High-dimensional regression over disease subgroups

Frank Dondelinger¹, Sach Mukherjee², and The Alzheimer’s
Disease Neuroimaging Initiative*

¹Lancaster Medical School, Lancaster University

²DZNE, Bonn

December 9, 2016

Abstract

We consider high-dimensional regression over subgroups of observations. Our work is motivated by biomedical problems, where disease subtypes, for example, may differ with respect to underlying regression models, but sample sizes at the subgroup-level may be limited. We focus on the case in which subgroup-specific models may be expected to be similar but not necessarily identical. Our approach is to treat subgroups as related problem instances and jointly estimate subgroup-specific regression coefficients. This is done in a penalized framework, combining an ℓ_1 term with an additional term that penalizes differences between subgroup-specific coefficients. This gives solutions that are globally sparse but that allow information-sharing between the subgroups. We present algorithms for estimation and empirical results on simulated data and using Alzheimer’s disease, amyotrophic lateral sclerosis and cancer datasets. These examples demonstrate the gains our approach can offer in terms of prediction and the ability to estimate subgroup-specific sparsity patterns.

1 Introduction

High-dimensional regression has been well studied in the case where all samples can reasonably be expected to follow the same model. However, in several

*Data used in preparation of this article were obtained from the Alzheimers 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

current and emerging applications, observations span multiple subgroups that may not be identical with respect to the underlying regression models. Examples abound, from disease subtypes in biomedicine to customer subsets in business applications. We are specifically motivated by biomedical problems, where sets of samples, such as disease subtypes, although related, may differ with respect to underlying biology and therefore have different relationships between covariates and a response of interest.

Thus, we focus on high-dimensional regression in group-structured settings. In particular, we consider linear regression in a commonly-encountered scenario in which the same set of p covariates or predictors is available in each of K subgroups. That is, we consider subgroup-specific linear regression problems indexed by k , each with subgroup-specific sample size n_k , a response vector y_k of length n_k , a $n_k \times p$ feature matrix X_k and a p -vector β_k of regression coefficients. The problem we address is estimating the regression coefficients $\beta_1 \dots \beta_K$.

We propose an approach to jointly estimate the regression coefficients that induces global sparsity and encourages similarity between subgroup-specific coefficients. We consider the following penalized formulation and its variants

$$\hat{B} = \arg \min_{B=[\beta_1 \dots \beta_K]} \sum_{k=1}^K \left\{ \|y_k - X_k \beta_k\|_2^2 + \lambda \|\beta_k\|_1 + \gamma \sum_{k' > k} \tau_{k,k'} \|\beta_k - \beta_{k'}\|_2^2 \right\}$$

where, $B = [\beta_1 \dots \beta_K]$ is a $p \times K$ matrix that collects together all the regression coefficients, $\|\cdot\|_q$ denotes the ℓ_q norm of its argument and λ, γ, τ are tuning parameters. The last term is a fusion-type penalty between subgroups; note that the difference is taken between entire vectors of subgroup-specific coefficients. An ℓ_2 fusion penalty is shown above, although other penalties may be used; in this manuscript, we also consider an ℓ_1 variant. The parameters $\tau_{k,k'}$ allow for the possibility of controlling the extent to which similarity is encouraged for specific pairs of subgroups.

Our proposal shares similarities with both the group lasso [Yuan and Lin, 2006] and the fused lasso [Tibshirani et al., 2005] but differs from both in important ways. In contrast to the group lasso, we consider subgroups of samples or observations rather than groups of coefficients and in contrast to the fused lasso, we consider fusion of entire (subgroup-specific) coefficient vectors, rather than successive coefficients under a pre-defined ordering. Obozinski et al. [2010] showed how the group lasso could be used in subgroup-structured settings, essentially by considering the global problem and defining groups (in the group lasso sense) corresponding to the same covariate across all subgroups. This means that each covariate tends either to be included in all subgroup-specific models or none. In contrast, our approach allows subgroups to have different sparsity patterns, whilst pulling subgroup-specific coefficients together and inducing global sparsity. Our work is also similar in spirit to recent work concerning joint estimation of graphical models over multiple problem instances [Danaher et al., 2014, Oates et al., 2014, 2015].

We are motivated by emerging problems in biomedical research and specifically in personalized medicine. High-dimensional regression problems are now becoming common in this area, with several high-dimensional data types already in mainstream use. In the personalized medicine setting, samples usually correspond to individuals and the subgroups k to e.g. diseases or disease subtypes. It is increasingly clear that many disease subtypes differ in their biology [see e.g. Weinstein et al., 2013, Akbani et al., 2014], suggesting that relationships between covariates and responses of interest may differ between them. However, sample sizes tend to be limited, especially at the subgroup level, posing problems for the subgroup-wise strategy of solving each problem separately. On the other hand, pooling all the data together into a single regression problem may lead to severe mis-specification if the underlying subgroup-specific models do indeed differ.

These issues may lead to losses in terms of predictive ability and perhaps just as important in the ability to efficiently estimate subgroup-specific influences that may themselves be of interest. In contrast to simple pooling, our approach allows subgroups to have different sparsity patterns and regression coefficients, but in contrast to the subgroup-wise approach it takes advantage of similarities between subgroups.

