

REGULARIZED 3D FUNCTIONAL REGRESSION FOR BRAIN IMAGE DATA VIA HAAR WAVELETS¹

BY XUEJING WANG², BIN NAN^{2,3}, JI ZHU^{2,4}, ROBERT KOEPPE²
AND FOR THE ALZHEIMER'S DISEASE NEUROIMAGING INITIATIVE

University of Michigan

The primary motivation and application in this article come from brain imaging studies on cognitive impairment in elderly subjects with brain disorders. We propose a regularized Haar wavelet-based approach for the analysis of three-dimensional brain image data in the framework of functional data analysis, which automatically takes into account the spatial information among neighboring voxels. We conduct extensive simulation studies to evaluate the prediction performance of the proposed approach and its ability to identify related regions to the outcome of interest, with the underlying assumption that only few relatively small subregions are truly predictive of the outcome of interest. We then apply the proposed approach to searching for brain subregions that are associated with cognition using PET images of patients with Alzheimer's disease, patients with mild cognitive impairment and normal controls.

1. Introduction. Alzheimer's disease (AD) is the most frequent cause of dementia in our increasingly aging societies, representing a significant impact on the US population with 10% prevalence in individuals aged above 70 years old [Plassman et al. (2007)]. Despite the prevalence, this disease remains quite a mystery; there is neither a cure nor a definite treatment to arrest its course and, currently, the only definite way to diagnose AD is to examine the brain tissue after death. According to recent studies [Leifer (2003)], early diagnosis of AD is of great value since new drug therapies can be used to potentially delay the progression of the disease. To this end, much progress has been made in assisting the early diagnosis of AD with neuroimaging techniques. One such widely used neuroimaging technique is positron emission tomography (PET) imaging, which is one of the most promising tools for the early diagnosis of AD, and it is of great scientific interest in understanding the association between PET images and cognitive impairment. In particular, the fluorodeoxyglucose (FDG) PET has been used to measure the cerebral glucose metabolic activity for over 20 years.

Received July 2012; revised February 2014.

¹Supported by NIH Grants P30 AG010129 and K01 AG030514.

²Supported in part by NIH Grant R01-AG036802.

³Supported in part by NSF Grant DMS-1007590.

⁴Supported in part by NSF Grant DMS-0748389.

Key words and phrases. Alzheimer's disease, brain imaging, functional data analysis, Haar wavelet, Lasso, PET image, variable selection.

FDG PET scans used in the preparation of this article were obtained from a large multi-center follow-up study on Alzheimer's disease and early dementia, the Alzheimer's Disease Neuroimaging Initiative (ADNI). A total of 403 FDG PET scans were acquired for this application, including 102 normal control (NC) subjects, 206 subjects with mild cognitive impairment (MCI) and 95 subjects diagnosed with AD. In this study, we consider the baseline FDG PET scans with a standard $160 \times 160 \times 96$ voxel image grid as the predictor to the cognitive performance as measured by the mini-mental state exam (MMSE), which is a questionnaire test that is used to screen for cognitive impairment [Cockrell and Folstein (1988)]. The maximum MMSE score is 30 and, on average, MMSE scores decline as the disease progresses. The goal of our study is to identify brain subregions that are most closely related to the prediction of MMSE scores.

Many methods have been developed for the analysis of brain image data in order to identify disease-related brain subregions. Most of these methods focus on region of interest (ROI) and voxel-based univariate analysis; see, for example, Luo and Nichols (2003), Grimmer et al. (2009) and Shin et al. (2010), among many others. For AD in particular, several studies have shown that reduced metabolic activity in some regions of the brain, such as the posterior cingulate and the temporal and parietal cortices, are associated with the progression of cognitive impairment [Foster et al. (1984), Minoshima et al. (1995, 1997)]. These methods are intended to provide statistics by doing a separate analysis for each ROI or voxel and then to draw inferences at the region- or voxel-level. As a result of testing millions of hypotheses, appropriate adjustments for multiple comparisons have to be considered. In the neuroimaging literature, a distinction is often drawn between such univariate analyses and an alternative, multiple covariate regression models that treat every voxel as a covariate. Since the number of voxels is much larger than the number of scans, the ordinary least squares for linear regression cannot be implemented without applying, for example, some dimension reduction techniques. Such analysis, however, may lead to difficulties in interpretations and practical implications. Both the traditional univariate and multiple covariate approaches (if applicable) have one major limitation in common: they are developed without considering the spatial information of the brain, possibly resulting in some loss of information. There is an emerging awareness of the importance of taking such information into account. For example, multiple covariate analysis can be conducted with a focus on extracting principal components from the images [Friston et al. (1996), Kerrouche et al. (2006)]. More recently, a variety of Bayesian spatial modeling approaches have been proposed to model the correlation between neighboring voxels, which need to carefully specify the prior distributions; see, for example, Bowman et al. (2008), Kang et al. (2011). Our way of addressing this issue is to treat the entire 3D image as a single functional input, which allows retaining all the information from the original image without modeling the spatial correlation between voxels explicitly. Specifically, in this article, we treat PET image data as the 3D functional observations, and propose a novel Haar wavelet-based regularized approach to analyze PET image data in the framework of functional data analysis.

Functional regression models are known as one of the standard techniques in functional data analysis. It is noted that the models can be defined as functional in one or both of two ways: the response variable is functional; at least one of the covariates is functional. In this article, we focus on the functional linear regression model with a scalar response variable and a single functional predictor. Using the 1D case as an illustration, the functional linear regression model relates a scalar response variable Y to a functional predictor $X(t)$ as follows:

$$(1.1) \quad Y_i = \beta_0 + \int_0^T X_i(t)\beta(t) dt + \varepsilon_i, \quad i = 1, \dots, n,$$

where $\beta(t)$ is the regression coefficient function and t refers to time or location. For the 3D case we consider later, t is replaced by the coordinate (u, v, w) . Regularization methods, such as the roughness penalty approach or using restricted basis functions [Ramsay and Silverman (2005)], can be implemented to produce an estimator of $\beta(t)$ that is meaningful in interpretation and useful in prediction.

For the functional linear regression model (1.1), James, Wang and Zhu (2009) proposed a regularized approach that focuses on producing sparse and highly interpretable estimates of the coefficient function $\beta(t)$. This approach involves first dividing the domain into a fine grid of points, and then using appropriate variable selection methods to determine whether the d th derivative of $\beta(t)$ is zero or not at each of the grid points, that is, $\beta^{(d)}(t) = 0$ for one or more values of $d \in \{0, 1, 2, \dots\}$. They proposed the Dantzig selector [Candes and Tao (2007)] and a Lasso-type approach for the estimation of $\beta(t)$ using piecewise constant basis, where the Dantzig selector seems to be more natural. Empirical results show that their methods perform well when p , the number of basis functions, is not too large. When functional data are measured over a very fine grid such as brain image data, the Dantzig selector faces the challenge of solving a huge linear programming problem and the Lasso-type algorithm can be extremely slow; note that for the latter the fast coordinate descent algorithm [Fu (1998), Daubechies, Defrise and De Mol (2004), Friedman et al. (2007), Wu and Lange (2008)] does not apply due to the penalty on derivatives. Without imposing sparsity, Reiss and Ogden (2010) considered the functional principal component regression for image data.

