34th Annual International Conference of the IEEE EMBS
San Diego, California USA, 28 August - 1 September, 2012

Brain Shape Regression Components

Jing Xie[1] and Owen Carmichael[2]

Abstract— **Identifying associations between the shape properties of brain regions, measured from magnetic resonance imaging (MRI), and numerical measures of neurodegenerative disease burden can clarify whether disease processes lead to distinctive spatial patterns of brain atrophy. However, prior methods for identifying such associations between shape and clinical variables either failed to summarize shape patterns into a concise set of summary measurements, or risked failing to discover such associations by extracting summary shape features blinded to the clinical variables. We present a method that overcomes these limitations by directly searching for a small set of linear shape features–shape regression components– that simultaneously account for a large amount of population shape variability and are highly correlated with a numerical clinical variable of interest. When applied to hippocampi of 299 Alzheimer's Disease Neuroimaging Initiative (ADNI) participants, the method identified correlations between hippocampal atrophy and markers of AD pathology and cogniton that were stronger than, and covered a more extended spatial region than, those identified by competing approaches.**

I. Introduction

Shape features of brain regions have the potential to be important magnetic resonance imaging (MRI)-based biomarkers of neurodegenerative disease processes because these processes often inflict a stereotypical spatial pattern of damage to the brain regions over time. For example, the hippocampus (HP), a brain region that is critical to a variety of cognitive functions including memory, experiences a characteristic spatial pattern of atrophy over the biological course of Alzheimer's Disease (AD). Computational shape features that quantify the degree to which brain regions such as the HP have undergone such a pattern of shape change [1][2] could be used to detect neurodegenerative diseases preclinically or to detect beneficial effects of disease modifying therapies; but discovering such atrophy patterns depends on solving the difficult problem of finding shape characteristics that are strongly associated with numerical clinical variables such as biochemical surrogate markers of the disease from cerebrospinal fluid (CSF) or blood. Identifying such shape characteristics amounts to solving a regression problem with shape features as predictors and clinical variables as outcomes or *vice versa.*

There are currently two general approaches for solving this shape regression problem. Pointwise linear regression [3][4] solves separate linear regression problems at a large set of points sampled from the brain region surface: at each point, a measure of local region thickness is regressed against the clinical variable. Intuitive 3D color maps display the strength of association between local thickness and the clinical variable across the region surface. However, correcting for the large number of statistical tests performed by this technique reduces its sensitivity, and its inability to aggregate pointwise associations over extended spatial neighborhoods makes it difficult to summarize its discovered shape patterns into a compact set of summary measurements. Linear subspace methods [5][6] first decompose the high dimensional vector of brain region surface point coordinates or local thicknesses into a set of linear components that each represent a prominent mode of shape deformation among the population of shapes. These linear shape deformation features are then correlated with the clinical variables *post hoc*. Because the shape features are determined without any reference to the clinical variables, there is no guarantee that a linear subspace method will identify any shape features that are strongly associated with the clinical variable even if such an associated shape feature is present in the data.

We address the drawbacks of existing linear subspace methods by directly searching for a small number of dominant linear shape patterns that account for as much of the shape variability across the population as possible, while simultaneously being as highly correlated with a clinical variable as possible. Applying the technique to a large set of HP surfaces of healthy elderly individuals and elders with mild cognitive impairment (MCI) and Alzheimer's disease (AD) suggests that the technique is superior to existing linear subspace methods and pointwise regression for sensitively identifying a concise set of shape patterns that are strongly correlated to relevant numerical variables such as age, CSF markers of AD pathology burden, and measures of cognitive function. In Section 2, we present the framework and implementation of our method. In Section 3, we show experimental results comparing the technique to principal component analysis (PCA) and pointwise regression on a large dataset.

