

Stable Feature Selection from Brain sMRI

Bo Xin*, Lingjing Hu†, Yizhou Wang* and Wen Gao*

*National Engineering Laboratory for Video Technology, Key Laboratory of Machine Perception,
School of EECS, Peking University, Beijing, 100871, China

†Yanjing Medical College, Capital Medical University, Beijing, 101300, China

Abstract

Neuroimage analysis usually involves learning thousands or even millions of variables using only a limited number of samples. In this regard, sparse models, e.g. the lasso, are applied to select the optimal features and achieve high diagnosis accuracy. The lasso, however, usually results in independent unstable features. Stability, a manifest of reproducibility of statistical results subject to reasonable perturbations to data and the model (Yu 2013), is an important focus in statistics, especially in the analysis of high dimensional data. In this paper, we explore a *nonnegative generalized fused lasso model* for stable feature selection in the diagnosis of Alzheimer’s disease. In addition to sparsity, our model incorporates two important pathological priors: the spatial cohesion of lesion voxels and the positive correlation between the features and the disease labels. To optimize the model, we propose an efficient algorithm by proving a novel link between total variation and fast network flow algorithms via conic duality. Experiments show that the proposed nonnegative model performs much better in exploring the intrinsic structure of data via selecting stable features compared with other state-of-the-arts.

Introduction

Neuroimage analysis is challenging due to its high feature dimensionality and data scarcity. Sparse models such as the lasso (Tibshirani 1996) have gained great reputation in statistics and machine learning, and they have been applied to the analysis of such high dimensional data by exploiting the sparsity property in the absence of abundant data. As a major result, automatic selection of relevant variables/features by such sparse formulation achieves promising performance. For example, in (Liu, Zhang, and Shen 2012), the lasso model was applied to the diagnosis of Alzheimer’s disease (AD) and showed better performance than the support vector machine (SVM), which is one of the state-of-the-arts in brain image classification. However, in statistics, it is known that the lasso does not always provide interpretable results because of its instability (Yu 2013). “Stability” here means the reproducibility of statistical results subject to reasonable perturbations to data and

the model. (These perturbations include the often used Jackknife, bootstrap and cross-validation.) This unstable behavior of the lasso model is critical in high dimensional data analysis. The resulting irreproducibility of the feature selection are especially undesirable in neuroimage analysis/diagnosis. However, unlike the problems such as registration and classification, the *stability issue* of feature selection is much less studied in this field.

In this paper we propose a model to induce more stable feature selection from high dimensional brain structural Magnetic Resonance Imaging (sMRI) images. Besides sparsity, the proposed model harnesses two important additional pathological priors in brain sMRI: (i) the spatial cohesion of lesion voxels (via inducing fusion terms) and (ii) the positive correlation between the features and the disease labels. The correlation prior is based on the observation that in many brain image analysis problems (such as AD, frontotemporal dementia, corticobasal degeneration, etc), there exist strong correlations between the features and the labels. For example, gray matter of AD is degenerated/atrophied. Therefore, the gray matter values (indicating the volume) are positively correlated with the cognitive scores or disease labels $\{-1, 1\}$. That is, the less gray matter, the lower the cognitive score. Accordingly, we propose nonnegative constraints on the variables to enforce the prior and name the model as “non-negative Generalized Fused Lasso” ($n^2\text{GFL}$). It extends the popular generalized fused lasso and enables it to explore the intrinsic structure of data via selecting stable features. To measure feature stability, we introduce the “Estimation Stability” recently proposed in (Yu 2013) and the (multi-set) Dice coefficient (Dice 1945). Experiments demonstrate that compared with existing models, our model selects much more stable (and pathological-prior consistent) voxels. It is worth mentioning that the non-negativeness per se is a very important prior of many practical problems, e.g. (Lee and Seung 1999). Although $n^2\text{GFL}$ is proposed to solve the diagnosis of AD in this work, the model can be applied to more general problems.

Incorporating these priors makes the problem novel w.r.t the lasso or generalized fused lasso from an optimization standpoint. Although off-the-shelf convex solvers such as CVX (Grant and Boyd 2013) can be applied to solve the optimization, it hardly scales to high-dimensional problems in feasible time. In this regard, we propose an efficient algo-

rithm that solves the $n^2\text{GFL}$ problem exactly. We generalize the proximal gradient methods (such as FISTA) (Beck and Teboulle 2009) to solve our constrained optimization and prove its convergence. We then show that by using an element-wise post-processing, the resulting proximal operator can be reduced to the total variation (TV) problem. It is known that TV can be solved by parametric flow algorithms (Chambolle and Darbon 2009; Xin et al. 2014). In the present study, we provide a novel equivalence via conic duality, which gives us a minimum quadratic cost flow formulation (Hochbaum and Hong 1995). Fast flow algorithms (including parametric flow) are then easily applied. In practice, our algorithm runs hundreds of times faster than CVX at the same precision and can scale to high-dimensional problems.