In this article, we choose the Haar wavelet basis instead of the piecewise constant basis for analyzing 3D image data and show that the Haar wavelet-based approach presents a number of advantages. First, it yields regional sparseness without imposing constraints on derivatives, which is needed in James, Wang and Zhu (2009). In other words, by shrinking corresponding wavelet coefficients to zero, the estimator of the regression coefficient function can be exactly zero over regions where no relationship to the response variable is present. Second, the Haar wavelet transform offers a way to overcome the issue of high multicollinearity caused by high neighboring spatial correlations. Third, our approach is flexible enough to allow the coefficient function to be estimated at different levels of smoothness through choosing different levels of the Haar wavelet decomposition. Fourth, the

Haar wavelet transform can be applied as a dimension reduction technique prior to model fitting for high-dimensional image data by setting a common set of close to zero wavelet coefficients of PET images to zero, which is an effective way of removing voxels outside the brain or in the ventricles. It should be noted that a recent article by Zhao, Todd Ogden and Reiss (2012) considered a general wavelet-based Lasso approach in functional linear regression, but only concerned 1D $\beta(t)$. The Haar wavelet transform is a useful tool for image and signal analysis and has many other applications. For example, Picart, Butenschön and Shutler (2012) and Lovejoy and Schertzer (2012) discussed the use of Haar wavelet transforms in geophysics and climate research.

The rest of this article is organized as follows. In Section 2 we review some background on wavelet decomposition and properties of Haar wavelet basis functions using a 1D functional linear regression model as an illustration and then propose the ℓ_1 regularized shrinkage estimation for general functional data, including both 1D and 3D cases. To evaluate the numerical performance of our approach, we conduct extensive simulations in Section 3. We present the analysis of ADNI 3D PET image data in Section 4 and make some concluding remarks in Section 5. We also show that the proposed method achieves the desirable nonasymptotic error bounds for prediction and estimation, the so-called oracle inequalities, meaning that the method performs equally well (up to a constant) as if the true subregions with nonzero regression coefficient were given. The theoretical results are provided in the online supplementary material [Wang et al. (2014)].

2. Regularized Haar wavelets method. For ease of presentation, we describe the proposed methodology starting with the 1D case given in (1.1), then extend it to the 3D case using a tensor product of three 1D wavelet expansions.

2.1. Choice of basis. Basis expansions are commonly used in analyzing functional data. Among a variety of choices of basis expansions, wavelets have the important ability to allow simultaneous time (or space in this article) and frequency localization. Unlike many other commonly used basis systems, wavelet transforms are highly adaptable to different levels of smoothness and more capable of capturing edges, spikes and other types of discontinuities, especially for wavelet transforms with relatively small support such as the Haar wavelets. Wavelet transforms also provide a powerful tool to compress the data. A compressed approximation of the signal can be achieved by penalizing the wavelet coefficients [Wand and Ormerod (2011)], which involves shrinking small coefficients to zero and possibly shrinking the large ones without affecting the main features of the data. Hence, it is advantageous to use wavelet transforms to decompose images as well as the regression coefficient function for estimation.

In many applications, it is often the case that the association between $X(t)$ and Y in model (1.1) is sparse and potentially discontinuous at the boundaries of subregions. In particular, only few brain subregions in the aforementioned PET images

are believed to be related to cognitive impairment. To better identify such patterns, we choose to use Haar wavelets. The Haar wavelet transform is easily calculated and affected less by discontinuities. In addition, sparsity of $\beta(t)$ can be recovered by shrinking its wavelet coefficients to zero. The scaling function (also called a father wavelet) ϕ and the mother wavelet ψ of Haar wavelets defined on $[0, 1)$ are given below:

$$\phi(t) = \begin{cases} 1, & \text{if } 0 \leq t < 1; \\ 0, & \text{otherwise;} \end{cases}$$

$$\psi(t) = \begin{cases} 1, & \text{if } 0 \leq t < 1/2; \\ -1, & \text{if } 1/2 \leq t < 1; \\ 0, & \text{otherwise.} \end{cases}$$

The Haar wavelet bases are then generated in the form of translations and dilations of the above father and mother wavelet functions as

$$\phi_{j,k}(t) = \sqrt{2^j} \phi(2^j t - k),$$

$$\psi_{j,k}(t) = \sqrt{2^j} \psi(2^j t - k),$$

where $j = 0, 1, \dots$ and $k = 0, 1, \dots, 2^j - 1$. The index j refers to dilations and k refers to translations and $\sqrt{2^j}$ is the normalizing factor. It is noted that the basis functions are orthonormal. Therefore, for a sufficiently fine resolution J , the coefficient function $\beta(t)$ in (1.1) defined on $[0, 1)$ can be expanded in a Haar wavelet series:

$$(2.1) \quad \beta(t) = \sum_{k=0}^{2^{j_0}-1} a_{j_0,k} \phi_{j_0,k}(t) + \sum_{j=j_0}^J \sum_{k=0}^{2^j-1} d_{j,k} \psi_{j,k}(t) + e(t),$$

where $a_{j_0,k} = \int_0^1 \beta(t) \phi_{j_0,k}(t) dt$ are the approximation coefficients at the coarsest resolution j_0 , $d_{j,k} = \int_0^1 \beta(t) \psi_{j,k}(t) dt$ are the detail coefficients that characterize the finer structures of $\beta(t)$ as j grows, and $e(t)$ is the approximation error that goes to zero as J goes to infinity. The Haar wavelet representation of a signal thus consists of approximations together with details that can provide the desirable frequencies. See, for example, Walker (2008) for more details about Haar wavelets.

2.2. Model estimation. Rewrite $\beta(t)$ in (2.1) by

$$(2.2) \quad \beta(t) = B(t)^T \eta + e(t),$$

where $B(t)$ denotes the collection of all $\phi_{j,k}(t)$ and $\psi_{j,k}(t)$ in the above Haar wavelet expansion, and η is the corresponding wavelet coefficient vector of length p . Plugging (2.2) into (1.1), we obtain

$$(2.3) \quad Y_i = \beta_0 + \int_0^1 X_i(t) B(t)^T \eta dt + \varepsilon_i^* = \beta_0 + C_i^T \eta + \varepsilon_i^*, \quad i = 1, \dots, n,$$

where $C_i = \int_0^1 X_i(t)B(t) dt$ and $\varepsilon_i^* = \varepsilon_i + \int_0^1 X_i(t)e(t) dt$. It should be noted that C_i is the wavelet coefficient vector of $X_i(t)$ when we decompose $X_i(t)$ using the same set of Haar wavelet basis functions as those in (2.2). Model (2.3) can then be rewritten as follows:

$$(2.4) \quad Y = \beta_0 + C\eta + \varepsilon^*,$$

where $C = [C_1, C_2, \dots, C_n]^T$ is an $n \times p$ design matrix in linear model (2.4). Once an estimator $\hat{\eta}$ is obtained from (2.4), $\beta(t)$ can then be estimated by $B(t)^T \hat{\eta}$.