II. Methods

A. Shape Components

We begin by reviewing the mathematical formulation for traditional linear subspace methods that identify shape patterns without reference to clinical variables. Let $\mathbf{v}_j \in \mathbb{R}^n$ denote a vector of brain region surface point coordinates, or the local region thicknesses at those points, for individual j in a population. The linear subspace methods provide $\mathbf{v}_j^k$, a kth-order approximation of $\mathbf{v}_j$:

*This work was supported by NIH grants AG010129, AG 030514, AG021028, and AG024904, and by a grant from the Dana Foundation

[1]Jing Xie is with the Graduate Group in Computer Science, University of California, Davis, Davis, CA, 95616 xie@ucdavis.edu

[2]Owen Carmichael is with the Department of Neurology, University of California, Davis, Davis, CA, 95616 ocarmichael@ucdavis.edu

$$\mathbf{v}_j \approx \mathbf{v}_j^k = \mathbf{e}_0 + \sum_{i=1}^{k} \alpha_{i,j}\mathbf{e}_i, \quad (1)$$

where $\mathbf{v}_j^k$ is the sum of the $\mathbf{e}_0$, the average shape over all $\mathbf{v}_j$ in the population, and a linear combination of k *shape components*, $\{\mathbf{e}_1, \cdots \mathbf{e}_k\}$, $\mathbf{e}_i \in \mathbb{R}^n$. Without loss of generality, we assume that the shape vectors $\mathbf{v}_j$ have been transformed such that $\mathbf{e}_0 = \mathbf{0}$. Equation (1) can be also represented in matrix form as:

$$V = BC, \quad (2)$$

where V is a matrix whose column vectors are $\mathbf{v}_j$, for $j = 1 \cdots m$, $B = (\mathbf{e}_1\ \mathbf{e}_2\ \cdots\ \mathbf{e}_k)$ is the basis matrix, and $C = \{\alpha_{i,j}\}$ is the coefficient matrix. In this paper, shape components are unit-norm basis vectors and orthogonal to each other, so $B^T B = I$ and $C = B^T V$. The coefficient vector for a shape $\mathbf{c}_j = [\alpha_{1,j}, \alpha_{2,j}, \ldots, \alpha_{k,j}]^T$ is a representation of $\mathbf{v}_j$ in terms of the shape components $\mathbf{e}_i$, which capture the salient shape characteristics of $\mathbf{v}_j$ in terms of how it is deformed away from the population mean. Differing methods for estimating the $\mathbf{e}_i$ have differing advantages in terms of conciseness (*i.e.,* how many $\mathbf{e}_i$ are required so that $\mathbf{v}_j^k$ accurately approximates $\mathbf{v}_j$) and interpretability (*i.e.,* how well $\mathbf{e}_i$ corresponds to an easy-to-explain aspect of brain anatomy) [7][6]. Typically, to associate region shape with clinical variables, we first estimate the $\mathbf{e}_i$ and coefficients $\alpha_{i,j}$ in a way that does not depend on the clinical variables. Then, we linearly regress the resulting coefficients $\alpha_{i,j}$ against the clinical variables.

B. Shape Regression Components

In this paper, rather than estimate $\mathbf{e}_i$ blinded to clinical variables, we estimate *shape regression components* that are designed to maximize associations with a clinical variable. Given m individuals and a vector $\mathbf{x} \in \mathbb{R}^m$ that contains one numerical clinical variable value per individual, we seek shape components $\mathbf{e}_i$ such that the *shape regression component coefficient* vector $\boldsymbol{\alpha}_i = [\alpha_{i,1}, \alpha_{i,2}, \cdots, \alpha_{i,m}]^T$ is strongly correlated with $\mathbf{x}$. Without loss of generality we assume that $\boldsymbol{\alpha}_i$ and $\mathbf{x}$ have been scaled such that they have unit norms. Let $\beta_i = \boldsymbol{\alpha}_i^T \mathbf{x}$ and $\boldsymbol{\beta} = C\mathbf{x} = [\beta_1, \beta_2, \cdots, \beta_k]^T$. Greater values of β_i^2 suggest stronger correlations between $\boldsymbol{\alpha}_i$ and $\mathbf{x}$; for this reason we refer to β_i^2 as the *regression power* of $\mathbf{e}_i$. Our goal is to find a concise set of $\mathbf{e}_i$ that have high regression power.

