




Generalized Scalar-on-Image Regression Models via Total Variation

Xiao Wang, Hongtu Zhu & for the Alzheimer's Disease Neuroimaging Initiative


To cite this article: Xiao Wang, Hongtu Zhu & for the Alzheimer's Disease Neuroimaging Initiative (2017) Generalized Scalar-on-Image Regression Models via Total Variation, Journal of the American Statistical Association, 112:519, 1156-1168, DOI: [10.1080/01621459.2016.1194846](https://doi.org/10.1080/01621459.2016.1194846)

To link to this article: <https://doi.org/10.1080/01621459.2016.1194846>


 View supplementary material [↗](#)

 Accepted author version posted online: 10 Jun 2016.
Published online: 13 Apr 2017.

 Submit your article to this journal [↗](#)

 Article views: 736

 View Crossmark data [↗](#)

 Citing articles: 3 View citing articles [↗](#)

Generalized Scalar-on-Image Regression Models via Total Variation

Xiao Wang^a and Hongtu Zhu^b for the Alzheimer's Disease Neuroimaging Initiative

^aDepartment of Statistics, Purdue University, West Lafayette, IN; ^bDepartment of Biostatistics, University of North Carolina, Chapel Hill, NC

ABSTRACT

The use of imaging markers to predict clinical outcomes can have a great impact in public health. The aim of this article is to develop a class of generalized scalar-on-image regression models via total variation (GSIRM-TV), in the sense of generalized linear models, for scalar response and imaging predictor with the presence of scalar covariates. A key novelty of GSIRM-TV is that it is assumed that the slope function (or image) of GSIRM-TV belongs to the space of bounded total variation to explicitly account for the piecewise smooth nature of most imaging data. We develop an efficient penalized total variation optimization to estimate the unknown slope function and other parameters. We also establish nonasymptotic error bounds on the excess risk. These bounds are explicitly specified in terms of sample size, image size, and image smoothness. Our simulations demonstrate a superior performance of GSIRM-TV against many existing approaches. We apply GSIRM-TV to the analysis of hippocampus data obtained from the Alzheimers Disease Neuroimaging Initiative (ADNI) dataset. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received September 2014
Revised April 2016

KEYWORDS

Excess risk; Functional regression; Generalized scalar-on-image regression; Prediction; Total variation

1. Introduction

The aim of this article is to develop generalized scalar-on-image regression models via total variation (GSIRM-TV) with scalar response and imaging and/or scalar predictors. This new development is motivated by studying the predictive value of ultra-high-dimensional imaging data and/or other scalar predictors (e.g., cognitive score) for clinical outcomes including diagnostic status and the response to treatment in the study of neurodegenerative and neuropsychiatric diseases, such as Alzheimer's disease (AD) (Mu and Gage 2011). For instance, the growing public threat of AD has raised the urgency to discover and validate prognostic biomarkers that may identify subjects at greatest risk for future cognitive decline and accelerate the testing of preventive strategies. In this regard, prior studies of subjects at risk for AD have examined the utility of various individual biomarkers, such as cognitive tests, fluid markers, imaging measurements, or some individual genetic markers (e.g., APOE4 gene), to capture the heterogeneity and multifactorial complexity of AD (reviewed in Weiner et al. 2012).

Our GSIRM-TV considers the use of imaging predictor X and/or scalar predictors Z to predict scalar response Y . In practice, imaging data are often represented in the form of two-dimensional matrix or three-dimensional array. Assume that $X \in \mathbb{R}^{N \times N}$ is a two-dimensional matrix of size $N \times N$ which is observed without error and $Z \in \mathbb{R}^p$ is a $p \times 1$ vector with the first component being constant 1. Our GSIRM-TV assumes that Y given (X, Z) follows:

$$Y|(X, Z) \sim \text{Exponential Family}(\mu, \phi) \quad \text{and} \\ g(\mu) = \theta_0^T Z + \langle X, \beta_0 \rangle, \quad (1)$$

where μ and ϕ are, respectively, canonical and scale parameters, $\langle U, V \rangle = \sum_{i,j} u_{i,j} v_{i,j}$ for $U = (u_{i,j}) \in \mathbb{R}^{N \times N}$ and $V = (v_{i,j}) \in \mathbb{R}^{N \times N}$, and $g(\cdot)$ is a known link function. Moreover, θ_0 and $\beta_0(\cdot)$ are unknown parameters of interest and $\beta_0(\cdot)$ is called the *coefficient image/function*. Throughout the article, assume that images are observed without error. We may deal with such measurement errors in images by applying some smoothing methods to reduce error in images (Li, Wang, and Carroll 2010).

GSIRM-TV can be regarded as an extension of the well-known functional linear model (FLM) and the high-dimensional linear model (HLM) that have been extensively studied in the literature. If we regard $\langle U, V \rangle$ as an approximation of a two-dimensional integral, then GSIRM-TV is an approximated version of FLM. The literature on FLM is too vast to summarize here. Please see the well-known monographs Ramsay and Silverman (2005) and Ferraty and Vieu (2006). The functional principal component analysis (fPCA) and various penalization methods have been developed to estimate the coefficient function. For example, the fPCA method has been discussed by James (2002), Müller and Stadtmüller (2005), Hall and Horowitz (2007), Reiss and Ogden (2007, Reiss and Ogden 2010), James, Wang, and Zhu (2009), and Goldsmith et al. (2010) and the penalized method has been studied by Crambes, Kneip, and Sarda (2009), Yuan and Cai (2010), and Du and Wang (2014). On the other hand, if we vectorize X and $\beta_0(\cdot)$ as $N^2 \times 1$ vectors, model (1) takes the form of the high-dimensional generalized linear regression. To achieve sparsity in β_0 , various penalization methods, such as Lasso or SCAD, have been developed. Please see Tibshirani (1996), Chen, Donoho, and Saunders (1998), Fan and Li (2001), and references therein.




Figure 1. Results from a simulated dataset. The top row includes the true 64×64 coefficient image β_0 in the left panel, one realization of a 64×64 image predictor X in the middle panel, and the responses Y from $n = 300$ in the right panel. The bottom row includes the estimated coefficient functions obtained from fPCA (left), Lasso (middle), and total variation (right).

Compared with FLM and HLM, a key novelty of GSIRM-TV is that the coefficient image $\beta_0(\cdot)$ in model (1) is assumed to be a piecewise smooth image with unknown jumps and edges. Such assumption not only has been widely used in the imaging literature, but also is critical for addressing various scientific questions, such as the identification of brain regions associated with AD. As an illustration, we consider a dataset with $n = 300$ subjects simulated from a functional linear model which is a special case of (1). The first row of Figure 1 presents the true 64×64 image matrix β_0 , X , and Y from the left to the right. We have vectorized X , used fPCA for FLM, Lasso for HLM, and GSIRM-TV to estimate the coefficient image and presented the estimated coefficient images in the second row of Figure 1. Unfortunately, both FLM and HLM fail to capture the main feature of the true coefficient image due to their key limitations. First, fPCA requires that β_0 be well presented by the eigenfunctions of X , whereas it is not the case according to Figure 1. Second, the existing regularization methods can have difficulty in recovering β_0 , since the true coefficient image is nonsparse. Moreover, most regularization methods for FLM assume that the unknown coefficient function is one-dimensional and belongs to a smoothed function space, such as the Sobolev space, and thus they will not be able to preserve edge and boundary information for the dataset presented in Figure 1. In contrast, our GSIRM-TV estimate developed in this article can truly preserve the sharp edge of the original image.

In this article, we make two important contributions including a new estimation method based on the total variation analysis and nonasymptotic error bounds on the risk under the framework of GSIRM-TV. The total variation analysis plays a fundamental role in various image analyses since the path-breaking works of Rudin and Osher (1994) and Rudin, Osher, and Fatemi (1992). The total variation penalty has been proved

to be quite efficient for preserving the boundaries and edges of images (Rudin, Osher, and Fatemi 1992). Michel et al. (2011) proposed a similar total variation method for image regression and image classification, but they focus on the development of different algorithms for the TV optimization problem. According to the best of our knowledge, this is the first article on the development of statistical analysis of the total variation method for GSIRM-TV. The fused lasso (Tibshirani et al. 2005; Friedman, Hastie, and Tibshirani 2007) uses a similar penalty function. But for the two-dimensional parameter, the fused lasso, and the TV penalty can be quite different. For example, the isotropic total variation penalty uses the Euclidean norm of the first differences of the parameter, rather than the sum of the absolute values of the first differences. There are a few papers on the use of two-dimensional or three-dimensional imaging predictors in FLM (Guillas and Lai 2010; Reiss and Ogden 2010; Zhou, Li, and Zhu 2013; James, Wang, and Zhu 2009; Goldsmith et al. 2010; Gertheiss, Maity, and Staicu 2013; Wang et al. 2014; Reiss et al. 2015), but none of them consider the piecewisely smoothed function with jumps and edges and the total variation analysis. We also derive nonasymptotic error bounds on the risk for the estimated coefficient image under the total variation penalty. We are able to obtain finite-sample bounds that are specified explicitly in terms of the sample size n , the image size $N \times N$, and the image smoothness.

The rest of the article is organized as follows. Section 2 considers linear scalar-on-image regression model and proposes the TV optimization framework to estimate the unknown coefficient image. We also establish the nonasymptotic error bound for the prediction error. Section 3 extends linear scalar-on-image regression model to generalized scalar-on-image regression models. Section 4 examines the finite-sample performance of GSIRM-TV and compares it with several state-of-the-art



Figure 2. Left: the Shepp-Logan phantom image; Middle and Right: the two components of the discrete gradient of the phantom image.

methods, such as regularized matrix regression (Zhou and Li 2014). Section 5 applies GSIRM-TV to the use of the hippocampus imaging data for a binary classification problem. Future research directions are discussed in Section 6. The technical proofs of main theorems are given in the Appendix.