In practice, $X(t)$ is observed on only a finite set of grid points $\{t_1, \dots, t_p\}$, which also determines the highest and yet practically meaningful level of decomposition for $\beta(t)$. For the discrete wavelet transform, p is required to be a power of 2. Using the usual terminology for Haar wavelets [see, e.g., Walker (2008) and that used in the MATLAB Wavelet Toolbox (2011b)], we define the level 1 Haar wavelet decomposition by computing the average and the difference on each consecutive pair of values, and the maximum level is $\log_2 p$. The level number is directly determined by the integer j_0 in (2.1). For any level of Haar wavelet decomposition, the total number of basis functions $\phi_{j,k}$ and $\psi_{j,k}$ is always p , and the collection of $\phi_{j,k}$ and $\psi_{j,k}$ then forms a set of p -dimensional orthonormal basis functions.

A key advantage of using Haar wavelets is as follows. When $\beta(t) = 0$ in large regions of $t \in [0, 1)$ (in the ADNI brain image analysis where t is 3D, this would correspond to that large regions in the brain are not associated with the cognitive performance measured by MMSE), the coefficient vector η in (2.2) should be sparse with $e(t) = 0$ for those regions, that is, $\beta(t)$ can be well approximated by an economical wavelet expansion with few nonzero coefficients. We consider the Lasso approach [Tibshirani (1996)] and implement the method with the fast coordinate descent algorithm to obtain a desirable sparse solution for the wavelet coefficients.

For a given j_0 , which corresponds to a specific level of Haar wavelet expansion, the Lasso estimator for η is given by

$$(2.5) \quad \hat{\eta} = \arg \min_{\eta} \left\{ \frac{1}{n} \|Y - \beta_0 - C\eta\|_2^2 + 2\lambda \|\eta\|_1 \right\},$$

where $\|\cdot\|_1$ and $\|\cdot\|_2$ denote the ℓ_1 and ℓ_2 norms, respectively, and $\lambda \geq 0$ is a tuning parameter. In our estimating procedure, j_0 is also a tuning parameter.

It should be noted that in general the Haar wavelet coefficients with large magnitudes are related to salient features. The magnitudes of detail coefficients should be proportional to the differences between every pair of values, that is, larger magnitudes indicate sharper changes at corresponding locations and zero magnitudes indicate no change. If both detail and approximation coefficients of the Haar wavelet transform are close to zero, then $\beta(t)$ is close to zero. Thus, we are able to obtain a sparse solution of $\beta(t)$ by shrinking its small wavelet coefficients to zero.

2.3. *Selection of tuning parameters.* In addition to the Lasso tuning parameter λ in (2.5), we also need to take into account the level of the Haar wavelet decomposition. There should exist an optimal level of decomposition for $\beta(t)$ in terms of certain criterion, such as AIC, BIC or cross-validation. If the length of observed $X_i(t)$ is p , then the maximum possible level of the discrete Haar wavelet transform is $\log_2 p$, which is relatively small. Moreover, lower levels are usually considered in real applications. Therefore, including two tuning parameters does not increase computational burden by much.

2.4. *3D case.* The ADNI’s FDG PET brain images are 3D. A 3D function can be decomposed using a tensor product of three 1D Haar wavelets. In particular, the 3D Haar wavelet transform can be considered as averaging and differencing operations [Muraki (1992)]. The averaging operation is constructed by the 3D scaling function below:

$$\phi_{j,\{k,l,m\}}(u, v, w) = \phi_{j,k}(u)\phi_{j,l}(v)\phi_{j,m}(w).$$

The differencing operation is taken in seven directions constructed by the 3D wavelet functions as follows:

$$\psi_{j,\{k,l,m\}}^1(u, v, w) = \phi_{j,k}(u)\phi_{j,l}(v)\psi_{j,m}(w),$$

$$\psi_{j,\{k,l,m\}}^2(u, v, w) = \phi_{j,k}(u)\psi_{j,l}(v)\phi_{j,m}(w),$$

$$\psi_{j,\{k,l,m\}}^3(u, v, w) = \phi_{j,k}(u)\psi_{j,l}(v)\psi_{j,m}(w),$$

$$\psi_{j,\{k,l,m\}}^4(u, v, w) = \psi_{j,k}(u)\phi_{j,l}(v)\phi_{j,m}(w),$$

$$\psi_{j,\{k,l,m\}}^5(u, v, w) = \psi_{j,k}(u)\phi_{j,l}(v)\psi_{j,m}(w),$$

$$\psi_{j,\{k,l,m\}}^6(u, v, w) = \psi_{j,k}(u)\psi_{j,l}(v)\phi_{j,m}(w),$$

$$\psi_{j,\{k,l,m\}}^7(u, v, w) = \psi_{j,k}(u)\psi_{j,l}(v)\psi_{j,m}(w).$$

Let the image $X_i(u, v, w)$ be a 3D functional predictor and Y_i be a scalar response variable (MMSE, e.g.) for subject $i, i = 1, \dots, n$. The 3D functional linear regression model can be written as

$$(2.6) \quad Y_i = \beta_0 + \int_0^{T_1} \int_0^{T_2} \int_0^{T_3} X_i(u, v, w)\beta(u, v, w) du dv dw + \varepsilon_i.$$

For a sufficiently fine resolution J , the 3D coefficient function $\beta(u, v, w)$ can be approximated by

$$(2.7) \quad \sum_{k,l,m=0}^{2^{j_0}-1} a_{j_0,\{k,l,m\}}\phi_{j_0,\{k,l,m\}}(u, v, w) + \sum_{j=j_0}^J \sum_{k,l,m=0}^{2^j-1} \sum_{q=1}^7 d_{j,\{k,l,m\}}^q \psi_{j,\{k,l,m\}}^q(u, v, w).$$

Denote the set of all basis functions $\phi_{j,\{k,l,m\}}$ and $\psi_{j,\{k,l,m\}}^q$ in (2.7) by $B(u, v, w)$ and the wavelet coefficients in (2.7) by η , then $\beta(u, v, w)$ can be written as

$$(2.8) \quad \beta(u, v, w) = B(u, v, w)^T \eta + e(u, v, w).$$

Plugging (2.8) into model (2.6), we obtain

$$(2.9) \quad \begin{aligned} Y_i &= \beta_0 + \int_0^{T_1} \int_0^{T_2} \int_0^{T_3} X_i(u, v, w) B(u, v, w)^T \eta \, du \, dv \, dw + \varepsilon_i^* \\ &= \beta_0 + C_i^T \eta + \varepsilon_i^*, \end{aligned}$$

where $C_i = \int_0^{T_1} \int_0^{T_2} \int_0^{T_3} X_i(u, v, w) B(u, v, w) \, du \, dv \, dw$, which is equivalent to the wavelet coefficient vector when we apply the 3D wavelet transform to $X_i(u, v, w)$, and $\varepsilon_i^* = \varepsilon_i + \int_0^{T_1} \int_0^{T_2} \int_0^{T_3} X_i(u, v, w) e(u, v, w) \, du \, dv \, dw$. Then the methodology proposed in the previous subsections for the 1D case can be applied directly.

Following the calculations of [Bickel, Ritov and Tsybakov \(2009\)](#), we can show that our proposed method also enjoys the nonasymptotic oracle inequalities similar to the linear model with high-dimensional covariates. All the theoretical results are provided in the online supplementary material [[Wang et al. \(2014\)](#)]. It should be noted that the results are derived using the 1D notation for the estimator $\hat{\beta}(t)$ for simplicity, but they hold exactly for the 3D case of $\hat{\beta}(u, v, w)$.