Related work. In addition to sparsity, people leverage underlying data structures and introduce stronger priors such as the structured sparsity (Jacob, Obozinski, and Vert 2009) to increase model stability. However, for voxel-based sMRI data analysis, handcrafted grouping of the voxels or sub-structures may not coincide with various pathological topology priors. Consequently, group lasso (with overlap) (Jacob, Obozinski, and Vert 2009; Jenatton et al. 2012; Rao et al. 2013) is not an ideal model to the problem. In contrast, the graph-based structured sparse models adapt better to such a situation. The most popular one is referred here as LapL¹, which adopts l_2 norm regularization of neighborhood variable difference (e.g. (Ng and Abugharbieh 2011; Grosenick et al. 2013)). However, as we will show in the experiments, these models select many more features than necessary. Very recently, generalized fused lasso or total variation has been successfully applied to brain image analysis problems inducing the l_1 difference (Gramfort, Thirion, and Varoquaux 2013; Xin et al. 2014). In the experiments, we show that by including an extra nonnegative constraint, the features selected by our model is much more stable than that of such unconstrained models. A very recent work (Avants et al. 2014) also explored this positive correlation (partially supporting our assumption), but the problem formulation was quite different: neither structural assumption was considered, nor the stability of feature selection was discussed. From the optimization standpoint, the applied framework is similar to that of (Xin et al. 2014) but two key differences exist: (1) the FISTA and soft-thresholding process applied in (Xin et al. 2014) do not generalize to constrained optimization problems, we show important modifications and provide theoretical proof; (2) we propose a novel understanding of TV's relation with flow problems via conic duality and prove that the minimum norm point problem of (Xin et al. 2014) is a special case of our framework.

The Proposed Method

Nonnegative Generalized Fused Lasso ($n^2\text{GFL}$)

Let $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$ be a set of samples, where $\mathbf{x}_i \in \mathbb{R}^d$ and $y_i \in \mathbb{R}$ are features and labels, respectively. Also, we denote

¹Although different names are given in e.g. (Ng and Abugharbieh 2011; Grosenick et al. 2013), they are in fact fundamentally applying the graph Laplacian smoothing.

by $\mathbf{X} \in \mathbb{R}^{d \times N}$ and $\mathbf{y} \in \mathbb{R}^N$ the concatenations of \mathbf{x}_i and y_i . Then, we consider the formulation

$$\begin{aligned} \min_{\beta \in \mathbb{R}^d} \quad & l(\beta; \mathbf{X}, \mathbf{y}) + \lambda_1 \sum_{i=1}^d |\beta_i| + \lambda_2 \sum_{(i,j) \in E} w_{ij} |\beta_i - \beta_j|, \\ \text{s.t.} \quad & \beta \geq \mathbf{0} \end{aligned} \quad (1)$$

where $\lambda_1, \lambda_2 \geq 0$ are tuning parameters. l is a loss term of variable β (assumed to be convex and smooth). w_{ij} s are pre-defined weights. Here, the variables (e.g. the sMRI voxel parameters in the AD problem) are supposed to have certain underlying structure represented by a graph $G = (V, E)$ with nodes V and edges E . Each variable corresponds to a node on the graph. As mentioned above, in many brain image analysis, there exist strong directional correlations (positive or negative) between the features and the labels, thus we assume $\beta \geq \mathbf{0}$ (or $\beta \leq \mathbf{0}$). Due to the l_1 penalties on each variable as well as each adjacent pair of variables in (1), solutions tend to be both sparse and smooth, i.e., adjacent variables tend to be similar and spatially coherent. Also because we have added the nonnegative constraints, the model will not select negatively correlated features as support. In practice, we notice that unconstrained models will systematically select many negatively correlated features. The nonnegative constraints greatly reduce these falsely recovered variables and encourage genuine disease-related features to be selected.

Efficient Optimization of $n^2\text{GFL}$

The optimization of $n^2\text{GFL}$ is convex and off-the-shelf solver such as CVX can be applied. However, this solution hardly scales to a problem sized of thousands (mainly due to its choice of general second order frameworks), see Table 1. In this regard, we propose certain modifications to scalable first order methods (e.g. accelerated proximal methods (Beck and Teboulle 2009)). This is done by exploring Lagrange multiplier method to deal with the constraints. From the optimization standpoint, these modifications are non-trivial and compose one major contribution of this work.

We first extend the (fast) iterative shrinkage thresholding algorithm (ISTA and FISTA) (Beck and Teboulle 2009) as follows.

Proposition 1. Let β^* be the optimal solution to (1) and β^k defined as follows

$$\begin{aligned} \beta^{k+1} = \min_{\beta \in \mathbb{R}^d} \quad & \frac{1}{2} \|\beta - \mathbf{z}^k\|_2^2 + \frac{\lambda_1}{L} \sum_{i=1}^d |\beta_i| + \\ & \frac{\lambda_2}{L} \sum_{(i,j) \in E} w_{ij} |\beta_i - \beta_j|, \quad \text{s.t. } \beta \geq \mathbf{0}, \end{aligned} \quad (2)$$

where $L > 0$ is the Lipschitz constant of $\nabla l(\cdot)$ and k is the iteration number.

If $\mathbf{z}^k = \beta^k - \frac{1}{L} \nabla l(\beta^k)$, then $F(\beta_k) - F(\beta^*) \leq \frac{\alpha L \|\beta^0 - \beta^*\|_2^2}{2k}$, where $F(\cdot)$ is the objective of (1). If $\mathbf{z}^k = \mathbf{y}^k - \frac{1}{L} \nabla l(\mathbf{y}^k)$ where $\mathbf{y}^k = \beta^k + \alpha^k (\beta^k - \beta^{k-1})$ with α controlling the momentum, we have $F(\beta_k) - F(\beta^*) \leq \frac{2\alpha L \|\beta^0 - \beta^*\|_2^2}{(k+1)^2}$

The proof can be viewed as an instantiation of the convex analysis introduced in (Nesterov and Nesterov 2004). We provide a rigorous proof in the supplementary file.