We show empirical results in the context of two neurodegenerative diseases - Alzheimer’s disease and amyotrophic lateral sclerosis (ALS) - and cancer (see below for full details of the applications and data). The responses concern disease progression in Alzheimer’s and ALS and therapeutic response in cancer cell lines. In the Alzheimer’s and ALS examples, subgroups are based on clinical factors, while in the cancer data they are based on the tissue type of the cell lines.

Across the three examples, data types include genetic, clinical and transcriptomic variables. We find that our approach can improve performance relative to pooling or subgroup-wise analysis. Importantly, in cases where pooling or subgroup-wise analyses do well (perhaps reflecting a lack of subgroup structure or insufficient similarity respectively) our approach remains competitive. This gives assurance that penalization is indeed able to share information appropriately in real-world examples. We emphasize that the goal of the empirical analyses we present is not to give the best predictions possible in these applications, but rather to explore joint estimation in group-structured biomedical problems.

2 Methods

2.1 Notation

Each subgroup $k \in \{1 \dots K\}$ has the same set of p covariates, but subgroup-specific sample size n_k . Total sample size is $n = \sum_{k=1}^K n_k$. For subgroup k , X_k is the $n_k \times p$ feature matrix and y_k the corresponding $n_k \times 1$ vector of observed responses. Subgroup-specific regression coefficients are $\beta_k \in \mathbb{R}^p$.

Where convenient we collect all regression coefficients together in a $p \times K$ matrix $B = [\beta_1 \dots \beta_K]$ and accordingly we use $\beta_{j,k}$ to denote the coefficient for covariate j in subgroup k .

2.2 Model Formulation

We seek to jointly estimate the regression coefficients $B = [\beta_1 \dots \beta_K]$ whilst ensuring global sparsity and encouraging agreement between subgroup-specific coefficients. We propose the criterion

$$\hat{B} = \arg \min_{B=[\beta_1 \dots \beta_K]} \sum_{k=1}^K \left\{ \|y_k - X_k \beta_k\|_2^2 + \lambda \|\beta_k\|_1 + \gamma \sum_{k' > k} \tau_{k,k'} \|\beta_k - \beta_{k'}\|_2^2 \right\} \quad (1)$$

and a variant with an ℓ_1 norm in the last term

$$\hat{B} = \arg \min_B \sum_{k=1}^K \left\{ \|y_k - X_k \beta_k\|_2^2 + \lambda \|\beta_k\|_1 + \gamma \sum_{k' > k} \tau_{k,k'} \|\beta_k - \beta_{k'}\|_1 \right\}. \quad (2)$$

Here, λ, γ, τ are tuning parameters. The role of the last term is to encourage similarity between subgroup-specific regression coefficients. The special case $K = 1$ recovers the classical lasso (applied to all data pooled together). The tuning parameters $\tau_{k,k'}$ give the possibility of controlling the extent of fusion between specific subgroups. By default all τ 's are set to unity ("unweighted fusion"), but they can also be set to specific values as discussed below ("weighted fusion"). In the above formulation, we assume that y_k and X_k have been standardized (at the subgroup level) so that no intercept terms are required. Note that the regularization parameters λ, γ are the same across subgroups.

The difference between the two variants is that the first, ℓ_2 fusion encourages similarity between subgroup-specific coefficients, while the second ℓ_1 version allows for exact equality. The ℓ_2 formulation has the computational advantage that the fusion part of the objective function becomes continuously differentiable, and the estimate of the objective function at each step can be obtained by soft-thresholding, analogously to coordinate descent for regular lasso problems. In the ℓ_1 formulation on the other hand, the fusion constraint is only piece-wise continuously differentiable, leading to a more difficult optimisation problem (see below).

2.3 Comparison with group and fused lasso

Our formulation resembles the group lasso and fused lasso, but differs from both in important ways. The original group lasso [Yuan and Lin, 2006] was designed to consider groups of covariates within a single regression problem. Let X be the feature matrix and y the vector of responses in a standard regression problem. Letting $l \in \{1 \dots L\}$ index groups of covariates, the group lasso criterion is

$$\hat{\beta} = \arg \min_{\beta} \|y - \sum_{l=1}^L X^{(l)} \beta^{(l)}\|_2^2 + \sum_{l=1}^L \lambda_l \|\beta^{(l)}\|_2 \quad (3)$$

where $X^{(l)}$ is the submatrix of X corresponding to the covariates in group l , $\beta^{(l)}$ the corresponding regression coefficients and λ_l 's tuning parameters. The penalty tends to include or exclude all members of a group from the model, i.e. all coefficients in a group may be set to zero giving groupwise sparsity.

In our setting, the subgroups are subsets of samples rather than covariates. Nevertheless, as shown in Obozinski et al. [2010], one could use a group lasso-like criterion for estimation in the multiple subgroup setting by forming groups l each comprising all the coefficients for a single covariate $j \in \{1 \dots p\}$ across all K regression problems. This encourages covariates to either be included in all the subgroup-specific models or none.

The fused lasso [Tibshirani et al., 2005] is also aimed at a single regression problem, but assumes that the covariates can be ordered in such a way that successive coefficients may be expected to be similar. This leads to the following criterion

$$\hat{\beta} = \arg \min_{\beta} \|y - X\beta\|_2^2 + \lambda \|\beta\|_1 + \gamma \sum_{i=1}^{p-1} \|\beta_i - \beta_{i+1}\|_1 \quad (4)$$

where λ, γ are tuning parameters and we have assumed that the covariates are in a suitable order. The final term encourages similarity between successive coefficients. Efficient solutions for various classes of this problem exist [e.g. Hoeffling, 2010, Liu et al., 2010, Ye and Xie, 2011].

