

Doubly structured sparsity for grouped multivariate responses with application to functional outcome score modeling

Jared D. Huling¹  | Jennifer P. Lundine^{2,3} | Julie C. Leonard^{4,5}

¹Division of Biostatistics, University of Minnesota, Minneapolis, Minnesota, USA

²Department of Speech and Hearing Science, The Ohio State University, Columbus, Ohio, USA

³Division of Clinical Therapies and Inpatient Rehabilitation Program, Nationwide Children's Hospital, Columbus, Ohio, USA

⁴Division of Emergency Medicine, Department of Pediatrics, The Ohio State University College of Medicine, Columbus, Ohio, USA

⁵Abigail Wexner Research Institute, Nationwide Children's Hospital, Columbus, Ohio, USA

Correspondence

Jared D. Huling, Division of Biostatistics, University of Minnesota, Minneapolis, MN, USA.

Email: huling@umn.edu

Funding information

Eunice Kennedy Shriver National Institute of Child Health and Human Development, Grant/Award Number: R03HD101083

Summary

This work is motivated by the need to accurately model a vector of responses related to pediatric functional status using administrative health data from inpatient rehabilitation visits. The components of the responses have known and structured interrelationships. To make use of these relationships in modeling, we develop a two-pronged regularization approach to borrow information across the responses. The first component of our approach encourages joint selection of the effects of each variable across possibly overlapping groups of related responses and the second component encourages shrinkage of effects towards each other for related responses. As the responses in our motivating study are not normally-distributed, our approach does not rely on an assumption of multivariate normality of the responses. We show that with an adaptive version of our penalty, our approach results in the same asymptotic distribution of estimates as if we had known in advance which variables have non-zero effects and which variables have the same effects across some outcomes. We demonstrate the performance of our method in extensive numerical studies and in an application in the prediction of functional status of pediatric patients using administrative health data in a population of children with neurological injury or illness at a large children's hospital.

KEYWORDS

fused lasso, hierarchical sparsity, high dimensional data, risk prediction, variable selection

1 | INTRODUCTION

Each year in the United States more than 20,000 children are diagnosed with an acute neurologic injury or illness that result in debilitating physical and cognitive complications, and reduced quality-of-life.¹⁻⁴ These events are costly. As an example, it is estimated that one billion dollars are spent annually on management of pediatric traumatic brain injury (TBI)-associated hospitalizations.⁵ Though neurologic illnesses and injuries as a category are a leading cause of morbidity in children, each individual diagnosis is uncommon. These low disease-specific incidences make the study of rehabilitation interventions for children with neurologic injuries or illnesses challenging. Further, since many of the

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2023 The Authors. *Statistics in Medicine* published by John Wiley & Sons Ltd.

material effects manifest in the long term and can change with later child development, it is important to be able to track outcomes over time.

The WeeFIM® is a validated scoring system for appraising functional ability in children and provides valuable information about a multitude of components of health in children with neurologic injury or illness. WeeFIM® is administered and scored by trained assessors and has been demonstrated to have high interrater reliability. WeeFIM® has been shown to be predictive of longitudinal functional recovery of children with neurological disorders and can be used for discharge planning, prediction of functional outcomes, and documentation of functional performance of children over time to assess recovery or decline. Yet, because the WeeFIM® system requires trained assessors, scores for all children are not always available. WeeFIM® training and subscription is expensive and time-intensive and thus there is significant interest in being able to assess broadly the functional ability of children across a health system without the need for explicit scoring by trained assessors. One avenue for doing so is to relate administrative data sources to WeeFIM® scores to facilitate care management and identify individuals who may require additional care or rehabilitation. Our interest is thus in building predictive models to understand functional ability of children using administrative health data in scenarios where WeeFIM® is not used. WeeFIM® scores have particular characteristics and organization. In this work we seek to leverage these characteristics to facilitate the building of accurate and interpretable WeeFIM® risk models. We now describe these characteristics and how we aim to utilize them in this paper.

WeeFIM® is comprised of 18 component scores, each of which is measured on a 7 point Likert scale with 7 indicating complete independence and 1 indicating complete dependence on others to perform various tasks. The WeeFIM® component scores are categorized into three main domains: mobility, cognition, and self-care, and represent three distinct clinical concepts related to functional outcomes. We describe the components of each domain in Section 4. In our dataset, the correlations of the WeeFIM® component scores align closely with these domains. We computed the distance correlations⁶ of all pairs of the 18 WeeFIM® components to form an 18 × 18 distance matrix and ran a hierarchical clustering algorithm on the resulting distance matrix. As can be seen in Figure 1, the scores in the cognition domain are immediately completely separated from the scores in mobility and self-care; self-care and mobility separate from each other later in the hierarchical clustering.

In our motivating study, we aim to use information commonly present in administrative health data including International Classification of Diseases (ICD)-9 and -10 codes, Current Procedural Terminology (CPT) codes, pharmaceutical codes, data indicating durable medical equipment use, and basic demographic information to predict and evaluate functional outcomes in children admitted to an inpatient rehabilitation unit with a diagnosis of neurologic injury or illness




FIGURE 1 Heatmap of the distance correlations of the WeeFIM® scores. The dendrogram from hierarchical clustering of the scores shows that they group naturally by the three major WeeFIM® domains: self care, mobility, and cognition; thus the pre-defined domains align well with the natural variation in the components in our motivating study.

at a large children's hospital. Our code information is high-dimensional, with well over 500 billing codes as potential predictors of a child's functional status.

To deal with high-dimensionality, one could fit lasso-penalized regression models⁷ for each of the 18 scores separately, however this would ignore the inherent relationships and similarities between the components. First, as some WeeFIM[®] components are highly related (eg, dressing lower and dressing upper, bowel management and bladder management, bathing and tub shower), it is plausible that some variables (eg, billing codes) may have the same or similar effects across some of the WeeFIM[®] components. Similarly, some variables may not have the same effect size across some WeeFIM[®] components, but are likely to be jointly predictive for multiple components, especially for related components, or even for all components. A complication is that while linearity of effects of covariates on the components may be reasonable, an assumption of normality of the responses or error terms is not, as the WeeFIM[®] component scores are not normally-distributed.

There exists some work for building models for multivariate responses using high-dimensional predictors. Rothman, et al⁸ introduced a lasso-based estimation procedure that incorporates variable selection and covariance estimation for multivariate normal responses and developed an iterative estimation procedure that makes use of the covariance structure among the responses by penalizing the inverse covariance matrix using the methods of Yuan and Lin.⁹ However, this approach does not make use of similar predictor-response relationships across related responses and only borrows information across responses through their covariance structure. Sofer et al¹⁰ extended these ideas to incorporate a wider variety of penalties such as the SCAD penalty¹¹ and the MC+ penalty¹² and developed new algorithms for estimation. Li et al¹³ incorporated group structure both in responses and among predictors using an overlapping group lasso penalty,¹⁴ which allows for borrowing of strength across responses through joint selection of variables across related responses.

In our work, we borrow strength across responses in two ways through two different structured sparsity inducing penalties. In particular, we first aim to incorporate joint selection and removal of variables across all responses *and* joint selection and removal of variables across related responses (ie, responses that are in the same functional domain). In our application, the group structure is the natural groups formed by the WeeFIM[®] components displayed in Figure 1. Thus, we require use of a group lasso¹⁵ with overlapping groups¹⁴ similar to Li et al.¹³ Second, we use a fused lasso¹⁶ for shrinkage of the effects of a variable across related outcomes to borrow strength more explicitly in estimation by partially collapsing models across different responses into a single model for individual variables separately. The fused lasso penalty allows a variable's effects across related responses to be estimated to be exactly the same. We prove that with adaptively chosen weights^{17,18} but for the overlapping group lasso as in Huling et al¹⁹ and fused lasso as in Viallon et al,²⁰ our doubly-structured sparsity inducing estimator has an oracle property. The theoretical results are general in that they allow for any arbitrary overlapping group structure for the group lasso penalty and any arbitrary fused lasso penalty. The oracle property we show demonstrates that estimation for the non-zero coefficients has the same asymptotic distribution as if we had known in advance both (i) which coefficients are non-zero and (ii) which coefficients are equal to each other and fit models with only the non-zero coefficients and forcing the truly equal coefficients to be equal to each other. We also show selection consistency for the non-zero coefficients and selection consistency for which coefficients are equal to each other. Neither our methodology nor our theory requires an assumption of normality or independence of the responses, as we work under a semiparametric multivariate linear model assumption, where we make no parametric distributional assumptions about the error terms. Further, our theory, methodology, and computational approach allows for general use of the overlapping group lasso in combination with the fused lasso beyond our application to multivariate response regression.

The remainder of our paper is organized as follows. In Section 2.1, we introduce the key definition of our methodology. In Section 2.5, we develop theory for our methodology using adaptive regularization under semiparametric multivariate linear models and in Section 3, we investigate the operating characteristics of our method in small samples using simulation experiments. In Section 4, we demonstrate the use of our methodology on our motivating application involving pediatric functional outcomes score modeling. Finally, we conclude with discussion.

2 | DOUBLY STRUCTURED SPARSITY FOR MULTIVARIATE RESPONSES

2.1 | The overlapping group + fused method

We assume that the observed multivariate response $\mathbf{y}_i = (y_{i1}, \dots, y_{iK})^\top$ follows the semiparametric linear model

$$\mathbf{y}_i = \boldsymbol{\beta}^{0\top} \mathbf{x}_i + \boldsymbol{\epsilon}_i \text{ for } i = 1, \dots, N, \quad (1)$$

where $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^\top$ is a length p vector of predictors, $\boldsymbol{\beta}^0$ is a $p \times K$ matrix of regression coefficients with k th column equal to $\boldsymbol{\beta}_{\cdot,k}^0 = (\beta_{1,k}^0, \dots, \beta_{p,k}^0)^\top$, and $\boldsymbol{\epsilon}_i = (\epsilon_{i1}, \dots, \epsilon_{iK})^\top$ is a vector of errors with mean zero and finite, positive definite variance-covariance matrix $\boldsymbol{\Sigma}$; no requirement that the response vector be continuous is required, as we do not assume a particular distribution for $\boldsymbol{\epsilon}_i$. Denote the j th variable's coefficients across the K response as $\boldsymbol{\beta}_j^0 = (\beta_{j,1}^0, \dots, \beta_{j,K}^0)^\top$. Let \mathbf{X} denote the $N \times p$ matrix of predictors with i th row as \mathbf{x}_i^\top and j th column as $\mathbf{x}_{\cdot,j}$ and let \mathbf{Y} denote the $N \times K$ response matrix with i th row as \mathbf{y}_i^\top and denote the k th column of \mathbf{Y} as $\mathbf{Y}_{\cdot,k}$. The typical least squares estimator is the solution of

$$\underset{\boldsymbol{\beta}}{\operatorname{argmin}} (2N)^{-1} \operatorname{tr} [(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})^\top],$$

which is simply $(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{Y}$, where for a matrix \mathbf{A} , $\operatorname{tr}(\mathbf{A})$ indicates its trace.