Now the key to solve (1) is how efficiently we solve (2). If there were no constraints, (2) is the fused lasso signal approximation proposed in (Friedman et al. 2007), where it was shown that by utilizing the separability of the l_1 norm, an element-wise soft-threshold technique can be applied to remove the sparse term. Since the constraints of (2) are also separable, we show how (2) can be further reduced likewise.

Proposition 2. *If we define*

$$\tilde{\beta} = \min_{\beta \in \mathbb{R}^d} \frac{1}{2} \|\beta - \mathbf{z}\|_2^2 + \frac{\lambda_2}{L} \sum_{(i,j) \in E} w_{ij} |\beta_i - \beta_j|, \quad (3)$$

then the optimal solution to (2) (denoted as β^) can be achieved by an element-wise post-processing to $\tilde{\beta}$ as follows*

$$\beta^* = \max(\text{sign}(\tilde{\beta}) \odot \max(|\tilde{\beta}| - \frac{\lambda_1}{L}, 0), 0); \quad (4)$$

where \odot is an element-wise product operator.

Proof. We define $\theta_i = \frac{\lambda_1}{L}$ and $\theta_{ij} = \frac{\lambda_2 w_{ij}}{L}$ respectively for all $i \in V$ and $(i, j) \in E$. We denote β' as the optimal solution of the unconstrained problem of (2). According to (Friedman et al. 2007), $\beta'_i = \text{sign}(\tilde{\beta}_i) \max(|\tilde{\beta}_i| - \theta_i, 0)$. We now consider the nonnegative constraints in (2). According to the Karush-Kuhn-Tucker (KKT) conditions, the necessary and sufficient conditions for $\beta_1^*, \dots, \beta_d^*$ are

$$\begin{aligned} Lg_i &= (\beta_i - z_i) + \theta_i s_i + \sum_{j:(i,j) \in E} \theta_{ij} t_{ij} - \\ &\sum_{j:(j,i) \in E} \theta_{ij} t_{ji} - \alpha_i = 0, \quad \text{s.t.} \quad \alpha_i \beta_i = 0, \end{aligned} \quad (5)$$

where $\alpha_i \geq 0$ are the Lagrange multipliers and \mathbf{s}, \mathbf{t} are subgradients: $s_i = \text{sign}(\beta_i)$ if $\beta_i \neq 0$ and $s_i \in [-1, 1]$ if $\beta_i = 0$; $t_{ij} = \text{sign}(\beta_i - \beta_j)$ for $\beta_i \neq \beta_j$ and $t_{ij} \in [-1, 1]$ if $\beta_i = \beta_j$. The objective equation in (5) is to set the derivative of the Lagrange function to zero and the constraint equations are obtained from the complementary slackness condition. We now consider two cases of β :

Case 1 $\beta'_i \geq 0$: Note that by setting $\alpha_i = 0$, $\beta'_i \geq 0$ satisfies the conditions in (5), thus $\beta_i^* = \beta'_i = \text{sign}(\tilde{\beta}_i) \max(|\tilde{\beta}_i| - \theta_i, 0)$ is the solution of (2).

Case 2 $\beta'_i < 0$: We can set $\beta_i^* = 0$ and $\alpha_i = -\beta'_i > 0$, then we have $\beta'_i = \beta_i^* - \alpha_i$.

$$\begin{aligned} Lg_i &= (\beta_i^* - z_i) + \theta_i s_i + \sum_{j:(i,j) \in E} \theta_{ij} t_{ij} - \sum_{j:(j,i) \in E} \theta_{ij} t_{ji} - \alpha_i \\ &= (\beta_i^* - \alpha_i - z_i) + \theta_i s_i + \sum_{j:(i,j) \in E} \theta_{ij} t_{ij} - \sum_{j:(j,i) \in E} \theta_{ij} t_{ji} \\ &= (\beta'_i - z_i) + \theta_i s_i + \sum_{j:(i,j) \in E} \theta_{ij} t_{ij} - \sum_{j:(j,i) \in E} \theta_{ij} t_{ji} = 0. \end{aligned}$$

Hence, in summary, we have

$$\beta^* = \max(\text{sign}(\tilde{\beta}) \odot \max(|\tilde{\beta}| - \frac{\lambda_1}{L}, 0), 0). \quad \square$$

Notice that, (3) is a (continuous) total variation problem, which is known can be efficiently solved by parametric flow algorithms in (Chambolle and Darbon 2009; Xin et al. 2014). Here we present a more general perspective of such an equivalence via conic duality, which gives us a natural and novel minimum quadratic cost flow formulation. Fast flow algorithms, such as but not limited to parametric flow (Gallo, Grigoriadis, and Tarja 1989) etc. are then easily applied. For example, we show that the minimum norm point problem solved by parametric flow in (Xin et al. 2014) can be viewed as a special case of the proposed dual.

Conic Dual to Total Variation

To solve (3), we apply generalized inequalities and its corresponding Lagrange duality introduced in (Boyd and Vandenberghe 2004). Specifically, we first define a set C such that $C = \{(\beta, \alpha) \in \mathbb{R}^{d+1} \mid \forall (i, j), |\beta_i - \beta_j| \leq \alpha\}$.