If the linear subspace spanned by B is fixed, it can be proven that $\|\boldsymbol{\beta}\|^2 = \sum_{i=1}^{k} \beta_i^2$ is a constant, which implies that the total amount of regression power over the $\mathbf{e}_i$ is constant. Our optimization seeks to compress this regression power into the smallest number of $\mathbf{e}_i$ possible; *i.e.,* we iteratively update $\mathbf{e}_i$ in order to set as many β_i to 0 as possible; because $\|\boldsymbol{\beta}\|^2$ is constant, this has the effect of creating a small number of large-magnitude β_i. To do so, we consider the set $[\frac{\beta_1^2}{\|\boldsymbol{\beta}\|^2}, \frac{\beta_2^2}{\|\boldsymbol{\beta}\|^2}, \cdots, \frac{\beta_k^2}{\|\boldsymbol{\beta}\|^2}]$ as a probability distribution that we want to sparsify. We update the $\mathbf{e}_i$ to minimize the entropy of this distribution, which effectively drives as many β_i to zero as possible while setting a small number of β_i to large values.

C. Implementation

We combine entropy minimization with a criterion that forces the $\mathbf{e}_i$ to account for the greatest amount of population shape variability possible. As in a prior approach [8], we initialize $\mathbf{e}_i$ to be the principal components of the $\mathbf{v}_j$, and we iteratively rotate all possible pairs of $(\mathbf{e}_i, \mathbf{e}_j)$ together in the plane they span by the same angle θ to minimize an energy function. The energy function is $E = (1-\lambda)E_1 + \lambda E_2$. E_1 is used to encourage the $\mathbf{e}_i$ to account for the greatest amount of population shape variability possible and E_2 encourages a low-entropy distribution of β_i as described above. The trade-off factor λ controls the relative contributions of these two competing terms. As described previously [8], E_1 is:

$$E_1 = \sum_{i=1}^{k} -\frac{\|\boldsymbol{\alpha}_i\|^2}{s_\alpha} log(\frac{\|\boldsymbol{\alpha}_i\|^2}{s_\alpha}),$$

where $s_\alpha = \sum_{i=1}^{k} \|\boldsymbol{\alpha}_i\|^2$. As discussed in section II-B, E_2 takes on lesser values when the entropy of the β_i values is minimized:

$$E_2 = \sum_{i=1}^{k} -\frac{\beta_i^2}{\|\boldsymbol{\beta}\|^2} log(\frac{\beta_i^2}{\|\boldsymbol{\beta}\|^2}).$$

III. Experiments

We identified baseline HP shape features associated with baseline clinical variables including age, CSF measures of AD pathology (amyloid, tau, and phosphorylated tau), and cognitive measures among 299 left HP from participants in the Alzheimer's Disease Neuroimaging Initiative (ADNI), 168 of whom had CSF measurements. MRI acquisition, HP delineation, and calculation of $\mathbf{v}_j$ representing local HP region thicknesses at corresponding surface points across subjects have been described previously [9]. We compared shape regression components with analogous pointwise linear regression models [3] and PCA by assessing the strengths of associations between the clinical variables and shape regression component coefficients, PCA coefficients, and local HP radial distances respectively. In all shape regression optimizations, we set λ to 0.9.

A. Shape regression components provide stronger associations with clinical variables

First, we ran shape regression component optimization with each optimization searching for shape regression components associated with one of the clinical variables. For each clinical variable we ran linear regression models that assessed the strengths of associations between the shape regression component coefficients and the clinical variable. We then ran PCA on the HP to provide shape components that accounted for the greatest amount of HP shape variability possible; we used linear regression to test the strengths of associations between each clinical variable and the PCA coefficients. Finally, we ran pointwise regression [3], *i.e.* we calculated a separate linear regression model at each HP surface point that associated the local HP radius there to each clinical variable.

