Title
Eigenanatomy: Sparse dimensionality reduction for multi-modal medical image analysis

Authors
Kandel, Benjamin M
Wang, Danny JJ
Gee, James C
et al.

Publication Date
2015-02-01

DOI
10.1016/j.ymeth.2014.10.016

Peer reviewed

Methods 73 (2015) 43–53

Contents lists available at ScienceDirect

Methods

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

Eigenanatomy: Sparse dimensionality reduction for multi-modal medical image analysis

Benjamin M. Kandel [a,b,*], Danny J.J. Wang [c], James C. Gee [a,d], Brian B. Avants [a,d]

[a] Penn Image Computing and Science Laboratory, University of Pennsylvania, Philadelphia, PA, United States
[b] Department of Bioengineering, University of Pennsylvania, Philadelphia, PA, United States
[c] Department of Neurology, University of California, Los Angeles, Los Angeles, CA, United States
[d] Department of Radiology, Hospital of the University of Pennsylvania, Philadelphia, PA, United States

ARTICLE INFO

Article history:
Received 9 April 2014
Received in revised form 2 October 2014
Accepted 15 October 2014
Available online 22 October 2014

Keywords:
Multi-modal
Magnetic resonance imaging
Sparse
Pediatric

ABSTRACT

Rigorous statistical analysis of multimodal imaging datasets is challenging. Mass-univariate methods for extracting correlations between image voxels and outcome measurements are not ideal for multimodal datasets, as they do not account for interactions between the different modalities. The extremely high dimensionality of medical images necessitates dimensionality reduction, such as principal component analysis (PCA) or independent component analysis (ICA). These dimensionality reduction techniques, however, consist of contributions from every region in the brain and are therefore difficult to interpret. Recent advances in sparse dimensionality reduction have enabled construction of a set of image regions that explain the variance of the images while still maintaining anatomical interpretability. The projections of the original data on the sparse eigenvectors, however, are highly collinear and therefore difficult to incorporate into multi-modal image analysis pipelines. We propose here a method for clustering sparse eigenvectors and selecting a subset of the eigenvectors to make interpretable predictions from a multi-modal dataset. Evaluation on a publicly available dataset shows that the proposed method outperforms PCA and ICA-based regressions while still maintaining anatomical meaning. To facilitate reproducibility, the complete dataset used and all source code is publicly available.

1. Introduction

Modern imaging datasets are increasingly multimodal. Virtually all modern large-scale imaging studies, even those that concentrate on a given modality, such as resting state fMRI [1], include a variety of imaging measures [2,3]. Although some groups have reported improvements in classification accuracy in Alzheimer's Disease when using multimodal data [4], others have claimed that multimodal classification does not tend to outperform a single sensitive test [5]. This trend towards multimodal data presents challenges in data processing, visualization, and statistical inference. In particular, the extremely high dimensionality of medical imaging data presents challenges to classical linear model-based statistical analyses, which assume that there are more subjects than measured variables ($n > p$). Several approaches exist to deal with the high-dimensional nature of medical imaging datasets.

1.1. Mass-univariate approaches

One of the most widely used methods to perform statistical analyses on medical images is to use voxel-based morphometry (VBM) [6]. VBM performs a statistical test on each voxel in the image, producing a spatial map that describes how closely the values at a given voxel are correlated with an outcome measure. The massive number of multiple comparisons conducted when using VBM necessitate appropriate corrections [7]. In addition, because brain function is spread over regions larger than a single voxel [8], multivariate approaches are more naturally suited to leveraging the spatially distributed information contained in medical imaging data [9].

When examining multimodal data, univariate approaches are further restricted because they do not provide insight into the relationships between the various modalities. One way of using univariate approaches to analyze multimodal data is to perform separate mass-univariate analyses on each modality and examine the degree of spatial overlap between the resulting statistical maps [10–12]. A drawback of this method is that spatial overlap alone does not give insight into the subject-wise interactions or

* Corresponding author at: Penn Image Computing and Science Laboratory, University of Pennsylvania, Philadelphia, PA.
E-mail address: bkandel@seas.upenn.edu (B.M. Kandel).

http://dx.doi.org/10.1016/j.ymeth.2014.10.016
1046-2023/

correlations of the various modalities. To take a somewhat extreme example, if half the experimental population have increased cortical thickness as compared to controls and the other half have increased BOLD activation, a spatial map may show overlapping significant areas, even though no individual subject actually has increased cortical thickness and increased BOLD activation. To provide greater insight into the biological mechanisms underlying observed changes, several studies have begun investigating multivariate approaches to multimodal data [11,13,14], looking at, for example, the correlation between cortical thickness and BOLD activation in a given region.

One challenge of integrating large multimodal datasets is the difficulty in visualizing and interpreting the results, especially when performing multivariate analyses of data. Interpretation of multivariate data is often made easier by sparse methods, which ensure that only a small part of the data set is used for predicting an outcome variable. Sparse methods have enjoyed a resurgence in popularity in recent years, with several groups proposing sparse methods tuned for neuroimaging data [15–24]. Applying sparse techniques to multi-modal data enables specific and biologically interpretable statements to be made about data; for example, "Decreased cortical thickness in the left parietal lobe is correlated with decreased perfusion in the left and right parietal lobes, and this network together predicts a decrease in verbal ability."

1.2. Data-driven dimensionality reduction

Many clinical studies using multimodal imaging data average image values over predefined regions of interest (ROI's) to reduce the dimensionality of the data so that it will be more amenable to standard statistical analyses. Although this approach may be ideal if the ROI's are already known and have anatomically meaningful boundaries, this is not ideal for exploratory analyses which have minimal prior knowledge. Traditionally, linear regression from a high-dimensional dataset is performed after a dimensionality reduction step, such as principal component analysis (PCA) [25]. However, PCA-derived eigenvectors have global support and therefore do not provide anatomical specificity. Sparse techniques can provide more local specificity. In particular, a recently introduced sparse dimensionality reduction technique, "eigenanatomy," has proven to provide greater power to detect group differences than either voxel-based morphometry (VBM) [26] or pre-defined ROI's [27] while maintaining anatomical interpretability. Here, we extend the eigenanatomy approach to a multi-modal setting. Although the sparse eigenvectors are orthogonal in the image space, orthogonality is not enforced on the low-dimensional coefficients generated by projecting the imaging data onto the sparse eigenvectors. Therefore, care must be taken to prevent excessive collinearity among the predictor variables. We demonstrate that even with collinearity in the predictor variables, our method of extending eigenanatomy to multi-modal datasets produces a more accurate prediction of age in a pediatric population than principal component regression, independent component regression, or regression on average values within regions defined by the AAL atlas.

The eigenanatomy objective function is not new to this work. Here, we focus on the practical challenges, including validation, interpretation, and visualization of predictive models, involved in multimodal data analysis, and demonstrate the advantages of the eigenanatomy framework for multi-modal neuroimaging studies as compared to either classical dimensionality reduction techniques or predefined regions of interest (ROI's). The release of all data and code used to generate the paper will facilitate the use of this technique as a template for future studies, as well as encourage reproduction of similar evaluations with different datasets.

2. Methods

2.1. Reproducibility

To facilitate the use of this study as a template for other multimodal population studies, we have attempted to make it as reproducible as possible. All the data is available from an open-access data repository. The paper itself is written using the R package `knitr` [28], which facilitates on-the-fly production of figures from data, enhancing reproducibility and documenting all data processing steps. The full code for producing the paper, including raw data and code for producing figures, is available from https://bitbucket.org/bkandel/multimodaleanat.

2.2. Dimensionality reduction techniques

Dimensionality reduction is a technique to reduce the complexity of input data into a relatively small number of summary measures. Linear dimensionality techniques can be written as a matrix factorization problem. We assume the input data is given in an $n \times p$ input matrix $\mathbf{X}$, where n is the number of observations or subjects and p is the number of variables associated with each observation. In the context of medical imaging, n typically ranges from a few tens to a few hundred, and p is on the order of 10^3–10^6, depending on the size of images. Dimensionality reduction seeks to find a factorization of $\mathbf{X}$ into an $n \times k$ *coefficient* matrix $\mathbf{U}$ and a $p \times k$ *loading* or eigenvector matrix $\mathbf{V}$ so that $\mathbf{X} \approx \mathbf{U}\mathbf{V}^{\mathrm{T}}$. The most well-established method for dimensionality reduction is principal component analysis (PCA), which finds an orthogonal matrix (i.e. $\mathbf{V}^{\mathrm{T}}\mathbf{V} = \mathbf{1}$) that projects the input matrix to a lower-dimensional subspace. More recently, independent component analysis (ICA, e.g. [29]) has become widely used in the neuroimaging community. ICA seeks a decomposition of $\mathbf{X}$ in which the components are independent, which is a stronger condition than orthogonality.