In our setting, the K outcomes are grouped into M different, possibly overlapping groupings that describe inherent relationships between the outcomes. In our motivating data, the WeeFIM outcomes are grouped into three pre-defined domains that represent different categories of functional status. In other settings, these domains may be further subcategorized. In this manner, the groupings generally correspond to hierarchically-organized categorizations of the outcomes. In this work, we make use of the natural grouping of the outcomes by adding regularizers that induce the coefficients of each different variable to be selected or removed jointly across an entire group of the responses. The j th grouping \mathcal{G}_j consists of a set of g_j groups $G_{j,1}, \dots, G_{j,g_j}$ such that the union of the groups includes all of the K outcomes. Generally, since the groupings will be hierarchically defined, g_j will be larger than g_{j+1} and groups in \mathcal{G}_j can be expressed as unions of groups in \mathcal{G}_{j+1} . More formally, the groupings are written as $\mathcal{G}_1 = \{G_{1,1}, \dots, G_{1,g_1}\}, \dots, \mathcal{G}_M = \{G_{M,1}, \dots, G_{M,g_M}\}$, where $G_{m,g} \in \{1, \dots, K\} \equiv \mathcal{K}$ and $\bigcup_{G \in \mathcal{G}_m} G = \mathcal{K}$ for $m = 1, \dots, M$. To encourage a variable to be selected jointly across all responses, we define the trivial group $\mathcal{G}_0 = \{\mathcal{K}\}$, thus $G_{0,1} = \mathcal{K} = \{1, \dots, K\}$. To allow variables to be selected or removed individually, we further include the group $\mathcal{G}_{M+1} = \{\{1\}, \dots, \{K\}\}$, thus $G_{M+1,k} = \{k\}$. The group structure for a toy example with $K = 8$ responses and $M = 1$ groupings is displayed in Figure 2. We note that in this toy example, the groups are overlapping (eg, $G_{1,1} \subset G_{0,1}$) because the groups are hierarchically structured. In our motivating application, the group structure we use involves such overlapping groups.

In our setting the outcome groupings are pre-defined, but in general when such a pre-defined set of groups are not available *a priori*, the M groupings can be formed by iteratively refined, hierarchical groupings of the responses, as illustrated in Figure 1; for example, the M groupings could be formed by clustering the responses at M different levels in a hierarchical clustering of the responses.

For a particular group $G \in \mathcal{K}$ with size $|G|$, define the length $|G|$ subvector of $\boldsymbol{\beta}_j$, limited to the outcomes in group G as $\boldsymbol{\beta}_{j,G}$. We propose to add to the above objective function two regularizers that induce sparsity in a structured manner. Let $\hat{\boldsymbol{\beta}}$ be the solution of the following problem

$$\underset{\boldsymbol{\beta}}{\operatorname{argmin}} (2N)^{-1} \operatorname{tr} [(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})^\top] + \lambda_1 P_1(\boldsymbol{\beta}) + \lambda_2 P_2(\boldsymbol{\beta}) \tag{2}$$




FIGURE 2 An illustration of the structure of the two components of the penalization methods used to induce group-wise selection of covariate effects and fuse them to be more similar. **Blue** indicates joint selection across all responses via the group \mathcal{G}_0 , **red** indicates joint selection by response group $\mathcal{G}_M = \mathcal{G}_1$, **green** indicates individual coefficient selection achieved via the group \mathcal{G}_{M+1} , and **black** indicates shrinkage towards a common effect within response group $G_{1,k}$.

where

$$P_1(\boldsymbol{\beta}) = \sum_{j=1}^p \sum_{m=0}^{M+1} \sum_{G \in \mathcal{G}_m} \lambda_{1,j,G} \|\boldsymbol{\beta}_{j,G}\|_2$$

and

$$P_2(\boldsymbol{\beta}) = \sum_{j=1}^p \sum_{G \in \mathcal{G}_M} \sum_{l,o \in G: l \neq o} \lambda_{2,j,l,o} |\beta_{j,l} - \beta_{j,o}|,$$

where the penalty $\lambda_{1,j,G} \|\boldsymbol{\beta}_{j,G}\|_2$ is an overlapping group lasso penalty encourages joint selection or removal of the effects of variable j across the responses in response group G . All possible zero patterns of coefficients can be represented as unions of groups due to the results of Jenatton et al,¹⁴ hence nonzero patterns of coefficients can be thought of as complements of unions of groups. Thus, the grouping \mathcal{G}_0 encourages joint selection and removal of effects for a variable across all outcomes simultaneously and grouping \mathcal{G}_{M+1} allows for individual selection or removal for each outcome separately, while the groupings $\mathcal{G}_1, \dots, \mathcal{G}_M$ allow for selection or removal of effects across related groups. The second penalty $\lambda_{2,j,l,o} |\beta_{j,l} - \beta_{j,o}|$ is a fused lasso penalty that encourages the effect of variable j on response l and on response o to be more similar. Thus, $P_1(\boldsymbol{\beta})$ and $P_2(\boldsymbol{\beta})$, respectively, help borrow strength across the K responses by (1) leveraging the natural groupings of the responses and taking advantage of joint significance of predictors across related responses and (2) encouraging the effect estimates for a particular variable to be similar across the most related responses according to the response groupings. These two regularizers incorporate structural knowledge about outcomes in two manners by utilizing two different types of structured sparsity inducing penalties. The overall structure of the regularizers are depicted in Figure 2 in a toy example with $K = 8$ responses. In our formulation, we only add a fused lasso penalty within the final grouping \mathcal{G}_M , reflective of the notion that with an iteratively refined grouping, the final groups are likely to have more similar outcomes within each group. However, the fused lasso penalty can be added within any grouping and can even be set so that fused lasso terms are included for every possible pair of responses; our asymptotic results cover any arbitrary set of fused pairs. In our application, the multivariate outcomes are measured on the same scale, but our approach can be adapted to outcomes that have different scales by standardizing the outcomes prior to analysis. Further, as we list below, our adaptive penalization allows for the fused lasso penalty to adapt to observed differences in effects across outcomes.

To perform adaptive penalization, we take $\lambda_{1,j,G} = \|\hat{\boldsymbol{\beta}}_{j,G}^{OLS}\|_2^{-\gamma_1}$ and $\lambda_{2,j,l,m} = |\hat{\beta}_{j,l}^{OLS} - \hat{\beta}_{j,m}^{OLS}|^{-\gamma_2}$, where $\gamma_1, \gamma_2 > 0$ and $\hat{\boldsymbol{\beta}}^{OLS} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{Y}$ is the ordinary least squares estimate of $\boldsymbol{\beta}$. We study this adaptive version of our estimator in Section 2.5. For non-adaptive penalization, we take $\lambda_{1,j,G} = |G|^{1/2}$, as is used in Jenatton et al¹⁴ and Huling et al¹⁹ and $\lambda_{2,j,l,m} = 1$, however for this choice selection consistency is not guaranteed without an irreproducible condition.¹⁴ In order to get reliable estimates of out-of-sample performance using cross validation when the group structure is derived in a data-driven manner, the procedure used to estimate the group structure should be applied in each cross validation fold instead of being fixed.

2.2 | Implementation details

In our implementation, we use a generic ADMM algorithm described in the next section for computation and feed into it the response vector $\tilde{\mathbf{Y}}$ and design matrix $\tilde{\mathbf{X}}$, where the latter is treated as a sparse matrix object that only stores the values and locations of the non-zero entries of $\tilde{\mathbf{X}}$, which both dramatically saves space, memory used, and reduces computation time. The sparse matrix object is of the type provided in the Eigen C++ linear algebra library²¹ with interface to an R package through RcppEigen,²² allowing for highly efficient sparse-matrix manipulations. Our R package implementation of our method is `groupFusedMulti` available in the open source repository <https://github.com/jaredhuling/groupFusedMulti>, which uses an interface similar to the interface used in `glmnet`.

For ease of tuning the tuning parameters, we utilize the following re-parameterization of λ_1 and λ_2 . Instead of using the penalty $\lambda_1 P_1(\boldsymbol{\beta}) + \lambda_2 P_2(\boldsymbol{\beta})$, we use $\lambda(1 - \alpha)P_1(\boldsymbol{\beta}) + \lambda\alpha P_2(\boldsymbol{\beta})$, where $\lambda \geq 0$ and $\alpha \in [0, 1]$ so that α controls the proportion of the total penalty that the fused lasso comprises.