3. Simulation studies. To investigate the performance of the proposed Haar wavelet-based approach, we have conducted extensive simulations for both 1D and 3D functional data. The results for 1D cases can be easily visualized, whereas the 3D case mimics the brain images more naturally.

3.1. *1D simulation.* We consider a variety of settings of $X(t)$ and $\beta(t)$. For $X(t) = X^*(t) + \mathcal{E}(t)$ defined on $0 \leq t \leq 1$, where $\mathcal{E}(t) \sim N(0, \sigma_{\mathcal{E}}^2)$ is the noise term independent of time t , we consider the following two scenarios:

- Fourier: $X^*(t) = a_0 + a_1 \sin(2\pi t) + a_2 \cos(2\pi t) + a_3 \sin(4\pi t) + a_4 \cos(4\pi t)$.
- *B*-splines: $X^*(t)$ is a linear combination of cubic *B*-splines with interior knots at $1/7, \dots, 6/7$ and coefficients a_i , that is, $X^*(t) = \sum a_i \phi_i(t)$, where $\phi_i(t)$ are the *B*-spline basis functions.

In both scenarios, the coefficients $a_i \sim N(0, 1)$. To assess the performance of the proposed approach in identifying continuous and discontinuous signals, we consider two cases of the regression coefficient function $\beta(t)$:

- *Case 1:* $\beta(t)$ is a smooth function,

$$\beta(t) = \begin{cases} 0.5(\sin(20t - \pi) + 1), & \text{if } \pi/8 \leq t < 9\pi/40, \\ 0, & \text{otherwise.} \end{cases}$$

– *Case 2:* $\beta(t)$ is piecewise constant,

$$\beta(t) = \begin{cases} 1, & \text{if } 0.2 \leq t < 0.3, \\ 0.5, & \text{if } 0.5 \leq t < 0.7, \\ 0, & \text{otherwise.} \end{cases}$$

For each curve $X^*(t)$, we record $p = 128$ equally spaced measurements for convenience. The variance of the noise term $\mathcal{E}(t)$ is set to be $\sigma_{\mathcal{E}}^2 = \frac{1}{p-1} \sum_{j=1}^p (X^*(t_j) - \bar{X}^*(t_j))^2$, where $\bar{X}^*(t_j)$ is the mean of $X^*(t_j)$. The error term ε in model (1.1) also follows a normal distribution $N(0, \sigma^2)$. The value of σ^2 is determined by the signal-to-noise ratio

$$(3.1) \quad \text{SNR} = \frac{\sigma_g^2}{\sigma^2},$$

where σ_g^2 is the sample variance of $g(X_i) = \int X_i(t)\beta(t) dt$. The simulation results presented in this article are under $\text{SNR} = 9$, which is also considered in Example 4 of Tibshirani (1996). For the Lasso method, Zou (2006) observed that smaller SNR usually yields smaller relative prediction error. For each of the settings, we use $n = 100$ training observations to fit the model. The optimal tuning parameter is selected by using one of the following methods: (i) validating by a separate validation (SV) data set of the same size; (ii) 5-fold cross-validation (CV); (iii) AIC and (iv) BIC [Zou, Hastie and Tibshirani (2007)] given below:

$$(3.2) \quad \text{AIC} = \frac{\|Y - \hat{g}(X)\|^2}{n\hat{\sigma}^2} + \frac{2}{n} \hat{d}f,$$

$$(3.3) \quad \text{BIC} = \frac{\|Y - \hat{g}(X)\|^2}{n\hat{\sigma}^2} + \frac{\log(n)}{n} \hat{d}f,$$

where $\hat{d}f$ is the number of nonzero elements of $\hat{\eta}$ in model (2.4). We estimate σ^2 by the refitted cross-validation method introduced in Fan, Guo and Hao (2012). We then generate $n = 10,000$ test observations to calculate the mean squared errors (MSEs) of the corresponding selected models. The procedure is repeated 100 times and the average MSEs and their standard errors (SE) for each of the models are presented in Table 1. We also report the percentages of correctly identified zero regions and nonzero regions in Table 1. We can see that all four methods perform reasonably well, while the nonpractical SV method performs the best. The CV method seems to have a nice trade-off between sparsity and prediction accuracy. Averages of the estimates of $\beta(t)$ using the CV method over 100 replications are shown in Figure 1.

We also conduct permutation tests to assess the significance of the regularized estimates of $\beta(t)$. For each of the training data sets, we generate 200 permutation data sets by randomly shuffling the response values. Using the same model selection technique for each of the 200 permutation data sets, 200 sets of $\hat{\beta}_{\text{perm}}(t)$ are

TABLE 1
Average MSEs with standard errors (SE, in parentheses) and average percentage of correctly identified nonzero and zero elements over 100 replications for 1D cases

Type	Method	Average MSE (SE) ($\times 10^{-3}$)		Average percentage (%)			
		Case 1	Case 2	Case 1		Case 2	
				Nonzero	Zero	Nonzero	Zero
B-spline	SV	0.11 (0.05)	0.19 (0.08)	84.30	69.20	96.00	57.26
	CV	0.15 (0.11)	0.23 (0.11)	82.95	69.68	95.03	58.90
	BIC	0.60 (1.96)	1.63 (3.10)	72.70	96.14	83.26	79.36
	AIC	0.56 (1.96)	1.56 (3.12)	75.80	93.80	82.51	82.27
Fourier	SV	0.65 (0.30)	1.20 (0.49)	84.00	70.59	95.87	58.93
	CV	0.92 (0.56)	1.46 (0.63)	82.30	71.39	95.56	55.76
	BIC	1.12 (0.86)	10.62 (20.69)	72.75	96.59	84.03	67.07
	AIC	1.05 (1.28)	10.29 (20.82)	75.80	93.64	83.85	69.01

obtained. At each t_j , $j = 1, \dots, p$, the two-sided critical values are set to be the 2.5th and 97.5th percentiles of $\hat{\beta}_{\text{perm}}(t_j)$ for the significance level of 0.05. Supposing the null hypothesis is $\beta(t_j) = 0$ at each t_j , we will reject the null hypothesis if $\hat{\beta}(t_j)$ is within the critical region. Repeating this permutation process 100 times, we can compute the percentages that we reject the null hypothesis at each t_j . The results of the permutation tests using the CV method are presented in Figure 2, which shows high rejection frequency in the regions where $\beta(t)$ is nonzero.

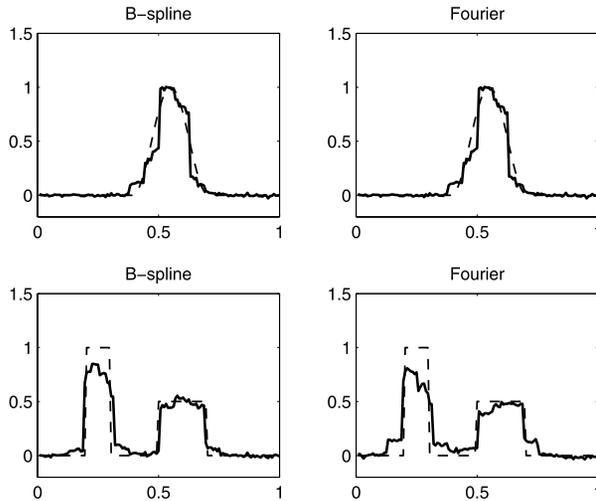


FIG. 1. Average of $\hat{\beta}(t)$ estimated using 5-fold cross-validation with 100 replications (solid line). The dashed line is the true $\beta(t)$. The top panel is for case 1, and the bottom panel is for case 2.