2. Linear Scalar-on-Image Regression Model

We start with considering a linear scalar-on-image regression model, which is the simplest case of GSIRM-TV (1), as follows:

$$Y = \langle X, \beta_0 \rangle + \epsilon, \quad (2)$$

where ϵ is the random error with $\mathbb{E}(\epsilon|X) = 0$ and $\mathbb{E}(\epsilon^2|X) = \sigma^2$, and without loss of generality, both X and Y are assumed to be centered with $\mathbb{E}(Y) = \mathbb{E}(X) = 0$. Model (2) may be treated as a special case of FLM since discrete images are isometric to the space of piecewise-constant functions defined as

$$\mathcal{X} = \left\{ x \in L^2(\Omega) : x(u, v) = NX_{jk}, \frac{j-1}{N} \leq u < \frac{j}{N}, \frac{k-1}{N} \leq v < \frac{k}{N} \text{ for } 1 \leq j, k \leq N \right\},$$

where X_{jk} is the (j, k) th pixel value of the image X and $\Omega = [0, 1]^2$. By treating β_0 as an integrable function in Ω , that is, $\beta_0 \in L^2(\Omega)$, model (2) can be rewritten as

$$Y = \int_0^1 \int_0^1 x(u, v) \beta_0(u, v) dudv + \epsilon.$$

2.1. The Space of Bounded Variation

Throughout the article, it is assumed that β_0 is a function of bounded variation in Ω if the total variation of β_0 in Ω , denoted by $\|\beta_0\|_{\text{TV}}$, is finite and defined as follows:

$$\|\beta_0\|_{\text{TV}} = \sup \left\{ \int_{\Omega} \beta_0(u, v) \operatorname{div} f(u, v) dudv : f \in C_c^\infty(\Omega; \mathbb{R}^2), |f|_\infty \leq 1 \right\},$$

where $|f|_\infty = \sup_{(u,v) \in \Omega} |f(u, v)|$ and $C_c^\infty(\Omega; \mathbb{R}^2)$ denotes the vector field with value in \mathbb{R}^2 , which is infinitely differentiable and has compact support in Ω . Moreover, $f(u, v) = (f_1(u, v), f_2(u, v))$ and $\operatorname{div} f(u, v) = \partial_u f_1(u, v) + \partial_v f_2(u, v)$, where $\partial_u = \partial/\partial u$ and $\partial_v = \partial/\partial v$. The vector space of functions of bounded variation in Ω is denoted by $\text{BV}(\Omega)$. For example, if β_0 is differentiable in Ω , then $\|\beta_0\|_{\text{TV}}$ reduces to

$\int_{\Omega} \sqrt{(\partial_u \beta_0)^2 + (\partial_v \beta_0)^2} dudv$. In this case, β_0 belongs to the Sobolev space $W^{1,1}(\mathcal{D})$, that is, functions with integrable first order partial derivatives. However, the power of total variation in image analysis arises exactly from the relaxation of such constraints. The $\text{BV}(\Omega)$ is much larger than $W^{1,1}(\mathcal{D})$ and contains many interesting piecewise continuous functions with jumps and edges. This is exactly the advantage of using TV regularization over other familiar regularization methods used in the non-parametric literature. For example, the smoothing spline penalty term is not sensitive enough to capture sharp edges and jumps.

There are at least two additional advantages of using bounded variation functions in model (2). First, many real images with edges have small total variation since image edges usually reside in a low-dimensional subset of pixels. As an illustration, in Figure 2, the left panel displays the Shepp-Logan phantom image, while the middle and right panels show the two components of the discrete gradient of the phantom image, which have obvious sparse patterns. Second, $\text{BV}(\Omega)$ is mathematically tractable even though it contains many more functions with edges and jumps compared with $W^{1,1}(\mathcal{D})$.

2.2. Estimation

On the basis of model (2) and $\text{BV}(\Omega)$, we propose to solve the following TV minimization:

$$\begin{aligned} & \text{minimize} \quad \|\beta\|_{\text{TV}} \\ & \text{subject to} \quad \sum_{i=1}^n (Y_i - \langle X_i, \beta \rangle)^2 \leq \lambda^2, \end{aligned} \quad (3)$$

where λ is a smoothing parameter, which controls the noise level. It is known that the above minimization problem is equivalent to the penalized optimization

$$\sum_{i=1}^n (Y_i - \langle X_i, \beta \rangle)^2 + \check{\lambda} \|\beta\|_{\text{TV}}, \quad (4)$$

where $\check{\lambda}$ is a different smoothing parameter. The TV optimization has been widely used to reconstruct images in the compressive sensing literature (see, e.g., Candès, Romberg, and Tao 2006a, 2006b; Needell and Ward 2013). Using the TV optimization for one-dimensional regression has been studied by Mammen and van de Geer (1997) and Tibshirani (2014). Michel et al. (2011) discussed some algorithms to solve a similar optimization problem. To the best of our knowledge, nothing has been done on the statistical properties of the TV estimator for scalar-on-image regression models.

To solve the TV minimization (3) (or (4)), we treat $\beta = (\beta_{jk}) \in \mathbb{R}^{N \times N}$ as an $N \times N$ block of pixels with β_{jk} as its (j, k) element. Then, we define the discrete total variation of $\beta = (\beta_{jk}) \in \mathbb{R}^{N \times N}$. For any $\beta \in \text{BV}(\Omega)$, the discrete gradient $\nabla : \text{BV}(\Omega) \rightarrow \mathbb{R}^{N \times N \times 2}$ is defined by

$$(\nabla \beta)_{jk} = \begin{cases} (\beta_{j+1,k} - \beta_{jk}, \beta_{j,k+1} - \beta_{jk}), & 1 \leq j, k \leq N - 1, \\ (0, \beta_{j,k+1} - \beta_{jk}), & j = N, 1 \leq k \leq N - 1, \\ (\beta_{j+1,k} - \beta_{jk}, 0), & 1 \leq j \leq N - 1, k = N, \\ (0, 0), & k = j = N. \end{cases}$$

Based on $(\nabla \beta)_{jk} = ((\nabla \beta)_{jk,1}, (\nabla \beta)_{jk,2})^T$, the anisotropic version of the total variation norm $\|\beta\|_{\text{TV}}$ can be rewritten as

$$\|\beta\|_{\text{TV}}^{\text{aniso}} = \|\nabla \beta\|_1 = \sum_{jk} \left\{ |(\nabla \beta)_{jk,1}| + |(\nabla \beta)_{jk,2}| \right\}.$$

On the other hand, its isotropic version is defined by

$$\|\beta\|_{\text{TV}}^{\text{iso}} = \sum_{jk} \|(\nabla \beta)_{jk}\|_2 = \sum_{jk} \sqrt{(\nabla \beta)_{jk,1}^2 + (\nabla \beta)_{jk,2}^2}.$$

The anisotropic and isotropic induced total variation norms are equivalent up to a factor of $\sqrt{2}$, that is,

$$\frac{1}{\sqrt{2}} \|\beta\|_{\text{TV}}^{\text{iso}} \leq \|\beta\|_{\text{TV}}^{\text{aniso}} \leq \sqrt{2} \|\beta\|_{\text{TV}}^{\text{iso}}.$$

We will write all results in terms of the anisotropic total variation seminorm, but our results also extend to the isotropic version.

Let A_X be an $n \times N^2$ design matrix such that the i th row is the vectorized X_i . With a slight abuse of notation, we use β to denote the coefficient matrix and its corresponding vector. We may rewrite (3) as the matrix form given by

$$\hat{\beta} = \arg \min \|\beta\|_{\text{TV}} \quad \text{subject to} \quad \|Y - A_X \beta\|_2 \leq \lambda. \quad (5)$$

We adapt an algorithm called TVAL3 based on the augmented Lagrangian method (Hestenes 1969; Powell 1969; Li 2013). Specifically, we solve an equivalent optimization problem given by

$$\min_{w, \beta} \sum_{l=1}^{N^2} \|w_l\|_1 \quad \text{subject to} \quad \|Y - A_X \beta\|_2 \leq \lambda \quad \text{and} \\ D_l \beta = w_l \quad \text{for all } l,$$

where D_l is an $2 \times N^2$ vector of constants associated with the discrete gradient. As an illustration, we consider a simple case with $N = 2$. In this case, we have $\beta = (\beta_{11}, \beta_{12}, \beta_{21}, \beta_{22})^T$. We may choose

$$D_1 = \begin{bmatrix} -1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 \end{bmatrix}, D_2 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 1 \end{bmatrix}, \\ D_3 = \begin{bmatrix} 0 & 0 & -1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad \text{and} \quad D_4 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

so that we have $D_1 \beta = (\nabla \beta)_{11}$, $D_2 \beta = (\nabla \beta)_{12}$, $D_3 \beta = (\nabla \beta)_{21}$, and $D_4 \beta = (\nabla \beta)_{22}$.

Its corresponding augmented Lagrangian function is given by

$$\mathcal{L}_A(w, \beta) = \sum_{l=1}^{N^2} \left\{ \|w_l\|_1 - \mathbf{v}_l^T (D_l \beta - w_l) + \frac{\alpha_l}{2} \|D_l \beta - w_l\|_2^2 + \frac{\gamma}{2} \|A_X \beta - Y\|_2^2 \right\},$$

where \mathbf{v}_l , α_l , and γ are tuning parameters. We may find the minimizer iteratively, and then the subproblem at each iteration of TVAL3 becomes $\min_{w_l, \beta} \mathcal{L}_A(w_l, \beta)$. In our algorithm, \mathbf{v}_l is updated at each iteration. Moreover, α_l 's and γ as smoothing parameters can be selected by using either the C_p criterion or the K -fold cross-validation (CV). However, its computational time can be long even under current computing facilities. In our numerical examples, we pre-fix the tuning parameters by setting $\alpha_l = 2^5$ for $l = 1, \dots, N^2$ and $\gamma = 2^8$. The simplest way to choose γ is to try different values from 2^4 up to 2^{13} and compare the recovered images. The value of α_l is much less sensitive to the choice of γ . We leave tuning parameter optimization for our future research topic.