One drawback of standard PCA and ICA is that the eigenvectors often cover the entire input matrix, meaning that each entry in the coefficient matrix is a weighted average of all the voxels in the image. This makes interpretation of the output difficult for two related reasons. First, the lack of spatial specificity of the eigenvectors makes it difficult to use the coefficients to investigate anatomically-informed biological hypotheses. For example, it is impossible to use the coefficients from a PCA decomposition to look at the relation between left precuneal atrophy and age. In addition, because the eigenvectors have both positive and negative components, interpreting the weights in a linear regression model that relates the coefficients to an outcome measure is not intuitive. This is for two reasons. First, the PCA eigenvectors can contain negative weights, even if the input data is strictly positive, as is the case in cortical thickness, perfusion, and fractional anisotropy (FA) images. If the weight of a given PCA coefficient in a linear model is positive but the corresponding entry in the eigenvector is negative, it follows that an *increase* in the input matrix corresponds to a *decrease* in the outcome variable. Second, this problem is compounded by the overlapping nature of the eigenvectors: Because a given brain region can contribute positively or negatively to each eigenvector, it is very difficult to go back from the coefficient weights to the biological meaning of the weight. Furthermore, interpreting the eigenvectors themselves without accounting for the weights is ill-advised [30], as the eigenvectors can be confounded by biases in the data in non-obvious ways [31]. Therefore, although PCA may be used for predicting age in unseen data, it is not as useful for testing biological hypotheses.

Sparse dimensionality reduction techniques [32,33] deal with the problems of global support of PCA eigenvectors by enforcing

a sparsity constraint on the matrix decomposition. The sparseness forces the eigenvectors to have only a few non-zero entries, thus making the eigenvectors more amenable to anatomically-specific hypothesis testing. Eigenanatomy [26,34] augments these constraints with smoothness and cluster thresholds, which ameliorate the inherent instability of sparse solutions [35]. In addition, eigenanatomy requires the eigenvectors to be unsigned, thereby eliminating the difficulty of interpreting linear model weights: A positive weight in a linear model means that an increase in the predictor variable corresponds to an increase in the outcome variable and vice versa.

Mathematically, the eigenanatomy objective function is

$$\begin{aligned} \underset{U,V}{\text{argmin}} \quad & \|\mathbf{X} - \mathbf{U}\mathbf{V}^\top\|_2^2 + \text{smooth}(G^*(\mathbf{V})) + \text{cluster}(\mathbf{V}) \\ \text{subject to} \quad & \|\mathbf{V}\|_0 < \gamma, \\ & \mathbf{V} \succeq 0 \end{aligned} \tag{1}$$

where γ is the desired level of sparseness, the ℓ_0 "norm" operator $\|\cdot\|_0$ is the number of non-zero entries in the argument, and G^* is the manifold in image space representing a geometric object [36]. Although optimizing over an ℓ_0 penalty is an NP-hard problem, thresholded gradient techniques have been proposed for sparsity-constrained optimization problems [37], and we adopt the conjugate gradient version used in [26]. We use the open-source eigenanatomy decomposition function `sparseDecom` as implemented in the `ANTsR` package (github.com/stnava/antsr).

One challenge with sparse dimensionality reduction techniques is that the projections on the sparse eigenvectors are not orthogonal [38], which can lead to collinearity in the regressors. To deal with this problem, we first perform an eigendecomposition using very sparse eigenvectors (50 eigenvectors, each covering 2% of the image), and then use hierarchical graph clustering to cluster eigenvectors into fewer, but larger, eigenvectors. This strategy of oversegmenting into "supervoxels" and then clustering has proven fruitful in image segmentation [39]. Eigenvectors which, when the population of images is projected onto them, have correlated values are clustered together into communities. This clustering reduces the degree of collinearity between the eigenvectors. The clustering stops when the graph density of the original eigenvectors reaches a pre-set threshold. We have found that a graph density of 0.10 is a good value for maintaining anatomical specificity and predictive accuracy. We used the `igraph` package for graph-based community construction [40]. For PCA, we used the `princomp` function in `R`, and we used the `fastICA` package [29] for ICA computations.

2.3. Multi-modal data visualization and integration

One key advantage of sparse dimensionality reduction techniques is that they facilitate incorporation of multiple modalities into a single interpretable multivariate model. In our experiments, we limit ourselves to linear and generalized linear models, as the coefficients and p-values of these models are well-established and can be interpreted using mature, standard statistical techniques. The predictor variables for the linear models are the coefficient scores for the various decomposition techniques, obtained by projecting the original image space data onto the eigenvectors to obtain an n by k matrix, where n is the number of subjects and k is the number of eigenvectors.

To choose which of the projections we include in the final linear model relating imaging data to an outcome, we use a model selection. The sparsity induced by model selection complements the sparsity of the individual eigenvectors, allowing multiple eigenvectors to enter the model if they each contribute to the outcome but still maintaining the anatomical interpretability afforded by the spatial sparsity of the eigenvectors. Because imaging data is generally highly correlated between regions, we use the VIF regression [41] model selection technique, which explicitly accounts for correlation among predictors. Because the focus of the study is on feature generation and not feature selection, we fix the model selection method and apply it to all decomposition techniques.

Visualization of multivariate regressions is challenging, especially when considering multiple input and multiple output variables. We use the Sankey diagrams from the `d3` library for visualizing correlations between input and output variables. These diagrams permit quick and intuitive visual exploration of the correlation between each input and output variable.

Because multivariate models are more prone to overfitting than univariate models, a rigorous training–testing split of data is necessary to assess generalizability of results. Two steps are necessary to interpret an optimized multivariate linear model. First, the regression weights can be interpreted. A cautionary note is in order here: Interpretation of regression coefficients in multivariate linear regression is not trivial and is subject to confounding factors and multicollinearity artifacts, even when using a model selection technique like VIF that attempts to minimize collinearity between predictors. Specifically, the regression coefficients from a multivariate linear model correspond to the effect of a given regressor *keeping all other regressors constant*. Second, statistical assessment of the significance of the model can be performed by applying the model to unseen testing data and evaluating the performance of the model. Note that because model selection is performed in the training data, naïvely looking at the p-value of the model in the training data is inappropriate and would constitute statistical "double-dipping" unless special care is taken [42]. The ability to evaluate a model in unseen testing data is a key advantage to predictive models as opposed to correlation-based approaches and is valuable even if predicting the outcome from neuroimaging data is not clinically useful.

2.4. Clinical data

We use the Pediatric Template of Brain Perfusion dataset, which contains multimodal data for normally developing pediatric subjects. The full pediatric dataset is available at http://figshare.com/articles/The_Pediatric_Template_of_Brain_Perfusion_PTBP_/923555. Briefly, the dataset consists of 133 subjects, 82 females and 83 males, ages 7.07–17.99, mean 12.48 years. Cohort selection aimed to match the demographic distribution of children in the United States, based on US census data, for race, ethnicity, gender, and family income. Distributions of age and income with respect to sex are shown in Fig. 1.

2.5. Image acquisition

All MRI experiments were performed on a Siemens 3T TIM Trio scanner. For the T1-weighted acquisition, Magnetization-Prepared Rapid Acquisition Gradient Echo (MPRAGE) images were acquired using a 3D inversion recovery sequence with TR/TE/TI = 2170/4.33/1100 ms. The resolution is 1x1x1mm^3 with a matrix size of $256 \times 256 \times 192$. The flip angle = 7° and total scan time was 8 min and 8 s. For perfusion imaging, pseudo continuous arterial spin labeled (pCASL) images were acquired using gradient-echo echo-planar imaging (EPI) with TR/TE = 4000/12 ms. The resolution was $3.125 \times 3.125 \times 6$ mm (5 mm thickness with a 1 mm gap) over a $64 \times 64 \times 24$ matrix. 40 label/control pairs were acquired. Generalized autocalibrating partially parallel acquisition (GRAPPA) was used with an acceleration factor of 2. Labeling duration was 1.5 s and the post-labeling delay was 1.2 s. Total imaging time was 5 min and 30 s. Diffusion weighted images were acquired with single-shot spin-echo EPI with TR/TE = 9500/87 ms. A single