2.3 | Computation via a multi-block ADMM algorithm

We utilize an alternating direction method of multipliers (ADMM)²³⁻²⁵ algorithm for optimization. The ADMM algorithm works by decomposing an objective function and solving the decomposed subproblems iteratively, where each subproblem has a simple and computationally tractable form. ADMM solves problems of the form

$$\text{minimize } f(\boldsymbol{\beta}) + P(\boldsymbol{\gamma}) \quad \text{subject to } \mathbf{A}\boldsymbol{\beta} + \mathbf{B}\boldsymbol{\gamma} = \mathbf{c}$$

where \mathbf{A} and \mathbf{B} are constraint matrices, both f and P are convex functions, $\boldsymbol{\beta}$ is the parameter vector of interest, and typically f represents some loss function and P represents a penalty. In the simplest case the constraint is of the form $\boldsymbol{\beta} = \boldsymbol{\gamma}$ and the purpose of the constraint is to find a new variable for which the penalty is equivalent, but separable across loss and penalty in the terms of the new variable. To optimize our objective, we require the following multi-block version of ADMM

$$\text{minimize } f(\boldsymbol{\beta}) + P_1(\boldsymbol{\gamma}) + P_2(\boldsymbol{\eta}) \quad \text{subject to } \mathbf{A}\boldsymbol{\beta} + \mathbf{B}\boldsymbol{\gamma} + \mathbf{C}\boldsymbol{\eta} = \mathbf{c},$$

where \mathbf{C} is an additional constraint matrix that allows the second penalty to have a decomposed form with loss and the two penalties all separable from each other. To solve the above problem, the augmented Lagrangian is formed as:

$$\begin{aligned} L_\rho(\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}, \boldsymbol{\nu}) &= f(\boldsymbol{\beta}) + P_1(\boldsymbol{\gamma}) + P_2(\boldsymbol{\eta}) + \boldsymbol{\nu}^\top (\mathbf{A}\boldsymbol{\beta} + \mathbf{B}\boldsymbol{\gamma} + \mathbf{C}\boldsymbol{\eta} - \mathbf{c}) \\ &\quad + (\rho/2) \|\mathbf{A}\boldsymbol{\beta} + \mathbf{B}\boldsymbol{\gamma} + \mathbf{C}\boldsymbol{\eta} - \mathbf{c}\|_2^2, \end{aligned}$$

where ρ is any strictly positive number. The multi-block ADMM algorithm iterates by alternatingly minimizing with respect to $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$, and $\boldsymbol{\eta}$ and following these minimizations, updating the Lagrangian parameter $\boldsymbol{\nu}$:

$$\boldsymbol{\beta}^{(t+1)} = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} L_\rho(\boldsymbol{\beta}, \boldsymbol{\gamma}^{(t)}, \boldsymbol{\eta}^{(t)}, \boldsymbol{\nu}^{(t)}) \quad (3)$$

$$\boldsymbol{\gamma}^{(t+1)} = \underset{\boldsymbol{\gamma}}{\operatorname{argmin}} L_\rho(\boldsymbol{\beta}^{(t+1)}, \boldsymbol{\gamma}, \boldsymbol{\eta}^{(t)}, \boldsymbol{\nu}^{(t)}) \quad (4)$$

$$\begin{aligned} \boldsymbol{\eta}^{(t+1)} &= \underset{\boldsymbol{\eta}}{\operatorname{argmin}} L_\rho(\boldsymbol{\beta}^{(t+1)}, \boldsymbol{\gamma}^{(t+1)}, \boldsymbol{\eta}, \boldsymbol{\nu}^{(t)}) \\ \boldsymbol{\nu}^{(t+1)} &= \boldsymbol{\nu}^{(t)} + \rho(\mathbf{A}\boldsymbol{\beta}^{(t+1)} + \mathbf{B}\boldsymbol{\gamma}^{(t+1)} + \mathbf{C}\boldsymbol{\eta}^{(t+1)} - \mathbf{c}) \end{aligned} \quad (5)$$

where t indexes the iteration number. The standard ADMM has been shown to converge for any $\rho > 0$ and the multi-block ADMM algorithm above has been shown to converge under certain conditions on \mathbf{A} , \mathbf{B} , and \mathbf{C} . These conditions, from Chen et al,²⁶ are met if either $\mathbf{A}^\top \mathbf{B} = \mathbf{0}$, $\mathbf{B}^\top \mathbf{C} = \mathbf{0}$, or $\mathbf{A}^\top \mathbf{C} = \mathbf{0}$.

The following describes the multi-block ADMM algorithm applied to the overlapping group lasso with fused lasso problem. Let $m = \sum_{G \in \mathcal{G}} |G|$, $g = |G|$, and suppose $\mathcal{G} = \{G_1, \dots, G_g\}$, let $\mathbf{F} = (\mathbf{F}_1, \dots, \mathbf{F}_g)$ be a matrix of dimension $m \times Kp$ where \mathbf{F}_l is a $|G_l| \times Kp$ matrix with (i,j) th entry equal to 1 if j is the i^{th} element of group G_l , and 0 otherwise, $\forall j = 1, \dots, g$. Then $\mathbf{F}\boldsymbol{\beta}$ is a vector of length m comprised components of $\boldsymbol{\beta}$ where each element of $\boldsymbol{\beta}$ appears in $\mathbf{F}\boldsymbol{\beta}$ the total number of times it appears in any group. For example, if $p=1$, $K=3$, $\boldsymbol{\beta} = (\beta_1, \beta_2, \beta_3)^\top$ and $\mathcal{G} = \{\{1, 2\}, \{2, 3\}\}$, then

$$\mathbf{F} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{F}\boldsymbol{\beta} = \begin{pmatrix} \beta_1 \\ \beta_2 \\ \beta_2 \\ \beta_3 \end{pmatrix}.$$

In this example, the penalty $P_1(\boldsymbol{\gamma})$ is $\lambda_1(\lambda_{1,G_1} \|(\gamma_1, \gamma_2)\|_2 + \lambda_{1,G_2} \|(\gamma_3, \gamma_4)\|_2)$, which is a standard group lasso penalty on the variable $\boldsymbol{\gamma}$. In general, the overlapping group lasso penalty $P_1(\boldsymbol{\gamma})$ can be written as $P_1(\boldsymbol{\gamma}) = \lambda_1 \sum_{l=1}^g \lambda_{1,G_l} \|\boldsymbol{\gamma}_l\|_2$, where $\boldsymbol{\gamma} = (\boldsymbol{\gamma}_1, \dots, \boldsymbol{\gamma}_g)$ and $\boldsymbol{\gamma}_l$ is a $|G_l|$ -dimensional vector. Note that for the nonoverlapping group lasso, $\mathbf{A} = \mathbf{I}_{Kp}$, where \mathbf{I}_{Kp} is

the identity matrix of dimension $Kp \times Kp$. To accommodate the fused lasso penalty, we allow for a more general penalty, the generalized fused lasso penalty. Denote the matrix \mathbf{D} with e rows and Kp columns where each row represents a pair of effects which have a fused lasso penalty applied. Thus e is the total number of fused lasso terms/pairs. For example, if the pair of effects $\beta_{j,\ell}$ and $\beta_{j,m}$ are have a fused lasso applied, then the row of \mathbf{D} corresponding to this pair will have ℓ th position taking the value 1 and the m th position taking the value -1 . In the toy example in Figure 2, if β_1 has a fused lasso penalty with β_2 and β_2 has a fused lasso penalty with β_3 , then

$$\mathbf{D} = \begin{pmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \end{pmatrix}, \text{ and } \mathbf{D}\boldsymbol{\beta} = \begin{pmatrix} \beta_1 - \beta_2 \\ \beta_2 - \beta_3 \end{pmatrix}.$$

In this example, the $P_2(\boldsymbol{\eta}) = \lambda_2(\lambda_{2,1}|\eta_1| + \lambda_{2,2}|\eta_2|)$, which is a standard lasso penalty on the variable $\boldsymbol{\eta}$. In general, the fused lasso penalty $P_2(\boldsymbol{\eta})$ can be written as $P_2(\boldsymbol{\beta}) = \lambda_2 \sum_{l=1}^e \lambda_{2,l} |\eta_l|$, where $\boldsymbol{\eta}$ is an e -dimensional vector. The ADMM algorithm for the overlapping group lasso plus the fused lasso is constructed by taking

$$\mathbf{A} = \begin{pmatrix} \mathbf{F} \\ \mathbf{D} \end{pmatrix}, \mathbf{B} = \begin{pmatrix} -\mathbf{I}_m \\ \mathbf{0} \end{pmatrix}, \mathbf{C} = \begin{pmatrix} \mathbf{0} \\ -\mathbf{I}_e \end{pmatrix}, \text{ and } \mathbf{c} = \mathbf{0},$$

which meets the condition that $\mathbf{B}^T \mathbf{C} = \mathbf{0}$, meaning a multi-block ADMM algorithm using this setup is valid and convergent. When $f(\boldsymbol{\beta}) = \frac{1}{2} \|\tilde{\mathbf{Y}} - \tilde{\mathbf{X}}\boldsymbol{\beta}\|_2^2$, step (3) for the overlapping group lasso is simply the solution of $(\tilde{\mathbf{X}}^T \tilde{\mathbf{X}} + \rho(\mathbf{F}^T \mathbf{F} + \mathbf{D}^T \mathbf{D}))\boldsymbol{\beta} = \tilde{\mathbf{X}}^T \tilde{\mathbf{Y}} + (\mathbf{F}^T, \mathbf{D}^T)\mathbf{v}^{(t)} + \rho(\mathbf{F}^T \boldsymbol{\gamma}^{(t)} + \mathbf{D}^T \boldsymbol{\eta}^{(t)})$. When $f(\boldsymbol{\beta})$ is the negative log-likelihood, step (3) can be carried out by Newton-Raphson or other standard optimization techniques. As step (4) is group-separable, it can be minimized by minimizing with respect to each group $\boldsymbol{\gamma}_l$ independently. This is achieved by the block soft-thresholding operator $S_{\lambda_1 \lambda_{G_l} / \rho}((\mathbf{F}\boldsymbol{\beta}^{(t+1)})_l - \mathbf{v}_l^{(t)} / \rho)$, where $S_\lambda(\mathbf{u}) = \mathbf{u} (1 - \lambda / \|\mathbf{u}\|_2)_+$ and $(\mathbf{F}\boldsymbol{\beta}^{(t+1)})_l$ and $\mathbf{v}_l^{(t)}$ are defined in the same way as $\boldsymbol{\gamma}_l$. Since (5) is separable and is equivalent to a type of lasso penalization problem, it can be similarly minimized simply via soft thresholding: $S_{\lambda_2 \lambda_j / \rho}((\mathbf{D}\boldsymbol{\beta}^{(t+1)})_j - \mathbf{v}_j^{(t)} / \rho)$, where with some abuse of notation $(\mathbf{D}\boldsymbol{\beta}^{(t+1)})_j$ is the j th element of $(\mathbf{D}\boldsymbol{\beta}^{(t+1)})$, $\mathbf{v}_j^{(t)}$ is the j th element of $\mathbf{v}^{(t)}$, and $S_\lambda(u) = \text{sign}(u) (|u| - \lambda)_+$, which is the equivalent of the block soft thresholding operator applied to a scalar. Our convergence criterion is the same as suggested in section 3.3.1 of Boyd et al²⁵ with $\epsilon^{\text{abs}} = \epsilon^{\text{rel}} = 10^{-5}$. The parameter ρ for the ADMM algorithm used is the adaptive scheme described in section 3.4.1 of Boyd et al.²⁵

2.4 | Connections with existing literature

The methods of Rothman et al⁸ and Li et al¹³ both also address scenarios where high dimensional covariates are used to predict a set of multivariate outcomes in a penalized regression framework. These two methods in addition to ours borrow strength in information across outcomes through penalization techniques, albeit in different ways.