We describe the complete algorithm as follows.

Step 1. Initialize $\beta^{(0)}$ and $\mathbf{v}_l^{(0)}$;

Step 2. Given $\beta = \beta^{(k)}$ and $\mathbf{v}_l = \mathbf{v}_l^{(k)}$, we solve for $\omega_l^{(k+1)}$, $l = 1, \dots, N^2$, by minimizing

$$\|w_l\|_1 - \mathbf{v}_l^T (D_l \beta - w_l) + \frac{\alpha_l}{2} \|D_l \beta - w_l\|_2^2.$$

The explicit solution (component-wise) is given by

$$\omega_l = \begin{cases} D_l \beta - \frac{\mathbf{v}_l + 1}{\alpha_l}, & \text{if } D_l \beta > \frac{\mathbf{v}_l + 1}{\alpha_l}; \\ 0, & \text{if } \frac{\mathbf{v}_l - 1}{\alpha_l} \leq D_l \beta \leq \frac{\mathbf{v}_l + 1}{\alpha_l}; \\ D_l \beta - \frac{\mathbf{v}_l - 1}{\alpha_l}, & \text{if } D_l \beta < \frac{\mathbf{v}_l - 1}{\alpha_l}. \end{cases}$$

Step 3. Given $\omega_l = \omega_l^{(k+1)}$ and $\mathbf{v}_l = \mathbf{v}_l^{(k)}$, $l = 1, \dots, N^2$, we solve for $\beta^{(k+1)}$ by minimizing

$$\sum_{l=1}^{N^2} \left\{ -\mathbf{v}_l^T D_l \beta + \frac{\alpha_l}{2} \|D_l \beta - \omega_l\|_2^2 + \frac{\gamma}{2} \|A_X \beta - Y\|_2^2 \right\}.$$

The explicit solution is given by

$$\beta^{(k+1)} = \left\{ \sum_{l=1}^{N^2} (\alpha_l D_l^T D_l + \gamma A_X^T A_X) \right\}^{-1} \\ \times \left\{ \sum_{l=1}^{N^2} (\mathbf{v}_l^T D_l + \alpha_l D_l^T \omega_l + \gamma A_X^T Y) \right\}.$$

Step 4. Given $\beta = \beta^{(k+1)}$, $\omega_l = \omega_l^{(k+1)}$, update $\mathbf{v}_l^{(k+1)}$ by using

$$\mathbf{v}_l^{(k+1)} = \mathbf{v}_l^{(k)} - \alpha_l (D_l \beta^{(k+1)} - \omega_l^{(k+1)}).$$

Step 5. Iterate Steps 2–4 until convergence.

2.3. The Error Bound

In this section, we establish the nonasymptotic error bound for the TV estimator $\hat{\beta}$ based on model (2). We consider two types

of distances to measure the error. The first one is a weighted L_2 distance such that

$$\|\hat{\beta} - \beta_0\|_{X,2} = \left\{ \mathbb{E}^* (\langle X_{n+1}, \hat{\beta} - \beta_0 \rangle^2) \right\}^{1/2},$$

where \mathbb{E}^* represents taking expectation with respect to (Y_{n+1}, X_{n+1}) only. The second one is the TV distance between $\hat{\beta}$ and β_0 , $\|\hat{\beta} - \beta_0\|_{\text{TV}}$.

We derive both error bounds by means of Haar wavelet basis. Various wavelet bases are commonly used to effectively represent images and the Haar wavelet is the simplest possible wavelet. The bivariate Haar wavelet basis for $L_2(\Omega)$ can be constructed as follows. Let $\phi^0(t) = I_{[0,1]}$ be the indicator function, and the mother wavelet $\phi^1(t) = 1$ for $t \in [0, 1/2)$ and -1 for $t \in [1/2, 1)$. Starting from the multivariate functions

$$\phi^d(s, t) = \phi^{d_1}(s)\phi^{d_2}(t), \quad d \in \{(0, 1), (1, 0), (1, 1)\},$$

the bivariate Haar basis functions include the indicator function $I_{[0,1]^2}$ and other functions

$$\begin{aligned} \phi_{j,k}^d(u, v) &= 2^j \phi^d(2^j x - k), \quad d \in \{(0, 1), (1, 0), (1, 1)\}, \\ x &= (u, v), \quad j \geq 0, \quad k \in \mathbb{Z}^2 \cap 2^j [0, 1]^2. \end{aligned}$$

The bivariate Haar wavelet basis is an orthonormal basis for $L_2[0, 1]^2$. Note that discrete images are isometric to the space $\mathcal{I}_N \subset L_2[0, 1]^2$ of piecewise constant functions $\mathcal{I}_N = \{f \in L_2[0, 1]^2 : f(s, t) = c_{jk}, \frac{j-1}{N} \leq s < \frac{j}{N}, \frac{k-1}{N} \leq t < \frac{k}{N}\}$ via the identification $c_{jk} = NX_{jk}$. Letting $N = 2^J$, the bivariate Haar basis restricted to the N^2 basis functions $\{I_{[0,1]^2}, \phi_{j,k}^d, j \leq J-1, d \in \{(0, 1), (1, 0), (1, 1)\}, k \in \mathbb{Z}^2 \cap 2^j [0, 1]^2\}$ forms an orthonormal basis for $\mathbb{R}^{N \times N}$. Denote by Φ the discrete bivariate Haar transformation and $\{\phi_l\}$ the Haar basis, in which $\Phi\beta \in \mathbb{R}^{N \times N}$ contains the bivariate Haar wavelet coefficients of β . Next, we review a theoretical result of Petrushev et al. (1999), who proved a deep and nontrivial result on $\text{BV}(\Omega)$. Specifically, it states that the Haar wavelet coefficients of $\beta_0 \in \text{BV}(\Omega)$ are in weak ℓ_1 . That is, if the Haar coefficients are sorted decreasingly according to their absolute values, then the l th rearranged coefficient is in absolute value less than $c\|\beta_0\|_{\text{BV}}/l$ with c being an absolute constant.

Invoking Haar wavelets is only for theoretical investigation and we do not estimate the Haar coefficients directly. We now introduce the main assumptions of this article:

- A1. Assume that the coefficient image β_0 in the space of $N \times N$ blocks of pixel values with bounded variation. Assume that the error ϵ is sub-Gaussian.
- A2. Assume that the discrete Haar representation of the image predictor X is $X = \sum_l \rho_l^{1/2} \xi_l \phi_l$, where ρ_l are positive constants and ξ_l are independently and identically distributed sub-Gaussian random variables with zero mean and unit variance.
- A3. For any $\beta \in \text{BV}(\Omega)$, write $\beta = \sum_l \gamma_l \phi_l$, where the γ_l are the Haar basis coefficients of β . We arrange γ_l in a decreasing order according to their absolute values and denote the sorted coefficients as $\gamma_{(l)}$. Assume that the corresponding sorted $\rho_{(l)}$ associated with the same basis function satisfies $c_1 s^{-2q} \leq \rho_{(s)} \leq c_2 s^{-2q}$ with $q > 1/2$ for each s and two positive constants c_1, c_2 .

Assumption A2 on the wavelet representation of X is reasonable because the discrete wavelet transformation approximately decorrelates or “whitens” data (Vidakovic 1999). Although we might use the Karhunen-Loève expansion of X , we do not adopt this approach to avoid additional complexity associated with the estimation of eigenfunctions. When we sort the Haar wavelet coefficients of both β and X , the corresponding basis functions may not follow the same order. Assumption A3 specifies the decay rate of the Haar wavelet coefficients of X . From A2, the predictor images X_i can be written as $X_i = \sum_l \rho_l^{1/2} \xi_{il} \phi_l$. Let \tilde{A} be an $n \times N^2$ matrix with the (i, l) th element being ξ_{il}/\sqrt{n} . It is well-known that \tilde{A} satisfies the restricted isometry property (RIP) with a large probability (Candès, Romberg, and Tao 2006a, 2006b). Specifically, if $n \geq C^{-2} s \log(N^2/s)$, then with probability exceeding $1 - 2e^{-Cn}$, we have

$$(1 - \delta)\|u\|_2^2 \leq \|\tilde{A}u\|_2^2 \leq (1 + \delta)\|u\|_2^2 \quad (6)$$

for all s -sparse vectors $u \in \mathbb{R}^{N^2}$ with a small RIP constant $\delta < C$.

Let $\{\hat{\gamma}_l\}$ and $\{\gamma_l\}$ be, respectively, the wavelet coefficients of $\hat{\beta}$ and β_0 . It turns out that $\|\hat{\beta} - \beta_0\|_{X,2} = \{\sum_l \rho_l (\hat{\gamma}_l - \gamma_l)^2\}^{1/2}$, which is the weighted L_2 -norm of the wavelet coefficient difference. On the other hand, since $\|\phi_l\|_{\text{TV}} \leq 8$ (Needell and Ward 2013),

$$\|\hat{\beta} - \beta_0\|_{\text{TV}} \leq \sum_l |\hat{\gamma}_l - \gamma_l| \|\phi_l\|_{\text{TV}} \leq 8\|\hat{\gamma} - \gamma\|_1,$$

which is bounded by the L_1 -norm of the wavelet coefficient difference. We obtain the following theorem, whose detailed proof can be found in the Appendix.