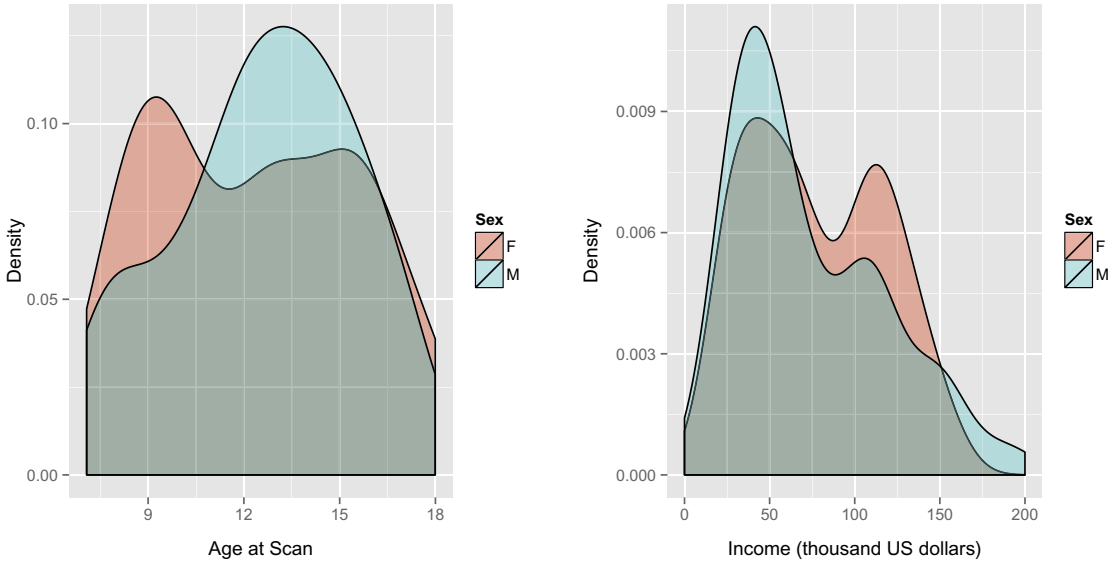

Fig. 1. Left: Distribution of ages among subjects. Right: Distribution of income among subjects.

$b = 0$ volume was acquired along diffusion weighted images for 30 directions with b-value = 1000. The resolution was $2 \times 2 \times 2$ mm with a matrix size of $128 \times 128 \times 75$ voxels, with a flip angle of 90°. Total acquisition time was 11 min with two acquisitions.

2.6. Image preprocessing

Our image processing pipeline is based on Advanced Normalization Tools (ANTs) and the `ANTsR` package, which provides an interface from the `R` computing environment to the ANTs image processing utilities. We used the ANTs template-building scripts to construct an unbiased template from the subject population [43]. Although template-building methods have traditionally been used for single-modality data, we leverage recent advances in template-building techniques to construct a multi-modal atlas [44,45]. AAL [46] and JHU [47] labels and priors for brain extraction and tissue segmentation are warped from pre-existing template space to the custom template using a multi-atlas label fusion (MALF) approach [48]. We generated a six-tissue segmentation: Cerebrospinal fluid (CSF), cortical gray matter (GM), white matter (WM), deep GM, brainstem, and cerebellum. Full details and code for constructing multi-modal templates, along with sample data, can be obtained from github.com/ntustison/TemplateBuildingExample.

For each subject's T1 image, we performed bias correction [49], brain extraction, normalization to the template [43], prior-based segmentation [50], and cortical thickness calculation [51]. The T1-based cortical thickness measurement pipeline is based on the `antsCorticalThickness.sh` script [52]. Diffusion processing, including fractional anisotropy (FA) calculation, is performed using Camino [53] as incorporated into ANTs [54]. pCASL images were processed using a robust regression approach [55,56] that regressed tag from untagged images. The M0 image was obtained by averaging the non-tag images, and was used as a reference for motion correction. Motion correction was performed using the `antsMotionCorr` function in ANTs, and motion parameters (three for rotation and three for translation) were regressed out as nuisance variables. Physiological nuisance parameters were estimated using the CompCor approach, which estimates physiological noise by computing an eigendecomposition of the high-variance regions in an image [57]. Full details are available in the open-source script at https://raw.github.com/stnava/ANTs/master/Scripts/antsASLProcessing.sh. The blood T1 value was adjusted for age and gender as $T1 = (2115.6 - 21.5 * \text{age} - 73.3 * \text{sex})$ ms, where female sex was set to 0 and male was set to 1, as suggested in [58].

3. Results

3.1. Template and average images

The population-specific template and the average of all images warped to the template are shown in Fig. 2. As expected, cortical thickness in the medial temporal and medial frontal is relatively high, and cortical thickness in the occipital lobe and motor cortex is relatively low [59]. Perfusion in the cortex is higher than in white matter or CSF, and the higher perfusion in the deep gray matter structures is clearly visible, consistent with existing literature [60]. The average FA image shows higher FA in the corpus callosum, with lower FA in the gray matter.

3.2. Eigenanatomy decompositions

We obtained eigenanatomy-derived eigenvectors for the populations of cortical thickness, CBF, and FA images. The correlation matrix of all the eigenvectors is shown in Fig. 4. Cortical thickness tended to be more correlated with CBF than with FA, and each modality showed strong intra-modality correlations. To help alleviate the collinearity of the predictors, we clustered the eigenvectors to create fewer, but larger, eigenvectors. The correlations between these clustered eigenvectors are shown in Fig. 4. Clustering significantly reduced the collinearity between the eigenvectors, especially for cortical thickness and FA. Because of the reduced collinearity, we used the clustered eigenvectors in our regressions.

A sample clustered eigenanatomy-derived eigenvector is shown in Fig. 5. The eigenanatomy eigenvector covers primarily the precuneus. On the other hand, the PCA and ICA vectors have global support. The eigenanatomy vector also has only positive values, whereas the PCA and ICA vectors each have positive and negative components.

3.3. Univariate data exploration

Before looking at the multivariate regression, it is helpful to look at univariate correlations between the eigenanatomy regions and the outcome variable. We use a Sankey diagram to show the

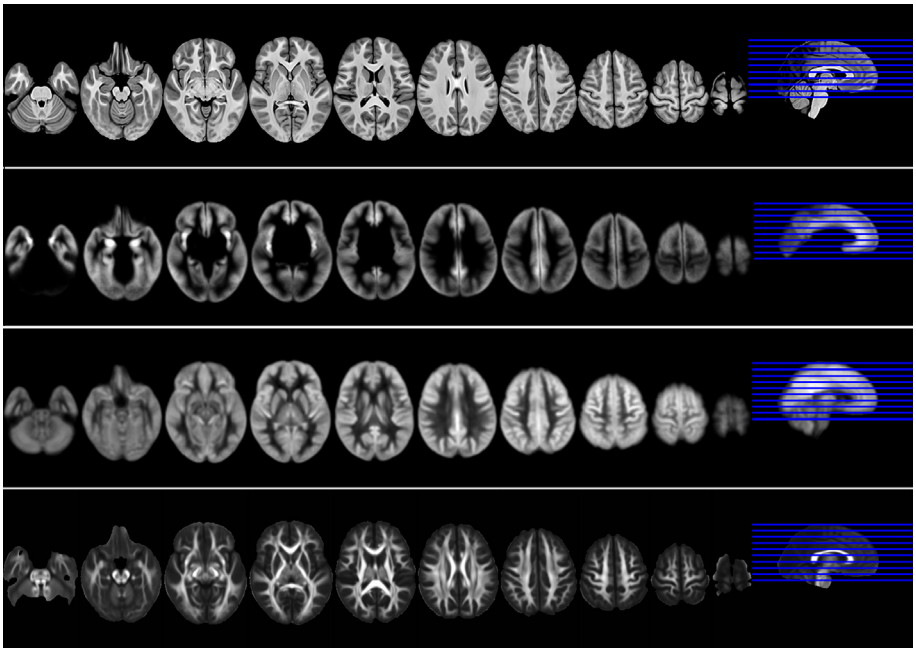

Fig. 2. From top to bottom: Population-specific template; average cortical thickness image for all subjects; average perfusion image; average fractional anisotropy (FA) image for all subjects.

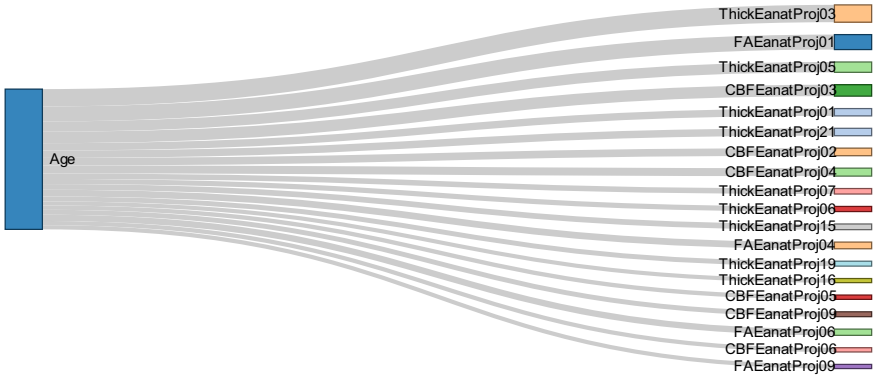

Fig. 3. A Sankey diagram showing the univariate relationships between image eigenvectors and age. Thicker links correspond to more highly significant correlations, and correlations with an FDR-corrected *p*-value of less than 0.05 are not shown.