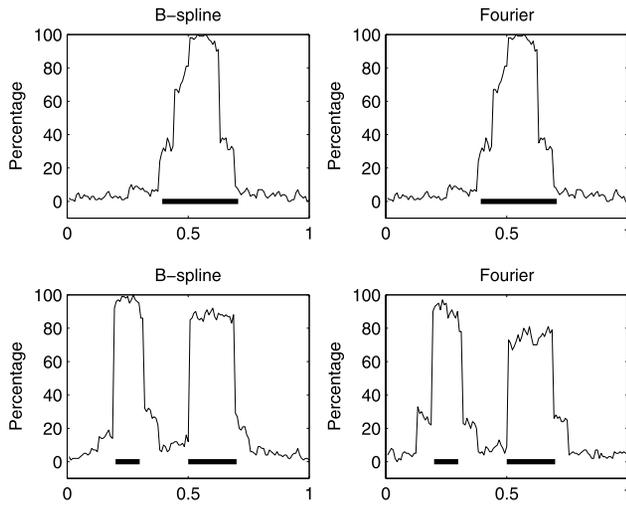


FIG. 2. Frequency of rejecting the null hypothesis $\beta(t) = 0$ using 5-fold cross-validation based on 100 permutation repetitions. The thick solid horizontal segments indicate the true nonzero regions. The top panel is for case 1, and the bottom panel is for case 2.

3.2. 3D simulation. For the 3D case, we generate the following type of images $X(u, v, w) = X^*(u, v, w) + \mathcal{E}(u, v, w)$ with

$$\begin{aligned}
 X^*(u, v, w) &= a_0 + a_1 \sin(2\pi u) + a_2 \cos(2\pi u) + a_3 \sin(2\pi v) \\
 &\quad + a_4 \cos(2\pi v) + a_5 \sin(2\pi w) + a_6 \cos(2\pi w), \quad 0 \leq u, v, w \leq 1,
 \end{aligned}$$

where $a_i \sim N(0, 1)$ and $\mathcal{E}(u, v, w) \sim N(0, \sigma_{\mathcal{E}}^2)$ with $\sigma_{\mathcal{E}}^2$ similarly defined as in the 1D case. For simplicity, we record $32 \times 32 \times 32$ equally spaced measurements in the unit cube. We define the coefficient function $\beta(u, v, w)$ as follows:

$$\beta(u, v, w) = \begin{cases} a(\sin(bu + c) + 1)(\sin(bv + c) + 1)(\sin(bw + c) + 1), & \text{if } (u - 7\pi/40)^2 + (v - 7\pi/40)^2 + (w - 7\pi/40)^2 \leq (3\pi/40)^2; \\ 0, & \text{otherwise,} \end{cases}$$

where $a = 1/8$, $b = 40/3$ and $c = \pi/6$. Note that $\beta(u, v, w)$ is zero outside a ball that is located in the center of the unit cube. The error term ε in model (2.6) also follows a normal distribution $N(0, \sigma^2)$ with $\text{SNR} = 9$. We generate 400 training images and apply the 3D Haar wavelet transform to decompose each image and obtain the wavelet coefficient matrix. Optimal tuning parameters are selected using the same procedures as for the 1D case. The results are summarized in Table 2. Figure 3 illustrates the comparison of the true $\beta(u, v, w)$ and the mean estimates of $\beta(u, v, w)$ over 100 replications at five different slices, which shows that our

TABLE 2
Average MSEs with standard errors (SE, in parentheses), and average percentages of correctly identified nonzero and zero elements over 100 replications for 3D case

Method	Average MSE (SE) ($\times 10^{-4}$)	Average percentage (%)	
		Nonzero	Zero
SV	0.97 (0.29)	77.15	61.97
CV	1.21 (0.51)	74.25	57.04
BIC	4.78 (1.52)	39.48	99.42
AIC	4.11 (2.13)	41.86	98.74

approach performs reasonably well in detecting signals based on visual inspection and on the high percentage of correctly identified nonzeros and zeros reported in Table 2.

4. ADNI PET analysis. The FDG PET data used in the preparation of this article were obtained from the ADNI database (adni.loni.ucla.edu). The ADNI was

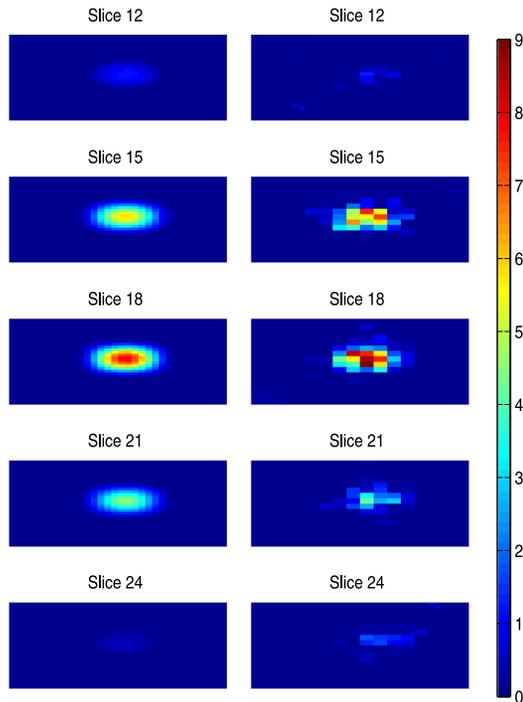


FIG. 3. *The left panel is true $\beta(u, v, w)$ at five selected slices and the right panel is the average of $\hat{\beta}(u, v, w)$ estimated using 5-fold cross-validation over 100 replications at the same five slices.*

launched in 2003 by NIA, NIBIB, FDA, private pharmaceutical companies and nonprofit organizations, as a \$60 million, 5-year public-private partnership. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), PET, other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of MCI and early AD. Determination of sensitive and specific markers of very early AD progression is intended to aid researchers and clinicians in developing new treatments, monitoring treatment effectiveness, and lessening the time and cost of clinical trials. The Principal Investigator of this initiative is Michael W. Weiner, MD, VA Medical Center and University of California, San Francisco. ADNI is the result of efforts of many co-investigators from a broad range of academic institutions and private corporations, and subjects have been recruited from over 50 sites across the U.S. and Canada. The initial goal of ADNI was to recruit 800 adults, ages 55 to 90, to participate in the research, approximately 200 cognitively normal older individuals to be followed for 3 years, 400 people with MCI to be followed for 3 years and 200 people with early AD to be followed for 2 years. For up-to-date information, see www.adni-info.org.