Theorem 2.1. Assumptions A1–A3 hold. Let C be an absolute constant and $\lambda = Cn^{1/2}$. If $n \geq Cs^{2q+1} \log(N^2/s^{2q+1})$ and $\delta < 1/3$ in (6), then with probability greater than $1 - 2\exp(-Cn)$, we have

$$\|\hat{\beta} - \beta_0\|_{X,2} \leq C \left\{ \sigma + \frac{1}{(s \log N)^{q+\frac{1}{2}}} \|\nabla \beta_0 - (\nabla \beta_0)_s\|_1 \right\}, \quad (7)$$

and

$$\begin{aligned} \|\hat{\beta} - \beta_0\|_{\text{TV}} &\leq C \log \left(\frac{N^2}{s} \right) \\ &\times \left\{ (s \log N)^{q+\frac{1}{2}} \sigma + \|\nabla \beta_0 - (\nabla \beta_0)_s\|_1 \right\}, \quad (8) \end{aligned}$$

where $(\nabla \beta_0)_s = \arg \min_{u:s\text{-sparse}} \|\nabla \beta_0 - u\|_1$ is the best s -sparse approximation to the discrete gradient $\nabla \beta_0$.

Theorem 2.1 provides nonasymptotic error bounds for $\|\hat{\beta} - \beta_0\|_{X,2}$ and $\|\hat{\beta} - \beta_0\|_{\text{TV}}$, which are specified explicitly in terms of sample size n and image size $N \times N$, and the underlying smoothness of the true coefficient image based on the discrete gradient.

Remark 2.1. We call a prediction “stable” if $\|\hat{\beta} - \beta_0\|_{X,2} \leq C\sigma$ holds with a high probability. Assume that the coefficient image has the sparse discrete gradient, that is, $\nabla \beta_0$ is supported on S_0 with $|S_0|_0 \leq s$. If $\lambda = Cn^{1/2}$, then **Theorem 2.1** shows that $\|\hat{\beta} - \beta_0\|_{X,2} \leq C\sigma$, which indicates that our prediction procedure is stable. Furthermore, for the extreme case with noiseless data, our prediction procedure is exact. The required sample size n is of

order $s^{2q+1} \log(N^2/s^{2q+1})$, which depends on the smoothness of the true coefficient image β_0 , the relative smoothness between β_0 and X , and the image size $N \times N$.

Remark 2.2. The parameter q characterizes the decay rate of the wavelet coefficients of X . The larger the q , the more the required sample size. [Theorem 2.1](#) also shows that the larger q is, the smaller the prediction error is. When $q = 0$, this gives the special case discussed in [Needell and Ward \(2013\)](#).

3. Generalized Scalar-on-Image Regression Models

In this section, we extend all developments for model (2) to GSIRM-(1). Given $X \in \mathbb{R}^{N \times N}$ and $Z \in \mathbb{R}^p$, the response variable Y is assumed to follow an exponential family distribution as

$$\exp \left(\{Y\eta(X, Z; \theta_0, \beta_0) - b(\eta(X, Z; \theta_0, \beta_0))\} / a(\psi) + c(y, \psi) \right), \tag{9}$$

where $a(\cdot)$, $b(\cdot)$, and $c(\cdot)$ are known functions, and ψ is either known or considered as a nuisance parameter. Our GSIRM-TV also assumes $\beta_0 \in \text{BV}(\Omega)$. It can be shown ([Nelder and Wedderburn 1972](#)) that

$$\begin{aligned} \mathbb{E}(Y|X) &= \mu(X, Z; \theta_0, \beta_0) = \dot{b}(\eta(X, Z; \theta_0, \beta_0)) \quad \text{and} \\ \text{var}(Y|X) &= a(\psi) \ddot{b}(\eta(X, Z; \theta_0, \beta_0)), \end{aligned}$$

where $\dot{b}(\eta)$ and $\ddot{b}(\eta)$ are, respectively, the first and second derivatives of $b(\eta)$ with respect to η . Moreover, $\eta(X, Z; \theta_0, \beta_0) = \dot{b}^{-1}(g^{-1}(\theta_0^T Z + \langle X, \beta_0 \rangle))$ is the canonical parameter of (9). A Gaussian distribution with variance σ^2 has $a(\psi) = \sigma^2$ and $b(\eta) = \eta^2/2$, a Bernoulli distribution has $a(\psi) = 1$ and $b(\eta) = \log(1 + e^\eta)$, and a Poisson distribution has $a(\psi) = 1$ and $b(\eta) = e^\eta$.

3.1. Estimation

Let $\xi = (\theta, \beta) \in \mathbb{R}^p \times \text{BV}(\Omega)$. Given the observed data, we propose to find estimates $\hat{\xi}$ by minimizing a penalized likelihood function given by

$$n^{-1} \sum_{i=1}^n \left\{ Y_i \eta(X_i, Z_i; \theta, \beta) - b(\eta(X_i, Z_i; \theta, \beta)) \right\} + \lambda \|\beta\|_{\text{TV}}. \tag{10}$$

We use an algorithm, which is a standard iteratively reweighted least squares for GLMs, modified to add a TV penalty, to calculate $\hat{\xi} = (\hat{\theta}, \hat{\beta})$. Given a trial estimate of ξ , denoted by $\hat{\xi}_t$, we introduce the iterative weights and the working dependent variable as

$$\hat{\omega}_{i,t} = \ddot{b}(\hat{\eta}_{i,t}) \quad \text{and} \quad \hat{Y}_{i,t} = g(\hat{\mu}_{i,t}) + (Y_i - \hat{\mu}_{i,t}) \dot{g}(\hat{\mu}_{i,t}), \tag{11}$$

where $\hat{\mu}_{i,t} = \mu(X_i, Z_i; \hat{\xi}_t)$, $\dot{g}(\mu) = dg(\mu)/d\mu$, and $\hat{\eta}_{i,t} = \eta(X_i, Z_i; \hat{\xi}_t)$. Then, we can calculate the next estimate of ξ , denoted by $\hat{\xi}_{t+1}$, by minimizing

$$\hat{\xi}_{t+1} = \underset{\xi}{\text{argmin}} \left\{ \sum_{i=1}^n \omega_{i,t} [\hat{Y}_{i,t} - \partial_\xi \mu_{i,t}(\hat{\xi}_t) \xi]^2 + \lambda \|\beta\|_{\text{TV}} \right\}, \tag{12}$$

where $\partial_\xi = \partial/\partial\xi$. The optimization in (12) can be effectively solved by using TVAL3 algorithm discussed in [Section 2](#). Finally, we can iteratively solve $\hat{\xi}_t$ until convergence.

We provide the complete algorithm as follows.

Step 1. Initialize $\xi^{(0)} = (\theta^{(0)}, \beta^{(0)})$.

Step 2. For each k , define the weights and the working dependent variable in (11), and define the objective function in (12). Use TVAL3 algorithm to solve for $\xi^{(k+1)} = (\theta^{(k+1)}, \beta^{(k+1)})$.

Step 3. Iterate Steps 2 and 3 until convergence.

We consider the logistic scalar-on-image regression model as an example. Specifically, Y_i given (X_i, Z_i) follows a Bernoulli distribution with the success probability p_i and $\text{logit}(p_i) = \langle X_i, \beta_0 \rangle + \theta_0^T Z_i$ for $i = 1, \dots, n$. Given the current estimate $\hat{\xi}_t$, it is easy to obtain the iterative weight and effective response variable, respectively, given by

$$\hat{\omega}_{i,t} = \frac{e^{\hat{\eta}_{i,t}}}{(1 + e^{\hat{\eta}_{i,t}})^2} \quad \text{and} \quad \hat{Y}_{i,t} = \hat{\eta}_{i,t} + \frac{Y_i - \hat{\mu}_{i,t}}{\hat{\mu}_{i,t}(1 - \hat{\mu}_{i,t})}.$$

Therefore, the estimate $\hat{\xi}_{t+1}$ can be obtained by solving a weighted penalized least squares in (12).

3.2. The Error Bound

We establish a nonasymptotic prediction error bound for GSIRM-TV. We need some additional assumptions as follows.

- B1. Assume $\eta(X, Z; \beta_0, \theta_0)$ is bounded almost surely. Given (X, Z) , the response Y is sub-Gaussian, that is, $\mathbb{E}\{\exp(t[Y - \dot{b}(\eta(X, Z; \beta_0, \theta_0))])|X\} \leq \exp(t^2 \bar{\sigma}^2/2)$ for some $\bar{\sigma}^2 > 0$ and all $t \in \mathbb{R}$.
- B2. The function $\dot{b}(\cdot)$ is monotonic with $\inf_t \ddot{b}(t) \geq c_3$ and $\sup_t \ddot{b}(t) \leq c_4$ for two positive constants c_3 and c_4 .

The sub-Gaussian assumption B1 holds for many well-known distributions, such as Gaussian. The assumption B2 requires that the second derivative of $b(\cdot)$ is bounded above and away from zero.