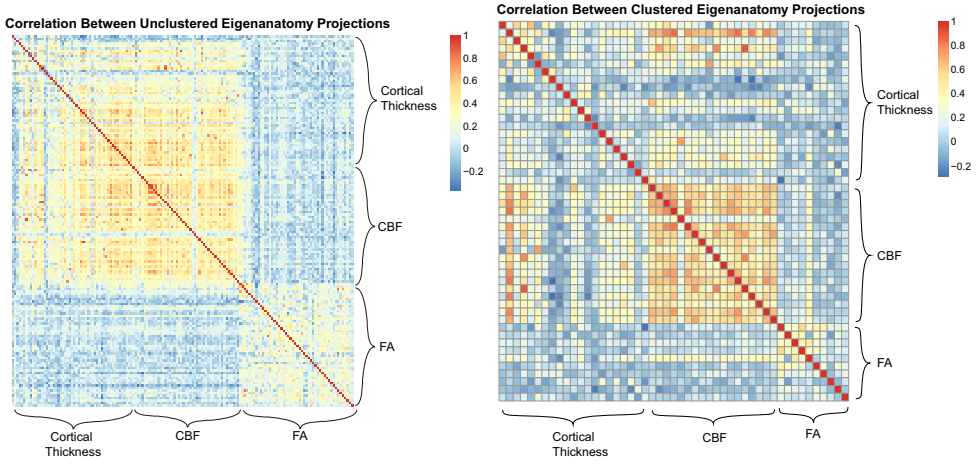

Fig. 4. Correlation heatmap between projections on all eigenanatomy eigenvectors for cortical thickness, CBF, and FA. Left: Correlations of projections on eigenvectors before the eigenvector clustering step. Right: Correlations after clustering eigenvectors. Clustering eigenvectors significantly reduces correlation between eigenvectors of cortical thickness and FA, but less so for CBF.

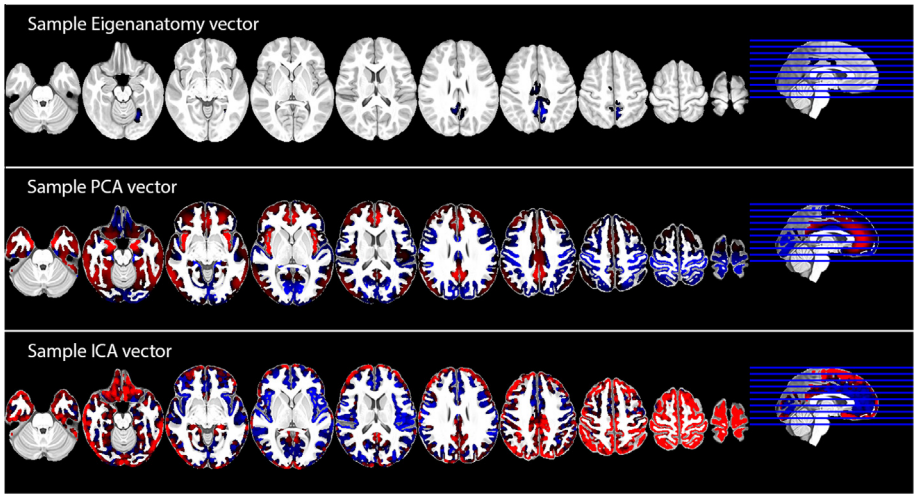

Fig. 5. Sample eigenanatomy, PCA, and ICA eigenvectors. Positive components are in blue, with negative components in red. Eigenanatomy returns unsigned components, facilitating interpretation of regression coefficients. The spatial localization of eigenanatomy enables testing spatially specific neuroanatomical hypotheses.

links between age and the various eigenanatomy regions (Fig. 3). Although many regions are correlated with age, not all survive FDR correction.

3.4. Regression analysis in training data

We ran a VIF regression-based model selection strategy to select projections for prediction of age from neuroimage data. The model selection, when performed on the decomposition-based methods (Eigenanatomy, PCA, and ICA), retained only thickness and FA projections, but did not retain CBF projections; for the pre-defined ROI's, thickness, CBF, and FA values were all retained. With all other regressors held constant, increases in thickness in bilateral anterior precuneus were weakly correlated with increased age, and decreases in thickness in bilateral posterior precuneus and cingulate were more strongly correlated with increases in age. Increases in FA in corpus callosum were correlated with increasing age, whereas with all other regressors held constant, decreases in FA in the inferior longitudinal fasciculus were correlated with increases in age. The eigenanatomy eigenvectors are rendered in Fig. 7. Anatomically specific descriptions and coefficient interpretation are not possible with PCA and ICA eigenvectors.

3.5. Evaluation on testing data

To verify the generalizability of the regression to unseen data, we predicted the age of half the data and evaluated the fit. The eigenanatomy predictors outperformed the PCA and ICA predictors, as well as the AAL/JHU regions (Fig. 6). Although all three methods produced a highly significant fit, the eigenanatomy projections gave greater correlation between true and predicted age, lower root mean squared error, a lower p-value for the correlation (Table 2) slope relating predicted to true age closer to unity.

3.6. Comparison to non-linear models

To examine the influence of regression method on prediction accuracy, we constructed a random forest predictor to relate eigenanatomy, PCA, and ICA projections, as well as AAL and JHU regions, to age. We compared the prediction using both the VIF regression-based model selection technique and using all of the projections. Random forests decreased the mean absolute error (MAE) and increased correlation for all methods, but eigenanatomy projections still outperformed PCA and ICA projections and predefined ROI's. Giving all the data to the random forest resulted in more accurate predictions than giving only the regions selected using the VIF model selection technique. Predicted vs. true age is plotted in Fig. 8, with quantitative values in Table 3.

4. Discussion

We have presented here extensions to a sparse dimensionality reduction method that enable biologically interpretable predictions from high-dimensional multi-modal data that requires minimal prior knowledge. We demonstrated how incorporating a variety of techniques, including over-segmentation and then clustering of eigenvectors and variance inflation factor regression, minimize the impact of collinearity of regressors on prediction accuracy, making these methods suited for multi-modal image analysis problems.

4.1. Interpretation of regression

One major advantage of sparse matrix decompositions for multi-modal image analysis is that the individual regressors in a linear model can be anatomically interpreted. Our results here for cortical thickness agree with standard results in developmental neuroscience that cortex generally thins throughout development [61]. The observation that increased FA in the corpus callosum is strongly correlated with increasing age is also in consonant with previous research [62]. Similarly, although the FA in the inferior longitudinal fasciculus (ILF) increases with age, it increases at a lower rate than the corpus callosum [62], and so when corrected for the corpus callosum, the regression coefficient becomes negative.

4.2. Collinearity of sparse eigenvectors

One confounding issue in sparse matrix decompositions is that projections onto eigenvectors are not generally constrained to be orthogonal [38]. This lack of orthogonality can complicate regression from sparse eigenvectors, as the collinearity of the regressors makes the design matrix ill-conditioned and can confound the regression coefficients. Variance inflation factor (VIF) regression [41] ameliorates this problem somewhat by explicitly choosing a regressor set that minimizes covariance, but cannot change the pre-existing collinearity among the regressors. Clustering the

