



# HHS Public Access

Author manuscript

*Biometrics*. Author manuscript; available in PMC 2023 June 25.

Published in final edited form as:

*Biometrics*. 2023 June ; 79(2): 826–840. doi:10.1111/biom.13632.

## A Joint Fairness Model with Applications to Risk Predictions for Under-represented Populations

Hyungrok Do<sup>1</sup>, Shinjini Nandi<sup>2</sup>, Preston Putzel<sup>3</sup>, Padhraic Smyth<sup>3</sup>, Judy Zhong<sup>1,\*</sup>

<sup>1</sup>Department of Population Health, NYU Grossman School of Medicine, New York, NY, 10016, USA

<sup>2</sup>Department of Mathematical Sciences, Montana State University, Bozeman, MT, 59717, USA

<sup>3</sup>Department of Computer Science, University of California, Irvine, CA, 92697, USA

### Summary:

In data collection for predictive modeling, under-representation of certain groups, based on gender, race/ethnicity, or age, may yield less-accurate predictions for these groups. Recently, this issue of fairness in predictions has attracted significant attention, as data-driven models are increasingly utilized to perform crucial decision-making tasks. Existing methods to achieve fairness in the machine learning literature typically build a single prediction model in a manner that encourages fair prediction performance for all groups. These approaches have two major limitations: i) fairness is often achieved by compromising accuracy for some groups; ii) the underlying relationship between dependent and independent variables may not be the same across groups. We propose a Joint Fairness Model (JFM) approach for logistic regression models for binary outcomes that estimates group-specific classifiers using a joint modeling objective function that incorporates fairness criteria for prediction. We introduce an Accelerated Smoothing Proximal Gradient Algorithm to solve the convex objective function, and present the key asymptotic properties of the JFM estimates. Through simulations, we demonstrate the efficacy of the JFM in achieving good prediction performance and across-group parity, in comparison with the single fairness model, group-separate model, and group-ignorant model, especially when the minority group's sample size is small. Finally, we demonstrate the utility of the JFM method in a real-world example to obtain fair risk predictions for under-represented older patients diagnosed with coronavirus disease 2019 (COVID-19).

### Keywords

algorithmic bias; algorithmic fairness; joint estimation; under-represented population

---

\* judy.zhong@nyulangone.org .

**Publisher's Disclaimer:** This article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: [10.1111/biom.13632](https://doi.org/10.1111/biom.13632)

Supporting Information

Web Appendices, Figures and Proofs referenced in Sections 2, 3, 4 and 5 are available with this paper at the Biometrics website on Wiley Online Library. The codes for the proposed JFM is available both on Wiley Online Library and on GitHub: <https://github.com/hyungrok-do/joint-fairness-model>.

# 1. Introduction

## 1.1 Applied Context

The issue of making fair predictions has attracted significant attention recently in machine learning as a critical issue in the application of data-driven models. Though machine learning models are increasingly utilized to perform crucial decision-making tasks, recent evidence reveals that many carefully designed algorithms can learn biases from the underlying data and exploit these inequities when making predictions. For example, large systematic biases in prediction performance have been detected for machine learning models in areas such as recidivism prediction relative to race (Angwin et al., 2016), ranking of job candidates relative to gender (Lahoti et al., 2018) and face recognition relative to both race and gender (Ryu et al., 2018; Buolamwini and Gebru, 2018).

There is an emerging recognition that such biases in data often leads to unfair performance from predictive models in healthcare for certain groups (Char et al., 2018), such as women (Larrazabal et al., 2020), ethnic and racial minorities (Seyyed-Kalantari et al., 2020; Chen et al., 2021), and individuals with public insurance (Seyyed-Kalantari et al., 2020; Chen et al., 2021). Biased representations of different populations in biomedical studies, and the under-performance for under-represented populations, limit the potential benefits that can be achieved for these communities. In particular, when model-based predictions are used to prioritize patients for rationed services (e.g. organ transplantation, care management programs, or ICU services), under-performance for the under-represented patient populations will lead to unfair treatments for these patients (Paulus and Kent, 2020).

A particularly motivating example that we use in this paper is the mortality prediction for patients infected with coronavirus disease 2019 (COVID-19). As of January 23 2021, COVID-19 has infected more than 96 million people globally, accounting for more than 2 million known deaths. Older patients are particularly vulnerable to severe outcomes and death due to COVID-19. The Centers for Disease Control and Prevention (CDC) reported that the fatality rate was 18.8% for patients older than 80 years whereas the overall fatality rate is estimated at 5% or less for all patients (Kompaniyets and Goodman, 2021). This difference in survival highlights an urgent need for risk stratification of older patients with COVID-19 based on routine clinical assessments. However, most COVID-19 studies have not been stratified by age groups (Tehrani et al., 2021). Thus, as an example, when applying a risk prediction equation generated from the general population to older patients with COVID-19, the model in Tehrani et al. (2021) predicts high-risk scores overall due to their older age, higher prevalence of comorbidities and more laboratory abnormalities. This resulted in insufficient and unfair risk stratification for these patients as not all older patients are at the same risk of death from COVID-19 (Tehrani et al., 2021).

## 1.2 Existing Approaches

Methods to address fairness in the machine learning literature typically begin with a formal probabilistic definition of fairness. In the context of risk prediction, predictive fairness at the group level means that a risk prediction model has performance characteristics (such as accuracy, ranking, or calibration) that are relatively independent of group membership. For

example, if the false positive rate for a classification model is defined as  $P(\hat{y} = 1 | y = 0)$ , where  $\hat{y}$  is the model's prediction, then enforcing equality can be stated as requiring that these distributions be as close as possible between groups. Other definitions include demographic parity (Calders et al., 2009), equalized odds or equal opportunity (Hardt et al., 2016), disparate treatment, and impact and mistreatment (Zafar et al., 2019, 2017a). It is well-recognized that there is no unique optimal way to define fairness, leading to trade-offs between different approaches (Zafar et al., 2017b).

Given a fairness criterion, the second component of a fairness strategy requires an algorithmic approach, typically consisting of either 1) pre-processing the data by mapping the training data to a transformed space where the dependencies between sensitive attributes and class labels disappear (Kamiran and Calders, 2012; Dwork et al., 2018); or 2) post-processing of a trained prediction model; for example, (Kamishima et al., 2012; Hardt et al., 2016) modify the probability of the decision being positive and negative predictions from an existing classifier to limit unfair discrimination; or 3) "in-process" methods, where fairness is accounted for during training of a model by adding a fairness constraint to the objective function. Examples of in-process methods include Zemel et al. (2013) who proposed to learn a fair representation of the data and classifier parameters by optimizing a non-convex function, and Zafar et al. (2017b) who defined a convex function as a measure of (un)fairness and suggested optimizing accuracy subject to the convex fairness constraints as well as their converse.

A key feature of nearly all existing approaches is that a single set of classifier parameters is estimated using fairness criteria that encourage fair prediction performance across all groups. This approach has two main limitations: i) fairness is often achieved by compromising accuracy of some groups; ii) the underlying relationship between dependent and independent variables may not be the same across groups, and the differences in predictive features may be of interest. In the example of predicting mortality risk for patients with COVID-19, while one would expect some features to have the same association with mortality for both older and younger patients, the associations between mortality and other features may be different between age groups. For instance, being overweight or obese (Body Mass Index [BMI]  $> 25 \text{ kg/m}^2$ ) increases the risk for mortality associated with COVID-19, particularly among adults aged  $< 65$  years (Kompaniyets and Goodman, 2021) However, geriatric BMI guidelines are different from those for younger adults. For older adults, higher BMIs are often associated with greater energy stores and a better nutritional state overall, which is beneficial for patients' survival outcomes when serious infections develop.