In the ADNI's FDG PET study, the injected dose of FDG was 5.0 ± 0.5 mCi, and subjects were scanned from 30 to 60 minutes post-injection acquiring 6 five-minute frames. The scans were preprocessed by the following steps: each frame was co-registered to the first frame of the raw image file; six co-registered frames were averaged to create a single 30-minute PET image; each subject's co-registered, averaged PET image from the baseline PET scan was reoriented into a standard $160 \times 160 \times 96$ voxel image grid with 1.5 mm cubic voxels and the anterior-posterior axis of the subject is parallel to a line connecting the anterior and posterior commissures (the AC-PC line). It should be noted that the number of voxels in each image is over 2.4 million, so the approach via linear programming, as in James, Wang and Zhu (2009), is too computationally expensive for this application. The data set consists of 403 scans, including 102 NCs, 206 subjects with MCI and 95 subjects diagnosed with AD. The demographic characteristics of the 403 subjects are described in Table 3. The goal of our analysis is to identify brain subregions that are most closely related to MMSE scores; we therefore choose not to adjust for age and other demographic variables. The summary of MMSE scores among the three groups of participants is given in Figure 4. We

TABLE 3
Demographics of ADNI participants (n = 403)

Category	Sex (% male)	Age (SD)	MMSE (SD)
NC (<i>n</i> = 102)	60.8%	80.9 (4.7)	28.9 (1.1)
MCI (<i>n</i> = 206)	67.0%	79.7 (7.3)	27.2 (1.7)
AD (<i>n</i> = 95)	58.9%	80.4 (7.5)	23.4 (2.1)

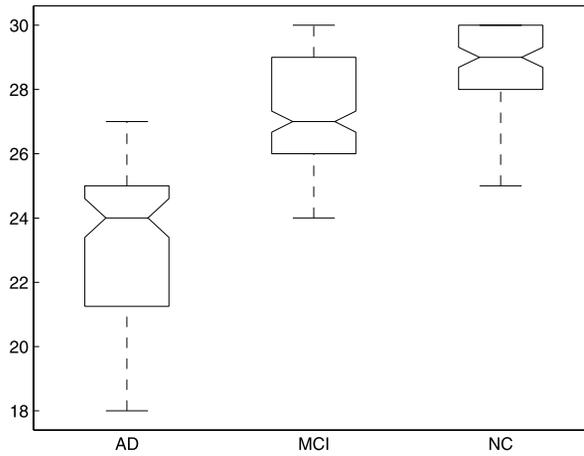


FIG. 4. Box plots of MMSE scores among AD, MCI and NC.

treat each PET image as a realization of the 3D functional predictor and then fit the 3D functional linear regression model (2.6). The voxel values outside the brain are set to zero prior to implementing the 3D Haar wavelet transform. We further reduce the computational cost by excluding those columns of the wavelet coefficient matrix where all the elements are zero.

In terms of applying the 3D Haar wavelet transforms to each subject's PET image data, we consider all the possible levels of the Haar wavelet decompositions. Two tuning parameters are therefore included in the model selection procedure: the level of the 3D Haar wavelet decomposition and the lasso regularization parameter.

First, we employ a 10-fold cross-validation to evaluate the predictive power of the proposed method. Specifically, for each set of 10% observations, we leave them out as a test set, use the remaining data as the training data to fit a model (including selecting the tuning parameters via 5-fold cross-validation) and compute the prediction error on the data points that have been left out. We aggregate these quantities by using the predictive R-square given by $1 - \sum (y_i - \hat{y}_{i,-i})^2 / \sum (y_i - \bar{y})^2$, where $\hat{y}_{i,-i}$ denotes the predicted value of y_i calculated by using the estimator obtained from the training data generated from the cross-validation. The result is 0.26 for the ADNI data set, whereas the standard R-square is 0.51, suggesting a moderate predictive power of the model.

Second, we investigate the voxels that are selected by our method. We use 5-fold cross-validation to the full data set to choose the optimal set of tuning parameters. The identified clusters of voxels [$\hat{\beta}(u, v, w) \neq 0$] are shown on selected axial slices in Figure 5, which are presented from the bottom of the brain to the top. The clusters of voxels with hot colors show a positive association to prediction of MMSE scores, whereas those with cold colors show a negative association. Each small square represents a small cluster of voxels. To assess the significance of the selected voxels, similar to what we have done in simulation studies, we permute the

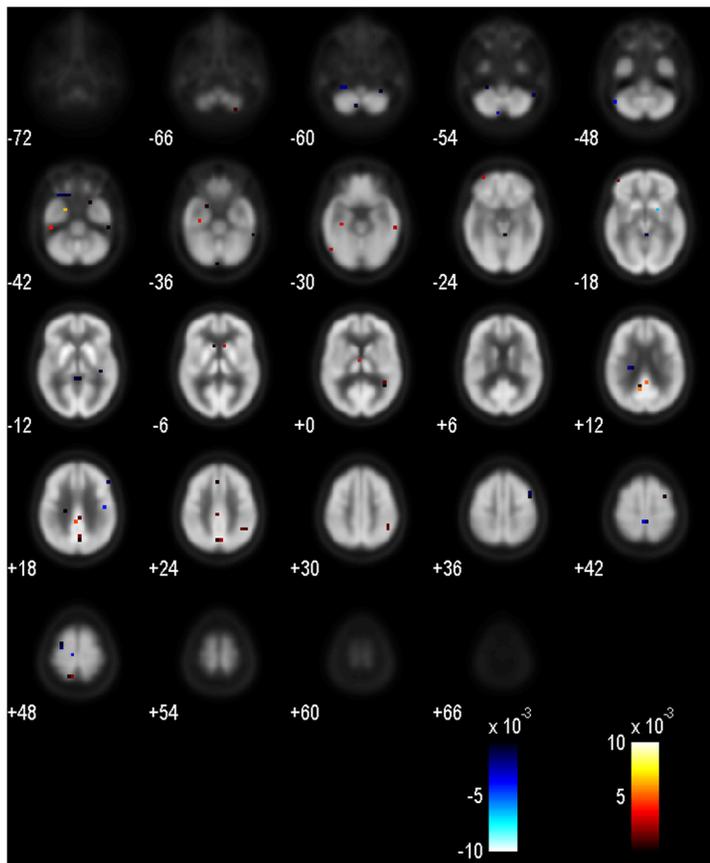


FIG. 5. Clusters of voxels identified using our approach for the ADNI data.

response variable MMSE score 200 times. It turns out that 95.3% of the selected voxels are significant at the 5% level. In addition to this pointwise testing, we also consider the global test described by [Nichols and Holmes \(2001\)](#), which provides a way to control the family-wise error rate by comparing $\hat{\beta}(t_j)$ to a “maximal statistic.” It turns out that only 15.6% of the selected voxels are significant at the 5% level, which is more conservative than the pointwise testing procedure. To further evaluate the stability of the selection, we generate 100 bootstrap samples and for each bootstrap sample, we apply our method including the tuning parameter selection via 5-fold cross-validation. Similar approaches also have been employed by other researchers, such as [Sauerbrei and Schumacher \(1992\)](#), [Royston and Sauerbrei \(2008\)](#) and [Meinshausen and Bühlmann \(2010\)](#). To summarize the results, we count the number of times that each voxel is selected over 100 bootstrap samples and denote it as the bootstrap inclusion frequency (BIF). The voxel BIFs are presented in Figure 6. The locations of these more frequently selected voxels are also presented in the 3D sagittal view in Figure 7 for ease of understanding. It