Lemma 3. *The set C is a proper cone.*

This can be easily shown by checking all the properties required by a proper cone. See the supplementary file for a proof. We now consider the following problem (we keep using $\theta_{ij} = \frac{\lambda_2 w_{ki}}{L}$ for $\frac{\lambda_2 w_{ki}}{L}$ in (3)):

$$\begin{aligned} \min_{\forall (i,j) \in E, \{\beta^{ij} \in \mathbb{R}^d, \alpha^{ij} \in \mathbb{R}\}} & \frac{1}{2} \|\beta - \mathbf{z}\|_2^2 + \sum_{(i,j) \in E} \theta_{ij} \alpha^{ij} \\ \text{s.t.} & (\beta^{ij}, \alpha^{ij}) \in C \text{ and } \beta_k^{ij} = \begin{cases} \beta_k & k = i, j \\ 0 & \text{else} \end{cases}. \end{aligned} \quad (6)$$

Since $(\beta^{ij}, \alpha^{ij}) \in C$, then $|\beta_i - \beta_j| \leq \alpha^{ij}$, therefore (6) is indeed equivalent to (3). Moreover, because C is a proper cone, we can rewrite (6) as follows,

$$\begin{aligned} \min_{\forall (i,j) \in E, \{\beta^{ij} \in \mathbb{R}^d, \alpha^{ij} \in \mathbb{R}\}} & \frac{1}{2} \|\beta - \mathbf{z}\|_2^2 + \sum_{(i,j) \in E} \theta_{ij} \alpha^{ij} \\ \text{s.t.} & \begin{bmatrix} \beta^{ij} \\ \alpha^{ij} \end{bmatrix} \succeq_C 0 \text{ and } \beta_k^{ij} = \begin{cases} \beta_k & k = i, j \\ 0 & \text{else} \end{cases}. \end{aligned} \quad (7)$$

where $\beta \succeq_C 0 \iff \beta \in C$ is defined as generalized inequality (Boyd and Vandenberghe 2004). We call (7) the primal problem (which equals to the original TV problem). Since the primal problem is both convex and satisfies Slater's condition, strong Lagrange duality holds (under generalized inequality). We define the Lagrange function as

$$\begin{aligned} L(\beta, \alpha, \xi, \tau) &= \frac{1}{2} \|\beta - \mathbf{z}\|_2^2 + \sum_{(i,j)} \theta_{ij} \alpha^{ij} - \sum_{ij} \begin{bmatrix} \xi^{ij} \\ \tau^{ij} \end{bmatrix}^T \begin{bmatrix} \beta^{ij} \\ \alpha^{ij} \end{bmatrix} \\ \text{s.t.} & \xi^{ij} \in \mathbb{R}^d : \begin{bmatrix} \xi^{ij} \\ \tau^{ij} \end{bmatrix} \succeq_{C^*} 0 \text{ and } \xi_k^{ij} = 0 \text{ if } k \neq i, j, \end{aligned} \quad (8)$$

where (ξ^{ij}, τ^{ij}) are the Lagrange multipliers and C^* is the dual cone of C , defined as $C^* = \{\mathbf{v} \mid \mathbf{w}^T \mathbf{v} \geq 0, \forall \mathbf{w} \in C\}$. To formulate the dual problem, we take the derivative of $L(\cdot)$ with respect to the primal variables (β, α) and we have

$$\beta - \mathbf{z} - \sum_{ij} \xi^{ij} = 0 \text{ and } \theta_{ij} - \tau^{ij} = 0. \quad (9)$$




Figure 1: Graph structure and the flows. Each source-to-node edge (s,i) has 0 cost and a maximum capacity γ_i and a minimum capacity γ_i , this ensures a flow γ_i ; each node-to-node edge (i,j) has 0 cost and a maximum capacity θ_{ij} ; each node-to-sink edge (i,t) has infinite capacity and a cost (in red) $\frac{1}{2}(y_i - (\bar{\xi}_i + \gamma_i))^2$, where $\bar{\xi}_i = \sum_{j, (i,j) \in E} \xi_i^{ij}$. All flows through node 2 is highlighted in blue for an example.

By applying (9) to (8), the dual problem is written as

$$\begin{aligned} \max_{\forall (i,j) \in E, \{\xi^{ij} \in \mathbb{R}^d\}} & -\frac{1}{2} \|\mathbf{z}\|_2^2 + \sum_{(i,j) \in E} \xi^{ij} \|\mathbf{z}\|_2^2 + \frac{1}{2} \|\mathbf{z}\|_2^2 \\ \text{s.t. } & (\xi^{ij}, \tau^{ij}) \in C^*; \xi_k^{ij} = 0, k \neq i, j. \end{aligned} \quad (10)$$

Proposition 4. $(\xi^{ij}, \tau^{ij}) \in C^*, \xi_k^{ij} = 0, k \neq i, j \iff \xi_i^{ij} + \xi_j^{ij} = 0, |\xi_i^{ij}| \leq \theta_{ij}$.

Please find the proof in the supplementary file. Accordingly, the dual problem becomes

$$\begin{aligned} \min_{\forall (i,j) \in E, \{\xi^{ij} \in \mathbb{R}^d\}} & \frac{1}{2} \|\mathbf{z} - \sum_{(i,j) \in E} \xi^{ij}\|_2^2 \\ \text{s.t. } & \xi_k^{ij} = 0, k \neq i, j, \xi_i^{ij} + \xi_j^{ij} = 0, |\xi_i^{ij}| \leq \theta_{ij}, \end{aligned} \quad (11)$$

where we omit $\|\mathbf{z}\|_2^2$ from the objective since it is a constant with respect to ξ s and we also changed the sign of all ξ s for better illustration of the flows.