Estimating separate prediction models for each group does not leverage potential similarities between the groups. Moreover, estimating a single prediction model, even while using fairness criteria, will likely result in sub-optimal estimation or prediction performance for one group in order to achieve fair performance with a single set of parameters shared across groups. Outside the context of algorithmic fairness, Danaher et al. (2014) proposed the joint graphical lasso method, a technique for jointly estimating multiple models corresponding to distinct but related conditions. Their approach is based upon a penalized log-likelihood approach, which penalizes the differences between parameter estimates across groups.

Penalized log-likelihood approaches have often been used by other authors like Yuan and Lin (2007), Friedman et al. (2007b) etc. for similar estimation purposes while minimizing the disparities in estimates across groups. In all such cases, however, prediction performance was not emphasized.

In this paper, we propose the joint fairness model, a technique for jointly estimating multiple logistic regression models corresponding to distinct but related groups, in order to achieve fair prediction performance across groups. The model parameters are estimated by encouraging prediction fairness, while simultaneously ensuring high predictive accuracy irrespective of heterogeneity across groups. The rest of this paper is organized as follows. In Section 2, we present the proposed joint fairness model. Section 3 describes the algorithm to find its optimal solution. In Section 4, we discuss asymptotic consistency of the estimators. We illustrate the performance of our proposed approach in simulation studies in Section 5; and apply our approach to the motivating example of predicting COVID-19 mortality outcomes for patients of different age groups in Section 6. Finally, we summarize and discuss our findings in Section 7.

## 2. Problem Formulation

For binary outcomes, consider  $K$  groups of datasets  $\mathcal{S}_k = \{(\mathbf{X}_{ki}, y_{ki}) \in \mathbb{R}^p \times \{0, 1\} : i = 1, \dots, n_k\}$  with  $K \geq 2$  representing group membership. Throughout the paper, group memberships are known and observed. Assuming that the  $n = \sum_{k=1}^K n_k$  observations are independently distributed, then  $y_{ki} \sim \text{Bernoulli}(p_{ki})$ , and  $\hat{y}_{ki} : \mathbb{R}^p \rightarrow \{0, 1\}$  is the predicted value based on predictor features  $\mathbf{X}_{ki}$ . We focus on the development of a fair prediction approach for the widely-used logistic regression model. The log-likelihood of the logistic model for the data from all groups takes the form

$$\sum_{k=1}^K \ell(\boldsymbol{\beta}_k; \mathbf{X}_k, \mathbf{y}_k) = \sum_{k=1}^K \sum_{i=1}^{n_k} (y_{ki} \mathbf{X}_{ki} \boldsymbol{\beta}_k - \log(1 + \exp(\mathbf{X}_{ki} \boldsymbol{\beta}_k))). \quad (1)$$

Define  $\boldsymbol{\beta} = (\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_K) \in \mathbb{R}^{pK}$ . Maximizing the likelihood function in (1) with respect to  $\boldsymbol{\beta}_k$  in each group separately yields the maximum likelihood estimates of parameters  $\hat{\boldsymbol{\beta}}_k$  for each group  $k$ , thus making separate predictions  $\hat{y}_k$  per group. If we ignore group memberships,  $\hat{\boldsymbol{\beta}}$  can be estimated by maximizing the likelihood function in equation (1) setting all  $\boldsymbol{\beta}_k$  equal to a single global parameter vector  $\hat{\boldsymbol{\beta}}$  and making predictions  $\hat{y}$  per individual (irrespective of group) using that global parameter vector.

If the  $K$  datasets correspond to observations collected from  $K$  distinct but related groups, then one might wish to borrow strength across the  $K$  groups to estimate  $\boldsymbol{\beta}$  and predict  $\hat{y}$ , rather than estimating parameters  $\boldsymbol{\beta}_k$  for each group separately, or estimating a single set of  $\boldsymbol{\beta}_k$  for all  $k$  which could lead to heterogeneous prediction performance across the groups. Therefore, instead of estimating  $\boldsymbol{\beta}$  by maximizing the likelihood in equation (1), we consider a penalized log-likelihood approach and jointly estimate  $\boldsymbol{\beta}$  by maximizing an objective

function of  $\sum_{k=1}^K \ell(\boldsymbol{\beta}_k; \mathbf{X}_k, \mathbf{y}_k)$  in equation (1) subject to constraints on (i) fairness,  $\mathcal{P}_F(\boldsymbol{\beta}; \mathbf{X}, \mathbf{y}, \lambda_F)$  (ii) parameter similarity,  $\mathcal{P}_{\text{sim}}(\boldsymbol{\beta}; \lambda_{\text{sim}})$ , and (iii) parameter sparsity,  $\mathcal{P}_{\text{sp}}(\boldsymbol{\beta}; \lambda_{\text{sp}})$ .

$$\begin{aligned} \text{minimize } F(\boldsymbol{\beta}) = & - \sum_{k=1}^K \frac{1}{n_k} \ell(\boldsymbol{\beta}_k; \mathbf{X}_k, \mathbf{y}_k) + \mathcal{P}_F(\boldsymbol{\beta}; \mathbf{X}, \mathbf{y}, \lambda_F) + \mathcal{P}_{\text{sim}}(\boldsymbol{\beta}; \lambda_{\text{sim}}) \\ & + \mathcal{P}_{\text{sp}}(\boldsymbol{\beta}; \lambda_{\text{sp}}). \end{aligned} \quad (2)$$

We propose a fairness penalty function  $\mathcal{P}_F(\boldsymbol{\beta}; \mathbf{X}, \mathbf{y}, \lambda_F)$  that encourages each group to have similar predictive performance. In this work we use equalized odds (Hardt et al., 2016) which encourages each group to have similar false positive rates (FPRs) and false negative rates (FNRs). Thus, we want to minimize the absolute difference between  $\text{FPR}_j$  and  $\text{FPR}_k |P(\hat{y}_j = 1 | y_j = 0) - P(\hat{y}_k = 1 | y_k = 0)|$ , and that between  $\text{FNR}_j$  and  $\text{FNR}_k |P(\hat{y}_j = 0 | y_j = 1) - P(\hat{y}_k = 0 | y_k = 1)|$ . Under the logistic regression model, the absolute differences of FPRs and FNRs are nonconvex due to the nonconvexity of the sigmoid function. We will instead minimize the absolute difference of the expected linear components of the two groups  $|\mathbb{E}[\mathbf{X}_j \boldsymbol{\beta}_j | y_j = 0] - \mathbb{E}[\mathbf{X}_k \boldsymbol{\beta}_k | y_k = 0]|$ . The proposition below shows that the absolute difference of the expected probabilities is upper-bounded by the absolute difference of the expected linear components. Thus, minimizing group differences in expected linear components guarantees that group difference of false predictions is minimized.

PROPOSITION 2.1: For any  $1 \leq j, k \leq K, j \neq k$ , and  $y \in \{0, 1\}$  the following inequality holds:

$$\begin{aligned} & \left| \mathbb{E} \left[ \frac{1}{1 + \exp(-\mathbf{X}_j \boldsymbol{\beta}_j)} \middle| y_j = y \right] - \mathbb{E} \left[ \frac{1}{1 + \exp(-\mathbf{X}_k \boldsymbol{\beta}_k)} \middle| y_k = y \right] \right| \\ & \leq \frac{1}{4} |\mathbb{E}[\mathbf{X}_j \boldsymbol{\beta}_j | y_j = y] - \mathbb{E}[\mathbf{X}_k \boldsymbol{\beta}_k | y_k = y]|. \end{aligned}$$

Proof: See Web Appendix 4.

