children diagnosed with Autism Spectrum Disorder (ASD). The dataset consisted of 22 typically developing controls (TDC) and 23 ASD patients. Whole brain HARDI was acquired using a Siemens 3T Verio MRI scanner using a spin-echo, echo-planar imaging sequence ($TR/TE=14.7s/110ms$, 2mm isotropic voxels, $b = 3000s/mm^2$, number of diffusion directions=64, 2 b_0 images). Total acquisition time was 18 minutes per subject.

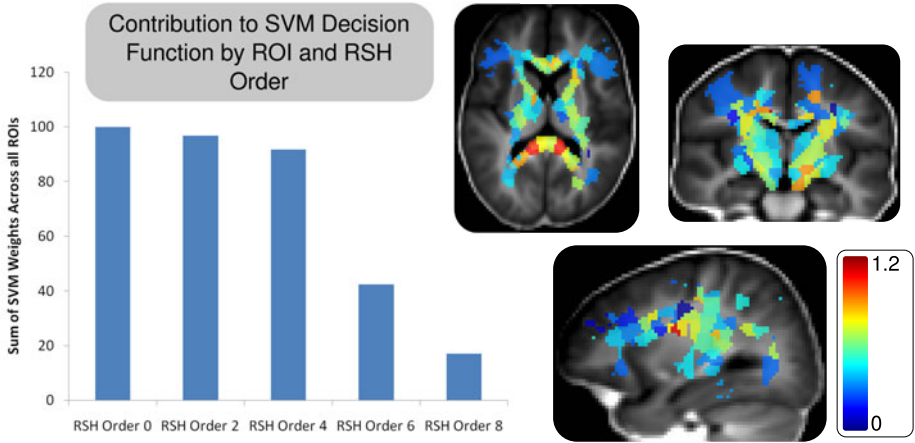


Fig. 3. The SVM weights are mapped into FOD feature space yielding the contribution of each feature to the classification score. The contribution of each RSH order (shown on Left) is obtained by summing the individual contributions across all of the ROIS. Similarly, by averaging the contributions across the rotation invariant features the importance of each spatial region to the classification score (Shown on Right) can be obtained. Higher contributions are indicative a larger group difference within that feature.

The diffusion weighted images (DWI) were first filtered using a joint linear minimum mean squared error filter for removal of Rician noise[13]. Eddy current correction was then performed using affine registration of each DWI volume to the unweighted b0 image [8]. The feature extraction process, described in section 2.1, was then applied using a 12-year old male TDC subject as the registration template as the age of this subject was closest to the population average, yielding a labeled dataset of 45 subjects each represented by a 4415 element feature vector.

The dataset was then divided into 5 folds and cross-validation was performed using the folds as test datasets and the remainder of the dataset for training. For each validation procedure (i.e. each fold) PCA was applied to the training dataset using a variance threshold of 90% and a linear SVM was trained on the resulting PCA coefficients. The SVM classifier was then applied to each subject in the test dataset yielding a classifier score. Additionally the SVM weights, which describe the contribution of each PCA coefficient to the classifier score, were computed and mapped back into the original feature space. In this way a classifier score and a predicted label for each subject in the dataset was obtained, as were 5 measures of the classifier weights mapped into the original feature space.

Figure 2 shows the classification results. The average classification accuracy was 78% (10 subjects misclassified) with a specificity of 77% (low type I error) and a sensitivity of 78% (low type II error). The receiver operating characteristic (ROC) curve shows a steep balanced curve with an area under the curve of 0.81.

These accuracy numbers are in-line with the existing DTI based classification methods [9,7,14,3], which achieve accuracies in the 70-95% range, where 95% mark was obtained using a hypothesis driven feature selection process designed specific for ASD[9], as opposed to the whole brain approach taken here. However such a comparison, based on published accuracies, is greatly hindered by the fact that each study utilizes a different dataset, with different patient characteristics, and in this case multiple diseases (ASD and schizophrenia).

In addition to computing accurate classifier scores our framework has the capability of determining the degree that each of the originally extracted features contribute to the classification score. By averaging the classifier weights in the original feature space, we obtained the mean contribution of each feature to the classification score. Our original feature space consisted of 5 orientation invariant features, one for each order of the RSH expansion, derived from 883 spatial ROIs. The relative contribution of each orientation invariant feature was determined by summing the contributions of that feature across all of the ROIs. Similarly the contribution of each ROI was determined by the sum across the orientation invariant features. These are shown in figure 3.