Figure 1 illustrates the key differences between PCA and shape regression components in terms of associations with

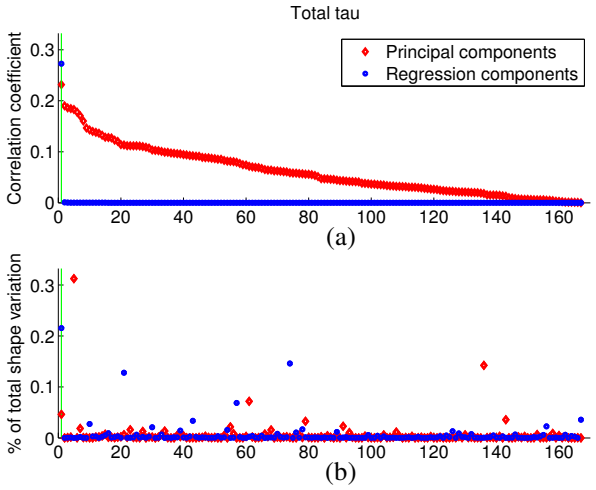

Fig. 1. (a) Plot of correlation coefficients between total tau and coefficients of principal components and shape regression components. Two sets of shape components were sorted by correlation coefficients in descending order respectively. (b) Plot of shape variability accounted for by the two sets of shape components in the same order as shown in (a). Green lines in both plots indicate the horizontal positions of the regression component and the principal component that have the strongest associations with total tau.

CSF total tau, a key biochemical marker of AD. Exactly one regression shape component has coefficients whose correlation with total tau is non-zero, while many principal component coefficients are correlated with total tau. In fact, the strength of association between the principal component coefficient and total tau decreases very gradually from component to component, making it unclear how many of these principal components should be analyzed in further detail by the end user. In addition, the one regression component whose coefficients were significantly correlated with total tau (green line) reflected a large amount of population variability in HP shape, while the principal component whose coefficients were maximally correlated with total tau reflected very little population shape variability. This suggests that while the principal component is associated with tau, the association may be irrelevant because it represents a relatively rare, under-represented shape feature in the population.

Figure 2 shows the p values of linear regressions between clinical variables and shape features, including principal component coefficients, shape regression component coefficients, and individual HP surface point radii. They are plotted on a logarithmic scale. For the shape regression components and principal components, the minimal p value over all such components is plotted. For pointwise regression, the average p value over all HP surface points with $p < .05$ is plotted. For all clinical variables, shape regression components had p values closer to zero, suggesting a superior ability to identify shape features that are the most strongly associated with clinical variables.

B. Shape regression components identify larger HP regions associated with clinical variables

The next experiment assessed whether there were HP regions whose significant associations with clinical variables were identified by shape regression components, but not

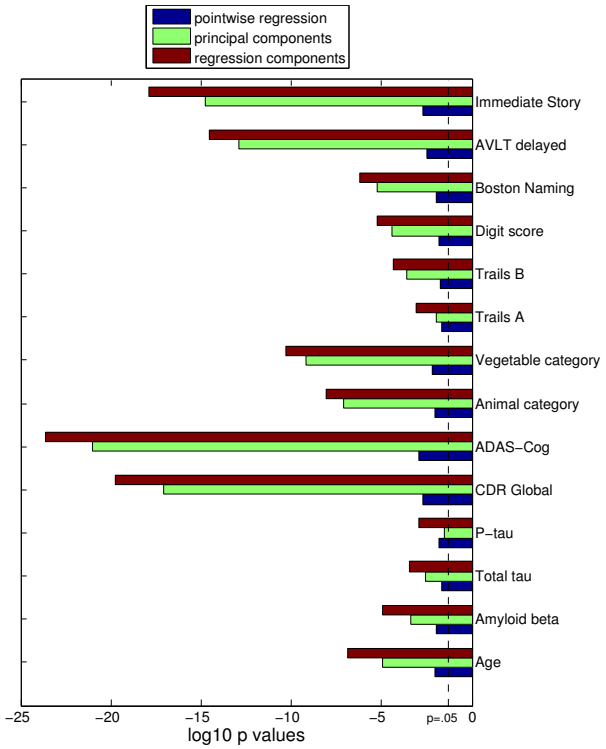

Fig. 2. P values of associations between clinical variables and shape features, including principal component coefficients, shape regression component coefficients, and HP surface point radii, plotted on a logarithmic scale. See Section III-A for details. The black dotted line represents the significance threshold 0.05.