Note that the empirical estimate of the expectation is

$$\mathbb{E}[\mathbf{X}_k \boldsymbol{\beta}_k | y_k = y] = \frac{1}{|S_{ky}|} \sum_{i \in S_{ky}} \mathbf{X}_i \boldsymbol{\beta}_k,$$

where  $S_{ky} = \{i: y_{ki} = y\}$  is a subgroup of subjects with a true response value  $y$  in group  $k$ , with  $y \in \{0, 1\}$ . Thus, our fairness penalty, that bridges the between-group gaps in the linear components of  $\text{FPR}_k$  and  $\text{FNR}_k$ , is defined as:

$$\begin{aligned} \mathcal{P}_F(\boldsymbol{\beta}; \mathbf{X}, \mathbf{y}, \lambda_F) &= \mathcal{P}_{\text{FPR}}(\boldsymbol{\beta}; \mathbf{X}, \mathbf{y}, \lambda_F) + \mathcal{P}_{\text{FNR}}(\boldsymbol{\beta}; \mathbf{X}, \mathbf{y}, \lambda_F) \\ &= \lambda_F \sum_{j < k} \left| \frac{1}{|S_{j0}|} \sum_{i \in S_{j0}} \mathbf{X}_i \boldsymbol{\beta}_j - \frac{1}{|S_{k0}|} \sum_{i \in S_{k0}} \mathbf{X}_i \boldsymbol{\beta}_k \right| \\ &+ \lambda_F \sum_{j < k} \left| \frac{1}{|S_{j1}|} \sum_{i \in S_{j1}} \mathbf{X}_i \boldsymbol{\beta}_j - \frac{1}{|S_{k1}|} \sum_{i \in S_{k1}} \mathbf{X}_i \boldsymbol{\beta}_k \right| \end{aligned} \quad (3)$$

where the summation  $\sum_{j < k}$  represents  $\sum_{k=2}^K \sum_{j=1}^{k-1}$  for notational simplicity.

The similarity penalty  $\mathcal{P}_{\text{sim}}(\boldsymbol{\beta}; \lambda_{\text{sim}})$  is chosen to encourage similarity across the  $K$  parameter estimates. Here we use the generalized fused lasso penalty (Hoefling, 2010; Danaher et al., 2014; Dondelinger et al., 2018) defined as

$$\mathcal{P}_{\text{sim}}(\boldsymbol{\beta}; \lambda_{\text{sim}}) = \lambda_{\text{sim}} \sum_{j < k} \|\boldsymbol{\beta}_j - \boldsymbol{\beta}_k\|_1, \quad (4)$$

herein referred to as the fusion-similarity penalty. The sparsity penalty  $\mathcal{P}_{\text{sp}}(\boldsymbol{\beta})$  is chosen to encourage sparse estimates and to avoid ill-defined maximum likelihood estimates when sample sizes  $n_k < p$ .

$$\mathcal{P}_{\text{sp}}(\boldsymbol{\beta}; \lambda_{\text{sp}}) = \sum_{k=1}^K \lambda_{\text{sp}_k} \|\boldsymbol{\beta}_k\|_1. \quad (5)$$

In the three penalty functions,  $\lambda_{\text{F}}$ ,  $\lambda_{\text{sim}}$ , and  $\lambda_{\text{sp}}$  are nonnegative hyperparameters. Here  $\mathcal{P}_{\text{sp}}(\boldsymbol{\beta}; \lambda_{\text{sp}})$ ,  $\mathcal{P}_{\text{F}}(\boldsymbol{\beta}; \mathbf{X}, \mathbf{y}, \lambda_{\text{F}})$ , and  $\mathcal{P}_{\text{sim}}(\boldsymbol{\beta}; \lambda_{\text{sim}})$  are convex penalty functions, so that the objective function in equation (2) is convex in  $\boldsymbol{\beta}$ . The proposed model jointly estimates  $\boldsymbol{\beta}$  to achieve fair performance across groups, herein referred to as the joint fairness model (JFM). In contrast, the dominant approach for fair predictions in the current literature is to estimate a single set of  $\boldsymbol{\beta}$  parameters with constraints on the quality of performance metrics across groups. In the context of logistic regression, Bechavod and Ligett (2017) proposed a Single Fairness Model (SFM) to minimize the following objective function.

$$\underset{\boldsymbol{\beta} \in \mathbb{R}^p}{\text{minimize}} - \sum_{k=1}^K \ell(\boldsymbol{\beta}; \mathbf{X}_k, \mathbf{y}_k) + \lambda_{\text{F}} \sum_{j < k} \sum_{y \in \{0, 1\}} \left| \frac{1}{|S_{jy}|} \sum_{i \in S_{jy}} \mathbf{X}_i \boldsymbol{\beta} - \frac{1}{|S_{ky}|} \sum_{i \in S_{ky}} \mathbf{X}_i \boldsymbol{\beta} \right| + \lambda_{\text{sp}} \|\boldsymbol{\beta}\|_1.$$

In contrast to the proposed JFM, it does not allow for group-specific different  $\boldsymbol{\beta}_k$ s, leading to two potential limitations: i) fairness is often achieved by compromising the likelihood ( $\ell(\boldsymbol{\beta}; \mathbf{X}_k, \mathbf{y}_k)$ ) of some groups; ii) inflexible model mis-specification when the  $\boldsymbol{\beta}_k$ s are different.

The proposed JFM objective function improves prediction parity through three components. First, it considers a weighted group total likelihood to upweight the groups with smaller sample sizes. Second,  $\mathcal{P}_{\text{F}}(\boldsymbol{\beta}; \mathbf{X}, \mathbf{y}, \lambda_{\text{F}})$  encourages estimates that achieve fair performance across groups. Third, the similarity penalty term improves estimation and prediction efficiencies when multiple subgroups are related but not identical (Hoefling, 2010; Danaher et al., 2014). Computationally, we found that multiple different combinations of  $\hat{\boldsymbol{\beta}}_k$  often result in very similar objective values; thus, the similarity term can help optimization algorithms for the JFM converge to one of the multiple combinations by favoring similar values of  $\hat{\boldsymbol{\beta}}_k$ . Other formats of the similarity penalty could be used in our proposed JFM framework. For example, the group lasso penalty (Yuan and Lin, 2006) has been shown to encourage similar sparsity patterns across groups (Obozinski et al., 2010; Danaher et al., 2014), while the fused lasso term is more aggressive in encouraging similar

$\hat{\beta}_k$  estimates. The penalty functions in (3), (4), and (5) are based on the  $L_1$  norm. They can be flexibly adapted to an  $L_2$  penalization or to a combination of  $L_1$  and  $L_2$  penalizations. The differences between  $L_1$  and  $L_2$  penalties have been well discussed in the literature (Tibshirani, 1996; Zou and Hastie, 2005). For the fairness penalty, Bechavod and Ligett (2017) showed that there are no significant differences in the empirical predictive performances between the  $L_1$  and  $L_2$  fairness penalty forms other than the  $L_2$  form of the similarity penalty penalizes large differences more aggressively so that models have less chance to obtain group-specific estimates.

### 3. Accelerated Smoothing Proximal Gradient Algorithm for JFM

In this section, we introduce an accelerated smoothing proximal gradient (ASPG) algorithm (Chen et al., 2012) to solve Problem (2) for the JFM. The objective function of (2) is convex in  $\beta$  so that a global optimal solution can be attained. However, conventional proximal gradient-based or coordinate descent approaches (generally used for lasso-like methods) cannot be directly applied to solve Problem (2) because there is no closed-form solution for the proximal operator associated with  $\mathcal{P}_{\text{FPR}}$  and  $\mathcal{P}_{\text{FNR}}$ .