Theorem 3.1. Assumptions A1–A3 and B1–B2 hold. Let $\lambda = Cn^{-1/2}$, where C is a positive constant. If $n \geq Cs^{2q+1} \log(N^2/s^{2q+1})$ and $\delta < 1/3$ in (6), with probability greater than $1 - 2 \exp(-Cn)$, we have $\|\hat{\theta} - \theta_0\|_2 \leq Cn^{-1/2}$,

$$\|\hat{\beta} - \beta_0\|_{X,2} \leq C \left\{ 1 + \frac{1}{(s \log N)^{q+\frac{1}{2}}} \|\nabla \beta_0 - (\nabla \beta_0)_s\|_1 \right\}, \tag{13}$$

and

$$\begin{aligned} \|\hat{\beta} - \beta_0\|_{\text{TV}} &\leq C \log \left(\frac{N^2}{s} \right) \\ &\quad \times \left\{ (s \log N)^{q+\frac{1}{2}} \sigma + \|\nabla \beta_0 - (\nabla \beta_0)_s\|_1 \right\}. \end{aligned} \tag{14}$$

The conditional mean of Y_{n+1} given X_{n+1} is $\dot{b}(\eta(X_{n+1}, \beta_0))$. We may measure the accuracy of $\hat{\beta}$ by $\mathbb{E}^* [\dot{b}(\eta(X_{n+1}, \hat{\beta})) - \dot{b}(\eta(X_{n+1}, \beta_0))]^2$. Under B1, this risk is bounded by $\|\hat{\beta} - \beta_0\|_{X,2}^2$ and thus, it is reasonable to study the nonasymptotic behavior of $\|\hat{\beta} - \beta_0\|_{X,2}$. [Theorem 2.1](#) is a special case of [Theorem 3.1](#) if it is

assumed in [Theorem 2.1](#) that responses follow a normal distribution. If assuming that the coefficient image has the sparse discrete gradient, [Theorem 3.1](#) shows that $\|\hat{\beta} - \beta_0\|_{X,2}$ is bounded by a constant, which is proportional to σ under the assumption of [Theorem 2.1](#). This shows that our prediction procedure is stable for GSIRM-TV.

4. Simulation Studies

In this section, we conducted a set of Monte Carlo simulations to examine the finite sample performance of the TV estimate $\hat{\beta}$ and compare it with five competing methods. The first approach (Lasso) is to calculate the Lasso estimates of β_0 . The second one (Lasso-Haar) is to calculate the Lasso estimates of the Haar coefficients of β_0 and use the inverse discrete wavelet transform to calculate the estimates of β_0 . The third one (Matrix-Reg) is to estimate β_0 by using a recent development called regularized matrix regression ([Zhou and Li 2014](#)), which treats the coefficient image as a matrix and penalizes the nuclear norm of this matrix. The fourth one (FPCR) is the functional principal component regression approach ([Reiss and Ogden 2007, 2010](#)) by using tensor product cubic B-splines to approximate the coefficient function. The fifth one (WNET) is to perform scalar-on-image regression in the wavelet domain by naive elastic net ([Zhao, Ogden, and Reiss 2014](#)). Among these six approaches, the TV, Lasso, Lasso-Haar, and Matrix-Reg methods have been implemented by Matlab and the FPCR and WNET methods have been implemented in the R packages “refund” and “refund.wave” (see [Reiss et al. 2015](#)), respectively. For the FPCR and WNET methods, we have used the default settings of both packages. The choice of wavelet basis in WNET is the Daubechies basis.

We present some results based on linear scalar-on-image regression model (2). Specifically, X_i were simulated from a 64×64 phantom map with $N = 64$ and 4096 pixels according to a spatially correlated random process $X_i = \sum_l l^{-q/2} \xi_{il} \phi_l$ with $q = 0, 0.5$, and 1, where the ξ_l are standard normal random variables and the ϕ_l are bivariate Haar wavelet basis functions. We consider four different β_0 images including triangle, oval, T-shape, and checkerboard shapes ([Figure 3](#)). Among them, the triangle and oval images are convex, while the other two are not. Errors ϵ_i were independently generated from $N(0, 1)$. We set $n_1 = 300$ for the training set and $n_2 = 100$ for the test set. We repeated each setting 100 times. We calculated the root mean squared prediction error (RMSPE) to compare the finite sample performance of the six different estimation methods. Let

$\hat{\beta}$ be the estimated coefficient image from the training set and $\hat{Y}_i = \langle \hat{\beta}, X_i \rangle$ be the predicted responses for the test set. For each test set, RMSPE is defined by

$$\text{RMSPE} = \sqrt{n_2^{-1} \sum_{i=1}^{n_2} (\hat{Y}_i - Y_i)^2}.$$

We also calculated the means and standard errors of RMSPEs for the 100 testing datasets.

[Figures 4 and 5](#) present the estimated β_0 from a randomly selected training dataset with $q = 0$ and $q = 0.5$, respectively, for the sample size $n = 300$. For all four different shapes, our TV estimates can capture the sharp boundaries of the underlying shapes. In contrast, the Lasso method fails for all shapes, since the predictor images X_i are highly correlated. The Lasso estimates of the Haar coefficients can roughly capture the true shapes. However, this method cannot faithfully recover the sharp boundaries of the triangle, oval, and T shapes, whereas it does work very well for the checkerboard shape, since this checkerboard shape is exactly one of the bivariate Haar wavelet basis functions. The matrix regression approach can roughly capture the true shapes when $q = 0$, and unfortunately this method fails for the case when $q = 0.5$, for which the entries of X are spatially correlated. The PCR approach uses splines to approximate the predictor images, and it cannot preserve the sharp edges of coefficient estimator for our examples. The WNET method fails for the case when $q = 0$ but it can capture the shapes of the true coefficient image when the predictors are more spatially correlated.

[Table 1](#) presents the RMSPEs of all six methods across all shapes. Overall, our TV method has significantly smaller prediction errors, in particular for $q = 0$. It is expected that the Lasso method leads to the largest prediction error. For all these methods, the larger q is, the smaller are their RMSPEs. For a larger q which means the predictor images are more spatially correlated, the performances of the TV, Lasso-Haar, FPCR, and WNET are similar to each other.

5. Real Data Analysis

To illustrate the usefulness of our proposed model, we consider anatomical MRI data collected at the baseline by the Alzheimer’s Disease Neuroimaging Initiative (ADNI) study, which is a large




Figure 3. The true coefficient images used for the simulation study.

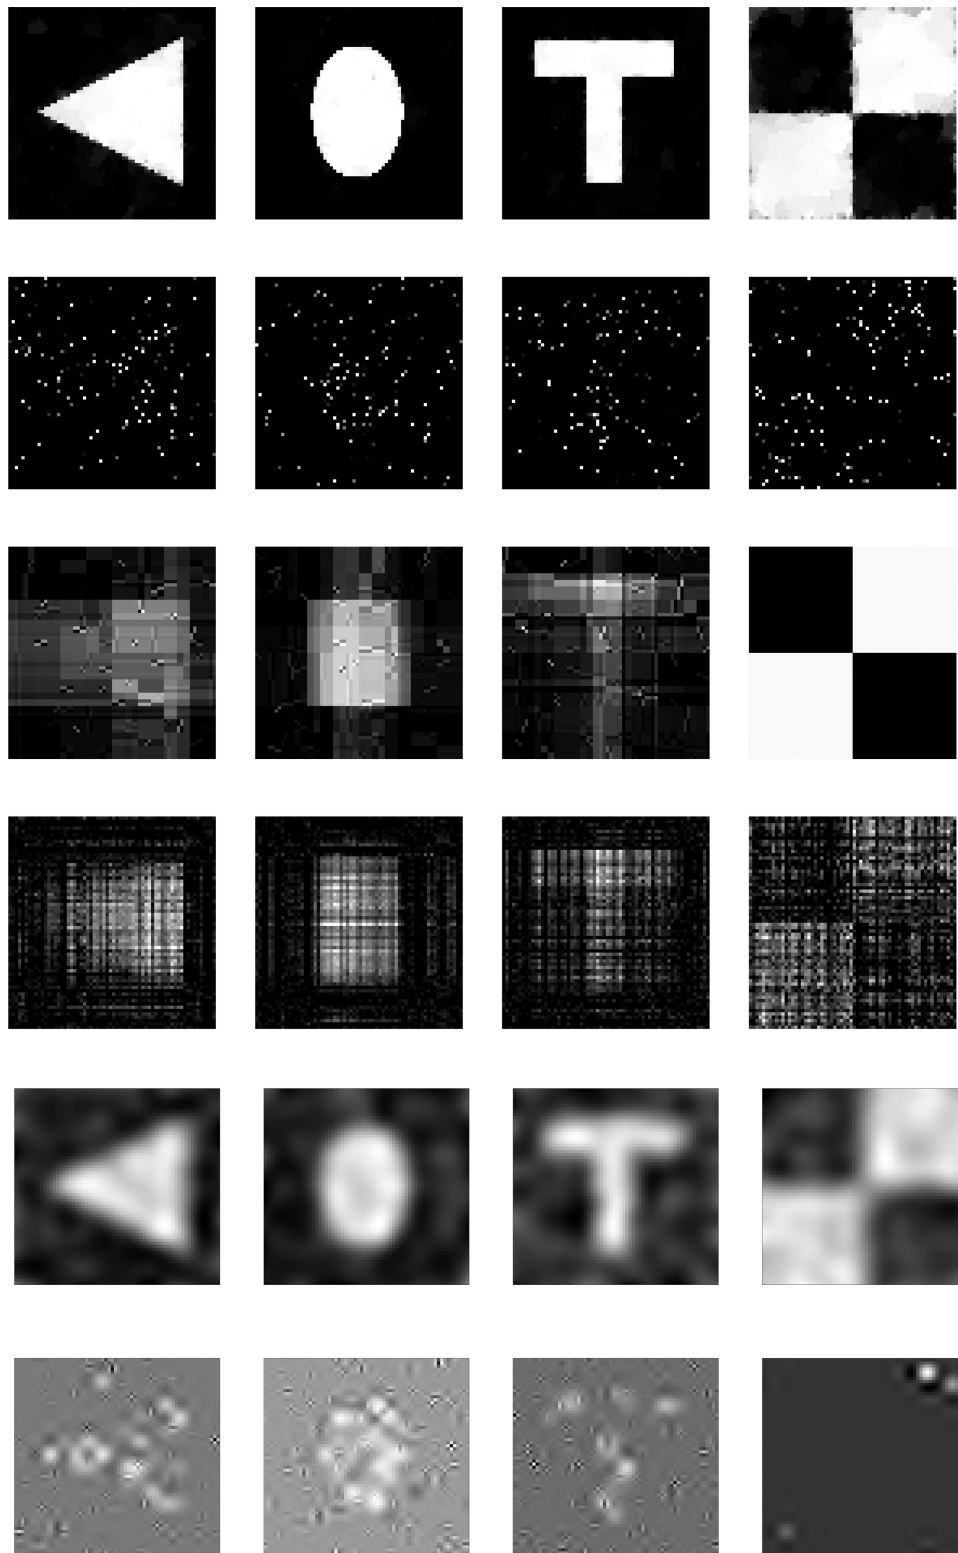


Figure 4. The estimated coefficient images from six estimation methods when $q = 0$ and $n = 300$: TV (top row); Lasso (second row); Lasso-Haar (third row); Matrix regression (fourth row); FPCR (fifth row); and WNET (sixth row).