Problem (11) can be viewed as the following minimum quadratic cost flow formulation,

$$\begin{aligned} \min_{\forall (i,j) \in E, \{\xi^{ij} \in \mathbb{R}^d\}} & \frac{1}{2} \|\mathbf{y} - (\sum_{(i,j) \in E} \xi^{ij} + \gamma)\|_2^2 \\ \text{s.t. } & \xi_k^{ij} = 0, k \neq i, j, \xi_i^{ij} + \xi_j^{ij} = 0, |\xi_i^{ij}| \leq \theta_{ij}, \end{aligned} \quad (12)$$

where we have induced γ ($\gamma_i = \max\{|z_i|, \sum_{j, (i,j) \in E} \theta_{ij}\}$) and denote $\mathbf{y} = \mathbf{z} + \gamma$ to ensure $\mathbf{y} \geq \mathbf{0}$ and $(\xi^{ij} + \gamma) \geq \mathbf{0}$. Thus each feasible ξ of (12) is a possible flow on graph $G=(V, E)$. Since $\xi_i^{ij} + \xi_j^{ij} = 0, |\xi_i^{ij}|$ can denote a flow on edge (i,j) such that $\xi_i^{ij} \geq 0$ denotes a flow coming into node i and $\xi_i^{ij} \leq 0$ denotes a flow leaving node i. Figure 1 illustrates such flows by taking node 2 as an example. Thus to minimize the objective of (12) is equivalent to computing a minimum cost flow on this graph. Since the cost is quadratic with respect to the flow, this problem is a minimum quadratic cost flow problem. According to (Hochbaum

Table 1: Runtime (in sec.) comparison of the proposed algorithm with CVX. d is the data dimensionality. * indicates the test did not finish within 24 hours.

d	400	900	2500	4900	10000
CVX	5.22	33.71	889.80	$1.08e^4$	*
Ours	0.87	2.39	15.95	64.33	321.93

and Hong 1995; Mairal et al. 2011), this type of problems can be efficiently solved via fast flow algorithms including but not limited to the parametric flow (Gallo, Grigoriadis, and Tarja 1989). Note that in (Xin et al. 2014), TV is shown equivalent to a minimum norm point (MNP) problem under submodular constraints which is solved via parametric flow. We now discuss the relation between the dual problem i.e. (12) or (11) and the MNP considered in (Xin et al. 2014).

Recall that the MNP problem is defined as follows

$$\min_{\mathbf{s} \in \mathbb{R}^d, \mathbf{s} \in B(\lambda f_c)} \|\mathbf{z} - \mathbf{s}\|_2^2, \quad (13)$$

where $f_c(S)$ is a cut function, defined as $f_c(S) = \sum_{i \in S, j \in V \setminus S} w_{ij}$ and $B(\cdot)$ is the base polyhedron of f_c .

Proposition 5. For any minimizer ξ^* of (11), define $\hat{\mathbf{s}}$ such that $\hat{\mathbf{s}} = \sum_{(i,j) \in E} \xi^{*ij}$, then $\hat{\mathbf{s}}$ is a minimizer of (13). For any minimizer \mathbf{s}^* of (13), there exists a decomposition such that $\mathbf{s}^* = \sum_{(i,j) \in E} \hat{\xi}^{ij}$, where $\hat{\xi}$ is one minimizer of (11).

According to Prop. 5, the MNP problem i.e. (13) can be viewed as a special case of (11) (the conic dual), where $\sum_{(i,j) \in E} \xi^{ij} = \mathbf{s}$. Moreover, since (11) has relatively “looser” constraints, it is possible to devise more efficient algorithms (than parametric flow) to solve (11) and thereafter TV. For example, in (Mairal et al. 2011), a faster (than parametric flow) flow algorithm is proposed to solve their specific minimum quadratic flow problem. Hence, the conic dual perspective opens a new opportunity to solve the famous TV problem more efficiently.

Optimization summary. In summary, by applying Prop. 1, we can solve n^2 GFL by iteratively solving (2). By applying Prop. 2, we further reduce (2) to the TV problem defined in (3), we then transform it to a minimum cost flow algorithm via conic duality and solve it by a fast flow algorithm.

In Tab. 1, we compare the proposed algorithm with an off-the-shelf solver on synthetic data. We generate a random $\beta \in \mathbb{R}^d$ and a 2D grid graph of d nodes with each node having four neighbors. We then generate $N = d/2$ samples: $\mathbf{x}_i \in \mathbb{R}^d$ and $y_i = \beta^T \mathbf{x}_i + 0.01n_i$, where \mathbf{x}_i and n_i are drawn from the standard normal distribution. All experiments are carried out on an Intel(R) Core(TM) i7-3770 CPU at 3.40GHz. The experiments show that the proposed optimization algorithm is more efficient and scalable.