identified by pointwise regression. This analysis was designed to assess how much value is added by shape regression components in terms of sensitivity to identify relevant HP regions that are larger than those pointwise regression provides. For each clinical variable we first identified HP surface points whose radii were associated with the clinical variable according to pointwise regression. These surface points were removed from the v_j vector, and PCA and shape regression component analyses were run on the vector of remaining points as in Section III-A. The significance of associations between clinical variables and the resulting principal components and shape regression components are shown in Figure 3. For each clinical variable, there was a shape regression component whose coefficient was significantly associated with the clinical variable, suggesting that shape regression components are sensitive enough to discover HP regions that pointwise regression have deemed irrelevant but are, in fact, associated with clinical variables. For many clinical variables, at least one principal component has coefficients associated with the clinical variable as well, but the strengths of these associations are generally weaker than those of the shape regression components.

C. Shape regression components are more specific to individual clinical variables

We then explored further whether the shape regression components identify unique spatial patterns of HP shape associated with unique clinical variables. To do so, we first recognized that increasing age is strongly associated with both HP shape variability and decreased cognitive function;

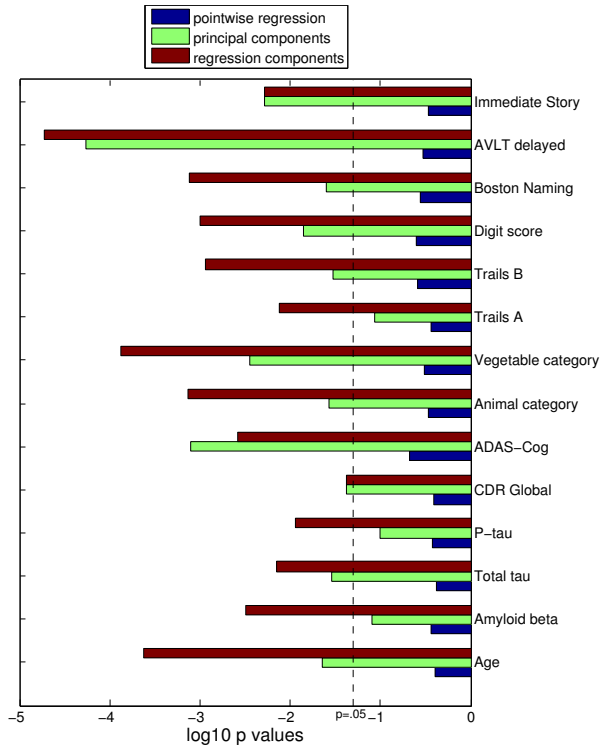

Fig. 3. P values of associations between clinical variables and shape features among HP regions that were deemed non-significantly associated by pointwise regression. See Section III-A for details. P values are plotted on a logarithmic scale. For pointwise regression, p values were the average over all non-significantly associated points, *i.e.* points with $p >= .05$. The black line represents the significance threshold 0.05.

this could lead to principal components that accidentally associate strongly with both age and cognitive function because they capture the large amount of HP shape variability that is associated with increasing age. Therefore, for each cognitive measure we generated a *residual cognitive measure* by performing linear regression with age as the sole predictor and the cognitive measure as the outcome, and using the signed residual of this regression as a measure of cognitive function that is dissociated from age. We then ran shape regression components analysis with each such residual cognitive measure, and correlated both the resulting coefficients and the principal component coefficients with the residual cognitive measures. A representative result is shown in Figure 4. The coefficients of one principal component were significantly associated with both age and a residual cognitive measure. In contrast, for age and the residual cognitive measure, there were highly distinct shape regression components whose coefficients were more strongly associated than the principal component coefficients were. Because the two significant shape regression components substantially differ, we conclude that individual clinical variables may be associated with individual, unique patterns of HP shape, and shape regression components may be better able to identify such unique spatial patterns compared to PCA.