If the multivariate outcomes are only correlated due to correlations between the outcomes themselves, then the multivariate regression with covariance estimation (MRCE) approach⁸ may be most appropriate as it directly incorporates the covariance of the responses via a penalized log-likelihood; here, the lasso penalty induces the estimated coefficients to be a function of the covariance of the responses, unlike the unpenalized maximum likelihood estimator. However, when many covariates are expected to have effects of similar magnitude across the multivariate outcomes, our additional fused lasso penalty is designed to make use of this by borrowing information across outcomes for such covariates at a more granular level. Further, when there is strong group structure of whether or not a given covariate has an effect across groups of outcomes, but the magnitudes of the effects are unrelated, then both our approach and the multinomial logistic regression with sparse group lasso (MSGSL) approach¹³ are appropriate options and their performances should be expected to be similar, despite our additional fused lasso penalty, whereas MRCE is not designed explicitly to make use of such information. However, if indeed there are many covariates with similar effects across outcomes and a group structure, our approach may be expected to perform well. If there is no relationship whatsoever across the multivariate outcomes, an approach of simply fitting a lasso-penalized model separately for each outcome may have reasonable performance.

2.5 | Asymptotic properties

We now present new asymptotic results for our penalization methods pertaining to semiparametric linear models. We show results for a more general, encompassing form for the overlapping group lasso penalty P_1 and the fused lasso penalty P_2 . We allow $P_1(\boldsymbol{\beta}) = \sum_{j=1}^p \sum_{g \in \mathcal{G}} \lambda_{1,j,G} \|\boldsymbol{\beta}_{j,G}\|_2$ for an arbitrary, potentially overlapping group structure \mathcal{G} and we allow $P_2(\boldsymbol{\beta}) = \sum_{j=1}^p \sum_{(l,o) \in \mathcal{F}} \lambda_{2,j,l,o} |\beta_{j,l} - \beta_{j,o}|$ for an arbitrary fusing set \mathcal{F} that contains pairs of indices of coefficients to be fused, where the elements of \mathcal{F} are of the form (l, o) with $l, o \in \mathcal{K}$. The penalty form introduced in Section 2.1 is a special case of the form we study in this section, as our Theorem 1 applies to any group structure, overlapping or not. We also denote the set of all pairs of coefficients which are equal and non-zero for the j th variable by $\mathcal{E}_j = \{(l, m) \in \mathcal{F} : \beta_{j,l}^0 = \beta_{j,m}^0 \neq 0\}$. Similarly, define $\hat{\mathcal{E}}_j = \{(l, m) : l, m \in \mathcal{F} \text{ and } \hat{\beta}_{j,l} = \hat{\beta}_{j,m} \neq 0\}$. Denote the union of these sets over all variables as $\mathcal{E} = \bigcup_{j=1}^p \mathcal{E}_j$ and $\hat{\mathcal{E}} = \bigcup_{j=1}^p \hat{\mathcal{E}}_j$.

To facilitate our explanation of the asymptotic results, for any vector $\mathbf{a}_{\cdot,k} \in \mathbb{R}^p$ associated with response k and a index set $\mathcal{I}_{\cdot,k} \subset \{1, \dots, p\}$ of size $|\mathcal{I}_{\cdot,k}|$ associated with response k , $\mathbf{a}_{\mathcal{I}_{\cdot,k}}$ represents a $|\mathcal{I}_{\cdot,k}|$ dimensional sub-vector with elements in $\mathbf{a}_{\cdot,k}$ indexed by $\mathcal{I}_{\cdot,k}$. Now let $\tilde{\mathbf{X}} = \mathbf{I}_K \otimes \mathbf{X} = \text{diag}(\mathbf{X}, \dots, \mathbf{X})$ be the block diagonal design matrix with K blocks, one for each outcome, with the k th block as the design matrix \mathbf{X} , where \otimes is the Kronecker product; similarly let $\tilde{\boldsymbol{\beta}}^0 = \text{vec}(\boldsymbol{\beta}^0) = (\boldsymbol{\beta}_{\cdot,1}^{0\top}, \dots, \boldsymbol{\beta}_{\cdot,K}^{0\top})^\top$ be the vectorization of the true coefficients for all K responses, let $\tilde{\mathbf{Y}}^\top = (\mathbf{Y}_{\cdot,1}^\top, \dots, \mathbf{Y}_{\cdot,K}^\top)$. Using this notation, we can re-express the model (1) as $\tilde{\mathbf{Y}} = \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}^0 + \tilde{\boldsymbol{\epsilon}}$, where $\tilde{\boldsymbol{\epsilon}}$ is the stacked error vector defined similarly as $\tilde{\mathbf{Y}}$.

Now let $\mathcal{J} = (\mathcal{J}_{\cdot,1}, \dots, \mathcal{J}_{\cdot,K}) \subseteq \{1, \dots, Kp\}$, where $\mathcal{J}_{\cdot,k} = \{j + (k-1)p : j \in \{1, \dots, p\} \text{ and } \beta_{j,k}^0 \neq 0\}$, to be the set of indices of all non-zero effects in $\tilde{\boldsymbol{\beta}}^0$ and let $\mathcal{H} = \text{Hull}(\mathcal{J}) = \{\bigcup_{G \in \mathcal{G}, G \cap \mathcal{J} = \emptyset} G\}^c$ be the hull of the nonzero pattern \mathcal{J} , where for a set S , S^c denotes its complement. The hull of the non-zero pattern is essentially the smallest set of groups in \mathcal{G} that contains all elements in \mathcal{J} . Similarly, denote $\hat{\mathcal{J}} = (\hat{\mathcal{J}}_{\cdot,1}, \dots, \hat{\mathcal{J}}_{\cdot,K})$ to be the set of indices of all nonzero variables in $\hat{\boldsymbol{\beta}}$. Note that by construction of $P_1(\cdot)$ due to each effect of each variable across the outcomes having its own group in \mathcal{G}_{M+1} , $\text{Hull}(\hat{\mathcal{J}}) = \hat{\mathcal{J}}$, however we present our results in terms of the hull so as to be fully general with respect to the group structure. We denote $\tilde{\mathbf{X}}_{\mathcal{H}}$ as the columns in $\tilde{\mathbf{X}}$ corresponding to variables in \mathcal{H} and similarly denote $\tilde{\boldsymbol{\beta}}_{\mathcal{H}}^0$ as the values in $\tilde{\boldsymbol{\beta}}^0$ corresponding to elements in \mathcal{H} . Then denote $\tilde{\mathbf{X}}_{\mathcal{H}}^*$ constructed by dropping and collapsing columns of $\tilde{\mathbf{X}}_{\mathcal{H}}$ corresponding to the *distinct*, non-zero values of $\tilde{\boldsymbol{\beta}}_{\mathcal{H}}^0$. Specifically, for each j , let $\{c_{j,1}, \dots, c_{j,L}\}$ for $L \leq K$ denote the unique *non-zero* values of $\{\hat{\beta}_{j,1}, \dots, \hat{\beta}_{j,K}\}$ and denote $C_{j,\ell} = \{k \in \mathcal{K} : \hat{\beta}_{j,k} = c_{j,\ell}\}$. Then for each $c_{j,\ell} \in \{c_{j,1}, \dots, c_{j,L}\}$ all columns $j + (k-1)p$ of $\tilde{\mathbf{X}}$ such that $k \in C_{j,\ell}$, $j \in \{1, \dots, p\}$ are collapsed and added together into a single column, that is, $\sum_{k \in C_{j,\ell}} \tilde{\mathbf{X}}_{j+(k-1)p}$. As an illustration, denote $\tilde{\mathbf{X}}_{j+(k-1)p}$ as the j th column of the design matrix for response k . Then if the coefficient for variable j is in \mathcal{H} for responses k, ℓ , and K and $\beta_{j,k}^0 = \beta_{j,\ell}^0 = \beta_{j,K}^0 \neq 0$, this results in the following to the corresponding columns in $\tilde{\mathbf{X}}$ in the formulation of $\mathbf{X}_{\mathcal{H}}^*$:

$$\begin{array}{l}
 \text{Block } k \rightarrow \\
 \text{Block } \ell \rightarrow \\
 \text{Block } K \rightarrow
 \end{array}
 \begin{array}{c}
 \tilde{\mathbf{X}}_{j+(k-1)p} \\
 \left(\begin{array}{c} \mathbf{0} \\ \vdots \\ \mathbf{x}_{\cdot,j} \\ \vdots \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{array} \right) \\
 + \\
 \tilde{\mathbf{X}}_{j+(\ell-1)p} \\
 \left(\begin{array}{c} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \vdots \\ \mathbf{x}_{\cdot,j} \\ \vdots \\ \mathbf{0} \end{array} \right) \\
 + \\
 \tilde{\mathbf{X}}_{j+(K-1)p} \\
 \left(\begin{array}{c} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \vdots \\ \mathbf{x}_{\cdot,j} \end{array} \right) \\
 \rightarrow \\
 \left(\begin{array}{c} \mathbf{0} \\ \vdots \\ \mathbf{x}_{\cdot,j} \\ \vdots \\ \mathbf{x}_{\cdot,j} \\ \vdots \\ \mathbf{x}_{\cdot,j} \end{array} \right)
 \end{array}$$

In other words, the columns in $\tilde{\mathbf{X}}$ corresponding to these three variables are collapsed and added together. Note that there exist matrices \mathbf{H} and \mathbf{E} such that $\tilde{\mathbf{X}}_{\mathcal{H}}^* = \tilde{\mathbf{X}}\mathbf{H}\mathbf{E}$, where \mathbf{H} is formed by removing all columns of positions in the indices of \mathcal{H} from the identity matrix of dimension $pK \times pK$ and \mathbf{E} is formed by taking the identity matrix of dimension $|\mathcal{H}| \times |\mathcal{H}|$ and collapsing and summing columns corresponding to non-zero coefficients that are equal to each other and are in \mathcal{F} , the fusing set used in P_2 .