Application to the Diagnosis of AD

In the diagnosis of AD, two fundamental issues are AD/NC (Normal Control) classification and MCI/NC (Mild Cognitive Impairment) classification. Let $\mathbf{x}_i \in \mathbb{R}^d$ be sub-

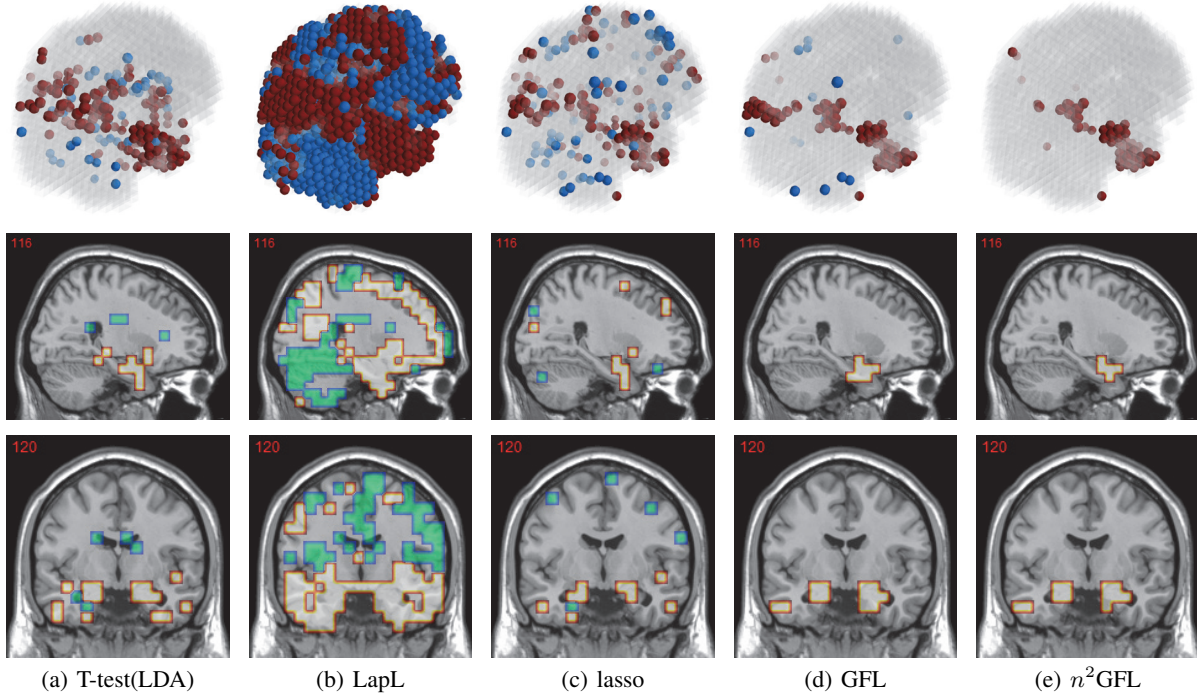


Figure 2: Feature selection by different models. The top row illustrates selected voxels in a 3D model (voxels with positive β are in brown and negative ones are in blue), the mid and bottom rows illustrate the corresponding projections on brain slices.

jects' sMRI voxels and $y_i = \{-1, 1\}$ be the disease status (AD/NC or MCI/NC). Since the problems are classifications, we use the logistic regression as the loss term

$$l(\beta) = \sum_{i=1}^N \log(1 + \exp(-y_i(\beta^T \mathbf{x}_i + c))), \quad (14)$$

where $c \in \mathbb{R}$ is the bias parameter (to be learned). For the graph structure, we define each voxel as a node and their spatial adjacency as the edges, i.e. $w_{ij}=1$ if voxels i and j are adjacent and 0 otherwise. The data are obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database². We split all the baseline data into 1.5T and 3.0T MRI scans datasets (named 15T and 30T). 64 AD patients, 90 NC and 208 MCI patients are included in our 15T dataset; 66 AD patients and 110 NC are included in our 30T dataset. (Most 30T MCI data are in an on-going phase and are not included). Data preprocessing follows the DARTEL VBM pipeline (Ashburner and others 2007) as commonly done in the literature. $2,527 \times 8 \times 8 \times 8 \text{ mm}^3$ size voxels that have values greater than 0.2 in the mean gray matter population template serve as the input features. We design experiments on three tasks, namely, 15ADNC, 30ADNC, 15MCINC.

Classification Accuracy. 10-fold cross-validation (CV) evaluation is applied and the classification accuracy for all tasks are summarized in Tab. 2. Under exactly the same experiment setup, we compare $n^2\text{GFL}$ with the state-of-the-art classifiers: logistic regression (LR), SVM, sparse models e.g. the lasso and its graph Laplacian structured variants, i.e.

²<http://adni.loni.ucla.edu>

Table 2: Classification accuracies.

	SVM	LR	MLDA	LapL	lasso	GFL	$n^2\text{GFL}$
15ADNC	83.1	83.1	83.1	84.4	85.7	85.1	86.4
30ADNC	87.5	86.9	83.5	87.5	88.6	85.8	90.3
15MCINC	71.1	70.5	59.4	70.1	70.8	69.8	71.8

the LapL, the unconstrained GFL (Xin et al. 2014), and the “MLDA” model (Dai et al. 2012), which applies a variant of Fisher Discriminant Analysis after univariate feature selection (via T-test). For each model, we used grid-search to find the optimal parameters respectively. Note that our accuracies may not be superior to the recent work (Liu et al. 2014), the main reason is that in (Liu et al. 2014), multi-modality data (including PET and sMRI data) are used. Nevertheless, Tab. 2 demonstrates that $n^2\text{GFL}$ outperforms all the other models using only voxel-based sMRI data.