#### 3.1 Nesterov smooth approximation

To overcome the difficulty originating from the non-differentiability of the fairness and similarity penalties, we decouple the terms into a linear combination of the decision variables via the dual norm and then apply the Nesterov smoothing approximation (Nesterov, 2005). We start with matrix representations of the fairness penalty terms  $\mathcal{P}_{\text{FPR}}(\beta; \mathbf{X}, \mathbf{y}, \lambda_{\text{F}}) = \lambda_{\text{F}} \|\mathbf{D}_0 \beta\|_1$  and  $\mathcal{P}_{\text{FNR}}(\beta; \mathbf{X}, \mathbf{y}, \lambda_{\text{F}}) = \lambda_{\text{F}} \|\mathbf{D}_1 \beta\|_1$ , where  $\mathbf{D}_y \in \mathbb{R}^{K(K-1)/2 \times pK}$  is defined as below. Similarly, we use the matrix representation of the similarity penalty  $\mathcal{P}_{\text{Sim}}(\beta; \lambda_{\text{Sim}}) = \lambda_{\text{Sim}} \|\mathbf{F} \beta\|_1$  with  $\mathbf{F}$  defined as below.

$$\mathbf{D}_y = \begin{pmatrix} \bar{\mathbf{x}}_{1y} & -\bar{\mathbf{x}}_{2y} & \mathbf{0} & \cdots & \mathbf{0} \\ & \vdots & & & \\ \mathbf{0} & \bar{\mathbf{x}}_{2y} & -\bar{\mathbf{x}}_{3y} & \cdots & \mathbf{0} \\ & & \vdots & & \end{pmatrix} \quad \mathbf{F} = \begin{pmatrix} \mathbf{I}_p & -\mathbf{I}_p & \mathbf{0} & \cdots & \mathbf{0} \\ & \vdots & & & \\ \mathbf{0} & \mathbf{I}_p & -\mathbf{I}_p & \cdots & \mathbf{0} \\ & & \vdots & & \end{pmatrix}$$

Here,  $\bar{\mathbf{x}}_{ky} = \frac{1}{|S_{ky}|} \sum_{i \in S_{ky}} \mathbf{X}_i$  is the average predictor vector for group  $k$  with true outcome  $y$ , and  $\mathbf{I}_p$  is the  $p$ -dimensional identity matrix. The matrix form of the fairness penalty term and the similarity penalty term is therefore defined as:

$$\mathcal{P}_{\text{F}}(\beta; \mathbf{X}, \mathbf{y}, \lambda_{\text{F}}) + \mathcal{P}_{\text{Sim}}(\beta; \lambda_{\text{Sim}}) = \left\| \begin{pmatrix} \lambda_{\text{F}} \mathbf{D}_0 \\ \lambda_{\text{F}} \mathbf{D}_1 \\ \lambda_{\text{Sim}} \mathbf{F} \end{pmatrix} \beta \right\|_1 = \|\mathbf{D}_{\lambda_{\text{F}}, \lambda_{\text{Sim}}} \beta\|_1.$$

Thus, the objective function (2) can be written in matrix form:

$$\underset{\beta}{\text{minimize}} - \sum_{k=1}^K \frac{1}{n_k} \ell(\beta_k; \mathbf{X}_k, \mathbf{y}_k) + \|\mathbf{D}_{\lambda_{\text{F}}, \lambda_{\text{Sim}}} \beta\|_1 + \sum_{k=1}^K \lambda_{\text{Spk}} \|\beta_k\|_1, \quad (6)$$

where the associated proximal operator of  $\|\mathbf{D}_{\lambda_F, \lambda_{\text{Sim}}}\boldsymbol{\beta}\|_1$  does not have a closed form solution. We apply the Nesterov smooth approximation to approximate  $\|\mathbf{D}_{\lambda_F, \lambda_{\text{Sim}}}\boldsymbol{\beta}\|_1$  by a smooth function

$$f_\mu(\boldsymbol{\beta}; \lambda_F, \lambda_{\text{Sim}}) = \sup\left\{\boldsymbol{\alpha}^T \mathbf{D}_{\lambda_F, \lambda_{\text{Sim}}}\boldsymbol{\beta} - \frac{\mu}{2} \|\boldsymbol{\alpha}\|_2^2 : \|\boldsymbol{\alpha}\|_\infty \leq 1\right\}. \quad (7)$$

Proposition A.1 in Web Appendix 1 provides the maximum gap between  $\|\mathbf{D}_{\lambda_F, \lambda_{\text{Sim}}}\boldsymbol{\beta}\|_1$  and its Nesterov approximation  $f_\mu(\boldsymbol{\beta}; \lambda_F, \lambda_{\text{Sim}})$ .

As demonstrated in Proposition A.2 in Web Appendix 1, for any  $\mu > 0$ ,  $f_\mu(\boldsymbol{\beta}; \lambda_F, \lambda_{\text{Sim}})$  is smooth and convex with respect to  $\boldsymbol{\beta}$ , whose gradient takes the following form:

$$\nabla f_\mu(\boldsymbol{\beta}; \lambda_F, \lambda_{\text{Sim}}) = \mathbf{D}_{\lambda_F, \lambda_{\text{Sim}}}^T \boldsymbol{\alpha}^*, \quad (8)$$

where  $\boldsymbol{\alpha}^* = \operatorname{argmax}\{\boldsymbol{\alpha}^T \mathbf{D}_{\lambda_F, \lambda_{\text{Sim}}}\boldsymbol{\beta} - \frac{\mu}{2} \|\boldsymbol{\alpha}\|_2^2 : \|\boldsymbol{\alpha}\|_\infty \leq 1\}$ . Moreover, the gradient is Lipschitz continuous with a Lipschitz constant  $L_\mu = \mu^{-1} \|\mathbf{D}_{\lambda_F, \lambda_{\text{Sim}}}\|_2^2$ , where  $\|\cdot\|_2$  denotes the matrix spectral norm (which is equivalent to the largest singular value of the matrix). We can show further that  $\boldsymbol{\alpha}^*$  can be calculated as  $\boldsymbol{\alpha}^* = \mathcal{S}_\infty(\mu^{-1} \mathbf{D}_{\lambda_F, \lambda_{\text{Sim}}}\boldsymbol{\beta})$ , where  $\mathcal{S}_\infty(\cdot)$  is the projection onto the unit  $L_\infty$  ball such that  $[\mathcal{S}_\infty(x)]_i = x_i \mathbb{1}_{[-1, 1]}(x_i) + \mathbb{1}_{(1, \infty)}(x_i) - \mathbb{1}_{(-\infty, -1)}(x_i)$ , where  $\mathbb{1}$  is the indicator function. Details are provided in Web Appendix 1.

### 3.2 Accelerated Smoothing Proximal Gradient Algorithm

With  $\|\mathbf{D}_{\lambda_F, \lambda_{\text{Sim}}}\boldsymbol{\beta}\|_1$  substituted by the Nesterov smooth approximation  $f_\mu(\boldsymbol{\beta}; \lambda_F, \lambda_{\text{Sim}})$ , problem (6) becomes

$$\underset{\boldsymbol{\beta}}{\text{minimize}} \tilde{F}(\boldsymbol{\beta}) = - \sum_{k=1}^K \frac{1}{n_k} \ell(\boldsymbol{\beta}_k; \mathbf{X}_k, \mathbf{y}_k) + f_\mu(\boldsymbol{\beta}; \lambda_F, \lambda_{\text{Sim}}) + \sum_{k=1}^K \lambda_{\text{Spk}} \|\boldsymbol{\beta}_k\|_1, \quad (9)$$

whose first two terms are convex smooth functions. Although the sparsity penalty term  $\sum_k \lambda_{\text{Spk}} \|\boldsymbol{\beta}_k\|_1$  is non-differentiable, it can be managed through the proximal gradient method using the soft-thresholding operator  $\mathcal{S}$  with a closed form solution (Friedman et al., 2007a).

Algorithm 1 presents the proposed ASPG algorithm, starting from parameter initialization, to gradient descent iterations with proximal and momentum steps, until convergence. Although Algorithm 1 minimizes the Nesterov smooth approximation, Theorem 3.1 proves that the solution can reach arbitrarily close to the global optimum of Problem (2).