We assume the following standard regularity conditions

- D.1 $\lim_{N \rightarrow \infty} \frac{1}{N} \mathbf{X}^\top \mathbf{X} \rightarrow \mathbf{Q}$ where \mathbf{Q} is positive definite
- D.2 The errors ϵ_i for all samples $i = 1 \dots, n$ are *i.i.d.* random vectors with mean zero and finite, positive definite variance-covariance matrix Σ .

The following result pertains to cases where the group structure has been misspecified. As such, we define β^{*0} to be the distinct values of $\beta_{\mathcal{H}}^0$ (ie, the union of the distinct values in $\{\beta_{j,k}^0, \beta_{j,\ell}^0 : (k, \ell) \in \mathcal{E} \text{ or } (\ell, k) \in \mathcal{E} \text{ and } (k, \ell) \in \mathcal{F} \text{ or } (\ell, k) \in \mathcal{F}\}$ and all remaining values $\beta_{j,k}^0 \neq 0$) appended with zeros corresponding to the elements of $\beta_{\mathcal{H}^c}^0$. Similarly, we define $\hat{\beta}^*$ to be the distinct values of $\hat{\beta}_{\mathcal{J}}$ appended with zeros corresponding to the elements of $\hat{\beta}_{\mathcal{J}^c}$.

Theorem 1. Assume the data are generated under the model described in (1) and that our estimator is given by (2). Furthermore, assume conditions (D.1) and (D.2) and let $\lambda_{1,j,G} = \|\hat{\beta}_{j,G}^{OLS}\|_2^{-\gamma_1}$ and $\lambda_{2,j,l,m} = |\hat{\beta}_{j,l}^{OLS} - \hat{\beta}_{j,m}^{OLS}|^{-\gamma_2}$, where $\gamma_1, \gamma_2 > 0$ such that $N^{(\gamma_\ell+1)/2} \lambda \rightarrow \infty$ for $\ell = 1, 2$. If $\sqrt{N} \lambda \rightarrow 0$, then we have the following:

$$P(\hat{\mathcal{J}}_{\cdot,j} = \mathcal{H}_{\cdot,j}) \rightarrow 1 \text{ as } N \rightarrow \infty, \tag{6}$$

$$P(\hat{\mathcal{E}}_{\cdot,j} = \mathcal{E}_{\cdot,j}) \rightarrow 1 \text{ as } N \rightarrow \infty, \tag{7}$$

for each $j = 1, \dots, p$ and

$$\sqrt{N}(\hat{\beta}^* - \beta^{*0}) \xrightarrow{d} \mathbf{Z}^*, \tag{8}$$

where $\mathbf{Z}^* = (\mathbf{Z}_{\mathcal{H}}^{*\top}, \mathbf{0}^\top)^\top$ and $\mathbf{Z}_{\mathcal{H}}^* \sim N_{|\mathcal{H}^*|}(0, \mathbf{Q}_{\mathcal{H}}^{*-1} \mathbf{V}_{\mathcal{H}}^* \mathbf{Q}_{\mathcal{H}}^{*-1})$, where $\mathbf{V}_{\mathcal{H}}^* = \mathbf{E}^\top \mathbf{H}^\top (\Sigma \otimes \mathbf{Q}) \mathbf{H} \mathbf{E}$, $\mathbf{Q}_{\mathcal{H}}^* = \mathbf{E}^\top \mathbf{H}^\top (\mathbf{I}_K \otimes \mathbf{Q}) \mathbf{H} \mathbf{E}$, and $|\mathcal{H}^*|$ denotes the number of distinct non-zero elements in the true coefficients β^{*0} .

The proof of Theorem 1 is provided in the Supplementary Material Appendix A. We reiterate that Theorem 1 applies to any group structure, including those with overlapping groups. We note that the variance-covariance matrix in the distribution of \mathbf{Z}^* does not simplify due to the allowance of selection of variables individually by outcome in combination with the Kronecker-product variance structure of the response vector $\tilde{\mathbf{Y}}$; without either it would simplify considerably. The terms \mathbf{H} and \mathbf{E} in the asymptotic variance of the estimates are due to selection of variables and collapsing of effects of a single variable across outcomes, respectively; without any selection these matrices are removed and the variance simplifies to the usual asymptotic variance of multivariate response regression via least squares. However, the asymptotic variance matches that of the estimator where all and only all truly non-zero variable effects are included in a model and all variables with equal effects across a subset of outcomes are accordingly collapsed. The result of Theorem 1 does not require that $\mathcal{E} \subseteq \mathcal{F}$, however, if indeed $\mathcal{E} \subseteq \mathcal{F}$, that is, if \mathcal{F} contains indices corresponding to all pairs of coefficients in $\beta_{\mathcal{H}}^0$ that are equal to each other, then due to (8), when using adaptive terms for the tuning parameters, our approach yields selection consistency for the hull of the non-zero terms; when the group structure has individual groups for each of the individual coefficients in the model, the convex hull of the non-zero coefficients is simply the set of non-zero coefficients. Thus, when such individual groups are included, we have selection consistency for the non-zero coefficients. Our theory also shows that with probability tending to one, our approach will fuse together the coefficients that are truly equal to each other. Finally, our theory shows that our estimate will converge to the same asymptotic distribution as if we had known which coefficients were non-zero and which coefficients were equal to each other.

3 | NUMERICAL EXPERIMENTS

3.1 | Estimators used

In this section, we conduct simulation experiments to assess the small sample operating characteristics of our proposed doubly-structured sparsity inducing estimation approach in comparison with several other state-of-the-art approaches for multivariate regression in high-dimensions. In our simulations, we vary both the dimensionality of the problem and

the sample size. We also consider data-generating mechanisms with varying degrees of sparsity of variable effects that align with grouping of the outcomes and further consider scenarios where the effects of each variable have either minimal similarity across outcomes or a varying degree of similarity across outcomes. This allows our studies to explore under what data-generating mechanisms our two penalties indeed help in estimation. We use adaptive and non adaptive versions of our method, denoted OGF_M(adapt) and OGF_M, respectively, where OGF_M indicates Overlapping Group + Fused Multivariate regression. For the adaptive version, we set $\gamma_1 = \gamma_2 = 0.5$ and for high dimensional settings with $p \geq n$, we used marginal regression estimates for the adaptive weights as in Huang et al.²⁷ We compare our approach with the approach of Li et al.¹³ (denoted as MSGL), which allows for overlapping group lasso penalties. For MSGL, we use the MSGLasso R package version 2.1 We further compare with a simple approach of fitting a separate lasso-penalized linear regression model for each outcome, with the tuning parameter chosen separately for each outcome (denoted as Sep-Lasso) and do so using the glmnet R package version 4.1-3. We also compare with the approach of Rothman et al.⁸ (denoted MRCE) implemented in the MRCE R package version 2.1. For all methods, the tuning parameters are chosen by 10-fold cross validation; MSGL has two tuning parameters (one for a lasso penalty, one for a group lasso penalty), MRCE has two tuning parameters (one for a lasso penalty for the coefficients for the variable effects on responses, another for a lasso penalty on the elements of the inverse covariance matrix of the residuals), Sep-Lasso, for separate lasso has a tuning parameter for each outcome for a lasso penalty on coefficients, and the OGF_M approaches have two tuning parameters (one for the overlapping group lasso penalty and another for the fused lasso penalty). For both the OGF_M approaches and MSGL, the group lasso penalty applied corresponds to the true underlying group structure of the data-generating mechanism.

3.2 | Data generation

For each replication of the simulation, we generate data under model (1), where \mathbf{x}_i are generated as i.i.d. multivariate normal random variables with mean vector $\mathbf{0}$ and covariance matrix $\Sigma_{\mathbf{X}} = [\sigma_{X_{jk}}]_{j,k=1}^p$, where $\sigma_{X_{jk}} = 0.5^{|j-k|}$, as was used in the simulations of Yuan and Lin⁹ and Rothman et al.⁸ The responses are generated according to model (1), where the error vectors ϵ_i are generated as i.i.d. multivariate normal random variables with mean vector $\mathbf{0}$ and covariance matrix $\Sigma_{\epsilon} = [\sigma_{\epsilon_{jk}}]_{j,k=1}^K$, where $\sigma_{\epsilon_{jk}} = 4(0.5^{|j-k|})$. In the Supplementary Material Appendix B.2, we explore a simulation setting with binary covariates and a Likert scale outcome similar to our motivating data. The dimensionality of the outcome/error vector is $K = 8$ and the outcomes form 3 groups, with the first three outcomes forming group 1, the fourth and fifth forming group 2, and the last three outcomes forming group 3. The variable effects are generated as

$$\beta^0 = \begin{pmatrix} \overbrace{\xi_{1,1} \xi_{1,1}^G \eta_{1,1} & \xi_{1,2} \xi_{1,1}^G \eta_{1,2} & \xi_{1,3} \xi_{1,1}^G \eta_{1,3} & \xi_{1,4} \xi_{1,2}^G \eta_{1,4} & \xi_{1,5} \xi_{1,2}^G \eta_{1,5} & \xi_{1,6} \xi_{1,3}^G \eta_{1,6} & \xi_{1,7} \xi_{1,3}^G \eta_{1,7} & \xi_{1,8} \xi_{1,3}^G \eta_{1,8}}^{G_{1,1}} & \vdots \\ \xi_{z,1} \xi_{z,1}^G \eta_{z,1} & \xi_{z,2} \xi_{z,1}^G \eta_{z,2} & \xi_{z,3} \xi_{z,1}^G \eta_{z,3} & \xi_{z,4} \xi_{z,2}^G \eta_{z,4} & \xi_{z,5} \xi_{z,2}^G \eta_{z,5} & \xi_{z,6} \xi_{z,3}^G \eta_{z,6} & \xi_{z,7} \xi_{z,3}^G \eta_{z,7} & \xi_{z,8} \xi_{z,3}^G \eta_{z,8} \\ 0 & \dots & & & & & \dots & 0 \\ \vdots & \ddots & & & & & & \vdots \\ 0 & \dots & & & & \dots & & 0 \end{pmatrix},$$

where the last $p - z$ rows of β^0 have all elements as 0, the terms $\xi_{j,k} \sim$ Bernoulli (0.9) induce sparsity at the individual effect level, the terms $\xi_{j,1}^G, \xi_{j,2}^G, \xi_{j,3}^G \sim$ Bernoulli $(1 - p_{\text{HS}})$ induce sparsity at the group level, the variable effect size terms $\eta_{j,k}$ are distributed i.i.d. uniformly from $\{-1, -0.5, -0.25, -0.125, 0.125, 0.25, 0.5, 1\}$ to create both small and large effects, and for each variable j separately the terms $\eta_{j,k}$ for k in the same group are set to be all equal to each other with probability $p_{\text{GE}}/2$ and independently, the terms $\eta_{j,k}$ for all outcomes $k = 1, \dots, 8$ are set to be all equal to $\eta_{j,1}$ with probability $p_{\text{GE}}/2$. The latter process induces effects for some variables within a group to be equal to each other and induces effects for some variables to be equal for all outcomes.