Our approach shares the use of a fusion-type penalty, but focuses on a different problem, namely that of jointly estimating regression coefficients across multiple, potentially non-identical, problems. Accordingly, our penalty encourages agreement between entire coefficient vectors from different subgroups and does not require any ordering of covariates.

2.4 Setting the tuning parameters τ

For weighted fusion, the parameters $\tau_{k,k'}$ could be set by cross-validation but this may be onerous in practice. As an alternative we consider setting $\tau_{k,k'}$ using a distance function $d(k, k')$ based on the covariates. The idea is to allow more fusion between subgroups that are similar with respect to d , while allowing the $\tau_{k,k'}$ to be set in advance of estimation proper. However, this assumes that similarity in the covariates reasonably reflects similarity between the underlying regression coefficients, which may or may not be the case in specific applications.

We consider two variants. The first sets $d(k, k') = \|\mu_k - \mu_{k'}\|_2$ where $\mu_k, \mu_{k'}$ are the sample means of the covariates in the subgroups k, k' respectively (we assume the data have been standardized). The second approach additionally takes the covariance structure into account by using the symmetrised Kullback-Leibler (KL) divergence, i.e. $d(k, k') = \frac{1}{2}(\text{KL}(\hat{p}_k \| \hat{p}_{k'}) + \text{KL}(\hat{p}_{k'} \| \hat{p}_k))$, where $\hat{p}_k, \hat{p}_{k'}$ are estimated marginal distributions over the covariates in the subgroups k, k' respectively and $\text{KL}(p \| q)$ is the KL-divergence between distributions p and q . In practice, this requires simplifying distributional assumptions. Below we use multivariate Normal models for this purpose, with the graphical lasso [Friedman et al., 2008] used to estimate the Σ_k 's. For both approaches, we set

$\tau_{k,k'} = 1 - d(k, k')/d_{max}$, with d_{max} the largest distance between any pair of groups k, k' (this scales τ to the unit interval).

2.5 Optimisation

We describe a coordinate descent approach for optimising equation (1). While it is possible to derive a block coordinate descent approach for equation (2) [e.g. following Friedman et al., 2007], this is generally inefficient for the high-dimensional problems that we consider. Instead, we will describe an optimization procedure based on a proximal gradient approximation derived in Chen et al. [2010].

2.5.1 Coordinate Descent for ℓ_2 Fusion

The ℓ_2 fusion penalty is continuously differentiable and we can obtain the optimal value for $\hat{\beta}_{j,k}$ in equation (1) at each step by first calculating optimal values without the lasso penalty:

$$\hat{\beta}_{k,j}^* = \frac{x_{j,k}^T (y_k - X_{-j,k} \beta_{-j,k}) + 2\gamma \sum_{k' \neq k} \tau_{k,k'} \beta_{j,k'}}{x_{j,k}^T x_{j,k} + 2\gamma \sum_{k' \neq k} \tau_{k,k'}} \quad (5)$$

Then $\hat{\beta}_{j,k}$ can be obtained by soft-thresholding on $\hat{\beta}_{k,j}^*$. The procedure is summarized in Algorithm 1.

Algorithm 1 Block Coordinate Descent

```

1: procedure BLOCKDESCENTL2( $n_{iter}, X, Y, \beta_{init}, \lambda, \gamma, \tau$ )
2:    $i \leftarrow 0$ 
3:    $\beta \leftarrow \beta_{init}$ 
4:   while not_converged AND  $i < n_{iter}$  do
5:     for all  $j$  in 1:P do
6:        $\beta_{j,1:K}^{temp} \leftarrow \text{DESCENTUPDATEL2}(X, Y, \beta, j, \gamma, \tau)$  using eq. (5)
7:        $\beta_{j,1:K} \leftarrow \text{sgn}(\beta_{j,1:K}^{temp}) * \max(\beta_{j,1:K}^{temp} - \lambda, 0)$ 
8:      $i \leftarrow i + 1$ 

```

While Algorithm 1 is easy to understand and implement, a naive implementation in most programming languages will be still be slow due to the need for an inner for-loop over p , where p can be in the tens of thousands for the kinds of problems we will consider. In order to efficiently optimize B , we reformulate (1) as a classical lasso problem and apply the **glmnet** software [Friedman et al., 2010]. We transform the sum in first part of the objective into matrix form $y_{flat} - X_{diag} b_{flat}$ by defining X_{diag} as a block-diagonal $n \times pK$ matrix with X_k along the diagonals. The vector b_{flat} is a flattened version of B with stacked β_k vectors, and similarly for y_{flat} . So we have:

$$X_{diag} = \begin{pmatrix} X_1 & & \\ & \ddots & \\ & & X_K \end{pmatrix} \quad b_{flat} = \begin{pmatrix} \beta_1 \\ \vdots \\ \beta_K \end{pmatrix} \quad y_{flat} = \begin{pmatrix} y_1 \\ \vdots \\ y_K \end{pmatrix}$$