**THEOREM 3.1:** *Let  $\{\boldsymbol{\beta}^{(t)} : t = 1, 2, \dots\}$  be a sequence generated by Algorithm 1. Then for any  $t \geq 1$  and desired  $\delta > 0$*

$$F(\boldsymbol{\beta}^{(t)}) - F(\boldsymbol{\beta}^{* *}) \leq \delta + \frac{2L \|\boldsymbol{\beta}^{(0)} - \boldsymbol{\beta}^*\|_2^2}{t^2},$$



where  $\beta^*$  and  $\beta^{**}$  are global minimizers of Problem (9) and Problem (2), respectively and  $L$  is the Lipschitz constant of the function in (7).

Proof: See Web Appendix 4.

Based on Theorem 3.1, the rate of the convergence of Algorithm 1 is  $\mathcal{O}\left(\sqrt{\frac{pK}{\delta(\varepsilon - \delta)}}\right)$ , given a desired accuracy  $\varepsilon > 0$ . The complexity of a single iteration of Algorithm 1 is  $\mathcal{O}((n+K^2)pK)$ . For additional details see Propositions A.4 and A.5 in Web Appendix 1. The proposed JFM Algorithm 1 is for the fusion-similarity term. The algorithm can be readily extended to include a group-similarity term as presented in Web Appendix 2.

#### 4. Asymptotic properties of the JFM estimates

We now present asymptotic results for the JFM parameter estimates  $\hat{\beta}$  obtained by solving Problem (2). We assume  $p$  remains constant and  $n$  increases to infinity.

Consider the following assumptions:

*Assumption 1.*  $\mathcal{I}(\beta_k)/n_k \rightarrow \mathbf{C}_k$ , where  $\mathbf{C}_k$  is a positive definite  $p \times p$  matrix, for  $k = 1, \dots, K$ , where  $\mathcal{I}(\beta_k)$  is the information matrix of size  $p \times p$ . For simplicity, we assume there are no intercept terms in  $\beta_k$ .

*Assumption 2.* As  $n = \min_{k=1, \dots, K} n_k \rightarrow \infty$ ,  $\max_{\hat{\beta}_k} \left\| \left( \mathcal{I}(\hat{\beta}_k)^{-\frac{1}{2}} \right) \mathcal{I}(\beta_k) \left( \mathcal{I}(\hat{\beta}_k)^{-\frac{1}{2}} \right)^T - \mathbf{I}_p \right\|_2 \rightarrow 0$ , where  $\mathcal{I}(\hat{\beta}_k)$  is the empirical information matrix, and  $\mathbf{I}_p$  is a  $p \times p$  identity matrix.

The following theorem proves  $\sqrt{n}$ -consistency for the estimators, complying with the fairness and similarity constraints between the two groups as well as the sparsity constraint. We note that the theorem holds even if the sample size of one group increases faster than the other group's.

**THEOREM 4.1:** Let  $\hat{\beta}_k$  for  $k = 1, \dots, K$ , minimize the loss function (2). If

$\lambda_F^{(0)}/\sqrt{n} \rightarrow \lambda_F^{(0)} \geq 0$ ,  $\lambda_{\text{Sim}}^{(0)}/\sqrt{n} \rightarrow \lambda_{\text{Sim}}^{(0)} \geq 0$ , and  $\lambda_{\text{Sp}}^{(0)}/\sqrt{n} \rightarrow \lambda_{\text{Sp}}^{(0)} \geq 0$ , then under the assumptions 1 and 2

$$\sqrt{n}(\hat{\beta}_k - \beta_k) \xrightarrow{d} \hat{\mathbf{u}}_k, \quad (10)$$

where  $(\hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_K) = \text{argmin } \mathcal{V}(\mathbf{u}_1, \dots, \mathbf{u}_K)$ ,

$$\begin{aligned} \mathcal{V}(\mathbf{u}_1, \dots, \mathbf{u}_K) &= \sum_{k=1}^K \mathbf{u}_k^T \mathbf{W}_k + \frac{1}{2} \sum_{k=1}^K \mathbf{u}_k^T \mathbf{C}_k \mathbf{u}_k \\ &+ \lambda_F^{(0)} \sum_{j < k} \sum_{y \in \{0, 1\}} \mathcal{F}(\bar{\mathbf{X}}_{jy} \mathbf{u}_j - \bar{\mathbf{X}}_{ky} \mathbf{u}_k, \bar{\mathbf{X}}_{jy} \beta_j - \bar{\mathbf{X}}_{ky} \beta_k) \\ &+ \lambda_{\text{Sim}}^{(0)} \sum_{j < k} \sum_{l=1}^p \mathcal{F}(u_{jl} - u_{kl}, \beta_{jl} - \beta_{kl}) + \lambda_{\text{Sp}}^{(0)} \sum_{k=1}^K \sum_{l=1}^p \mathcal{F}(u_{kl}, \beta_{kl}). \end{aligned}$$

Here  $\mathcal{T}(u, \beta) = u \cdot \text{sign}(\beta) \cdot \mathbb{1}(\beta \neq 0) + |u| \cdot \mathbb{1}(\beta = 0)$ , and  $\mathbf{W}_k \sim \mathcal{N}_p(\mathbf{0}, \mathbf{C}_k)$ , where  $\mathbf{C}_k = \lim_{n_k \rightarrow \infty} \frac{1}{n_k} \sum_{i=1}^{n_k} \mathbf{X}_{ki} \mathbf{X}_{ki}^T$ , and  $\frac{1}{|S_{ky}|} \sum_{i \in S_{ky}} \mathbf{X}_i = \bar{\mathbf{X}}_{ky}$  for  $k = 1, \dots, K$  and  $y = 0, 1$ .

Proof: See Web Appendix 4.

## 5. Simulation Study

We performed a series of simulations to evaluate the proposed JFM, and compared it with the approaches of a group-separate individual logistic regression model, a group-ignorant vanilla logistic regression model, and the SFM method (Bechavod and Ligett, 2017) implemented using an SFM-ASPG algorithm (see Web Appendix 3). When applying the group-separate model, regression coefficients were estimated for each group separately with an  $L_1$  penalty. The group-ignorant model estimates a single logistic regression with group membership as an additional covariate with an  $L_1$  penalty.

### 5.1 Simulation Setup

We consider a two-group problem ( $K = 2$ ) for simplicity, with group 1 as the over-represented group and group 2 as the under-represented group with respect to the sample sizes. The training samples were simulated as follows. The predictor matrix  $\mathbf{X}_k$  was independently generated from a standard normal distribution. The binary outcome  $y_{ki}$  was then simulated from Bernoulli( $\pi_i(\mathbf{x}_{ki})$ ), where  $\pi_i(\mathbf{x}_{ki}) = \frac{\exp(\mathbf{x}_{ki} \boldsymbol{\beta}_k)}{1 + \exp(\mathbf{x}_{ki} \boldsymbol{\beta}_k)}$ . Out of the total number of features, 40% in each group had non-zero coefficients ( $\boldsymbol{\beta}$ 's). The non-zero coefficients were each set to the value 3. The simulations were conducted under three scenarios to investigate performances at various levels of shared parameters, sample sizes and dimensionalities.

- **Scenario 1 (Difference in True Model):** The proportion of shared features between the two groups ranged from 0% to 100% of features with non-zero coefficients. The intercepts were selected so that the baseline event prevalences were at 30% and 50% for the under- and over-represented groups. The sample sizes were set at 500 and 200 for group 1 and 2 respectively. The number of features was set to  $p = 100$ . Note that when small proportions of features are shared between groups, the fairness and similarity penalty terms are misspecified given that  $\mathbf{X} \boldsymbol{\beta}_k$  and  $\boldsymbol{\beta}_k$  are different between groups. As a result, performances in these settings allow us to investigate the robustness of the JFM approach to misspecification of the fairness and similarity penalty terms.
- **Scenario 2 (Difference in Sample Size):** The sample size of the under-represented group (group 2) ranged from 50 to 300 while the sample size of group 1 was fixed at 500. The number of features was set to  $p = 100$ . Half of the features with non-zero coefficients were shared between the groups.
- **Scenario 3 (High Dimensionality):** The number of features  $p$  ranged from 50 to 2,000. Sample sizes were 500 and 200 for group 1 and 2 respectively. For each value of  $p$ , 40 features had non-zero coefficients, with half of the non-zero features being shared.