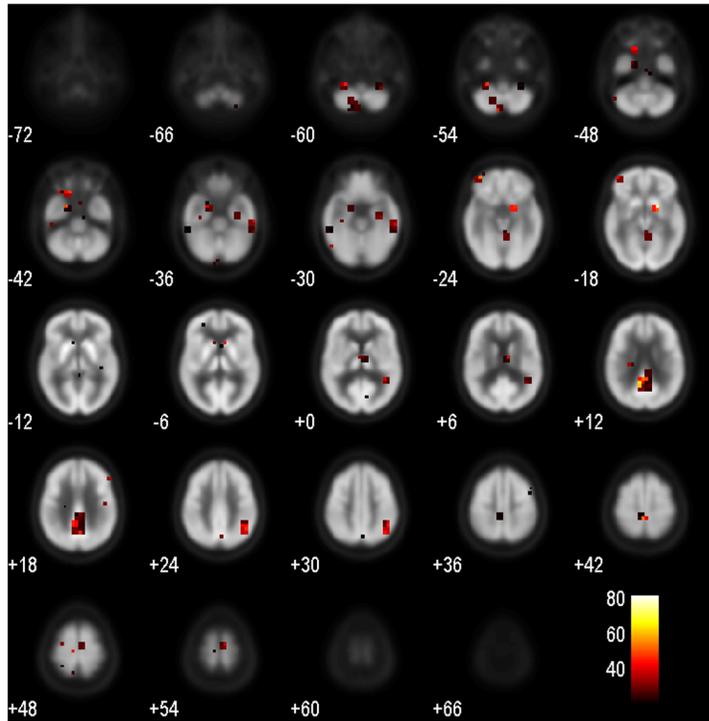


FIG. 6. *Bootstrap inclusion frequencies of the voxels over 100 bootstrap samples.*

can be seen that the highly selected brain regions agree well with the results in Figure 5. We note that the clusters of voxels identified in our analysis shown in Figures 5 and 6 reveal high associations of the expected anatomical regions with cognitive deficits. For example, the orange ones on slices “+12” and “+18” in Figure 5 and the big cluster on the same slices in Figure 6 indicate that the posterior cingulate/precuneus cortex is significantly related to cognitive impairment;

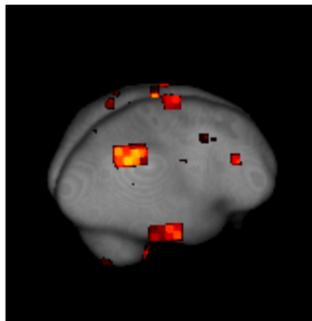


FIG. 7. *Locations of the frequently selected voxels in the 3D sagittal view.*

the blue ones on slices “-60,” “-54” and “-48” in Figure 5 and the clusters on the same slices in Figure 6 suggest that the medial temporal/hippocampal cortex is also closely involved; the red ones on slices “-42,” “-36” and “-30” in Figure 5 and the corresponding clusters on the same slices in Figure 6 correspond to the lateral temporal cortex. Many studies have demonstrated that the most prominent metabolic abnormalities are found in these regions; see, for example, Foster et al. (1984), Minoshima et al. (1995, 1997), Mueller et al. (2005). In our study, we have particularly found the most predictive voxels of the cognitive impairment in these regions. Other involved regions include the superior lateral parietal cortex and the frontal cortex, which are all known to be related to the progression of Alzheimer’s disease.

5. Discussion. In this article we propose a highly effective Haar wavelet-based regularization approach that can be easily applied to analyzing multidimensional functional data. Analysis of the PET image data demonstrates that our approach is useful in finding brain subregions that are most responsible for cognitive impairment in elderly people. It has great potential to efficiently assist the diagnosis of disease in neuroimaging studies, yielding easily interpretable results. Our approach is also computationally fast because of the implementation of the coordinate descent algorithm with the MATLAB glmnet package. For example, the real data analysis of 403 subjects’ PET image data can be finished in less than two hours on a 64-bit Intel Xeon 3.33 GHz server with about 35 GB of RAM, including the selection of tuning parameters. We should note that another practical advantage of our approach is that the wavelet transform itself can reduce the dimensionality of the large volume of brain image data. As a result, we can then apply the proposed approach on reduced data sets. In such situations, although the resolution of the original PET images is decreased, the results remain largely the same since the related subregions are usually not comprised of a single voxel but of a cluster of voxels.

Acknowledgments. Data used in preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (adni.loni.ucla.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf.

Data collection and sharing for this project were funded by the Alzheimer’s Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: Abbott; Alzheimer’s Association; Alzheimer’s Drug

Discovery Foundation; Amorfix Life Sciences Ltd.; AstraZeneca; Bayer Health-Care; BioClinica, Inc.; Biogen Idec Inc.; Bristol-Myers Squibb Company; Eisai Inc.; Elan Pharmaceuticals Inc.; Eli Lilly and Company; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; GE Healthcare; Innogenetics, N.V.; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Medpace, Inc.; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Servier; Synarc Inc.; and Takeda Pharmaceutical Company. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer's Disease Cooperative Study at the University of California, San Diego. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of California, Los Angeles.

The authors would like to thank Professor Karen Kafadar, the Associate Editor and two referees for their helpful and constructive comments.

SUPPLEMENTARY MATERIAL

Appendix (DOI: [10.1214/14-AOAS736SUPP](https://doi.org/10.1214/14-AOAS736SUPP); .pdf). The online supplementary material contains the technical appendix showing the theoretical results of the proposed approach and an illustrative example showing the desirable feature of Haar wavelets.

REFERENCES

- BICKEL, P. J., RITOV, Y. and TSYBAKOV, A. B. (2009). Simultaneous analysis of lasso and Dantzig selector. *Ann. Statist.* **37** 1705–1732. [MR2533469](#)
- BOWMAN, F. D., CAFFO, B., BASSETT, S. S. and KILTS, C. (2008). A Bayesian hierarchical framework for spatial modeling of fMRI data. *NeuroImage* **39** 146–156.
- CANDES, E. and TAO, T. (2007). The Dantzig selector: Statistical estimation when p is much larger than n . *Ann. Statist.* **35** 2313–2351. [MR2382644](#)
- COCKRELL, J. R. and FOLSTEIN, M. F. (1988). Mini-Mental State Examination (MMSE). *Psychopharmacol. Bull.* **24** 689–692.
- DAUBECHIES, I., DEFRISE, M. and DE MOL, C. (2004). An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Comm. Pure Appl. Math.* **57** 1413–1457. [MR2077704](#)
- FAN, J., GUO, S. and HAO, N. (2012). Variance estimation using refitted cross-validation in ultrahigh-dimensional regression. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **74** 37–65. [MR2885839](#)
- FOSTER, N. L., CHASE, T. N., MANSI, L., BROOKS, R., FEDIO, P., PATRONAS, N. J. and CHIRO, G. D. (1984). Cortical abnormalities in Alzheimer's disease. *Ann. Neurol.* **16** 649–654.
- FRIEDMAN, J., HASTIE, T., HÖFLING, H. and TIBSHIRANI, R. (2007). Pathwise coordinate optimization. *Ann. Appl. Stat.* **1** 302–332. [MR2415737](#)
- FRISTON, K. J., POLINE, J. B., HOLMES, A. P., FRITH, C. and FRACKOWIAK, R. S. J. (1996). A multivariate analysis of PET activation studies. *Hum. Brain Mapp.* **4** 140–151.