Now we move the ℓ_2 fusion penalty into the first squared term by defining the augmented matrix X_{diag}^{aug} , and augmented vector y_{flat}^{aug} , such that

$$\hat{b}_{flat} = \arg \min_{b_{flat}} \|y_{flat}^{aug} - X_{diag}^{aug} b_{flat}\|_2^2 + \lambda \|b_{flat}\|_1 \quad (6)$$

where

$$X_{diag}^{aug} = \begin{pmatrix} X_{diag} \\ \Gamma \end{pmatrix} \quad y_{flat}^{aug} = \begin{pmatrix} y_{flat} \\ \vec{0} \end{pmatrix}$$

with Γ a $pK(K-1)/2 \times pk$ matrix encoding the pair-wise fusion constraints, and $\vec{0}$ a $pK(K-1)/2 \times 1$ vector of zeros. Each block $\Gamma_{k,k'}$, $k, k' \in [1, K]$, $k < k'$ of p rows of Γ corresponds to the fusion constraint between two coefficient vectors β_k and $\beta_{k'}$, with:

$$\Gamma_{k,k'}(l, m) = \begin{cases} \gamma\tau_{k,k'} & \text{if } l = p(k-1) + m \\ -\gamma\tau_{k,k'} & \text{if } l = p(k'-1) + m \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

We can see that (6) is a classical lasso problem, to which **glmnet** can be directly applied.

2.5.2 Proximal-Gradient Approach for Fused L1 Penalty

Optimising equation (2) by block gradient descent, while possible, is highly inefficient due to having to deal with the discontinuities in the objective function space. In Chen et al. [2010], the authors describe a proximal relaxation of this problem that introduces additional smoothing to turn the objective function $f_{L1}(B)$ into a continuously differentiable function $f_{L2}^\mu(B)$. Chen et al. deal with the multi-task regression setting (with common X for each task); it is straightforward to adapt their procedure for the subgroup regression setting with different X_k per subgroup.

It is notationally convenient to first introduce a graph formulation of the fusion penalties. We will think of the fusion constraints in terms of an undirected graph $G = (V, E)$ with vertex set $V = \{1 \dots K\}$ corresponding to the subgroups and edges between all vertices. Then the ℓ_1 penalised objective function can be written as:

$$f_{L1}(B) = \sum_k \{\|y_k - X_k \beta_k\|_2^2\} + \|BC\|_1 \quad (8)$$

where the last term includes both sparsity and fusion penalties, via the matrix $C = (\lambda I_K, \gamma H)$, with I_K the identity matrix of size K , C a $K \times |E|$ matrix ($|E| = \binom{K}{2}$ in this case) and $H_{k,e} = \tau_{k,l}$ if $e = (m, l)$ and $k = l$, $H_{k,e} = -\tau_{m,k}$

if $e = (m, l)$ and $k = l$. Note that unlike in Chen et al. [2010], we require the explicit sum over k in the objective to account for different sample sizes N_k in different groups¹.

The graph formulation allows for zero edges by setting $\tau_{k,k'}$ to zero. We have implicitly assumed in the formulation of (1) and (2) that the relationship between subgroups is represented by an undirected graph. However, (8) is completely general, and it would be straightforward to incorporate a directed graph in our model. We have not pursued this avenue here, as there is no reason to suspect directionality in the subgroup relationships for the applications we consider below, and including directionality would double the number of tuning parameters $\tau_{k,k'}$ that need to be considered.

Following Chen et al. [2010], we can introduce an auxiliary matrix $A \in \mathcal{Q} = \{A' \mid \|A'\|_\infty \leq 1, A' \in \mathcal{R}^{p \times (K+|E|)}\}$. Because of duality between ℓ_1 and ℓ_∞ , we can write $\|BC\|_1 = \max_{\|A\|_\infty \leq 1} \langle A, BC \rangle$. A smooth approximation of $\|BC\|_1$ is then obtained by writing:

$$f_\mu(B) = \max_{\|A\|_\infty \leq 1} \langle A, BC \rangle - \mu d(A) \quad (9)$$

where μ is a positive smoothness parameter, and $d(A) \equiv \frac{1}{2} \|A\|_F^2$, with $\|\cdot\|_F$ the Frobenius norm. They show that for a desired accuracy ϵ , we need to set $\mu = \frac{\epsilon}{p(K+|E|)}$. Theorem 1 in Chen et al. [2010] gives the gradient of $f_\mu(B)$ as $\Delta f_\mu(B) = A^* C^T$, where A^* is the optimal solution of (9). Replacing $\|BC\|_1$ by $f_\mu(B)$ in equation (8), we obtain

$$\tilde{f}_{L1}(B) = \sum_k \{\|y_k - X_k \beta_k\|_2^2\} + f_\mu(B) \quad (10)$$

which is now continuously differentiable with gradient

$$\Delta \tilde{f}_{L1}(B) = \sum_k \{X_k^T (X_k \beta_k - y_k)\} + f_\mu(B). \quad (11)$$

Chen et al. further show that $A^* = S(BC/\mu)$ where function S truncates each entry of A^* to the range $[-1, 1]$ to ensure that $A^* \in \mathcal{Q}$. An upper bound L_U of the Lipschitz constant L can be derived as:

$$L_U = \max_k (\lambda_{\max}(X_k^T X_k)) + \frac{\lambda^2 + 2 * \gamma^2 \max_{k \in V} d_k}{\mu} \quad (12)$$

where $\lambda_{\max}(M)$ is the largest eigenvalue of M and $d_k = \sum_{k'}^K \tau_{k,k'}$.

With the derivation of the gradient in (11) and the Lipschitz bound in (12), we can now apply Nesterov's method [Nesterov, 2005] for optimizing (10). The procedure is summarized in Algorithm 2. For more details on the proximal approach see Chen et al. [2010].