We explore dimensions of $p = 50, 100$, and 200 and set $z = 25, 50, 50$ for each dimension setting, respectively. We explore hierarchical sparsity parameters from $p_{\text{HS}} = 0, 0.25$ and 0.5 and fusing probabilities $p_{\text{GE}} = 0, 0.5, 0.75$ and 0.95 ;

when $p_{HS} = 0$, there is no group-wise sparsity and thus any group lasso penalty applied is superfluous. For each replication of the simulation, we additionally generate an independent test set of size 10,000 for use in evaluation of predictive performance.

3.3 | Performance evaluation

We evaluate methods in terms of three metrics: root mean squared error of predictions (RMSE) on a large, independent test set of size 10,000 generated anew for each simulation replication, model error defined below, and balanced accuracy which is the average of the true positive rate (TPR) and true negative rate (TNR), described in detail in the Supplementary Material Appendix B.1 and not to be confused with standard classification accuracy. As our primary interest is performance of predictions on validation data, it is our primary focus of evaluation. For the RMSE, we compute the RMSE for each outcome separately and then average the RMSEs across the outcomes. Model error for a given estimate is defined as $ME(\hat{\beta}, \beta^0) = \text{tr}[(\hat{\beta} - \beta^0)^T \Sigma_X (\hat{\beta} - \beta^0)]$.

3.4 | Results

The RMSE results fixing the parameter at $p_{HS} = 0$ are displayed in Figure 3 and the RMSE results for $p_{HS} = 0.25$ are displayed in Figure 4. When the sample size is small ($n = 100$) and the hierarchical sparsity probability is 0 (Figure 3), our proposed approach OGFM results in the smallest validation RMSE on average across the replications across all settings including moderate dimensions ($p = 50, 100$) and high dimensions ($p = 200$), with MSGL as second best across the majority of settings and OGFM(adapt) a close third. The separate lasso tends to perform worse than all competing approaches, except for in some settings (eg, when $n = 400$) where it performs marginally better than MRCE. We note that the code for




FIGURE 3 Validation RMSEs for all methods across 100 replications of the simulation experiment holding the parameter $p_{HS} = 0$ that there is no group-level sparsity.

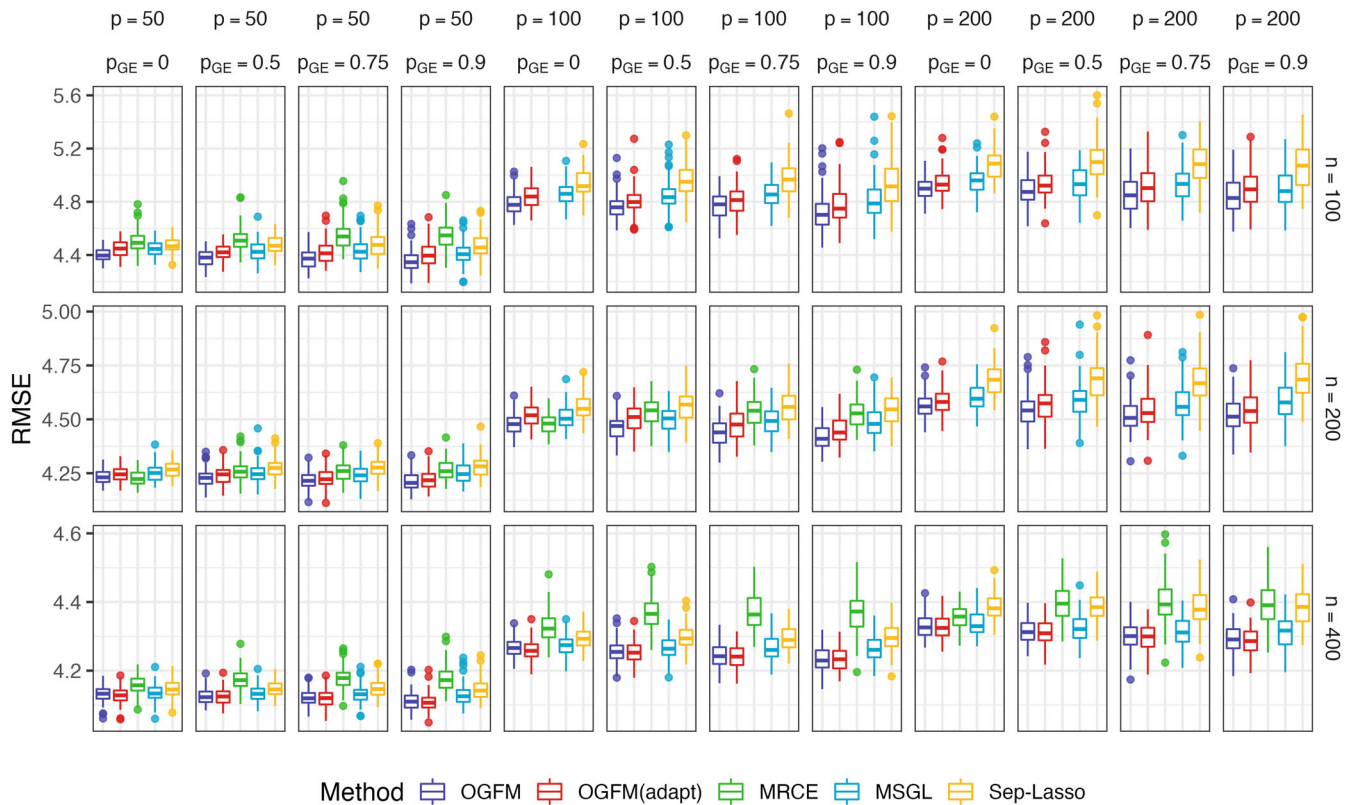


FIGURE 4 Validation RMSEs for all methods across 100 replications of the simulation experiment holding the parameter $p_{HS} = 0.25$ so that there is moderate group-level sparsity.

MRCE throws an error when $p \geq n$, although in principle MRCE can work for high dimensional settings. In the same settings but with $n = 200$, OGFM still tends to work the best across all settings, however when the probability of coefficients being equal to each other is zero ($p_{GE} = 0$), and the dimensionality is moderate, the MRCE approach performs nearly as well. In the high dimensional setting ($p = 200$), MSGL performs on par with OGFM and sometimes better, however this benefit slightly attenuates for the larger sample size setting ($n = 400$). In general the adaptive version of our approach, OGFM(adapt) performs worse than OGFM in small sample size settings and performs better with larger sample sizes. The results with the hierarchical sparsity probability set to 0.25 (Figure 4) roughly mirror the results with no hierarchical sparsity, however the performance for all methods is slightly better in terms of RMSE as there are fewer overall non-zero coefficients in the data-generating process. The model error results largely track with the validation RMSE results, albeit on a different scale, and are thus shown in the Supplementary Material Appendix B.1.

Figure 5 shows the same results as Figures 3 and 4, but is re-organized to focus on the effect of the fused coefficients probability (p_{GE}) and the hierarchical sparsity probability (p_{HS}). In this figure, we fix the sample size to be 200. In general, we can see that as the probability of coefficients being equal/fused within groups increases, the performance of OGFM and OGFM(adapt) relative to the Sep-Lasso, MSGL, and MRCE tends to improve, with the trend being pronounced in moderate dimensional settings. Only the OGFM methods improve as the probability of coefficients being equal increases, as they are the only approaches which explicitly allow for shrinkage of effect sizes to each other across outcomes. The relative performance of all methods tends to stay relatively consistent as p_{HS} is varied. However, we note that MRCE tends to perform best in the setting with the fewest true non-zero coefficients ($p_{HS} = 0.25$) and smallest amount of truly equal coefficients ($p_{GE} = 0$).

In the Supplementary Material Appendix B.1, we show simulation results in terms of the average of the TPR and TNR. From these results, it can be seen that while OGFM performs best in terms of prediction performance, it tends to over-select terms, resulting in a high TPR but low TNR and thus lower average TPR and TNR value. On the other hand, for larger sample sizes OGFM(adapt) performs very well in terms of the average TPR and TNR as it selects fewer variables with a higher proportion of selected variables being ones with truly non-zero coefficients. Overall, MSGL performs best in terms of the average TPR and TNR and MRCE worst, the latter trend having been concurrently observed in the original




FIGURE 5 Validation RMSEs for all methods across 100 replications of the simulation experiment holding fixing the sample size $n = 200$ and varying all other simulation parameters. The points are the average RMSEs across the 100 replications and error bars are plus and minus 1 standard deviation of this average.

work of Rothman et al,⁸ who noted that MRCE tends to result in better model error but with less benefit in terms of TPR and TNR. In the Supplementary Material Appendix B.1 we also present computation times for all methods. The OGFM approaches are highly competitive computationally as the sample size increases, however their performance deteriorates as the dimensionality increases.