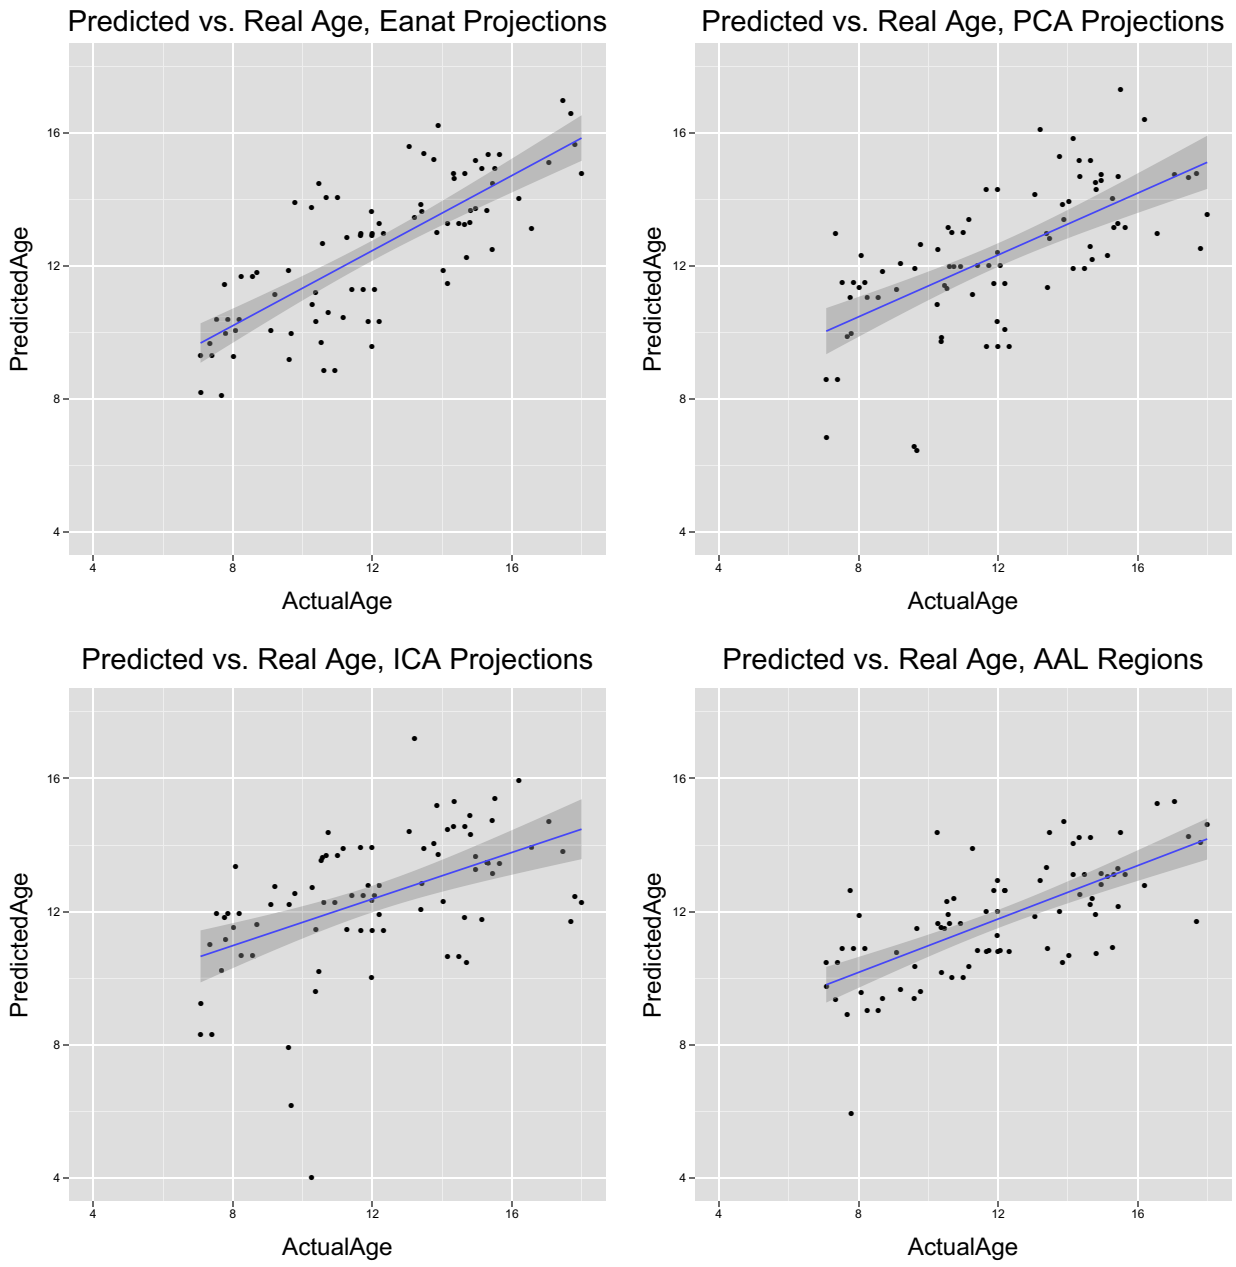

Fig. 6. Age prediction on testing data using eigenanatomy, PCA, and ICA projections, and using AAL regions. Quantitative statistics can be found in Table 2.

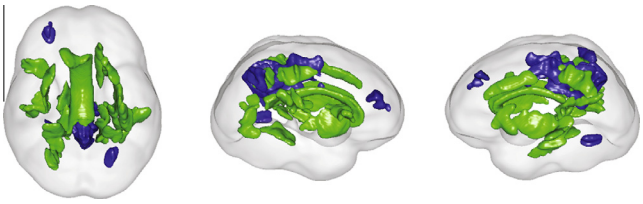

Fig. 7. Eigenanatomy regions retained for prediction of age. Thickness predictors are blue; FA predictors are green.

sparse eigenvectors based on their correlation further helps to minimize this issue, but does not completely resolve it. In particular, the CBF values appeared to be highly collinear, even after clustering the eigenvectors. This collinearity is a result of the greater degree of collinearity in the data, and resolving this issue for CBF specifically is an important issue. Because CBF is generally so correlated across the brain, several groups have divided each subject's CBF values by the mean of the whole brain CBF to create a relative CBF measure [63]. Although relative CBF may not be as predictive of age as raw CBF is, it may give better spatial discrimination of developmental trends and so may generate more meaningful spatial information.

4.3. Sparsity in multi-modal image analysis

Although sparsity can be a useful penalty in medical image analysis, it can also be overused and misinterpreted. For example, if a signal is not sparse and is instead distributed throughout the brain, enforcing sparsity on a solution can be misleading by implying that only certain parts of the brain contain signal. This is due to the fundamental nature of sparse solutions: By discarding redundant information, they generate a minimal set of information that can be used to predict the outcome – even if other parts of the

Table 1
Regression analysis for prediction of age in training data using eigenanatomy, PCA, and ICA projections.

Method	Projection	Estimate	Std. error	*t* value	Pr(>—*t*—)	Main positions
Eigenanatomy	(Intercept)	10.508	4.288E+00	2.45	1.657E−02	
	ThickEanatProj01	0.547	6.618E−01	0.827	4.110E−01	Bilateral anterior precuneus
	ThickEanatProj03	−2.729	6.399E−01	−4.26	5.711E−05	Bilateral posterior precuneus
	ThickEanatProj05	−1.890	8.070E−01	−2.34	2.180E−02	Cingulate
	FAEanatProj01	48.323	6.869E+00	7.03	7.559E−10	Corpus callosum
	FAEanatProj05	−14.994	4.105E+00	−3.65	4.742E−04	Inferior longitudinal fasciculus
PCA	(Intercept)	12.659	2.439E−01	51.9	1.126E−61	
	ThickPCAProj03	−0.023	5.759E−03	−3.93	1.838E−04	
	ThickPCAProj04	−0.022	5.962E−03	−3.74	3.495E−04	
	FAPCAProj02	0.286	6.414E−02	4.46	2.799E−05	
	FAPCAProj05	−0.275	7.749E−02	−3.55	6.690E−04	
ICA	(Intercept)	1.753	1.531E+00	1.14	2.557E−01	
	ThickICAProj10	−0.036	8.731E−03	−4.11	9.447E−05	
	FAICAProj08	−0.290	7.515E−02	−3.86	2.286E−04	
AAL/JHU	(Intercept)	11.715	4.685E+00	2.5	1.452E−02	
	ThickAALLabel11	−2.295	8.341E−01	−2.75	7.398E−03	Left lateral frontal
	CBF.AALLabel10	0.005	1.549E−02	0.354	7.242E−01	Right orbitofrontal
	CBF.AALLabel28	−0.083	2.298E−02	−3.63	5.118E−04	Right gyrus rectus
	FA.JHULabel17	31.026	6.622E+00	4.69	1.181E−05	Internal capsule

Table 2
Prediction performance (correlation, mean absolute error, and *p*-value for correlation) in testing data for predicting age from eigenanatomy, PCA, and ICA projections, and AAL/JHU regions. Predictions using multimodal data outperformed predictions using one modality in every instance, and eigenanatomy projections performed the best overall.

	Method	Correlation	MAE	*p*-Value
1	Eigenanatomy	0.772	1.59	1.40e−17
2	PCA	0.650	1.89	2.98e−11
3	ICA	0.496	2.16	1.84e−06
4	AAL/JHU Labels	0.691	1.74	5.07e−13
5	Eigenanatomy, Thickness	0.598	1.88	2.29e−09
6	PCA, Thickness	0.585	1.96	6.17e−09
7	ICA, Thickness	0.406	2.29	1.40e−04
8	AAL Labels, Thickness	0.221	2.50	4.50e−02
9	Eigenanatomy, CBF	0.529	2.08	2.66e−07
10	PCA, CBF	0.585	1.96	6.17e−09
11	ICA, CBF	0.178	2.52	1.08e−01
12	AAL Labels, CBF	0.494	2.04	2.09e−06
13	Eigenanatomy, FA	0.611	2.03	8.32e−10
14	PCA, FA	0.669	1.97	4.84e−12
15	ICA, FA	0.565	2.04	2.60e−08
16	JHU Labels, FA	0.617	1.84	5.16e−10

brain also contain similar information. Put another way, there is currently no standard method for setting the sparsity parameter in a sparse decomposition. The dependence of the solution on the sparsity penalty has long been known [64]. Although we do not solve this issue here, clustering the sparse eigenvectors allows the eigenvectors to adapt naturally to a grouping that allows different levels of sparsity for different regions, thereby making the sparsity penalty less arbitrary.