¹It would be possible to reformulate the first part of the objective in matrix form $Y_{diag} - X_{diag} B_{diag}$ by defining X_{diag} as a block-diagonal matrix as in Section 2.5.1, defining B_{diag} as a $pK \times K$ matrix with β_k along the diagonals and similarly y_{diag} as an $n \times K$ matrix with

Algorithm 2 Proximal Gradient Optimization

```

1: procedure PROXIMAL( $n_{iter}, X, Y, B_{init}, \lambda, \gamma, \tau, L_U, \mu$ )
2:    $i \leftarrow 0$ 
3:    $W^0 \leftarrow B_{init}$ 
4:   while not_converged AND  $i < n_{iter}$  do
5:     Compute  $\Delta \tilde{f}_{L1}(W^i)$  according to (11).
6:      $B^i \leftarrow W^i - \frac{1}{L_U} \Delta \tilde{f}_{L1}(W^i)$ 
7:      $Z^i \leftarrow -\frac{1}{L_U} \sum_{j=0}^i \frac{j+1}{2} \Delta \tilde{f}_{L1}(W^j)$ 
8:      $W^{i+1} \leftarrow \frac{i+1}{i+3} B^i + \frac{2}{i+3} Z^i$ 
9:      $i \leftarrow i + 1$ 

```

3 Simulation Study

To test the performance of the proposed approach, we simulated data from a model based on characteristics of a recent cancer dataset, the Cancer Cell Line Encyclopedia [CCLE; Barretina et al., 2012]. We treat cancer types as subgroups. To simulate data, we first estimated means and covariance matrices μ_k, Σ_k for each of $K = 9$ subgroups (the eight cancer types with the latest sample sizes in CCLE plus a ninth for all other cancer types; covariances were estimated using the graphical lasso). For each group k , we then sampled covariates from the multivariate normal $\mathcal{N}(\mu_k, \Sigma_k)$. For a given total sample size n , subgroup sizes were consistent with those in the original data. We used a random subset of 200 gene expression levels (i.e. the dimensionality was fixed at $p = 200$). This parametric approach allowed us to vary sample sizes freely, including the case of total n larger than in the original dataset. The set-up is intended to roughly reflect the correlation structure of the covariates, but we do not expect it to capture all aspects of the real data.

We are interested in the situation in which it may be useful to share information between subgroups. But we are also interested in investigating performance in settings that do not agree with our model formulation (the extreme cases being where subgroups are either entirely dissimilar or identical). Let $V = \{1 \dots K\}$ be the set of subgroup indices (here, $K = 9$). We set regression coefficients to be identical in a subset $V_0 \subseteq V$ of the subgroups, such that the size $K_0 = |V_0|$ of the subset governs the extent to which fusion could be useful. Specifically, if $K_0 = K$, all subgroups have the same regression coefficients (i.e. favoring a pooled analysis using a single regression model) and at the other extreme if $K_0 = 1$ all groups have differently drawn coefficients. Intermediate values of K_0 give differing levels of similarity.

For a given value K_0 , we defined membership of V_0 by considering the differences between the subgroup-specific models for the covariates. Specifically, we choose the K_0 groups that minimized the sum of symmetrised KL divergences between subgroup-specific models. A coefficient vector was then drawn

y_k along the diagonals; however, this formulation is neither practical nor intuitive, and the gain in notational simplicity is negligible.

separately for each subgroup $k \notin V_0$ and one, shared coefficient vector drawn for all $k \in V_0$. Each draw was done as follows. We first sampled a binary vector b of length p from a Bernoulli, i.e. $b_i \sim \text{Bernoulli}(0.1)$. Then we drew $\beta_i \sim \mathcal{N}_{trunc}(0, 1)$ if $b_i = 1$ and set $\beta_i = 0$ otherwise, where $\mathcal{N}_{trunc}(0, 1)$ denotes a standard Normal with the interval $(-0.1, 0.1)$ excluded (this is to ensure non-zero coefficients are not very small in magnitude). Note that in the case of $K_0 = 1$, all groups have separately drawn coefficients and the between-subgroup KL divergence plays no role.

We compare our approaches with pooled and subgroup-wise analyses. These are performed using classical lasso (we use the `glmnet` implementation) on respectively the whole dataset or each subgroup separately.

Figure 1 shows performance when varying the number K_0 of subgroups with shared coefficients, with the total number of samples fixed at $n = 250$. Here, a smaller value of K_0 corresponds to less similarity between subgroup-specific coefficients in the underlying models. At intermediate values of K_0 the fusion approaches offer gains over pooled and subgroup-wise analyses. This is because the pooled analyses are mis-specified due to the inhomogeneity of the data, while the subgroup-wise analyses, although correctly specified, must confront limited sample sizes since they analyze each subgroup entirely separately. In contrast, the fusion approaches are able to pool information across subgroups, but also allow for subgroup-specific coefficients. Importantly, even at the extremes of $K_0 = 1$ (separately drawn coefficients for each subgroup) and $K_0 = 9$ (all subgroups have exactly the same coefficients), the fusion approaches perform well. This demonstrates their flexibility in adapting the degree of fusion.