scale multi-site study collecting clinical, imaging, and laboratory data at multiple time points from healthy controls, individuals with amnesic mild cognitive impairment, and subjects with Alzheimer’s disease (AD). “Data used in the preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). The ADNI was launched in 2003 by the National Institute on Aging (NIA),

the National Institute of Biomedical Imaging and Bioengineering (NIBIB), the Food and Drug Administration (FDA), private pharmaceutical companies, and nonprofit organizations, as a \$60 million, 5-year publicprivate partnership. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological

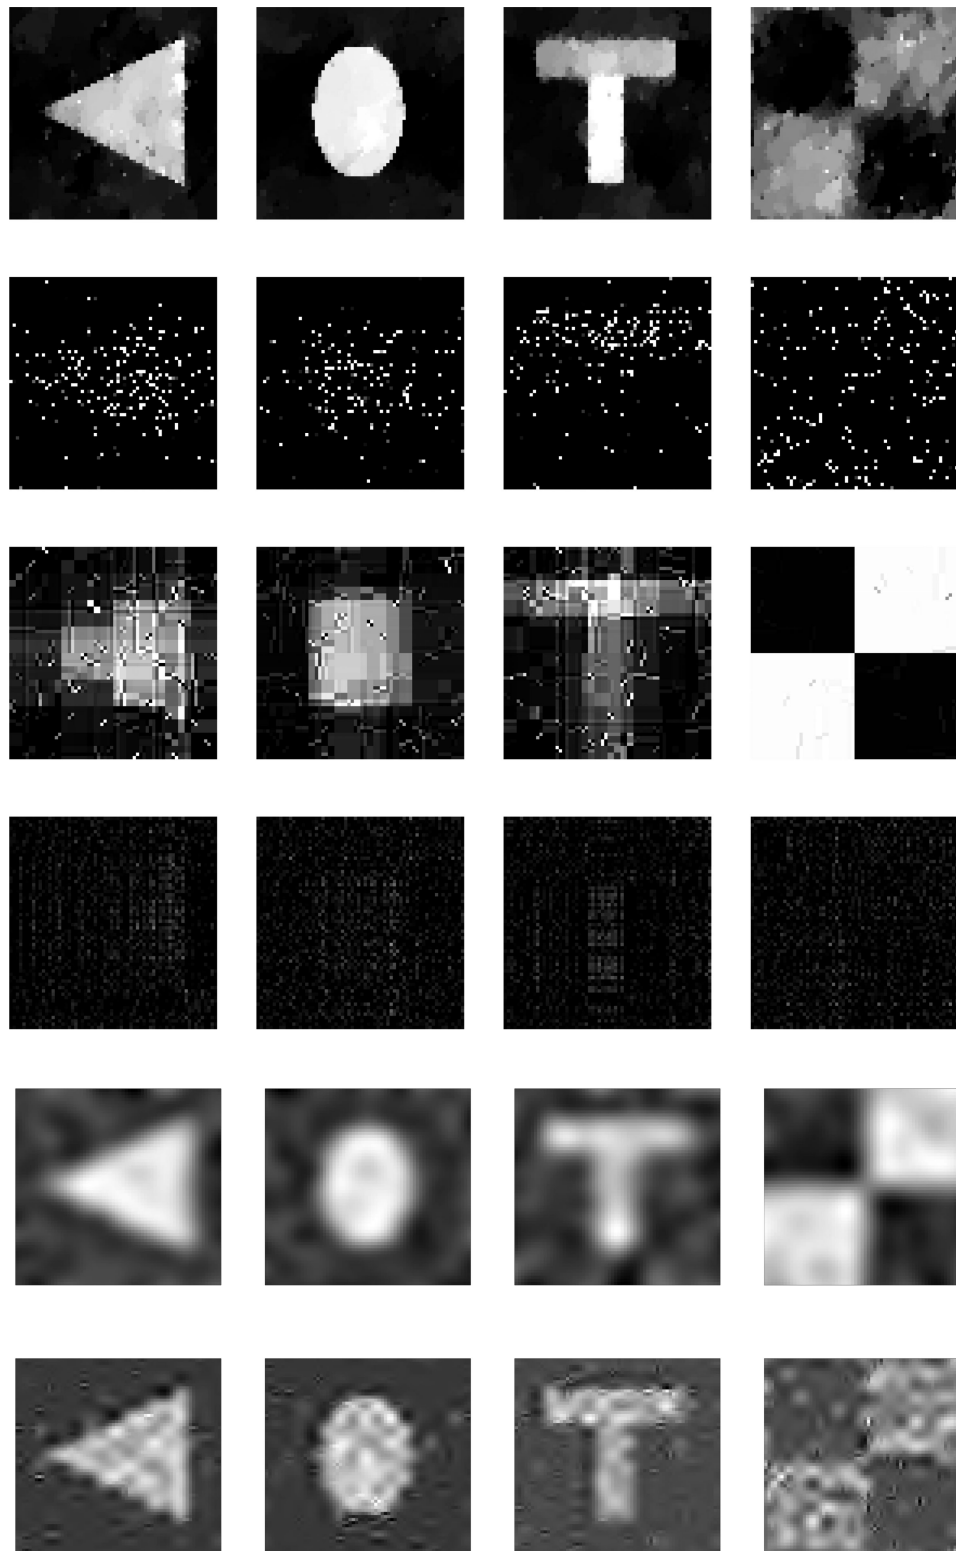


Figure 5. The estimated coefficient images from six methods when $q = 0.5$ and $n = 300$: TV (top row); Lasso (second row); Lasso-Haar (third row); Matrix regression (fourth row); FPCR (fifth row); and WNET (sixth row).

assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer's disease (AD). Determination of sensitive and specific markers of very early AD progression is intended to aid researchers and clinicians to develop new treatments and monitor their effectiveness, as well as lessen 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 coinvestigators 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 subjects but ADNI has been followed by ADNI-GO and ADNI-2. To date these three protocols have recruited over 1500 adults,




Figure 6. Observed left hippocampus images.

ages 55–90, to participate in the research, consisting of cognitively normal older individuals, people with early or late MCI, and people with early AD. The follow up duration of each group is specified in the protocols for ADNI-1, ADNI-2, and ADNI-GO. Subjects originally recruited for ADNI-1 and ADNI-GO had the option to be followed in ADNI-2. For up-to-date information, see www.adni-info.org.

Alzheimer’s disease as an age-related neurodegenerative brain disorder is often characterized by progressive loss in memory and deterioration of cognitive functions (De La Torre 2010; Weiner et al. 2012). Important neuropathological hallmarks of

AD are the gradual intraneuronal accumulation of neurofibrillary tangles formed as a result of abnormal hyperphosphorylation of cytoskeletal tau protein, extracellular deposition of amyloid- β ($A\beta$) protein as senile plaques, and massive neuronal death. These pathologies are evident in the hippocampus, which is located in the medial temporal lobe underneath the cortical surface, and other vulnerable brain areas. The hippocampus belongs to the limbic system and plays important roles in the consolidation of information from short-term memory to long-term memory and spatial navigation (Colom et al. 2013; Fennema-Notestine et al. 2009; Luders et al. 2013).