By examining the contributions of each orientation invariant feature, figure 3-left, we see that the first 3 RSH orders (0, 2 and 4) are predominant in determining the classification score. This suggests that the higher angular frequency information, contained in the higher RSH orders, is perhaps more variable across the population or inherently less reliable. The regional contributions to the classification score shows large contributions from portions of internal capsule (figure 3-right) as well as from the splenium of the corpus callosum, regions that have been previously implicated in ASD. While these results suggest that regional contributions may be useful in localizing WM areas that are affected in ASD a full investigation of these results is beyond the scope of this paper.

4 Conclusion

We have presented a classification methodology based on regional measures of white matter architecture and fidelity as derived from the FOD diffusion data model. The FOD model when coupled with an atlas-free parcellation algorithm yields a physiologically interpretable feature representation, as it contains the orientation and relative proportion of the underlying anatomical fibers in homogeneous WM regions. We demonstrate that this feature representation, when coupled with PCA and a linear SVM yields robust and accurate classification results in an ASD population. In addition to producing a classification score, which may aid in diagnosis, this framework provides the SVM feature weightings which identify the spatial regions and the FOD features contribute to the score. The feature weights elucidate the regions that are affected in the patient population, providing possible insight into the pathology as well as suggesting future directions for the development of hypothesis based classifiers and studies.

While classification based on DTI features has been attempted this is the first classification work that utilizes features derived from HARDI data models to perform patient classification. The high specificity, sensitivity and accuracy demonstrate the feasibility of HARDI based patient classification.

References

1. Bloy, L., Verma, R.: Demons registration of high angular resolution diffusion images. In: 2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pp. 1013–1016 (2010)
2. Bloy, L., Ingalhalikar, M., Verma, R.: Neuronal white matter parcellation using spatially coherent normalized cuts. In: 2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro (2011)
3. Caan, M.W.A., Vermeer, K.A., van Vliet, L.J., Majoie, C.B.L.M., Peters, B.D., den Heeten, G.J., Vos, F.M.: Shaving diffusion tensor images in discriminant analysis: a study into schizophrenia. *Med. Image Anal.* 10(6), 841–849 (2006)
4. Ecker, C., Rocha-Rego, V., Johnston, P., Mourao-Miranda, J., Marquand, A., Daly, E.M., Brammer, M.J., Murphy, C., Murphy, D.G., Consortium, M.R.C.A.: Investigating the predictive value of whole-brain structural mr scans in autism: a pattern classification approach. *Neuroimage* 49(1), 44–56 (2010)
5. Fan, Y., Shen, D., Gur, R.C., Gur, R.E., Davatzikos, C.: Compare: classification of morphological patterns using adaptive regional elements. *IEEE Trans. Med. Imaging* 26(1), 93–105 (2007)
6. Frank, L.R.: Characterization of anisotropy in high angular resolution diffusion-weighted mri. *Magnetic Resonance in Medicine* 47(6), 1083–1099 (2002)
7. Ingalhalikar, M., Kanterakis, S., Gur, R., Roberts, T.P.L., Verma, R.: DTI based diagnostic prediction of a disease via pattern classification. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 558–565. Springer, Heidelberg (2010)
8. Jezzard, P., Barnett, A.S., Pierpaoli, C.: Characterization of and correction for eddy current artifacts in echo planar diffusion imaging. *Magn. Reson. Med.* 39(5), 801–812 (1998)
9. Lange, N., Dubray, M.B., Lee, J.E., Froimowitz, M.P., Froehlich, A., Adluru, N., Wright, B., Ravichandran, C., Fletcher, P.T., Bigler, E.D., Alexander, A.L., Lainhart, J.E.: Atypical diffusion tensor hemispheric asymmetry in autism. *Autism Res.* (December 2010)
10. Platt, J.: Probabilistic outputs for support vector machines and comparison to regularized likelihood methods. In: *Advances in Large Margin Classifiers* (1999)
11. Schnell, S., Saur, D., Kreher, B.W., Hennig, J., Burkhardt, H., Kiselev, V.G.: Fully automated classification of hardi in vivo data using a support vector machine. *Neuroimage* 46(3), 642–651 (2009)
12. Tournier, J.D., Calamante, F., Connelly, A.: Robust determination of the fibre orientation distribution in diffusion mri: Non-negativity constrained super-resolved spherical deconvolution. *NeuroImage* 35(4), 1459–1472 (2007)
13. Tristán-Vega, A., Aja-Fernández, S.: Dwi filtering using joint information for dti and hardi. *Med. Image Anal.* 14(2), 205–218 (2010)
14. Wang, P., Verma, R.: On classifying disease-induced patterns in the brain using diffusion tensor images. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 908–916. Springer, Heidelberg (2008)