With this caveat, though, sparsity is a useful tool for multimodal image analysis problems. Because information is virtually always shared between different modalities, the reason to incorporate more than one modality into an analysis problem is to highlight the unique information given by one modality that is not given by another modality, ignoring the shared information between the modalities. This emphasis on extricating the unique information from the various modalities while ignoring shared information between the modalities naturally leads to sparse methods as a way to produce predictive models. In particular, while cortical thickness may measure the number of cells in the cortex, CBF measures cortical perfusion, while FA is a measure of white matter tract integrity. Understanding the unique components of these measures may give greater insight into how different neurobiological mechanisms develop and relate to each other.

4.4. Linear vs. non-linear models

One important tradeoff when constructing predictive models is the expressive power of the model vs. the interpretability of the output of the model. In particular, non-linear predictive models such as random forests [65], while achieving state-of-the-art classification accuracy, do not produce directly interpretable models. Although some attempts have been made at interpreting the prediction model, including a variable importance score [65], the model still does not have the direct interpretation and statistical significance theory attached to linear regression coefficients. The method proposed for selecting a subset of eigenanatomy regions and then feeding those projections into a random forest attains a compromise solution: The predictors used by the random forest correspond to discrete anatomical regions, but the way those regions are combined does not have a simple interpretation. Still, this level of interpretability may be sufficient for some research questions, and gives more insight into the prediction mechanism than feeding all brain regions into a random forest.

4.5. Comparison to state-of-the-art results

Prediction of age from medical imaging data has drawn a significant amount of attention in recent years [66–70]. Most of these methods use smooth, non-linear models to predict age from imaging data, so the results from the random forest predictor are most relevant to comparing against this work. Comparing directly against previous work is complicated by a variety of factors. The age range of the subjects can influence the accuracy of the prediction, as the prediction is generally more accurate for younger subjects. In addition, different methods use different combinations of modalities. Furthermore, the studies that produced the highest accuracy used training datasets of several hundred to nearly a thousand subjects [67,68], whereas our database included a training set of only 82 subjects. Even with these caveats, though, the random forest achieved a mean absolute error (MAE) of 1.27, which is comparable to the MAE of 1.1 reported in other studies.

4.6. Drawbacks of the framework

Although the framework for multi-modal image analysis shows promise as compared to standard dimensionality reduction techniques, it still has some drawbacks. Even with well-behaved data, interpreting coefficients of multivariate linear regressions is not

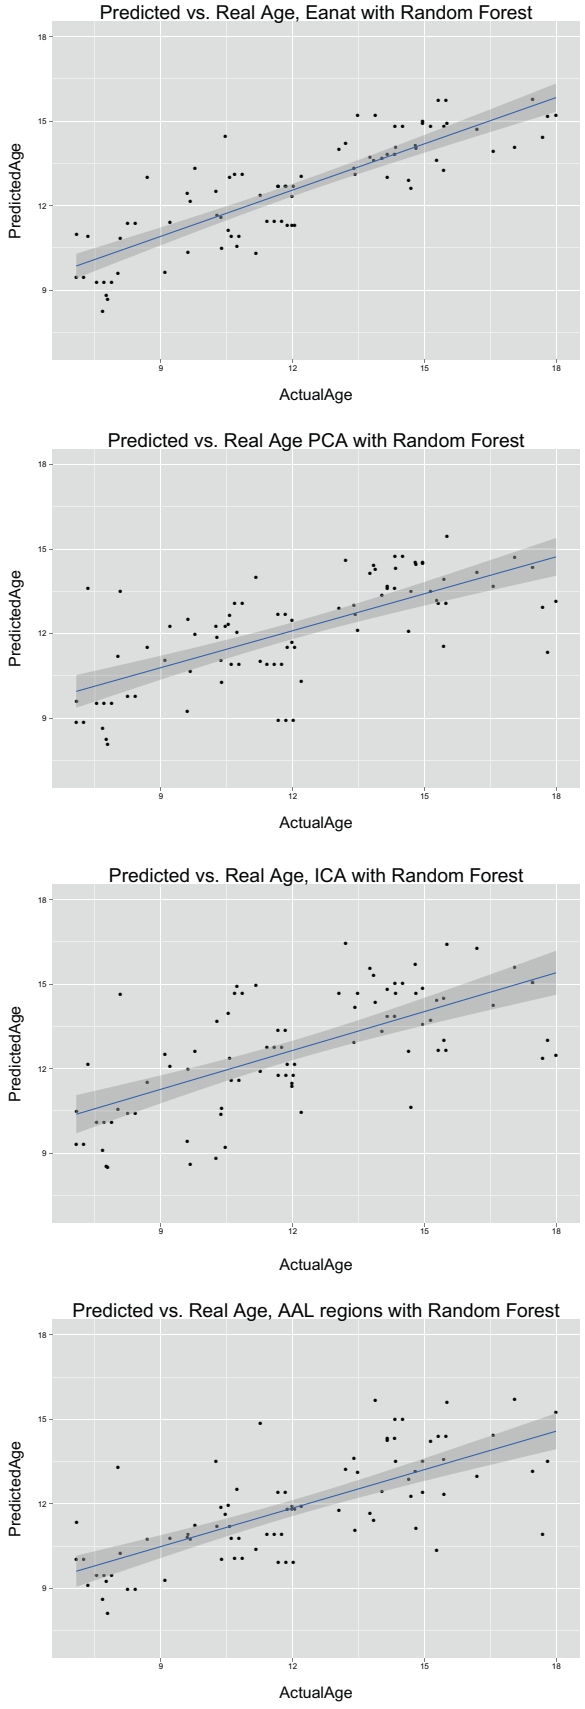

Fig. 8. Age prediction on testing data using eigenanatomy, PCA, and ICA projections, and AAL regions as input to random forests. Random forests achieved lower error than linear models at the expense of model interpretability. Quantitative descriptions of the predictions are found in Table 3.

Table 3
Prediction performance (correlation, mean absolute error, and *p*-value for correlation) in testing data, using random forests, for eigenanatomy, PCA, and ICA projections, and AAL regions. Predictions using all data were compared to predictions using the variance inflation factor (VIF) regression-selected predictors.

	Method	Correlation	MAE	*p*-Value
1	Eigenanatomy, VIF-selected predictors	0.849	1.36	<1e−16
2	Eigenanatomy, all data	0.867	1.27	<1e−16
3	PCA, VIF-selected predictors	0.693	1.62	4.1e−13
4	PCA, all data	0.784	1.44	<1e−16
5	ICA, VIF-selected predictors	0.654	1.82	2e−11
6	ICA, all data	0.832	1.36	<1e−16
7	AAL/JHU regions, VIF-selected predictors	0.724	1.56	1.1e−14
8	AAL/JHU regions, all data	0.849	1.29	<1e−16

as straightforward as interpreting coefficients of univariate regressions. In particular, because the coefficients of a multivariate regression are corrected for the effects of all the other covariates, careful attention must be paid to the whole model, and each individual regressor cannot be viewed in isolation.

Another drawback of sparse decomposition techniques is the instability of the solutions, as demonstrated by the different regions selected by the AAL/JHU regions and the eigenanatomy decomposition (Table 1). Instability of sparse learning techniques is a fundamental characteristic of sparsity in general [35]. Although bootstrapping-based techniques have been used to estimate the stability of sparse regression methods [20], extending this technique to achieve stable multi-component sparse dimensionality reduction is a significant open problem in the field.

5. Conclusion

We have proposed a method for incorporating eigenanatomy decompositions into a multi-modal data analysis pipeline. We have proposed methods for dealing with collinearity of sparse predictors. In addition, we have demonstrated that the sparse eigendecomposition method outperforms traditional PCA and ICA-based dimensionality reduction methods for predicting age from imaging data, while maintaining anatomical interpretability of the regression. This method is suited to extracting the relevant and unique information from multi-modal imaging datasets.

Acknowledgments

The authors would like to acknowledge the support of the following grants: R01-MH080892, R01-NS081077, R01-EB014922 (D.W.), T32-EB009384, P30-NS045839, R01-NS065347, R01-EB006266, R01-DA022807 (J.C.G. and B.M.K.).