Table 1. The RMSPEs of six methods including TV, Lasso, Lasso-Haar, Matrix-Reg, FPCR, and WNET for four different shapes: the numbers in brackets are the corresponding standard errors of those RMSPEs.

q	TV			Lasso			Lasso-Haar		
	0	0.5	1	0	0.5	1	0	0.5	1
Triangle	3.20 (0.70)	1.53 (0.12)	1.18 (0.08)	34.02 (2.25)	13.12 (0.95)	3.96 (0.32)	18.76 (1.43)	5.08 (0.52)	1.98 (0.17)
Oval	1.69 (0.22)	1.49 (0.13)	1.22 (0.09)	31.16 (2.00)	11.83 (0.95)	3.57 (0.25)	14.99 (1.24)	3.82 (0.28)	1.67 (0.15)
T-shape	2.04 (0.51)	1.47 (0.09)	1.23 (0.10)	30.10 (2.41)	18.81 (0.81)	3.34 (0.28)	19.65 (1.84)	4.62 (0.46)	1.69 (0.16)
Checkerboard	6.27 (1.09)	1.45 (0.11)	1.10 (0.08)	49.00 (3.54)	25.46 (1.65)	9.08 (0.83)	1.43 (0.14)	1.06 (0.08)	1.05 (0.07)
q	Matrix-Reg			FPCR			WNET		
	0	0.5	1	0	0.5	1	0	0.5	1
Triangle	23.13 (1.78)	6.46 (0.55)	4.00 (0.34)	10.89 (0.76)	2.27 (0.18)	1.19 (0.09)	28.42 (2.45)	2.10 (0.17)	1.28 (0.10)
Oval	18.81 (1.58)	5.58 (0.37)	3.41 (0.30)	9.28 (0.68)	2.12 (0.15)	1.19 (0.09)	24.74 (2.40)	2.03 (0.16)	1.26 (0.10)
T-shape	20.97 (1.45)	5.49 (0.38)	3.57 (0.33)	11.69 (0.82)	3.42 (0.26)	1.64 (0.12)	24.41 (2.18)	2.52 (0.21)	1.46 (0.13)
Checkerboard	35.33 (2.89)	12.80 (0.95)	6.35 (0.38)	10.70 (0.83)	3.05 (0.22)	1.14 (0.10)	44.24 (3.23)	5.41 (0.55)	2.03 (0.15)

Given the MRI scans, hippocampal substructures were segmented with FSL FIRST (Patenaude et al. 2011) and hippocampal surfaces were automatically reconstructed with the marching cube method (Lorensen and Cline 1987). We adopted a surface fluid registration based hippocampal subregional analysis package (Shi et al. 2013), which uses isothermal coordinates and fluid registration to generate one-to-one hippocampal surface registration for surface statistics computation. It introduced two cuts on a hippocampal surface to convert it into a genus zero surface with two open boundaries. The locations of the two cuts were at the front and back of the hippocampal surface. By using conformal parameterization, it essentially converts a three-dimensional surface registration problem into a two-dimensional image registration problem. The flow induced in the parameter domain establishes high-order correspondences between three-dimensional surfaces. Finally, various surface statistics were computed on the registered surface, such as multivariate tensor-based morphometry (mTBM) statistics (Wang et al. 2010), which retain the full tensor information of the deformation Jacobian matrix, together with the radial distance (Pizer et al. 1999). This software package and associated image processing methods have been adopted and described by various studies (Shi et al. 2014).

We applied GSIRM-TV to the hippocampus dataset calculated from ADNI. The sample in our investigation includes $n = 403$ subjects: 223 healthy controls (HC) (107 females and 116 males) and 180 individuals with AD (87 females and 93 males). We consider binary disease status with 0 being HC and 1 being AD as responses. The image predictor X_i is the two-dimensional representation of left hippocampus. Figure 6 displays some observed left hippocampus images. The covariate vector Z_i includes constant(=1), gender (Female = 0 and Male = 1), age (55–92), and behavior score (1–36). Given

Table 2. ADNI hippocampus dataset: the estimated coefficients of the four scalar covariates and their standard deviations in parentheses.

	Intercept	Sex	Age	Behavior score
$\hat{\theta}$	−1.807 (3.186)	−0.533 (0.590)	−0.093 (0.043)	0.869 (0.111)

(X_i, Z_i) , Y_i is assumed to follow a Bernoulli distribution with the success probability p_i satisfying

$$\text{logit}(p_i) = \langle X_i, \beta_0 \rangle + \theta_0^T Z_i \quad \text{for } i = 1, \dots, n.$$

We used the iterative reweighted algorithm described above to estimate the unknown parameters.

Table 2 presents the estimates of θ_0 and their corresponding standard deviations, which were calculated by using the bootstrap method. Figure 7 shows the estimated coefficient images by using the five estimation methods. The effects around pixels (5, 40), (40, 40), (95, 40) seem to be captured well by our TV estimate. The confidence band for the coefficient image can also be obtained by using the bootstrap method. We randomly partitioned the hippocampus dataset into a training set with $n_1 = 203$ and a test set with $n_2 = 200$. We repeated this random partition for 100 times and computes 100 classification errors. The average classification error of TV is 8.13% with a standard error 1.56%. We also obtain the average classification errors for other five methods. The average classification errors are 12.23%(7.36%), 21.65%(15.56%), 12.03%(11.55%), 17.13%(3.27%), 16.45%(15.57%), respectively, for Lasso, Lasso-Haar, matrix regression, FPCR, and WNET. For the WNET method, since the R code requires the image size to be a power of 2, we have added zeros to make the image size of 256×256 as suggested by one of the referees. Inspecting Table 2 reveals that sex and age are not significant in GSIRM-TV. We run the same procedure without sex and age and




Figure 7. Estimated coefficient images for hippocampus data based four methods: the two-dimensional representation of TV estimator (a) and the surface representation of TV estimator (b), Lasso estimator (c), Lasso-wavelet estimator (d), matrix regression estimator (e), FPCR estimator (f), and WNET estimator (g).

obtained a similar classification result as the full model, which is omitted from the article.

6. Conclusion

We have developed a class of GSIRM-TV's for scalar response and imaging and/or scalar predictors, while explicitly assuming that its slope function belongs to $BV(\Omega)$. We have developed an efficient penalized total variation minimization to estimate the coefficient image. We have used simulations and real data analysis to show that GSIRM-TV is quite efficient for estimating the slope function, while preserving its edges and jumps. We have established the nonasymptotic error bound of the TV estimate for the excess risk.

It is known that many image data have small total variation and are compressible with respect to wavelet transform. Therefore, we may generalize our approach to include both, total variation penalty and Lasso penalty on the wavelet coefficients. Specifically, let Φ be the wavelet transformation operator and γ be the wavelet coefficients of the coefficient image β_0 . We may calculate γ by minimizing

$$\sum_{i=1}^n \left(Y_i - \langle X_i, \Phi^{-1} \gamma \rangle \right)^2 + \lambda_1 \|\Phi^{-1} \gamma\|_{TV} + \lambda_2 \|\gamma\|_1, \quad (15)$$

where Φ^{-1} is the inverse discrete wavelet transform, and $\beta = \Phi^{-1} \gamma$. In (15), there are two smoothing parameters λ_1 and λ_2 which need to be selected. Efficient algorithm is also needed to be developed to solve (15). We leave this as further research work.

We have so far focused on two-dimensional images. It would be interesting to extend our method to analyze k -dimensional (k -D) images for $k \geq 2$ (Zhou, Li, and Zhu 2013; Zhu, Fan, and Kong 2014). For example, consider a 3-D image $f \in \mathbb{R}^{N^3}$, where $f = (f_e)$, in which $e = (e_1, e_2, e_3) \in \{1, 2, 3\}^3$. The inner product can be defined as

$$\langle f, g \rangle = \sum_{e \in \{1, 2, 3\}^3} f_e \cdot g_e.$$

For $\ell = 1, 2$, and 3, the discrete derivative of f in the direction of r_ℓ is $f_{r_\ell} \in \mathbb{R}^{N^{\ell-1} \times N \times N^{3-\ell}}$,

$$(f_1)_e = f_{(e_1+1, e_2, e_3)} - f_{(e_1, e_2, e_3)}, (f_2)_e = f_{(e_1, e_2+1, e_3)} - f_{(e_1, e_2, e_3)}, \\ (f_3)_e = f_{(e_1, e_2, e_3+1)} - f_{(e_1, e_2, e_3)},$$

and the three-dimensional discrete gradient is $(\nabla f)_e = (f_{r_\ell})_e$ for $e_\ell \leq N - 1$ and zero elsewhere. Hence, the three-dimensional anisotropic and isotropic total variation seminorm can be defined similarly. We may consider a similar total variation optimization (4) to estimate the three-dimensional coefficient image. This research is currently under investigation and will be presented in another report.

Supplementary Materials

The online supplementary materials contain the appendices for the article.

Acknowledgments

Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.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. The authors would like to thank Dr. Yalin Wang for sharing the processed data. We would also like to thank the Editor, Associate Editor, and referees for constructive comments.

Funding

The research of Xiao Wang is supported by NSF grants DMS1042967 and CMMI1030246. The research of Hongtu Zhu is partially supported by NIH grants MH086633 and 1UL1TR001111, NSF grants SES-1357666 and DMS-1407655, and a grant from Cancer Prevention Research Institute of Texas. This material was based upon work partially supported by the NSF grant DMS-1127914 to the Statistical and Applied Mathematical Science Institute.

References

Candès, W., Romberg, J., and Tao, T. (2006a), "Stable Signal Recovery From Incomplete and Inaccurate Measurements," *Communications on Pure and Applied Mathematics*, 59, 1027–1023. [1158,1160]

— (2006b), "Robust Uncertainty Principles: Exact Signal Reconstruction From Highly Incomplete Frequency Information," *IEEE Transactions on Information Theory*, 52, 489–509. [1158,1160]

Chen, S., Donoho, D. L., and Saunders, M. (1998), "Atomic Decomposition for Basis Pursuit," *SIAM Journal on Scientific Computing*, 20, 33–61. [1156]

Colom, R., Stein, J. L., Rajagopalan, P., Martínez, K., Hermel, D., Wang, Y., Alvarez Linera, J., Burgaleta, M., A., Quiroga, M., Shih, P. C., and Thompson, P. M. (2013), "Hippocampal Structure and Human Cognition: Key Role of Spatial Processing and Evidence Supporting the Efficiency Hypothesis in Females," *Intelligence*, 41, 129–140. [1165]

Crambes, C., Kneip, A., and Sarda, P. (2009), "Smoothing Splines Estimators for Functional Linear Regression," *Annals of Statistics*, 37, 35–72. [1156]

De La Torre, J. C. (2010), "Alzheimer's Disease is Incurable But Preventable," *Journal of Alzheimer's Disease*, 20, 861–870. [1165]

Du, P., and Wang, X. (2014), "Penalized Likelihood Functional Regression," *Statistica Sinica*, 24, 1017–1041. [1156]

Fan, J., and Li, R. (2001), "Variable Selection via Nonconcave Penalized Likelihood and Its Oracle Properties," *Journal of the American Statistical Association*, 96, 1348–1360. [1156]