- FU, W. J. (1998). Penalized regressions: The bridge versus the lasso. *J. Comput. Graph. Statist.* **7** 397–416. [MR1646710](#)
- GRIMMER, T., HENRIKSEN, G., WESTER, H.-J., FÖRSTL, H., KLUNK, W. E., MATHIS, C. A., KURZ, A. and DRZEZGA, A. (2009). Clinical severity of Alzheimer’s disease is associated with PIB uptake in PET. *Neurobiol. Aging* **30** 1902–1909.
- JAMES, G. M., WANG, J. and ZHU, J. (2009). Functional linear regression that’s interpretable. *Ann. Statist.* **37** 2083–2108. [MR2543686](#)
- KANG, J., JOHNSON, T. D., NICHOLS, T. E. and WAGER, T. D. (2011). Meta analysis of functional neuroimaging data via Bayesian spatial point processes. *J. Amer. Statist. Assoc.* **106** 124–134. [MR2816707](#)
- KERROUCHE, N., HERHOLZ, K., MIELKE, R., HOLTHOFF, V. and BARON, J.-C. (2006). 18FDG PET in vascular dementia: Differentiation from Alzheimer’s disease using voxel-based multivariate analysis. *J. Cereb. Blood Flow Metab.* **26** 1213–1221.
- LEIFER, B. P. (2003). Early diagnosis of Alzheimer’s disease: Clinical and economic benefits. *J. Am. Geriatr. Soc.* **51** S281–S288.
- LOVEJOY, S. and SCHERTZER, D. (2012). Haar wavelets, fluctuations and structure functions: Convenient choices for geophysics. *Nonlinear Process. Geophys.* **19** 513–527.
- LUO, W.-L. and NICHOLS, T. E. (2003). Diagnosis and exploration of massively univariate neuroimaging models. *NeuroImage* **19** 1014–1032.
- MEINSHAUSEN, N. and BÜHLMANN, P. (2010). Stability selection. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **72** 417–473. [MR2758523](#)
- MINOSHIMA, S., FREY, K. A., KOEPPE, R. A., FOSTER, N. L. and KUHL, D. E. (1995). A diagnostic approach in Alzheimer’s disease using three-dimensional stereotactic surface projections of fluORine-18-FDG PET. *J. Nucl. Med.* **36** 1238–1248.
- MINOSHIMA, S., GIORDANI, B., BERENT, S., FREY, K. A., FOSTER, N. L. and KUHL, D. E. (1997). Metabolic reduction in the posterior cingulate cortex in very early Alzheimer’s disease. *Ann. Neurol.* **42** 85–94.
- MUELLER, S. G., WEINER, M. W., THAL, L. J., PETERSEN, R. C., JACK, C. R., JAQUEST, W., TROJANOWSKI, J. Q., TOGA, A. W. and BECKETT, L. (2005). Ways toward an early diagnosis in Alzheimer’s disease: The Alzheimer’s Disease Neuroimaging Initiative (ADNI). *Alzheimer’s and Dementia: The Journal of the Alzheimer’s Association* **1** 55–66.
- MURAKI, S. (1992). Approximation and rendering of volume data using wavelet transforms. In *Proceedings of Visualization 1992* 21–28. IEEE.
- NICHOLS, T. E. and HOLMES, A. P. (2001). Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Hum. Brain Mapp.* **15** 1–25.
- PICART, S. S., BUTENSCHÖN, M. and SHUTLER, J. D. (2012). Wavelet-based spatial comparison technique for analysing and evaluating two-dimensional geophysical model fields. *Geosci. Model Dev.* **5** 223–230.
- PLASSMAN, B. L., LANGA, K. M., FISHER, G. G., HEERINGA, S. G., WEIR, D. R., OFSTEDAL, M. B., BURKE, J. R., HURD, M. D., POTTER, G. G., RODGERS, W. L., STEFFENS, D. C., WILLIS, R. J. and WALLACE, R. B. (2007). Prevalence of dementia in the United States: The aging, demographics, and memory study. *Neuroepidemiology* **29** 125–132.
- RAMSAY, J. O. and SILVERMAN, B. W. (2005). *Functional Data Analysis*, 2nd ed. Springer, New York. [MR2168993](#)
- REISS, P. T. and OGDEN, R. T. (2010). Functional generalized linear models with images as predictors. *Biometrics* **66** 61–69. [MR2756691](#)
- ROYSTON, P. and SAUERBREI, W. (2008). *Multivariable Model-Building: A Pragmatic Approach to Regression Analysis Based on Fractional Polynomials for Modelling Continuous Variables*, 1st ed. Wiley, Chichester. [MR2426993](#)
- SAUERBREI, W. and SCHUMACHER, M. (1992). A bootstrap resampling procedure for model building: Application to the Cox regression model. *Stat. Med.* **11** 2093–2109.

- SHIN, J., LEE, S. Y., KIM, S. J., KIM, S. H., CHO, S. J. and KIM, Y. B. (2010). Voxel-based analysis of Alzheimer's disease PET imaging using a triplet of radiotracers: PIB, FDDNP, and FDG. *NeuroImage* **52** 488–496.
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B* **58** 267–288. [MR1379242](#)
- WALKER, J. S. (2008). *A Primer on Wavelets and Their Scientific Applications*, 2nd ed. Chapman & Hall/CRC, Boca Raton, FL. [MR2400818](#)
- WAND, M. P. and ORMEROD, J. T. (2011). Penalized wavelets: Embedding wavelets into semiparametric regression. *Electron. J. Stat.* **5** 1654–1717. [MR2870147](#)
- WANG, X., NAN, B., ZHU, J., KOEPPE, R. and FOR THE ALZHEIMER'S DISEASE NEUROIMAGING INITIATIVE (2014). Supplement to "Regularized 3D functional regression for brain image data via Haar wavelets." DOI:[10.1214/14-AOAS736SUPP](#).
- WU, T. T. and LANGE, K. (2008). Coordinate descent algorithms for lasso penalized regression. *Ann. Appl. Stat.* **2** 224–244. [MR2415601](#)
- ZHAO, Y., TODD OGDEN, R. and REISS, P. T. (2012). Wavelet-based LASSO in functional linear regression. *J. Comput. Graph. Statist.* **21** 600–617. [MR2970910](#)
- ZOU, H. (2006). The adaptive lasso and its oracle properties. *J. Amer. Statist. Assoc.* **101** 1418–1429. [MR2279469](#)
- ZOU, H., HASTIE, T. and TIBSHIRANI, R. (2007). On the "degrees of freedom" of the lasso. *Ann. Statist.* **35** 2173–2192. [MR2363967](#)

X. WANG
 B. NAN
 DEPARTMENT OF BIostatISTICS
 UNIVERSITY OF MICHIGAN
 1415 WASHINGTON HEIGHTS
 ANN ARBOR, MICHIGAN 48109-2029
 USA
 E-MAIL: xuejwang@umich.edu
bnan@umich.edu

J. ZHU
 DEPARTMENT OF STATISTICS
 UNIVERSITY OF MICHIGAN
 1085 SOUTH UNIVERSITY AVE
 ANN ARBOR, MICHIGAN 48109-1107
 USA
 E-MAIL: jizhu@umich.edu

R. KOEPPE
 DEPARTMENT OF RADIOLOGY
 UNIVERSITY OF MICHIGAN
 1500 E. MEDICAL CENTER DRIVE
 ANN ARBOR, MICHIGAN 48109-5030
 USA
 E-MAIL: koeppe@umich.edu