Feature selection. For each task, the selected features are those whose β are not zero. In Figure 2, the result of 30ADNC is used to illustrate the feature selection by different models (using the parameters at their best accuracy). As shown, the selected voxels by both GFL and $n^2\text{GFL}$ cluster into several spatially connected regions, whereas those of lasso and T-test/MLDA scatter around. Also, as mentioned before, the LapL tends to select much more voxels than necessary due to the l_2 regularization. Moreover, the selected voxels by GFL and $n^2\text{GFL}$ are concentrated in Hippocampus, ParaHippocampal gyrus (which are believed to be the

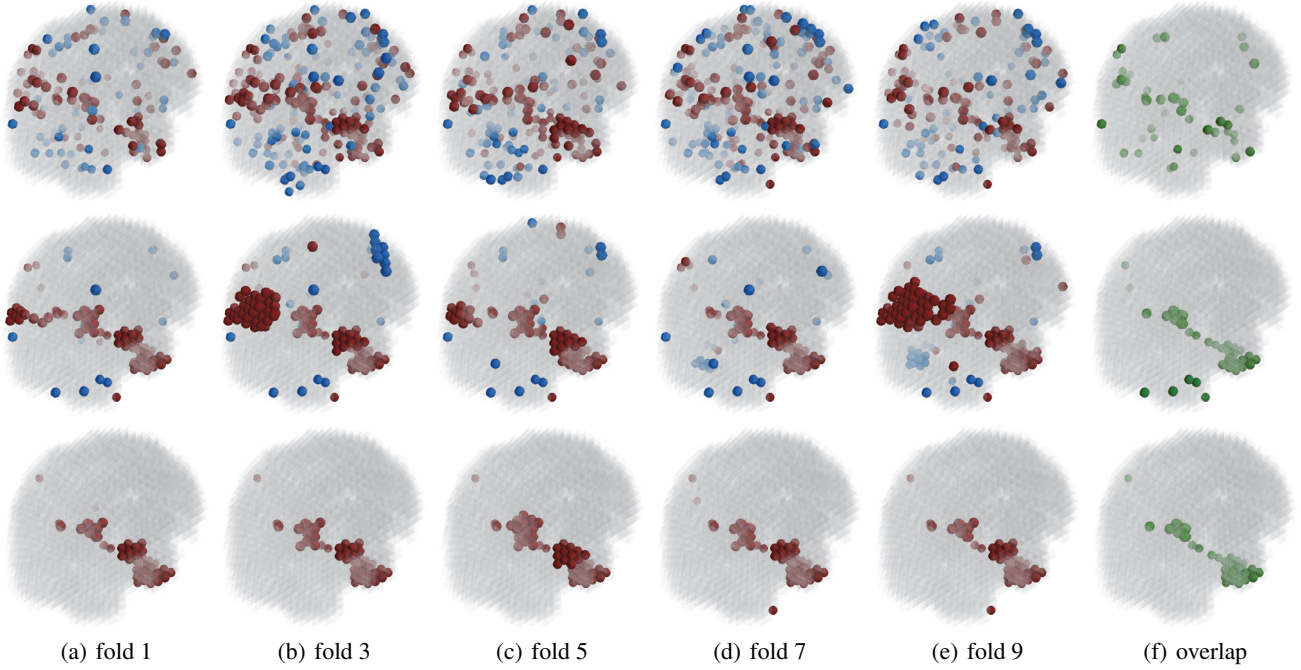


Figure 3: Stability of selected voxels across different folds of the cross-validation. The results of 5 different folds are shown in (a)-(e). The voxels with positive β are in brown, negative ones are in blue. The common/overlapped voxels selected in all 10 folds are shown in green (f). The top row illustrates voxels selected by the lasso model, the mid row illustrates those of GFL and the bottom row shows those of n^2 GFL.

early damaged regions). On the other hand, the lasso and T-test/MLDA either select less lesion voxels or select probably noisy voxels not in the early damaged regions.

Feature Stability. In Figure 3, we show the selected voxels across different folds of CV³. As shown, the selected voxels by lasso vary much across different folds, whereas the selected voxels by GFL are more stable. However, by assuming the positive correlation between the features and the disease labels in n^2 GFL, we further increase the stability. To quantitatively evaluate the stability gain, we denote the variables of the k th fold of CV as $\beta(k)$. We introduce two measurements here. In (Yu 2013), the **Estimation Stability (ES)** is proposed to measure the stability of the estimation

$$ES = \sum_{k=1}^K \|\mathbf{X}\beta(k) - \mathbf{X}\bar{\beta}\|_2^2 / K \|\bar{\beta}\|_2^2, \quad (15)$$

where $\bar{\beta} = \sum_{k=1}^K \beta(k) / K$. It is shown in (Yu 2013) that ES is a fair measurement of the estimation stability. To further understand the stability of feature selection, we also extend the Dice coefficient (Dice 1945) to multiple sets and apply the **multi-set Dice Coefficient (mDC)** as a measurement. We denote set $S(k) = \{i : \beta_i(k) \neq 0\}$ and define mDC as

$$mDC = K \#(\cap_{k=1}^K S(k)) / \sum_{k=1}^K \#(S(k)), \quad (16)$$

where $\#$ is the number of elements in a set. In Tab. 3, both measurements quantitatively suggest n^2 GFL obtains much

³Here, parameters were determined by accuracy. Similar results were observed using parameters producing same level of sparsity.

Table 3: Stability comparison of the models.

	lasso	GFL	n^2 GFL
ES (smaller is better)	0.035	0.033	0.022
mDC (larger is better)	0.267	0.374	0.644

more stable voxels due to the consideration of the correlation between the features and the disease labels⁴.