We evaluated the methods on independent test datasets with large sample sizes ( $n = 1,000$  for both groups) under the same simulation setups. The Area under the Receiver Operating Characteristic curve (AUC) was used to assess the predictive ability of each model. Prediction unfairness was assessed by the group difference in AUCs. In addition, Mean Squared Errors (MSEs) for all  $\hat{\beta}$ s were examined to assess parameter estimation performance, and TPRs of the associated features and TNRs of the non-associated features were used to assess variable selection performance. Medians and interquartile ranges (IQRs) of the assessment metrics were generated from 20 replicates for each experiment. Predictive performances and their unfairness in terms of FPR and FNR were calculated with a predicted probability cutoff of 0.5, as presented in Web Appendix 8. We present additional simulation scenarios in Web Appendix 9.

## 5.2 Choice of the Hyperparameters

The group-ignorant model, group-separate model, SFM, and JFM contain 1,  $K$ , 2, and  $K + 2$  hyperparameters respectively. For every method, five-fold cross-validation on the training dataset was used to determine the hyperparameters. For the vanilla models (group-separate and group-ignorant), the lasso penalty term was selected by optimizing cross-validation AUCs. For the fairness-aware models (SFM and JFM), we investigated and compared a series of evaluation metrics for selecting hyperparameters using cross-validation. The metrics include overall AUCs and Brier scores (defined as  $\sum (\hat{p}_{ki} - y_{ki})^2$ ) on all samples (ignoring group memberships), group average of AUCs or Brier scores, and the group average of AUCs or Brier scores subtracting their disparities. Web Appendix 7, Web Figures 2, 3, and 4 demonstrate the empirical performance of the hyperparameters selected by the various strategies in the test datasets for the simulation scenarios. In summary, we find that the hyperparameters optimizing cross-validated group-average metrics showed better performance than sample-average or average metrics subtracting disparities. The hyperparameters optimizing AUCs in general generated the best AUC performances, while the hyperparameters optimizing Brier scores generated the best MSEs from the perspective of parameter estimation. We find that both are better empirically than threshold-based metrics such as classification accuracy. Lastly, the group average calculated by the harmonic mean is more robust than the arithmetic mean when the group sample sizes are unbalanced.

## 5.3 Simulation Results

For Scenario 1, Figure 1(a) displays the estimated AUC for the under-represented group versus the proportion of shared features in the two groups. The AUCs of the under-represented group from the JFM, SFM, and group-ignorant models improved as the proportion of shared features increased. The SFM and group-ignorant models were highly sensitive to the percentage of shared non-zero features as they both estimate a single set of parameters for both groups. In contrast, the JFM showed consistently higher AUC than the other three methods. When the proportion of shared features is high, the JFM estimated higher AUCs and smaller variances than those from the group-separate model. The JFM's overall AUC performance was similar to that of the SFM and the group-ignorant model. When the proportion of shared features is low, the JFM estimated higher AUCs than the SFM and the group-ignorant model, and showed similar AUC to the group-separate model.

Figure 1(b) displays the estimated AUC for the majority group against the proportion of shared features in the two groups. The JFM was robust in achieving comparable AUC to that of the group-separate model. The SFM and group-ignorant models were highly sensitive to the percentage of shared features for the majority group, with lower AUCs when the proportion of shared parameters is low.

Figure 1(c) displays the estimated overall AUCs, and Figure 1(d) displays the group disparity of the AUCs from the four approaches. In particular, although both the SFM and JFM achieved smaller disparities than the group-ignorant model in Figure 1(d), Figure 1(a–b) underscores the limitation of the SFM with no flexibility of group-specific parameters. When the true models of the two groups are different, the SFM often achieved better parity by compromising the performance for the majority group. These figures together demonstrate that the JFM achieves fair prediction performance robustly across the range of possible proportions of shared features between groups, by training the classifiers jointly with a flexible parameterization. Web Figure 5(a) through 5(d) compares the average of prediction TPR and TNR and disparity in TPR and TNR differences of the four methods. The patterns are similar to those observed for AUCs. Coefficient MSEs are presented in Figure 4(a–b) for the minority and majority groups respectively. Under varying proportions of the shared parameters, JFM achieves the smallest MSEs across different methods for the minority group and similar MSEs to the group-separate model for the majority group, demonstrating its improvement for parameter estimation especially for the minority group. Web Figure 14(a) through 14(f) compares the variable selection TPRs and TNRs of the four methods. The patterns are similar to MSEs.

Figure 2 displays the performance of the four methods as a function of the sample size of the under-represented group, with other settings fixed. In Figure 2(a), the AUCs of the under-represented group from all models improved as sample size of that group increased. The JFM showed consistently higher AUCs and smaller variances than those from all the other models. JFM outperforms the other models the most when the minority group's sample size is small, showing the benefits of borrowing information between groups in situations with unbalanced sample sizes. Figure 2(b) illustrates that the AUC of the majority group was not impacted for the JFM and group-separate methods. However, the AUC of the majority group decreased as the sample size of the under-represented group increased for the SFM and the group-ignorant models. This decrease highlights an undesirable aspect of these two methods, namely, they compromise accuracy by estimating a single set of classifier parameters. Figure 2(c) and 2(d) illustrates that the JFM achieves overall satisfactory AUCs and parity between groups across varying sample sizes of the under-represented group. Web Figure 6 compares the average of TPR and TNR and disparity of TPR and TNR of the four methods. In addition, the JFM substantially reduced coefficient MSEs (Figure 4(c)) under all simulated sample sizes for the minority group compared to all competing methods. The MSEs were similar between the JFM and the group separate model for the majority group, and both were lower than the MSEs of the group ignorant and SFM models.

Figure 3 displays the performance of the four methods while varying the number of features from 200 to 2000, and holding the number of associated features constant at 40. It demonstrates that the JFM method in going from low dimensional to high dimensional

settings can maintain overall satisfactory prediction performances and parity between groups. Web Figure 11 displays the performance of the four methods while varying the number of features from 200 to 2000, and setting the number of associated features to a fixed proportion of the total number of features. The resultant patterns are similar to Figure 3. The JFM consistently had the smallest MSEs under all simulated numbers of covariates for both the minority and majority groups. In particular, the JFM showed the largest reduction of MSEs with small number of covariates.

We investigated the empirical computational complexity of the JFM approach as a function of increasing numbers of features and sample sizes in Web Appendix 3. Web Figure 1 shows that the JFM computation time is approximately  $\mathcal{O}(p^{1.5})$  and  $\mathcal{O}(n)$ . This matches our theoretical analysis on complexity since the per-iteration complexity is  $\mathcal{O}(np)$  and the rate of convergence is  $\mathcal{O}(p^{0.5})$ . Details are presented in Web Appendix 6. We also implemented an alternative JFM algorithm using a group lasso similarity penalty (referred to as JFM-Group, Web Figure 15), and compared its performance with JFM-Fusion (presented above). The results showed largely similar performances. The JFM-Fusion showed slightly better predictive performance than the JFM-Group when the two groups share more than 80% of the features, which is in line with previous reports that the fusion penalty term more aggressively enforces similarities across groups (Yuan and Lin, 2006; Obozinski et al., 2010; Danaher et al., 2014).