4 | ANALYSIS OF PEDIATRIC FUNCTIONAL STATUS DATA

Our motivating study consists of 1897 children, adolescents, and young adults who were admitted to an inpatient rehabilitation unit with a diagnosis of neurologic injury or illness at a major Midwestern children’s hospital for neurological injury or illness between the years 2000 and August of 2020. The sample consisted of individuals between the ages of 6 months to 32 years with a median age of 12 years and first and third quartiles of 6 and 16 years, respectively. The covariates used as predictors are (numbering 608 in total) included basic sociodemographic information including age and sex, and billing codes including CPT codes, ICD-9 and -10 codes, pharmaceutical codes, and data indicating durable medical equipment (DME) use. The billing codes represent information from a single episode of care involving an admission to an inpatient rehabilitation unit. The billing code information pertain only to information collected prior to the WeeFIM assessment. CPT and all ICD-9 and -10 codes were linked together to unique concept unique identifiers (CUIs) using the UMLS metathesaurus mapping system.²⁸ Concept Unique Identifiers (CUIs) are distinct medical concepts (codes, diseases, etc.) identified by the UMLS metathesaurus. By linking all billing codes to CUIs, we are able to in many cases handle the switch-over from ICD-9 to ICD-10 codes. An added benefit of CUIs is that they can be linked to other coding systems. We utilize all 18 WeeFIM® components as responses; these 18 components are divided into three domains. The

TABLE 1 Mean squared errors (MSEs) on the 18 WeeFIM components in the validation data and average MSE across the 18 components.

Method	Domain	Self-care WeeFIM responses								
	Average	Ea	G	B	D-U	D-L	T-g	Bl-M	Bo-M	
OGFM	2.153	2.074	1.816	2.034	1.673	2.099	2.052	3.533	3.492	
OGFM(adapt)	2.165	2.070	1.824	2.032	1.663	2.112	2.165	3.573	3.521	
MRCE	2.185	2.155	1.864	2.093	1.674	2.117	2.093	3.735	3.638	
MSGSL	2.222	2.145	1.860	2.149	1.712	2.218	2.190	3.729	3.617	
Sep-Lasso	2.213	2.142	1.860	2.189	1.672	2.225	2.148	3.788	3.641	

Method	Domain	Mobility WeeFIM responses					Cognition WeeFIM responses				
		C-W	T	T-S	W-W	S	C	Ex	S-I	P-S	M
OGFM		1.483	1.770	1.811	1.344	2.169	2.140	2.165	2.159	2.537	2.395
OGFM(adapt)		1.492	1.792	1.808	1.336	2.187	2.155	2.159	2.180	2.495	2.407
MRCE		1.571	1.811	1.889	1.331	2.164	2.126	2.081	2.140	2.446	2.399
MSGSL		1.494	1.857	1.790	1.305	2.182	2.156	2.185	2.264	2.685	2.453
Sep-Lasso		1.539	1.809	1.915	1.286	2.186	2.178	2.152	2.174	2.491	2.432

Note: Bold indicates the smallest MSE across the different methods for a particular component.

Abbreviations: B, bathing; Bl-M, bladder management; Bo-M, bowel management; C, comprehension; C-W, chair wheelchair; D-L, dressing lower; D-U, dressing upper; Ea, eating; Ex, expression; G, grooming; M, memory; P-S, problem solving; S, stairs; S-I, social interaction; T, toilet; T-g, toileting; T-S, tub shower; W-W, walk wheelchair.

self-care domain describes how well a child is able to feed themselves, groom, bath, dress, and complete toileting tasks including the management of their bowel and bladder. The mobility domain describes how well a child is able to transfer on and off a toilet, in and out of a bathtub or shower, and in and out of a chair or wheelchair. The mobility domain also describes a child's ability to walk, crawl, or use a wheelchair, and to move up or down stairs. The cognition domain describes how well a child can express themselves, understand information, interact with peers, solve daily problems, and recall information. Together, they describe the ability of children to function in routine and important aspects of daily life.

As our motivating use case of developing a model for the WeeFIM[®] components is to apply it to future data to assess functional ability across a health system population, we validate all developed models by splitting data into training and validation sub-datasets, where the training dataset is from years prior to data from the validation dataset. We use data prior to 2018 for training and data from 2018 to 2020 as validation data, leaving 1592 observations for training and 305 observations for validation. For validation, we compare methods in terms of the response-specific mean squared error for each response and the average validated mean squared error across all 18 responses. We also compute the corresponding validation R-squared values, that is, the proportion of the validation responses explained by the out-of-sample predictions. We exclude any covariates that have no variation in either the training or validation datasets. After this screening, the combined dimensionality of all the predictors was 608, a high-dimensional scenario given the sample size. We apply all methods described in Section 3.1 and use the approaches described therein for selection of tuning parameters. For OGFM(adapt), we use $\gamma_1 = \gamma_2 = 0.5$, as use of $\gamma_1 = \gamma_2 = 1$ resulted in some extreme weights given the high dimensionality of the data. For both the OGFM approaches and MSGSL, the group lasso penalty applied corresponds to the three pre-defined domains of the WeeFIM[®] components (self-care, mobility, and cognition).

The validation MSEs and R-squared values are displayed in Tables 1 and 2, respectively. Our OGFM approach has the best performance on the validation data on average across the 18 responses and also most often performs best for the individual responses, with the adaptive version of OGFM with the second lowest validation MSE on average. For responses where OGFM does not perform best, most often its performance does not differ much from the best MSE across the remaining methods, excluding OGFM(adapt). When OGFM performs best for particular responses, its improvement in terms of RMSE and R-squared is often large, as is depicted in Figure 6, which shows the difference between OGFM and competing approaches in terms of MSE on the validation data.

For both OGFM approaches, the fused lasso mixing parameter α was chosen by cross-validation to be 1×10^{-5} , indicating only a small amount of fused lasso was required by the data. However, this small amount of fused lasso penalty

TABLE 2 R-squared values on the 18 WeeFIM components in the validation data and average R-squared value across the 18 components.

Method	Domain	Self-care WeeFIM responses							
	Average	Ea	G	B	D-U	D-L	T-g	BI-M	Bo-M
OGFM	0.432	0.526	0.536	0.476	0.446	0.447	0.498	0.449	0.429
OGFM(adapt)	0.429	0.527	0.534	0.477	0.449	0.444	0.470	0.443	0.425
MRCE	0.425	0.507	0.524	0.461	0.446	0.442	0.488	0.418	0.405
MSGL	0.416	0.509	0.525	0.447	0.433	0.416	0.464	0.419	0.409
Sep-Lasso	0.419	0.510	0.525	0.437	0.446	0.414	0.474	0.410	0.405

Method	Domain	Mobility WeeFIM responses					Cognition WeeFIM responses				
		C-W	T	T-S	W-W	S	C	Ex	S-I	P-S	M
OGFM		0.433	0.522	0.447	0.287	0.298	0.410	0.405	0.400	0.375	0.384
OGFM(adapt)		0.430	0.516	0.448	0.290	0.292	0.406	0.406	0.394	0.385	0.381
MRCE		0.400	0.510	0.423	0.293	0.299	0.414	0.428	0.405	0.397	0.383
MSGL		0.429	0.498	0.454	0.307	0.294	0.406	0.399	0.371	0.338	0.369
Sep-Lasso		0.412	0.511	0.415	0.317	0.292	0.400	0.408	0.396	0.386	0.374

Note: Bold indicates the larger R-squared across the different methods for a particular components.

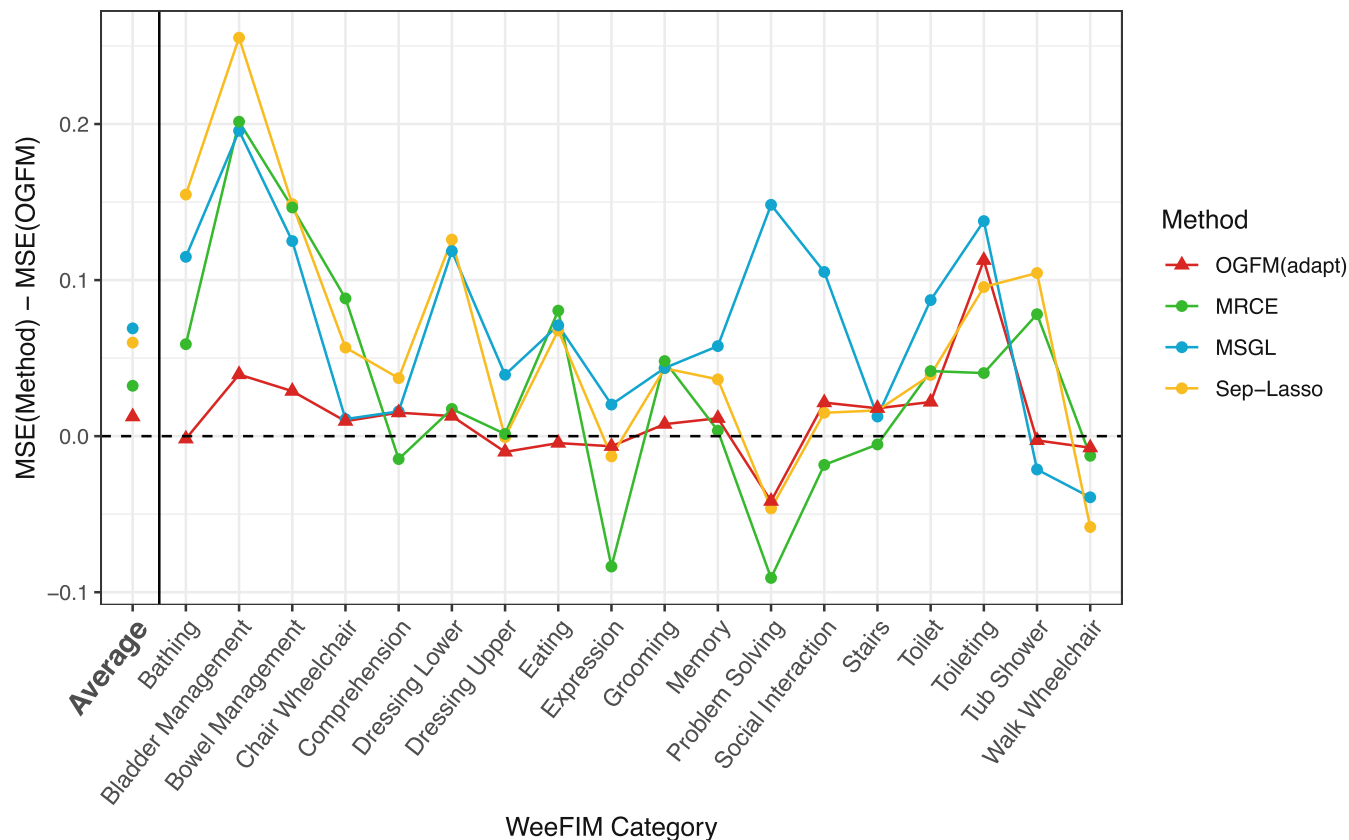


FIGURE 6 Comparisons of the validation MSE for each of the 18 WeeFIM components and on average (left) across the 18 components.