Figure 2 shows performance as a function of total sample size n . Here, the number of subgroups with identical coefficients is fixed at $K_0 = 4$. This gives a relatively weak opportunity for information sharing, since 5/9 groups have separately drawn coefficients. Since the true β_k 's are not identical, the pooled analysis is mis-specified and accordingly even at large sample sizes, it does not catch up with the other approaches. As expected, subgroup-wise analyses perform increasingly well at larger sample sizes. However, at smaller sample sizes the fusion approaches show some gains.

The ℓ_1 and ℓ_2 fusion approaches seem similar in performance. Our ℓ_2 implementation leverages the `glmnet` package and is more computationally efficient than the ℓ_1 approach. For computational convenience, in examples below we show results from the ℓ_2 approach only.

4 Alzheimer's disease: prediction of cognitive scores

Here, we use data from the Alzheimer's Disease Neuroimaging Initiative (ADNI) [Mueller et al., 2005] to explore the ability of fusion approaches to estimate regression models linking clinical and genetic covariates to disease progression, as captured by cognitive test scores.

In 2014, ADNI made a subset of its data available for a DREAM challenge [Allen et al., 2016] and we use these data here. The dataset consists of a total

Figure 1: Simulated data, performance for varying K_0 . Here, the number of subgroups is fixed at $K = 9$ of which K_0 have shared coefficients in the underlying data-generating model (see text for details of simulation set-up). A smaller K_0 corresponds to less similarity between underlying subgroup-specific models, with $K_0 = 1$ representing the case where all subgroups have separately drawn coefficients while $K_0 = 9$ represents an entirely homogenous model in which each subgroup has exactly the same regression coefficients. The total sample size is fixed at $n = 250$. Upper panel: Weighted root mean squared error (RMSE). RMSE is weighted by subgroup sizes. Lower panel: Area under the ROC curve (AUROC; with respect to the true sets of active variables with non-zero coefficients).

Figure 2: Simulated data, performance at varying sample sizes. Here, the number of subgroups is fixed at $K = 9$ of which $K_0 = 4$ have shared coefficients in the underlying data-generating model (see text for details of simulation set-up). Top left: Root mean squared error (RMSE; weighted by subgroup sizes). Top right: Area under the ROC curve (AUROC; with respect to the true sets of active variables with non-zero coefficients). Bottom: Computational time taken in log seconds.

of $n = 767$ individuals who were followed up over at least 24 months. Cognitive function was evaluated using the mini-mental state examination (MMSE). At baseline, individuals were classified as either cognitively normal (CN), early mild cognitive impairment (EMCI), late mild cognitive impairment (LMCI) or diagnosed with Alzheimer's disease (AD). These form clinically-defined subgroups for our analysis. For the present analysis, we use only genetic data (single nucleotide polymorphisms or SNPs) and clinical profile as covariates and disregard the neuroimaging data.

The task is to predict the slope of MMSE scores over a 24-month period. The total number of SNPs available is $\sim 10^7$. Filtering by linkage disequilibrium reduces this to $\sim 2 \times 10^6$. For computational ease, we pre-selected 20,000 of this latter group that gave the smallest residuals when regressed with the clinical variables against responses in the training set. We note that this biases our analyses, but we emphasize that our goal in this section is not biomarker

discovery but comparison between approaches all using the same (pre-selected) covariates.

Figure 3 shows root mean squared error (RMSE) separately for each of the four subgroups. The fusion approaches offer substantial gains compared with pooled and subgroup-wise analyses (the latter performed very badly and are not shown in the figure). The biggest gain is for the AD subgroup. For the weighted fusion analysis the tuning parameters $\tau_{k,k'}$ were set using the distance between the means of each subgroup (in the space of genetic and clinical variables). Weighting did not appear to improve performance.

Figure 4 shows scatter plots of predicted MMSE slopes versus the true slopes. The predictions shown were obtained in a held-out fashion via 10-fold cross-validation (CV), as were the RMSE and Pearson correlations shown.

Figure 5 shows a comparison of the estimated regression coefficients themselves. The subgroup-wise approach is much sparser than the other methods, likely due to the fact that it must operate entirely separately on each (relatively small-sample) subgroup. The pooled approach finds more influential variables but obviously there is no subgroup-specificity. The fused approach selects more variables than the subgroup-wise analysis, but there are many instances of subgroup-specificity in the estimates.

Figure 3: Alzheimer's disease prediction results, ADNI data. Box plots showing difference in RMSE of fused methods compared to the pooled linear regression model (higher values indicate better performance by the fused methods). [Subgroup-wise analysis performed less well than pooled and is not shown; box-plots are over 10-fold cross-validation.]

Figure 4: Alzheimers disease data, predicted vs. observed responses. Scatter plots show predicted and observed 24-month slopes for each of the standard and fused linear regression models. All predictions were obtained via 10-fold cross-validation.

Figure 5: Alzheimers disease data, estimated regression coefficients. Heatmap showing estimated regression coefficients for the clinical variables and a representative subsample of the SNPs. Absolute coefficients are thresholded at e^{-2} to improve readability.

5 Prediction of therapeutic response in cancer cell lines

The Cancer Cell Line Encyclopedia [CCLE, Barretina et al., 2012] is a panel of 947 cancer cell lines with associated molecular measurements and responses to 24 anti-cancer agents. Here, we use these data to explore group-structured regression. We treat the area above the dose-response curve as the response and use expression levels of $\sim 20,000$ human genes as covariates. We treat the cancer types as subgroups k . After discarding cell lines with missing values, we arrive at $n \sim 500$ samples.