## 6. COVID-19 Risk Prediction Case Study

We applied the JFM, in comparison with other methods, to predict mortality related to COVID-19 from patients' routine ambulatory encounters and laboratory records prior to COVID-19 infection, with the goal of better stratification of patient risk for clinical management. We used a retrospective EHR dataset of 11,594 patients of age 50+ with laboratory-confirmed COVID-19 at New York University Langone Health (NYULH) from March 2020 to February 2021. Among the 11,594 patients, 1,242 (10.7%) died of COVID-19. The patients were divided into four groups by their age at the time of COVID-19 diagnosis: 50-64, 65-74, 75-84, and 85+ with 5,905 (50.9%), 2,946 (25.4%), 1,814 (15.6%), and 929 (8.0%) patients, respectively. The observed mortality rates were 4.44%, 11.17%, 18.96% and 33.05%, respectively. Candidate features ( $p = 82$ ) included demographic variables, such as age, sex, race/ethnicity, smoking status, body mass index (BMI); common chronic disease history such as diabetes, dementia, chronic kidney diseases (CKD); Myocardial Infarction (MI) & Atrial Fibrillation (AF); and routinely collected laboratory markers, such as lipid panels, blood panels, albumin, creatinine, aspartate aminotransferase (AST) etc. obtained from patients routine ambulatory histories before their COVID-19 infections. To build the prediction models, we randomly split the dataset into training ( $n = 8,115$ , 70%) and testing ( $n = 3,479$ , 30%) sets. We first standardized all features to zero-mean and unit variance. Five-fold cross-validation was conducted on the training set to determine the hyperparameters for each model. Hyperparameters for the group-separate and group-ignorant models were selected to maximize the groupwise AUCs and the overall AUC, respectively, while those for the SFM and JFM were determined to maximize the harmonic mean of groupwise AUCs. Subsequently, we trained the final models with the optimal hyperparameters using the entire training set and applied the final

models to the testing dataset to demonstrate their predictive performance. We repeated the training/testing split 10 times and averaged the performances across the 10 splits. Table 1 presents the AUCs and the averages of TPR and TNR of the four methods for each age group. The JFM performed better across all age groups than the separate model did, demonstrating that joint modeling yields higher efficiency. Compared with the group ignorant model, the JFM performed better in the three older age groups, with comparable AUC for the 50-64 age group, which resulted in smaller disparities in prediction performance overall. This phenomenon supports the observed pattern in simulation studies that the JFM reduced disparities in prediction performances without impacting those from the majority groups. In contrast, the SFM tended to reduce prediction disparities by lowering the performances for the majority groups. Figure 5 presents the boxplots of odds ratios (ORs) of selected demographic and clinical features estimated by the JFM. These results support the hypothesis that some features have common associations between groups, and some have group-specific ORs. For example, the decreasing OR estimates of BMI along age-groups confirmed the prior hypothesis that the association between BMI and COVID-19 mortality is heterogeneous between age-groups. In the JFM estimates, BMI is positively associated with higher risks of COVID-19 mortality for patients younger than 75, but with smaller and even reversed ORs in the oldest age groups. For older adults, higher BMIs are often associated with greater energy stores and a better nutritional state overall, which is beneficial for patients' survival outcomes when infected by COVID-19. The proportion of underweight patients ( $BMI < 18$ ) increased from 0.6% in the age group 50-64 to 5.5% in the age group 85+. The underweight status, often a proxy of frailness, has been repeatedly reported as a strong risk factor of COVID-19-induced multi-organ failure and mortality in older patients (Tehrani et al., 2021). On the other hand, the JFM can improve efficiencies for covariates with rare prevalence in a subgroup. For instance, dementia has been reported as a risk factor with COVID-19 mortality. In the group-separate model, dementia was insignificant in patients aged 50-64, mainly due to its low prevalence in this group (0.6%). In contrast, dementia was significantly associated with mortality in all age groups with similar ORs in the JFM estimates.

## 7. Conclusions and Discussion

In this study we introduced a joint fairness model for jointly estimating sparse parameters, on the basis of observations drawn from distinct but related groups, with the goal of achieving fair performances across groups. We employ an efficient accelerated smoothing proximal gradient algorithm to solve the joint fair objective function, which has convex penalty functions. Our algorithm is computationally tractable for high-dimensional datasets. Further, we presented the asymptotic distributions of  $\hat{\beta}_k$ . Our JFM predictions outperform competing approaches over a range of simulations and in an example application dataset.

We note that the JFM relies on separate hyperparameters ( $K + 2$  hyperparameters) to control sparsity, fairness and similarity. This reliance can be viewed as a strength rather than a drawback because one can vary separately the amount of similarity, sparsity and fairness to enforce in the group-specific estimates. In situations with many groups, further assumptions

can be made to reduce the number of sparsity hyperparameters (i.e.  $\lambda_{Spk} = c_k \lambda_{Sp}$ ). Possible choices of  $c_k$  include  $\frac{1}{\sqrt{m_k}}$  so that sparsity is inversely proportional to the number of samples.

While nearly all existing fairness-aware prediction approaches estimate a single set of classifier parameters, one exception is a recent study that proposes using multi-task learning (MTL) to improve algorithm fairness (Oneto et al., 2019). However, MTL research focuses on joint architecture, optimization, and task relationship learning, which is a different emphasis from the proposed JFM approach to improve risk prediction performance for underrepresented populations.

Theorem 4.1 established consistency of the JFM estimators. Its asymptotic distribution needs to be further investigated to lay the foundation of its inference. Moving forward, the proposed JFM framework can be extended for time-to-event outcomes by using similar constraints to those proposed here. It can also in principle be extended to non-linear models by adding a suitable fairness penalty term to the objective function.

Given the increasing ability to subclassify diseases according to their molecular features and the recognition that substantial heterogeneity exists in many molecular subtypes, most diseases will be eventually classified into a collection of multiple subtypes with unbalanced sample sizes. Therefore, the proposed JFM has wide application potential to improve prediction efficiencies and reduce subgroup prediction disparities beyond applications addressing gender, race/ethnicity and age disparities.

A Python package implementing the JFM is available at <https://github.com/hyungrok-do/joint-fairness-model>.

## Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

## Data Availability Statement

The data that support the findings in this paper are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

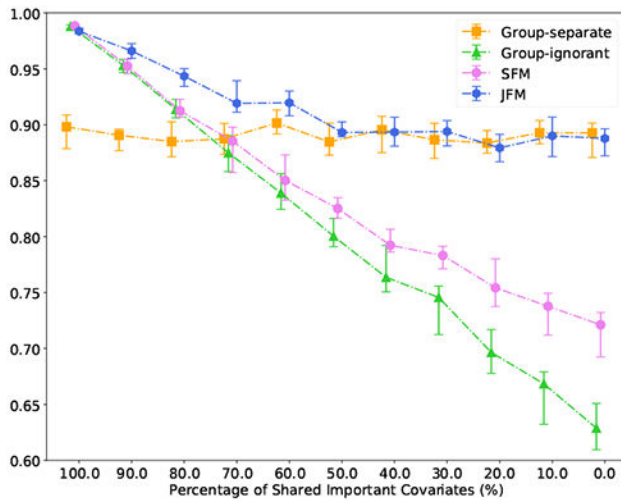
## References

- Angwin J, Larson J, Mattu S, and Kirchner L (2016). Machine bias. ProPublica, May, 23:2016.
- Bechavod Y and Ligett K (2017). Penalizing unfairness in binary classification. arXiv preprint arXiv:1707.00044
- Buolamwini J and Gebru T (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. In Proceedings of the 1st Conference on Fairness, Accountability and Transparency, pages 77–91.
- Calders T, Kamiran F, and Pechenizkiy M (2009). Building classifiers with independency constraints. In 2009 IEEE International Conference on Data Mining Workshops, pages 13–18.
- Char DS, Shah NH, and Magnus D (2018). Implementing machine learning in health care—addressing ethical challenges. The New England journal of medicine, 378(11):981. [PubMed: 29539284]