FIGURE 7 Coefficient path plots across all 18 WeeFIM components for six variables from the pediatric functional status data. The dashed vertical line indicates the value of λ that minimizes the cross-validation error.

had the effect of shrinking a large number of coefficients either close to each other or exactly to each other. In Figure 7, we display the coefficient paths versus λ (described in Section 2.2) for 6 variables, where each plot is the coefficient path for a particular variable across the 18 outcomes. For each of the coefficient path plots, we fix the mixing tuning parameter α at the value that minimizes the cross validation error. The patterns of these coefficient paths demonstrate that our approach fuses coefficients for a variable to be the same across multiple responses when appropriate, allows effects to differentiate when warranted by the data, an allows for joint selection of effects across an entire domain of related responses. For the ICD code “shock, unspecified”, the fused lasso grouped the effects for all cognition outcomes together along the entire path, all mobility coefficients together, and grouped the self-care coefficients into two groups. As can be seen for other variables, different amounts fusing occurred for different groups of outcomes, and for some outcomes effects were shrunk towards each other but without exact fusing.

From the coefficient paths in Figure 7, our approach yields benefit in terms of predictive performance and also results in clinically sensible estimated coefficients. We chose these 6 variables to demonstrate effects of several classes of billing codes in the data: ICD codes, DME codes, pharmaceutical codes, and CPT codes. The three figures on the left of Figure 7 represent examples of diagnostic codes from billing data. The top panel represents the functional domains' coefficients path plot for children coded with autistic disorder. Children with autism spectrum disorders are diagnosed based on deficits in areas of social communication and behavior. The degree of impairment in social skills as well as comorbid neuro developmental disorders often leads to functional deficits in self-care. Mobility is typically unaffected in children with autism. Another example is the diagnosis code for shock. Shock is a life-threatening condition in which the body is not getting enough blood flow, resulting in end-organ damage, including damage to the brain. Consistent with this, in the figure we see from the coefficient path plots that the diagnosis code of shock is most related with the functional domain of cognition/communication as compared to mobility and self-care skills. Hypothyroidism is a disorder caused when the thyroid gland does not make enough hormone. In children, it has numerous etiologies including autoimmune reactions, brain injury, or radiation treatment. Hypothyroidism causes slow growth, decreased strength, and impaired cognition. The associated plot shows some relationship to the self-care and mobility domains, but a stronger relationship to the cognition/communication domain. The three figures in the right of the panel represent billing codes for a medication, a durable good, and a procedure. Methylphenidate Hydrochloride, a stimulant medication used to treat attention deficit hyperactivity disorder, showed minimal relationships with mobility and self-care. However a pronounced relationship with cognition/communication is observed, as would be expected in a child with being treated with this medication. An example of durable medical equipment are enteral feeding supplies. These are ordered for patients who require feeding through a tube inserted into the stomach. This plot shows a clear relationship between feeding supplies and eating outcomes, but markedly less strong relationship with other functional outcome domains. Mechanical chest wall oscillation is used in patients who have an impaired lung function such as in children with cystic fibrosis or other diseases which lead to general physical debilitation or loss of strength. The plot demonstrates that this procedure code holds associations with mobility and self-care limitations as would be expected, but does not relate to cognition/communication.

5 | DISCUSSION

In this paper, we introduced a doubly structured sparsity regression approach for modeling multivariate outcomes with known groupings to borrow information across related outcomes. Our approach allows for overlapping groupings of the responses and borrows strength by joint selection of the effects of variables across groups of related outcomes and encourages shrinkage of the effects of variables to be more similar across related outcomes. We prove an oracle property for an adaptive version of our penalty, showing that our approach results in estimates with the same asymptotic distribution as if the true non-zero coefficients and *which* variable effects are truly equal across related responses were both known in advance. The results pertaining to our motivating study indicate our approach yields benefit in terms of predictive performance and yields estimated coefficients that are clinically sensible, as was demonstrated in our analysis of the pediatric functional status data. Our work is motivated by modeling functional status scores in a pediatric population however the methods are applicable more generally. Since our approach does not rely on normality of the multivariate outcomes, it can be straightforwardly extended to more general models, such as generalized linear models both in theory and computationally. Incorporation of covariance information in addition to borrowing strength through joint selection and effect shrinkage may be a promising direction for additionally borrowing strength of information across outcomes. However, in this work we have not focused on this as it would introduce another tuning parameter and another layer of complexity to computation, which may limit the applicability of such an approach to large scale and high dimensional settings.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their careful reading of our manuscript and their helpful comments and suggestions. Research reported in this publication was supported by the Eunice Kennedy Shriver National Institute of Child Health & Human Development of the National Institutes of Health under Award Number R03HD101083. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

DATA AVAILABILITY STATEMENT

The data are not publicly available due to privacy restrictions.

ORCID

Jared D. Huling  <https://orcid.org/0000-0003-0670-4845>

REFERENCES

1. Lo W, Stephens J, Fernandez S. Pediatric stroke in the United States and the impact of risk factors. *J Child Neurol*. 2009;24(2):194-203.
2. Patel S, Bhatnagar A, Wear C, et al. Are pediatric brain tumors on the rise in the USA? Significant incidence and survival findings from the SEER database analysis. *Child's Nerv Syst*. 2014;30(1):147-154.
3. Taylor CA, Bell JM, Breiding MJ, Xu L. Traumatic brain injury-related emergency department visits, hospitalizations, and deaths—United States, 2007 and 2013. *MMWR Surveill Summar*. 2017;66(9):1.
4. Dhillon JK, Shi J, Janezic A, Wheeler KK, Xiang H, Leonard JC. US estimates of pediatric spinal cord injury: implications for clinical care and research planning. *J Neurotrauma*. 2017;34(12):2019-2026.
5. Schneier AJ, Shields BJ, Hostetler SG, Xiang H, Smith GA. Incidence of pediatric traumatic brain injury and associated hospital resource utilization in the United States. *Pediatrics*. 2006;118(2):483-492.
6. Székely GJ, Rizzo ML, Bakirov NK. Measuring and testing dependence by correlation of distances. *Ann Stat*. 2007;35(6):2769-2794.
7. Tibshirani R. Regression shrinkage and selection via the lasso. *J R Stat Soc. Ser B (Methodol)*. 1996;58:267-288.
8. Rothman AJ, Levina E, Zhu J. Sparse multivariate regression with covariance estimation. *J Comput Graph Stat*. 2010;19(4):947-962.
9. Yuan M, Lin Y. Model selection and estimation in the Gaussian graphical model. *Biometrika*. 2007;94(1):19-35.
10. Sofer T, Dicker L, Lin X. Variable selection for high dimensional multivariate outcomes. *Stat Sin*. 2014;24(4):1633.
11. Fan J, Li R. Variable selection via nonconcave penalized likelihood and its oracle properties. *J Am Stat Assoc*. 2001;96(456):1348-1360.
12. Zhang CH. Nearly unbiased variable selection under minimax concave penalty. *Ann Stat*. 2010;38(2):894-942.
13. Li Y, Nan B, Zhu J. Multivariate sparse group lasso for the multivariate multiple linear regression with an arbitrary group structure. *Biometrics*. 2015;71(2):354-363.
14. Jenatton R, Audibert JY, Bach F. Structured variable selection with sparsity-inducing norms. *J Mach Learn Res*. 2011;12:2777-2824.
15. Yuan M, Lin Y. Model selection and estimation in regression with grouped variables. *J R Stat Soc: Ser B (Stat Methodol)*. 2006;68(1):49-67.
16. Tibshirani R, Saunders M, Rosset S, Zhu J, Knight K. Sparsity and smoothness via the fused lasso. *J R Stat Soc: Ser B (Stat Methodol)*. 2005;67(1):91-108.
17. Zou H. The adaptive lasso and its oracle properties. *J Am Stat Assoc*. 2006;101(476):1418-1429.
18. Wang H, Leng C. A note on adaptive group lasso. *Comput Stat Data Anal*. 2008;52(12):5277-5286.
19. Huling JD, Yu M, Liang M, Smith M. Risk prediction for heterogeneous populations with application to hospital admission prediction. *Biometrics*. 2018;74(2):557-565.
20. Viallon V, Lambert-Lacroix S, Höfling H, Picard F. Adaptive generalized fused-lasso: asymptotic properties and applications. *Tech Rep*. 2013.
21. Guennebaud G, Jacob B. Eigen v3. 2010 <http://eigen.tuxfamily.org>
22. Bates D, Eddelbuettel D. Fast and elegant numerical linear algebra using the rcppeigen package. *J Stat Softw*. 2013;52(5):1-24.
23. Glowinski R, Marroco A. Sur l'approximation, par éléments finis d'ordre un, et la résolution, par pénalisation-dualité d'une classe de problèmes de Dirichlet non linéaires. *Rev Fr Autom Inform Rech Opér Anal Numér*. 1975;9(2):41-76.
24. Gabay D, Mercier B. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Comput Math Appl*. 1976;2(1):17-40.
25. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found Trends® Mach Learn*. 2011;3(1):1-122.
26. Chen C, He B, Ye Y, Yuan X. The direct extension of ADMM for multi-block convex minimization problems is not necessarily convergent. *Math Program*. 2016;155(1-2):57-79.
27. Huang J, Ma S, Zhang CH. Adaptive Lasso for sparse high-dimensional regression models. *Stat Sin*. 2008;18(4):1603-1618.
28. Bodenreider O. The unified medical language system (UMLS): integrating biomedical terminology. *Nucleic Acids Res*. 2004;32(1):D267-D270.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Huling JD, Lundine JP, Leonard JC. Doubly structured sparsity for grouped multivariate responses with application to functional outcome score modeling. *Statistics in Medicine*. 2023;42(15):2619-2636. doi: 10.1002/sim.9740