Figure 6: Cancer cell line therapeutic response prediction. Difference in weighted RMSE between L2 fusion approach and a pooled analysis. Results shown over 24 responses (anti-cancer agents) using data from the Cancer Cell Line Encyclopedia (CCLE); the dashed vertical line at zero indicates no difference, boxplots to the right indicate improvement (lower RMSE) over pooled.

Figure 6 shows results over all 24 responses (anti-cancer agents). We observe that for most responses the L2 fusion approach shows either improved or similar

prediction performance to pooled in terms of RMSE (weighted by subgroup size). Weighted fusion shows a similar performance to unweighted fusion.

Figure 7: Cancer cell line therapeutic response prediction, broken down by subgroup (cancer type) for agents PD-0332991 and PLX4720.

Figure 7 shows results broken down by subgroup for two examples (responses PD-0332991 and PLX4720). In the former case, the fusion approaches largely outperform pooled and subgroup-wise analyses. In the second, pooled is the best performer, although the fusion approaches are similar in most subgroups.

6 ALS: prediction of disease progression

Amyotrophic lateral sclerosis (ALS) is an incurable neurodegenerative disease that can lead to death within three to four years of onset. However, about ten percent of patients survive more than 10 years. Prediction of disease progression remains an open question. We use data from the PROACT database, specifically data that were used in the 2015 DREAM ALS Stratification Prize4Life Challenge (data were retrieved from the PROACT database on 22/06/2015). Our aim is not to optimize predictive performance *per se* but rather to provide a case study exploring the use of fusion approaches in a moderate-dimensional, clinical data setting. In contrast to the Alzheimers example above, here the data are less high-dimensional and the subgrouping less clear cut (see below).

The data consist of observations from $n = 2,393$ patients. Each patient was enrolled in a clinical trial and followed up for a minimum of 12 months after the start of the trial. Disease progression is captured by a clinical scale called the ALS Functional Rating Scale (ALSFRS). The task is to predict the slope of the ALSFRS score from 3 to 12 months (after the start of the trial). For each patient, available covariates include ALSFRS scores for the 0-3 month period, demographic information and longitudinal measurements of clinical variables. We follow the featurization and imputation procedures devised by Mackay [see Küffner et al., 2015] and obtain a total of $p = 615$ covariates.

Subgroups were defined as follows. The first subgroup consists of patients with disease onset before the start of the trial. The second subgroup consists of patients for whom onset was after the start of the trial and who have negative

ALSFRS slope. The third subgroup of patients also had onset after the start of the trial but positive ALSFRS slope. Thus, the subgroups reflect severity of onset. As we believe that the pre-trial onset group are likely to differ most from the others, we manually set the distance between groups 1 and the other two groups to 1 (the maximum), and set the distance between groups 2 and 3 to 0.1.

Figure 8 shows (held-out) RMSEs by subgroup; we see that the largest improvement in prediction performance is in subgroup 1. The fusion approach leads to a modest improvement. The difference between weighted and un-weighted fusion is negligible².

Figure 8: ALS prediction results. Box plots showing difference in RMSE of fused methods compared with the pooled linear regression model (higher values indicate better performance by the fused methods). [Subgroup-wise analysis performed less well than pooled and is not shown; boxplots are over 10-fold cross-validation.]

7 Conclusions

Many biomedical datasets are heterogenous, spanning multiple disease types (or other biological contexts) that are related but also expected to have specific underlying biology. This means that large datasets are often usefully thought of as comprising several smaller datasets, that have similarities but that cannot be assumed to be identically distributed. Statistically efficient regression in these group-structured settings requires ways to pool information where useful to do

²This dataset is different from the one reported in [Küffner et al., 2015], with larger variance in the slopes, and so RMSE values are not directly comparable; however, we note that performance for our methods compares favorably with that reported in the reference.

so, whilst retaining the possibility of subgroup-specific parameters and sparsity patterns. We proposed a penalized likelihood approach for high-dimensional regression in the group-structured setting that provides group-specific estimates with global sparsity and that allows for information sharing between groups.

In any given application, even when there are good scientific reasons to suspect differences in regression models between subgroups, it is hard to know in advance whether the nature of any differences is such that a specific kind of joint estimation would be beneficial. For example, if sample sizes are small and groups only slightly different, pooling may be more effective, or if the groups are entirely different, fusion of the kind we consider may not be useful. This means that in practice, either simple pooling or subgroup-wise analysis may be more effective than fusion. In our approach, the tuning parameter γ (set by cross-validation) determines the extent of fusion in a data-adaptive manner, and we saw in several examples that this appears successful in giving results that are at worst close to the best of pooling and subgroup-wise analyses. For settings with widely divergent n_k 's, it may be important to allow tuning parameters to depend on n_k (we did not do so) and to consider alternative formulations that allow for asymmetric fusion.

An appealing feature of our approach is that it allows for subgroup-specific sparsity patterns and parameter estimates that may themselves be of scientific interest. We discussed point estimation, but did not discuss uncertainty quantification for these subgroup-specific estimates. A number of recent papers have discussed significance testing for lasso-type models [see e.g. Wasserman and Roeder, 2009, Lockhart et al., 2014, Städler and Mukherjee, 2016] and we think some of these ideas could be used with the models proposed here.