Fennema-Notestine, C., Hagler, D. J. Jr, McEvoy, L. K., Fleisher, A. S., Wu, E. H., Karow, D. S., Dale, A. M.; Alzheimer's Disease Neuroimaging Initiative. (2009), "Structural MRI Biomarkers for Preclinical and Mild Alzheimer's Disease," *Human Brain Mapping*, 30, 3238–3253. [1165]

Ferraty, F., and Vieu, P. (2006), *Nonparametric Functional Data Analysis: Theory and Practice*, New York: Springer-Verlag Inc. [1156]

Friedman, J. T., Hastie, H., and Tibshirani, R. (2007), "Pathwise Coordinate Optimization," *Annals of Applied Statistics*, 1, 302–332. [1157]

Gertheiss, J., Maity, A., and Staicu, A. M. (2013), "Variable Selection in Generalized Functional Linear Model," *Stat*, 2, 86–101. [1157]

Goldsmith, J., Bobb, J., Crainiceanu, C. M., Caffo, B., and Reich, D. (2010), "Penalized Functional Regression," *Journal of Computational and Graphical Statistics*, 20, 830–851. [1156,1157]

Guillas, S., and Lai, M. J. (2010), "Bivariate Splines for Spatial Functional Regression Models," *Journal of Nonparametric Statistics*, 22, 477–497. [1157]

Hall, P., and Horowitz, J. L. (2007), "Methodology and Convergence Rates for Functional Linear Regression," *Annals of Statistics*, 35, 70–91. [1156]

James, G. M. (2002), "Generalized Linear Models With Functional Predictors," *Journal of the Royal Statistical Society, Series B*, 64, 411–432. [1156]

- James, G. M., Wang, J., and Zhu, J. (2009), “Functional Linear Regression That’s Interpretable,” *Annals of Statistics*, 37, 2083–2108. [1156,1157]
- Hestenes, M. R. (1969), “Multiplier and Gradient Methods, Journal of Optimization Theory and Applications,” in *Computing Methods in Optimization Problems* (vol. 4), eds. L. A. Zadeh, L. W. Neustadt, and A. V. Balakrishnan, New York: Academic Press, pp. 303–320 [1159]
- Li, Y., Wang, N., and Carroll, R. J. (2010), “Generalized Functional Linear Models With Semi Parametric Single-Index Interactions,” *Journal of the American Statistical Association*, 105, 621–633. [1156]
- Li, C. (2013), “Compressive Sensing for 3D Data Processing Tasks: Applications, Model and Algorithms,” Ph.D. Dissertation, Rice University. [1159]
- Lorensen, W. E., and Cline, H. E. (1987), “Marching Cubes: A High Resolution 3D Surface Construction Algorithm,” in *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 87*, pp. 163–169.0 [1166]
- Luders, E., Thompson, P. M., Kurth, F., Hong, J. Y., Phillips, O. R., Wang, Y., Gutman, B. A., Chou, Y. Y., Narr, K. L., and Toga, A. W. (2013), “Global and Regional Alterations of Hippocampal Anatomy in Long-Term Meditation Practitioners,” *Human Brain Mapping*, 34, 3369–3375. [1165]
- Mammen, E., and van de Geer, S. (1997), “Locally Adaptive Regression Splines,” *Annals of Statistics*, 25, 387–413. [1158]
- Michel, V., Gramfort, A., Varoquaux, G., Eger, E., and Thirion, B. (2011), “Total Variation Regularization for fMRI-Based Prediction of Behavior,” *IEEE Transactions on Medical Imaging*, 30, 1328–1340. [1157,1158]
- Mu, Y., and Gage, F. (2011), “Adult Hippocampal Neurogenesis and Its Role in Alzheimer’s Disease,” *Molecular Neurodegeneration*, 6, 85. [1156]
- Müller, H. G., and Stadtmüller, U. (2005), “Generalized Functional Linear Models,” *Annals of Statistics*, 33, 774–805. [1156]
- Needell, D., and Ward, R. (2013), “Stable Image Reconstruction Using Total Variation Minimization,” *SIAM Journal of Imaging Sciences*, 6, 1035–1058. [1158,1160,1161]
- Nelder, J., and Wedderburn, R. (1972), “Generalized Linear Models,” *Journal of the Royal Statistical Society, Series A*, 135, 370–384. [1161]
- Patenaude, B., Smith, S. M., Kennedy, D. N., and Jenkinson, M. (2011), “A Bayesian Model of Shape and Appearance for Subcortical Brain Segmentation,” *NeuroImage*, 56, 907–922. [1166]
- Petrushev, P. P., Cohen, A., Xu, H., and DeVore, R. (1999), “Nonlinear Approximation and the Space $BV(R^2)$,” *American Journal of Mathematics*, 121, 587–628. [1160]
- Pizer, S., Fritsch, D., Yushkevich, P., Johnson, C., and Chaney, E. (1999), “Segmentation, Registration, and Measurement of Shape Variation via Image Object Shape,” *IEEE Transactions on Medical Imaging*, 18, 851–865. [1166]
- Powell, M. J. D. (1969), “A Method for Nonlinear Constraints in Minimization Problems,” in *Optimization*, ed. R. Fletcher, London, New York: Academic Press, pp. 283–298. [1159]
- Ramsay, J. O., and Silverman, B. W. (2005), *Functional Data Analysis*, New York: Springer-Verlag Inc. [1156]
- Reiss, P. T., Huo, L., Zhao, Y., Kelly, C., and Ogden, R. T. (2015), “Wavelet-Domain Regression and Predictive Inference in Psychiatric Neuroimaging,” *Annals of Applied Statistics*, 9, 1076–1101. [1157,1162]
- Reiss, P. T., and Ogden, R. T. (2007), “Functional Principal Component Regression and Functional Partial Least Squares,” *Journal of the American Statistical Association*, 102, 984–996. [1156,1162]
- (2010), “Functional Generalized Linear Models With Images as Predictors,” *Biometrics*, 66, 61–69. [1156,1157,1162]
- Rudin, L. I., and Osher, S. (1994), “Total Variation Based Image Restoration With Free Local Constraints,” *Proceedings of the 1st IEEE ICIP*, 1, 31–35. [1157]
- Rudin, L. I., Osher, S., and Fatemi, E. (1992), “Nonlinear Total Variation Noise Removal Algorithm,” *Physica D*, 60, 259–268. [1157]
- Shi, J., Lepore, N., Gutman, B., Thompson, P. M., Baxter, L., Caselli, R. J., and Wang, Y. (2014), “Genetic Influence of APOE4 Genotype on Hippocampal Morphometry—An N=725 Surface-Based ADNI Study,” *Human Brain Mapping*, 35, 3902–3918. [1166]
- Shi, J., Thompson, P. M., Gutman, B., and Wang, Y. (2013), “Surface Fluid Registration of Conformal Representation: Application to Detect Disease Burden and Genetic Influence on Hippocampus,” *NeuroImage*, 78, 111–134. [1166]
- Tibshirani, R. (1996), “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [1156]
- (2014), “Adaptive Piecewise Polynomial Estimation via Trend Filtering,” *Annals of Statistics*, 42, 285–323. [1158]
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. (2005), “Sparsity and Smoothness via the Fused Lasso,” *Journal of the Royal Statistical Society, Series B*, 67, 91–108. [1157]
- Vidakovic, B. (1999), *Statistical Modeling by Wavelets*, New York: Wiley. [1160]
- Wang, X., Nan, B., Zhu, J., Koppe, R., and ADNI (2014), “Regularized 3D Functional Regression for Brain Image Data via Haar Wavelets,” *Annals of Applied Statistics*, 8, 1045–1064. [1157]
- Wang, Y., Zhang, J., Gutman, B., Chan, T. F., Becker, J. T., Aizenstein, H. J., Lopez, O. L., Tamburo, R. J., Toga, A. W., and Thompson, P. M. (2010), “Multivariate Tensor Based Morphometry on Surfaces: Application to Mapping Ventricular Abnormalities in HIV/AIDS,” *NeuroImage*, 49, 2141–2157. [1166]
- Weiner, M. W., Veitcha, D. P., Aisen, P. S., Beckett, L. A., Cairns, N. J., Green, R. C., Harvey, D., Jack, C. R., Jagust, W., Liu, E., Morris, J. C., Petersen, R. C., Saykino, A. J., Schmidt, M. E., Shaw, L., Siu-ciak, J. A., Soares, H., Toga, A. W., Trojanowski, J. Q., and ADNI (2012), “The Alzheimer’s Disease Neuroimaging Initiative: A Review of Papers Published Since Its Inception,” *Alzheimers and Dementia*, 8, S1–S68. [1156,1165]
- Yuan, M., and Cai, T. T. (2010), “A Reproducing Kernel Hilbert Space Approach to Functional Linear Regression,” *Annals of Statistics*, 38, 3412–3444. [1156]
- Zhao, Y., Ogden, R. T., and Reiss, P. T. (2014), “Wavelet-Based LASSO in Functional Linear Regression,” *Journal of Computational and Graphical Statistics*, 21, 600–617. [1162]
- Zhou, H., and Li, L. (2014), “Regularized Matrix Regression,” *Journal of Royal Statistical Society, Series B*, 76, 463–483. [1158,1162]
- Zhou, H., Li, L., and Zhu, H. (2013), “Tensor Regression with Applications in Neuroimaging Data Analysis,” *Journal of the American Statistical Association*, 108, 540–552. [1157,1167]
- Zhu, H., Fan, J., and Kong, L. (2014), “Spatially Varying Coefficient Model for Neuroimaging Data With Jump Discontinuities,” *Journal of the American Statistical Association*, 109, 977–990. [1167]