References

[1] B.B. Biswal, M. Mennes, X.-N. Zuo, S. Gohel, C. Kelly, S.M. Smith, C.F. Beckmann, J.S. Adelstein, R.L. Buckner, S. Colcombe, A.-M. Dogonowski, M. Ernst, D. Fair, M. Hampson, M.J. Hoptman, J.S. Hyde, V.J. Kiviniemi, R. Ktter, S.-J. Li, C.-P. Lin, M.J. Lowe, C. Mackay, D.J. Madden, K.H. Madsen, D.S. Margulies, H.S. Mayberg, K. McMahon, C.S. Monk, S.H. Mostofsky, B.J. Nagel, J.J. Pekar, S.J. Peltier, S.E. Petersen, V. Riedl, S.A.R.B. Rombouts, B. Rypma, B.L. Schlaggar, S. Schmidt, R.D. Seidler, G.J. Siegle, C. Sorg, G.-J. Teng, J. Veijola, A. Villringer, M. Walter, L. Wang, X.-C. Weng, S. Whitfield-Gabrieli, P. Williamson, C. Windischberger, Y.-F. Zang, H.-Y. Zhang, F.X. Castellanos, M.P. Milham, Proc. Natl. Acad. Sci. 107 (10) (2010) 4734–4739 [Online]. Available: http://www.pnas.org/content/107/10/4734.
[2] S.G. Mueller, M.W. Weiner, L.J. Thal, R.C. Petersen, C.R. Jack, W. Jagust, J.Q. Trojanowski, A.W. Toga, L. Beckett, Alzheimers Dement. 1 (1) (2005) 55–66.
[3] T.D. Satterthwaite, M.A. Elliott, K. Ruparel, J. Loughead, K. Prabhakaran, M.E. Calkins, R. Hopson, C. Jackson, J. Keefe, M. Riley, F.D. Mentch, P. Sleiman, R. Verma, C. Davatzikos, H. Hakonarson, R.C. Gur, R.E. Gur, NeuroImage (2014) [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811913008331.

[4] D. Zhang, D. Shen, Lecture Notes in Computer Science, vol. 7012, Springer, Berlin/Heidelberg, 2011, pp. 60–67 [Online]. Available: http://www.springerlink.com/content/r00q76440l888871/abstract/.

[5] S.M. Landau, D. Harvey, C.M. Madison, E.M. Reiman, N.L. Foster, P.S. Aisen, R.C. Petersen, L.M. Shaw, J.Q. Trojanowski, C.R. Jack, M.W. Weiner, W.J. Jagust, Neurology 75 (3) (2010) 230–238 [Online]. Available: http://www.neurology.org/content/75/3/230.

[6] J. Ashburner, K.J. Friston, NeuroImage 11 (6 Pt. 1) (2000) 805–821 [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/10860804.

[7] T. Nichols, S. Hayasaka, Stat. Methods Med. Res. 12 (5) (Oct. 2003) 419–446.

[8] J.V. Haxby, M.I. Gobbini, M.L. Furey, A. Ishai, J.L. Schouten, P. Pietrini, Science 293 (5539) (2001) 2425–2430 [Online]. Available: http://www.sciencemag.org/content/293/5539/2425.

[9] C. Davatzikos, NeuroImage 23 (1) (2004) 17–20 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811904002861.

[10] J.J. Chen, H.D. Rosas, D.H. Salat, NeuroImage 55 (2) (2011) 468–478 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811910016162.

[11] D. Tosun, H. Rosen, B.L. Miller, M.W. Weiner, N. Schuff, NeuroImage 59 (3) (2012) 098–2109 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811911011943.

[12] R.J. Haier, S. Karama, L. Leyba, R.E. Jung, BMC Res. Notes 2 (1) (2009) 174 [Online]. Available: http://www.biomedcentral.com/1756-0500/2/174/abstract.

[13] I. Anurova, L.A. Renier, A.G.D. Volder, S. Carlson, J.P. Rauschecker, Cerebral Cortex (2014) bhu009 [Online]. Available: http://cercor.oxfordjournals.org/content/early/2014/02/10/cercor.bhu009.

[14] M.G. Naylor, V.A. Cardenas, D. Tosun, N. Schuff, M. Weiner, A. Schwartzman, Human Brain Mapp. (2013). pp. n/a–n/a. [Online]. Available: http://onlinelibrary.wiley.com/doi/10.1002/hbm.22217/abstract.

[15] V. Michel, A. Gramfort, G. Varoquaux, E. Eger, C. Keribin, B. Thirion, Pattern Recognit. 45 (6) (2012) 2041–2049 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0031320311001439.

[16] L. Grosenick, S. Greer, B. Knutson, IEEE Trans. Neural Syst. Rehabil. Eng. 16 (6) (2008) 539–548.

[17] L. Grosenick, B. Klingenberg, K. Katovich, B. Knutson, J.E. Taylor, NeuroImage 72 (2013) 304–321 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811912012487.

[18] L. Grosenick, B. Klingenberg, S. Greer, J. Taylor, B. Knutson, NeuroImage 47 (Suppl. 1) (2009) S58 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811909702320.

[19] B. Kandel, D. Wolk, J.C. Gee, B. Avants, Predicting Cognitive Data from Medical Images Using Sparse Linear Regression, Asilomar, NY, 2013.

[20] G. Varoquaux, A. Gramfort, B. Thirion, Small-sample Brain Mapping: Sparse Recovery on Spatially Correlated Designs with Randomization and Clustering, 2012, pp. 1375–1382. [Online]. Available: http://icml.cc/discuss/2012/688.html.

[21] M. Sabuncu, K. VanLeemput, IEEE Trans. Med. Imag. 31 (12) (2012) 2290–2306.

[22] N. Batmanghelich, B. Taskar, C. Davatzikos, Inf. Proc. Med. Imag. 21 (2009) 423–434.

[23] N. Batmanghelich, A. Dong, B. Taskar, C. Davatzikos, in: Medical Image Computing and Computer-assisted Intervention: MICCAI. International Conference on Medical Image Computing and Computer-Assisted Intervention, vol. 14, Pt. 3, pp. 17–24, 2011.

[24] N. Batmanghelich, B. Taskar, C. Davatzikos, IEEE Trans. Med. Imag. 31 (1) (2012) 51–69.

[25] H. Hotelling, Br. J. Stat. Psychol. 10 (2) (1957) 69–79 [Online]. Available: http://onlinelibrary.wiley.com/doi/10.1111/j.2044-8317.1957.tb00179.x/abstract.

[26] B. Avants, P. Dhillon, B.M. Kandel, P.A. Cook, C.T. McMillan, M. Grossman, J.C. Gee, in: Medical Image Computing and Computer-assisted Intervention: MICCAI. International Conference on Medical Image Computing and Computer-Assisted Intervention, vol. 15, Pt. 3, pp. 206–213, 2012.

[27] C.T. McMillan, B.B. Avants, P. Cook, L. Ungar, J.Q. Trojanowski, M. Grossman, Human Brain Map. (2014). pp. n/a–n/a. [Online]. Available: http://onlinelibrary.wiley.com/doi/10.1002/hbm.22515/abstract.

[28] Y. Xie, Dynamic Documents with R and knitr, CRC Press, 2013 [Online]. Available: http://books.google.com/books?hl=en&lr=&id=QZwAAAAAQBAJ&oi=fnd&pg=PP1&dq=Dynamic+Documents+with+R+and+knitr&ots=4VEuoaqpkh&sig=pmsSSsr-FmofufnoYKEthsdf6Pk.

[29] A. Hyvarinen, IEEE Trans. Neural Networks 10 (3) (May 1999) 626–634.

[30] F.Bießmann, S. Dhne, F.C. Meinecke, B. Blankertz, K. Grgen, K.-R. Müller, S. Haufe, in: 2nd NIPS Workshop on Machine Learning and Interpretation in NeuroImaging, 2012. [Online]. Available: http://www.user.tu-berlin.de/felix.biessmann/pub/2012-MLINIFiltersandPatterns.pdf.

[31] S. Haufe, F. Meinecke, K. Grgen, S. Dhne, J.-D. Haynes, B. Blankertz, F. Biemann, NeuroImage (2013) [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811913010914.

[32] H. Zou, T. Hastie, R. Tibshirani, J. Comput. Graph. Stat. 15 (2) (2006) 265–286.

[33] D.M. Witten, R. Tibshirani, T. Hastie, Biostatistics (Oxford, England) 10 (3) (2009) 515–534 [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/19377034.

[34] P.S. Dhillon, D.A. Wolk, S.R. Das, L.H. Ungar, J.C. Gee, B.B. Avants, NeuroImage 99 (2014) 14–27 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811914003917.