8 Software Availability

The R code used for the experiments in this paper has been made available as R package `fuser` on GitHub: <https://github.com/FrankD/fuser>. Scripts for reproducing the results in this paper can be obtained at: <http://fhm-chicas-code.lancs.ac.uk/dondelin/SubgroupFusionPrediction>.

9 Acknowledgements

Data collection and sharing for the Alzheimer's data application was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimers Association; Alzheimers Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-

La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimers Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

References

- Rehan Akbani, Patrick Kwok Shing Ng, Henrica MJ Werner, Maria Shahmoradgoli, Fan Zhang, Zhenlin Ju, Wenbin Liu, Ji-Yeon Yang, Kosuke Yoshihara, Jun Li, et al. A pan-cancer proteomic perspective on the cancer genome atlas. *Nature Communications*, 5, 2014.
- Genevra I Allen, Nicola Amoroso, Catalina Anghel, Venkat Balagurusamy, Christopher J Bare, Derek Beaton, Roberto Bellotti, David A Bennett, Kevin L Boehme, Paul C Boutros, et al. Crowdsourced estimation of cognitive decline and resilience in alzheimer’s disease. *Alzheimer’s & Dementia*, 12 (6):645–653, 2016.
- Jordi Barretina, Giordano Caponigro, Nicolas Stransky, Kavitha Venkatesan, Adam A Margolin, Sungjoon Kim, Christopher J Wilson, Joseph Lehár, Gregory V Kryukov, Dmitriy Sonkin, et al. The cancer cell line encyclopedia enables predictive modelling of anticancer drug sensitivity. *Nature*, 483(7391): 603–607, 2012.
- Xi Chen, Seyoung Kim, Qihang Lin, Jaime G. Carbonell, and Eric P. Xing. Graph-structured multi-task regression and an efficient optimization method for general fused lasso. *arXiv:1005.3579 [cs, math, stat]*, May 2010.
- Patrick Danaher, Pei Wang, and Daniela M Witten. The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(2):373–397, 2014.
- Jerome Friedman, Trevor Hastie, Holger Hoefling, and Robert Tibshirani. Pathwise coordinate optimization. *The Annals of Applied Statistics*, 1(2):302–332, December 2007. ISSN 1932-6157, 1941-7330. doi: 10.1214/07-AOAS131. URL <http://projecteuclid.org/euclid.aoas/1196438020>.

- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. 9(3):432–441, 2008. ISSN 1465-4644, 1468-4357. doi: 10.1093/biostatistics/kxm045. URL <http://biostatistics.oxfordjournals.org/content/9/3/432>.
- Jerome Friedman, Trevor Hastie, and Rob Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1, 2010.
- Holger Hoefling. A path algorithm for the fused lasso signal approximator. *Journal of Computational and Graphical Statistics*, 19(4):984–1006, 2010.
- Robert Küffner, Neta Zach, Raquel Norel, Johann Hawe, David Schoenfeld, Liuxia Wang, Guang Li, Lilly Fang, Lester Mackey, Orla Hardiman, et al. Crowdsourced analysis of clinical trial data to predict amyotrophic lateral sclerosis progression. *Nature Biotechnology*, 33(1):51–57, 2015.
- Jun Liu, Lei Yuan, and Jieping Ye. An efficient algorithm for a class of fused lasso problems. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 323–332. ACM, 2010. URL <http://dl.acm.org/citation.cfm?id=1835847>.
- Richard Lockhart, Jonathan Taylor, Ryan J Tibshirani, and Robert Tibshirani. A significance test for the lasso. *Annals of Statistics*, 42(2):413, 2014.
- Susanne G Mueller, Michael W Weiner, Leon J Thal, Ronald C Petersen, Clifford R Jack, William Jagust, John Q Trojanowski, Arthur W Toga, and Laurel Beckett. Ways toward an early diagnosis in alzheimer’s disease: the alzheimer’s disease neuroimaging initiative (adni). *Alzheimer’s & Dementia*, 1(1):55–66, 2005.
- Yu Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.
- Chris J Oates, Jim Korkola, Joe W Gray, and Sach Mukherjee. Joint estimation of multiple related biological networks. *The Annals of Applied Statistics*, 8(3):1892–1919, 2014.
- Chris J Oates, Jim Q Smith, Sach Mukherjee, and James Cussens. Exact estimation of multiple directed acyclic graphs. *Statistics and Computing*, pages 1–15, 2015.
- Guillaume Obozinski, Ben Taskar, and Michael I Jordan. Joint covariate selection and joint subspace selection for multiple classification problems. *Statistics and Computing*, 20(2):231–252, 2010.
- Nicolas Städler and Sach Mukherjee. Two-sample testing in high-dimensional models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) (To appear)*, 2016.

- Robert Tibshirani, Michael Saunders, Saharon Rosset, Ji Zhu, and Keith Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108, 2005.
- Larry Wasserman and Kathryn Roeder. High dimensional variable selection. *Annals of Statistics*, 37(5A):2178, 2009.
- John N Weinstein, Eric A Collisson, Gordon B Mills, Kenna R Mills Shaw, Brad A Ozenberger, Kyle Ellrott, Ilya Shmulevich, Chris Sander, Joshua M Stuart, Cancer Genome Atlas Research Network, et al. The cancer genome atlas pan-cancer analysis project. *Nature Genetics*, 45(10):1113–1120, 2013.
- Gui-Bo Ye and Xiaohui Xie. Split bregman method for large scale fused lasso. *Computational Statistics & Data Analysis*, 55(4):1552–1569, 2011.
- Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.