Conclusions

In this paper, we explore the nonnegative generalized fused lasso model to address an important problem of neuroimage analysis, i.e. the stability of feature selection. Experiments show that our model greatly improves the stabilities of feature selection over existing methods for brain image analysis. Although n^2 GFL is applied to the diagnosis of AD problem, it can be applied to solve more general problems. Moreover, we believe that the theoretical points made here e.g. nonnegative FISTA, soft-thresholding and the conic dual of TV, provide motivation for future work of general interest.

⁴We notice that, in (Xin et al. 2014), the stability is computed using the top 50 positive voxels because these voxels are believed to be the most atrophied ones. By computing the stability of all non-zero voxels, the mDC of GFL drops around 30%. This clearly shows that the instability is caused largely by the undesirable voxels that disagree with the correlation prior (those scattered blue voxels in the mid row).

Acknowledgments

This work was supported in part by Natural Science Foundation of China (NSFC) grants 973-2015CB351800, NSFC-61272027, NSFC-61231010, NSFC-61121002 and NSFC-61210005.

References

- Ashburner, J., et al. 2007. A fast diffeomorphic image registration algorithm. *Neuroimage* 38(1):95–113.
- Avants, B. B.; Libon, D. J.; Rascovsky, K.; Boller, A.; McMillan, C. T.; Massimo, L.; Coslett, H.; Chatterjee, A.; Gross, R. G.; and Grossman, M. 2014. Sparse canonical correlation analysis relates network-level atrophy to multi-variate cognitive measures in a neurodegenerative population. *NeuroImage* 84:698–711.
- Beck, A., and Teboulle, M. 2009. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences* 2(1):183–202.
- Boyd, S. P., and Vandenberghe, L. 2004. *Convex optimization*. Cambridge university press.
- Chambolle, A., and Darbon, J. 2009. On total variation minimization and surface evolution using parametric maximum flows. *International journal of computer vision* 84(3):288–307.
- Dai, Z.; Yan, C.; Wang, Z.; Wang, J.; Xia, M.; Li, K.; and He, Y. 2012. Discriminative analysis of early alzheimer’s disease using multi-modal imaging and multi-level characterization with multi-classifier (m3). *Neuroimage* 59(3):2187–2195.
- Dice, L. R. 1945. Measures of the amount of ecologic association between species. *Ecology* 26(3):297–302.
- Friedman, J.; Hastie, T.; Höfling, H.; Tibshirani, R.; et al. 2007. Pathwise coordinate optimization. *The Annals of Applied Statistics* 1(2):302–332.
- Gallo, G.; Grigoriadis, M.; and Tarja, R. 1989. A fast parametric maximum flow algorithm and applications. *SIAM Journal of Computing* 18(1):30–55.
- Gramfort, A.; Thirion, B.; and Varoquaux, G. 2013. Identifying predictive regions from fmri with tv-l1 prior. In *Pattern Recognition in Neuroimaging (PRNI), 2013 International Workshop on*, 17–20. IEEE.
- Grant, M., and Boyd, S. 2013. CVX: Matlab software for disciplined convex programming, version 2.0 beta. <http://cvxr.com/cvx>.
- Grosenick, L.; Klingenberg, B.; Katovich, K.; Knutson, B.; and Taylor, J. E. 2013. Interpretable whole-brain prediction analysis with graphnet. *NeuroImage* 72:304–321.
- Hochbaum, D. S., and Hong, S. 1995. About strongly polynomial time algorithms for quadratic optimization over submodular constraints. *Mathematical programming* 69(1-3):269–309.
- Jacob, L.; Obozinski, G.; and Vert, J.-P. 2009. Group lasso with overlap and graph lasso. In *Proceedings of the 26th Annual International Conference on Machine Learning*, 433–440. ACM.
- Jenatton, R.; Gramfort, A.; Michel, V.; Obozinski, G.; Eger, E.; Bach, F.; and Thirion, B. 2012. Multiscale mining of fmri data with hierarchical structured sparsity. *SIAM Journal on Imaging Sciences* 5(3):835–856.
- Lee, D. D., and Seung, H. S. 1999. Learning the parts of objects by non-negative matrix factorization. *Nature* 401(6755):788–791.
- Liu, F.; Wee, C.-Y.; Chen, H.; and Shen, D. 2014. Inter-modality relationship constrained multi-modality multi-task feature selection for alzheimer’s disease and mild cognitive impairment identification. *NeuroImage* 84:466–475.
- Liu, M.; Zhang, D.; and Shen, D. 2012. Ensemble sparse classification of alzheimer’s disease. *NeuroImage* 60(2):1106–1116.
- Mairal, J.; Jenatton, R.; Obozinski, G.; and Bach, F. 2011. Convex and network flow optimization for structured sparsity. *The Journal of Machine Learning Research* 12:2681–2720.
- Nesterov, Y., and Nesterov, I. 2004. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer.
- Ng, B., and Abugharbieh, R. 2011. Generalized sparse regularization with application to fmri brain decoding. In *Information Processing in Medical Imaging*, 612–623. Springer.
- Rao, N.; Cox, C.; Nowak, R.; and Rogers, T. T. 2013. Sparse overlapping sets lasso for multitask learning and its application to fmri analysis. In *Advances in Neural Information Processing Systems*, 2202–2210.
- Tibshirani, R. 1996. Regression shrinkage and selection via the Lasso. *Journal of Royal Statistical Society B* 58(1):267–288.
- Xin, B.; Kawahara, Y.; Wang, Y.; and Gao, W. 2014. Efficient generalized fused lasso and its application to the diagnosis of alzheimer’s disease. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.
- Yu, B. 2013. Stability. *Bernoulli* 19(4):1484–1500.