IV. Discussion

In this paper, we showed that directly optimizing for linear shape patterns that both account for population shape variability and associate strongly with a clinical variable

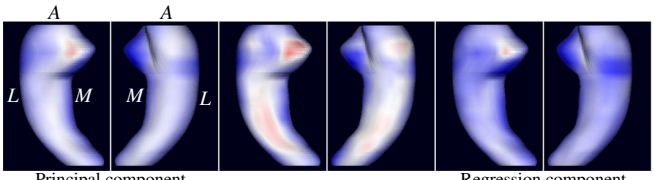

Principal component associated with age and immediate story

Regression component associated with age

Regression component associated with immediate story

Fig. 4. Color-coded maps of the principal component and regression components associated with age and residual immediate story recall. See Section III-C for details. Red points indicate $\mathbf{e}_i$ entries greater than 0, and blue indicates less than 0. Odd and even columns render the HP from superior and inferior viewpoints respectively (anterior, medial and lateral directions are marked with *A*, *M* and *L*).

can identify variable-specific patterns of HP shape better than PCA and pointwise linear regression can. Future work should explore whether we are able to further refine these shape regression components by encouraging them to take on additional desirable properties such as spatial locality or discrimination of coefficient values between clinically defined groups. In addition, as suggested by our analysis of residualized variables, our method could be extended in a factor analytic framework to directly ascertain whether there exist aspects of shape that are jointly associated with several clinical variables, as opposed to uniquely associated with a single, individual clinical variable. Finally, future work could more widely apply the method to other brain regions and clinical variables beyond those relevant to AD.

References

[1] P. Thompson, K. Hayashi, G. de Zubicaray, A. Janke, S. Rose, J. Semple, M. Hong, D. Herman, D. Gravano, D. Doddrell, and A. Toga, "Mapping hippocampal and ventricular change in Alzheimer's disease," *NeuroImage*, vol. 22, no. 4, pp. 1754–1766, August 2004.

[2] L. Wang, S. Joshi, M. Miller, and J. Csernansky, "Statistical analysis of hippocampal asymmetry in schizophrenia," *NeuroImage*, vol. 14, no. 3, pp. 531–545, September 2001.

[3] L. Apostolova, K. Hwang, J. Andrawis, A. Green, S. Babakchanian, J. Morra, J. Cummings, A. Toga, J. Trojanowski, L. Shaw, C. Jack, R. Petersen, P. Aisen, W. Jagust, R. Koeppe, C. Mathis, M. Weiner, P. Thompson, and the Alzheimers Disease Neuroimaging Initiative, "3d pib and csf biomarker associations with hippocampal atrophy in adni subjects," *Neurobiol. Aging*, vol. 31, no. 8, pp. 1284–1303, August 2010.

[4] J. H. Morra, Z. Tu, L. G. Apostolova, A. E. Green, C. Avedissian, S. K. Madsen, N. Parikshak, X. Hua, A. W. Toga, C. R. Jack, N. Schuff, M. W. Weiner, and P. M. T. andthe Alzheimers Disease Neuroimaging Initiative, "Automated 3d mapping of hippocampal atrophy and its clinical correlates in 400 subjects with alzheimers disease, mild cognitive impairment, and elderly controls," *Human Brain Mapping*, vol. 30, no. 9, pp. 2766–2788, September 2009.

[5] L. Wang, J. S. Swank, I. E. Glick, M. H. Gado, M. I. Miller, J. C. Morris, and J. G. Csernansky, "Changes in hippocampal volume and shape across time distinguish dementia of the alzheimer type from healthy aging," *NeuroImage*, vol. 20, no. 2, pp. 667–682, October 2003.

[6] J. Xie, D. Alcantara, N. Amenta, E. Fletcher, O. Martinez, M. Persianinova, C. DeCarli, and O. Carmichael, "Spatially localized hippocampal shape analysis in late-life cognitive decline," *Hippocampus*, vol. 19, no. 6, pp. 526–532, June 2009.

[7] I. Jolliffe, *Principal Component Analysis*. Springer, 2002.

[8] D. Alcantara, O. Carmichael, E. Delson, W. Harcourt-Smith, K. Sterner, S. Frost, R. Dutton, P. Thompson, H. Aizenstein, O. Lopez, J. Becker, and N. Amenta, "Localized components analysis," in *Proc. IPMI*, 2007.

[9] O. Carmichael, J. Xie, E. Fletcher, B. Singh, C. DeCarli, and the Alzheimers Disease Neuroimaging Initiative, "Localized hippocampus measures are associated with Alzheimer pathology and cognition independent of total hipocampal volume," *Neurobiol. Aging*, 2011.