[35] H. Xu, C. Caramanis, S. Mannor, IEEE Trans. Pattern Anal. Mach. Intell. 34 (1) (2012) 187–193.

[36] B. Avants, J. Gee, in: Information Processing in Medical Imaging, in: C. Taylor, J.A. Noble, (Eds.), Lecture Notes in Computer Science, vol. 2732, Springer, Berlin Heidelberg, 2003, pp. 101–113. [Online]. Available: http://link.springer.com/chapter/10.1007/978-3-540-45087-09.

[37] D.L. Donoho, IEEE Trans. Inf. Theory 41 (3) (1995) 613–627.

[38] H. Shen, J. Huang, J. Multivar. Anal. 99 (6) (2008) 1015–1034.

[39] O. Veksler, Y. Boykov, P. Mehrani, in: Computer Vision ECCV 2010, in: K. Daniilidis, P. Maragos, N. Paragios, (Eds.), Lecture Notes in Computer Science, vol. 6315, Springer, Berlin Heidelberg, 2010, pp. 211–224. [Online]. Available: http://link.springer.com/chapter/10.1007/978-3-642-15555-016.

[40] G. Csardi, T. Nepusz, InterJournal 99 (1695) (2006) 1–9.

[41] D. Lin, D.P. Foster, L.H. Ungar, J. Am. Stat. Assoc. 106 (493) (2011) 232–247 [Online]. Available: http://amstat.tandfonline.com/doi/abs/10.1198/jasa.2011.tm10113.

[42] R. Berk, L. Brown, A. Buja, K. Zhang, L. Zhao, Ann. Stat. 41 (2) (2013) 802–837. zentralblatt MATH identifier1267.62080, Mathematical Reviews number (MathSciNet) MR3099122. [Online]. Available: http://projecteuclid.org/euclid.aos/1369836961.

[43] B. Avants, C. Epstein, M. Grossman, J. Gee, Med. Image Anal. 12 (1) (2008) 26–41.

[44] B.B. Avants, N.J. Tustison, G. Song, P.A. Cook, A. Klein, J.C. Gee, NeuroImage 54 (3) (2011) 2033–2044.

[45] V. Jain, J. Duda, B. Avants, M. Giannetta, S.X. Xie, T. Roberts, J.A. Detre, H. Hurt, F.W. Wehrli, D.J.J. Wang, Radiology 263 (2) (2012) 527–536.

[46] N. Tzourio-Mazoyer, B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, B. Mazoyer, M. Joliot, NeuroImage 15 (1) (2002) 273–289 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811901909784.

[47] S. Wakana, H. Jiang, L.M. Nagae-Poetscher, P.C.M. van Zijl, S. Mori, Radiology 230 (1) (2004) 77–87 [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/14645885.

[48] H. Wang, J.W. Suh, S.R. Das, J. Pluta, C. Craige, P.A. Yushkevich, IEEE Trans. Pattern Anal. Mach. Intell. (2012).

[49] N. Tustison, B. Avants, P. Cook, Y. Zheng, A. Egan, P. Yushkevich, J. Gee, IEEE Trans. Med. Imag. 29 (6) (Jun. 2010) 1310–1320.

[50] B. Avants, N. Tustison, J. Wu, P. Cook, J. Gee, Neuroinformatics 9 (4) (2011) 381–400.

[51] S.R. Das, B.B. Avants, M. Grossman, J.C. Gee, NeuroImage 45 (3) (2009) 867–879 [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/19150502.

[52] N.J. Tustison, B.B. Avants, P.A. Cook, G. Song, S. Das, N. van Strien, J.R. Stone, J.C. Gee, 8672 (2013) 86720K–86720K–4. [Online]. Available: http://dx.doi.org/10.1117/12.2007128.

[53] P.A. Cook, Y. Bai, S. Nedjati-Gilani, K.K. Seunarine, M.G. Hall, G.J. Parker, D.C. Alexander, [Online]. Available: http://cds.ismrm.org/ismrm-2006/files/02759.pdf.

[54] N.J. Tustison, B.B. Avants, P.A. Cook, J. Kim, J. Whyte, J.C. Gee, J.R. Stone, Human Brain Map. (2012) [Online]. Available: http://proxy.library.upenn.edu:2170/doi/10.1002/hbm.22211/abstract.

[55] C. Maumet, P. Maurel, J.-C. Ferr, C. Barillot, Magnet. Reson. Imag. (2014) [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0730725X14000204.

[56] B.B. Avants, S.K. Lakshmikanth, J.T. Duda, J.A. Detre, M. Grossman, in: Proceedings of Perfusion MRI: Standardization, Beyond CBF and Everyday Clinical Applications, International Society for Magnetic Resonance in Medicine Scientific Workshop, Amsterdam, 2012, p. 21.

[57] Y. Behzadi, K. Restom, J. Liau, T.T. Liu, NeuroImage 37 (1) (2007) 90–101 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811907003837.

[58] W.-C. Wu, V. Jain, C. Li, M. Giannetta, H. Hurt, F.W. Wehrli, D.J.J. Wang, Magnet. Reson. Med. 64 (4) (2010) 1140–1147 [Online]. Available: http://onlinelibrary.wiley.com/doi/10.1002/mrm.22484/abstract.

[59] B. Fischl, A.M. Dale, Proc. Natl. Acad. Sci. 97 (20) (2000) 11050–11055 [Online]. Available: http://www.pnas.org/content/97/20/11050.

[60] D.A. Roberts, J.A. Detre, L. Bolinger, E.K. Insko, J.S. Leigh Jr., Proc. Natl. Acad. Sci. USA 91 (1) (1994) 33–37.

[61] N. Gogtay, J.N. Giedd, L. Lusk, K.M. Hayashi, D. Greenstein, A.C. Vaituzis, T.F. Nugent, D.H. Herman, L.S. Clasen, A.W. Toga, J.L. Rapoport, P.M. Thompson, Proc. Natl. Acad. Sci. 101 (21) (2004) 8174–8179 [Online]. Available: http://www.pnas.org/content/101/21/8174.

[62] C. Lebel, L. Walker, A. Leemans, L. Phillips, C. Beaulieu, NeuroImage 40 (3) (2008) 1044–1055.

[63] S. Aslan, H. Lu, Magnet. Reson. Imag. 28 (7) (2010) 928–935 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0730725X10001001.

[64] R.J. Tibshirani, Ann. Stat. 39 (3) (2011) 1335–1371. zentralblatt MATH identifier: 05947535; Mathematical Reviews number (MathSciNet): MR2850205. [Online]. Available: http://projecteuclid.org/euclid.aos/1304514656.

[65] L. Breiman, Mach. Learn. 45 (1) (2001) 5–32 [Online]. Available: http://link.springer.com/article/10.1023/A%3A1010933404324.

[66] K. Franke, G. Ziegler, S. Klppel, C. Gaser, NeuroImage 50 (3) (2010) 883–892 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811910000108.

[67] K. Franke, E. Luders, A. May, M. Wilke, C. Gaser, NeuroImage 63 (3) (2012) 1305–1312 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S105381191200794X.

[68] T. Brown, J. Kuperman, Y. Chung, M. Erhart, C. McCabe, D. Hagler Jr., V. Venkatraman, N. Akshoomoff, D. Amaral, C. Bloss, B.J. Casey, L. Chang, T. Ernst, J. Frazier, J. Gruen, W. Kaufmann, T. Kenet, D. Kennedy, S. Murray, E. Sowell, T. Jernigan, A. Dale, Curr. Biol. 22 (18) (2012) 1693–1698 [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0960982212007932.

[69] A.M. Fjell, K.B. Walhovd, T.T. Brown, J.M. Kuperman, Y. Chung, D.J. Hagler Jr, V. Venkatraman, J.C. Roddey, M. Erhart, C. McCabe, N. Akshoomoff, D.G. Amaral, C.S. Bloss, O. Libiger, B.F. Darst, N.J. Schork, B.J. Casey, L. Chang, T.M. Ernst, J.R. Gruen, W.E. Kaufmann, T. Kenet, J. Frazier, S.S. Murray, E.R. Sowell, P. van Zijl, S. Mostofsky, T.L. Jernigan, A.M. Dale, Proc. Natl. Acad. Sci. USA 109 (48) (2012) 19620–19625.

[70] N.U.F. Dosenbach, B. Nardos, A.L. Cohen, D.A. Fair, J.D. Power, J.A. Church, S.M. Nelson, G.S. Wig, A.C. Vogel, C.N. Lessov-Schlaggar, K.A. Barnes, J.W. Dubis, E. Feczko, R.S. Coalson, J.R. Pruett, D.M. Barch, S.E. Petersen, B.L. Schlaggar, Science 329 (5997) (2010) 1358–1361 [Online]. Available: http://www.sciencemag.org/content/329/5997/1358.