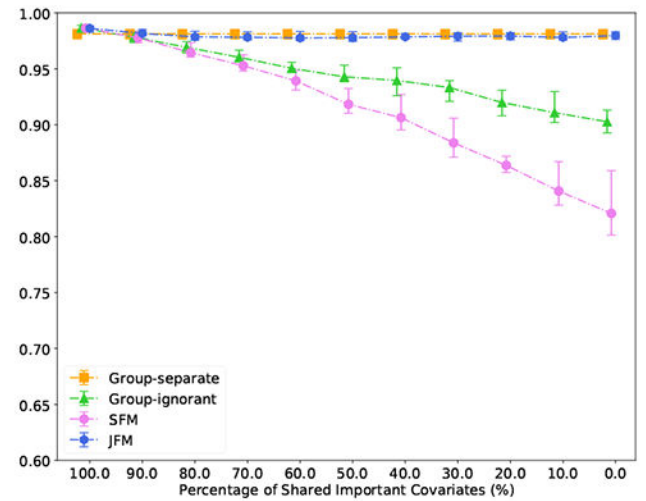
- Chen IY, Pierson E, Rose S, Joshi S, Ferryman K, and Ghassemi M (2021). Ethical machine learning in healthcare. *Review of Biomedical Data Science*, 4(1):123–144.
- Chen X, Lin Q, Kim S, Carbonell JG, Xing EP, et al. (2012). Smoothing proximal gradient method for general structured sparse regression. *The Annals of Applied Statistics*, 6(2):719–752.
- Danaher P, Wang P, and Witten DM (2014). 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. [PubMed: 24817823]
- Dondelinger F, Mukherjee S, and Initiative TADN (2018). The joint lasso: high-dimensional regression for group structured data. *Biostatistics*, 21(2):219–235.
- Dwork C, Immorlica N, Kalai AT, and Leiserson M (2018). Decoupled classifiers for group-fair and efficient machine learning. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, pages 119–133.
- Friedman J, Hastie T, Höfling H, Tibshirani R, et al. (2007a). Pathwise coordinate optimization. *The annals of applied statistics*, 1(2):302–332.
- Friedman J, Hastie T, and Tibshirani R (2007b). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441. [PubMed: 18079126]
- Hardt M, Price E, and Srebro N (2016). Equality of opportunity in supervised learning. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, page 3323–3331, Red Hook, NY, USA. Curran Associates Inc.
- Hoefling H (2010). A path algorithm for the fused lasso signal approximator. *Journal of Computational and Graphical Statistics*, 19(4):984–1006.
- Kamiran F and Calders T (2012). Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1):1–33.
- Kamishima T, Akaho S, Asoh H, and Sakuma J (2012). Fairness-aware classifier with prejudice remover regularizer. pages 35–50.
- Kompaniyets L and Goodman AB (2021). Body mass index and risk for covid-19–related hospitalization, intensive care unit admission, invasive mechanical ventilation, and death. *MMWR Morb Mortal Wkly Rep* 2021, 70:355–361.
- Lahoti P, Weikum G, and Gummadi KP (2018). ifair: Learning individually fair data representations for algorithmic decision making. *CoRR*, abs/1806.01059.
- Larrazabal AJ, Nieto N, Peterson V, Milone DH, and Ferrante E (2020). Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. *Proceedings of the National Academy of Sciences*, 117(23):12592–12594.
- Nesterov Y (2005). Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152.
- Obozinski G, Taskar B, and Jordan MI (2010). Joint covariate selection and joint subspace selection for multiple classification problems. *Statistics and Computing*, 20(2):231–252.
- Oneto L, Doninini M, Elders A, and Pontil M (2019). Taking advantage of multitask learning for fair classification. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, page 227–237.
- Paulus JK and Kent DM (2020). Predictably unequal: understanding and addressing concerns that algorithmic clinical prediction may increase health disparities. *NPJ digital medicine*, 3(1):1–8. [PubMed: 31934645]
- Ryu HJ, Adam H, and Mitchell M (2018). Inclusivefacenet: Improving face attribute detection with race and gender diversity. In *Fairness, Accountability, and Transparency in Machine Learning (FAT/ML)*. Fairness, Accountability, and Transparency in Machine Learning (FAT/ML).
- Seyyed-Kalantari L, Liu G, McDermott M, Chen IY, and Ghassemi M (2020). Chexclusion: Fairness gaps in deep chest x-ray classifiers. In *BIOCOMPUTING 2021: Proceedings of the Pacific Symposium*, pages 232–243. World Scientific.
- Tehrani S, Killander A, Åstrand P, Jakobsson J, and Gille-Johnson P (2021). Risk factors for death in adult covid-19 patients: Frailty predicts fatal outcome in older patients. *International Journal of Infectious Diseases*, 102:415–421. [PubMed: 33130213]
- Tibshirani R (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288.



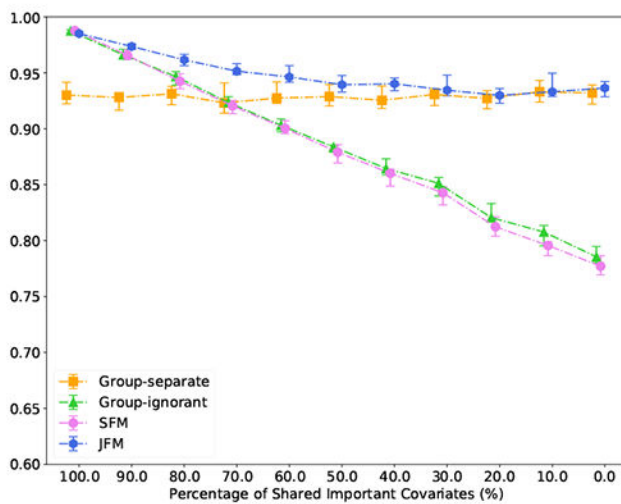
- Yuan M and Lin Y (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67.
- Yuan M and Lin Y (2007). Model selection and estimation in the gaussian graphical model. *Biometrika*, 94(1):19–35.
- Zafar MB, Valera I, Gomez Rodriguez M, and Gummadi KP (2017a). Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment.
- Zafar MB, Valera I, Gomez-Rodriguez M, and Gummadi KP (2019). Fairness constraints: A flexible approach for fair classification. *Journal of Machine Learning Research*, 20(75):1–42.
- Zafar MB, Valera I, Rogriguez MG, and Gummadi KP (2017b). Fairness constraints: Mechanisms for fair classification. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54, pages 962–970.
- Zemel R, Wu Y, Swersky K, Pitassi T, and Dwork C (2013). Learning fair representations. volume 28, pages 325–333.
- Zou H and Hastie T (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 67(2):301–320.



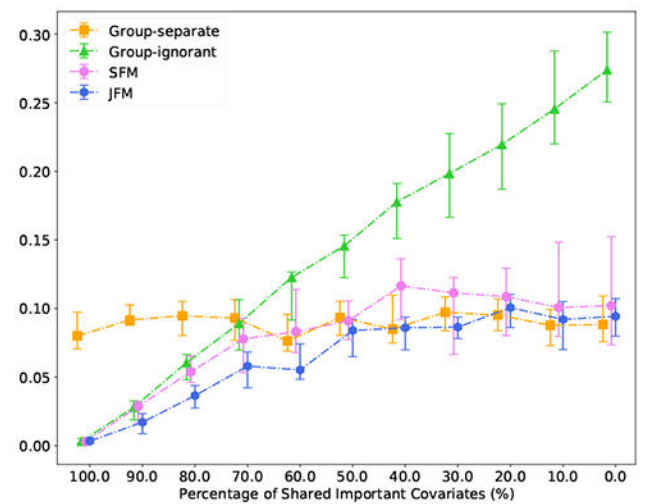
(a) AUC of the Under-represented Group



(b) AUC of the Over-represented Group

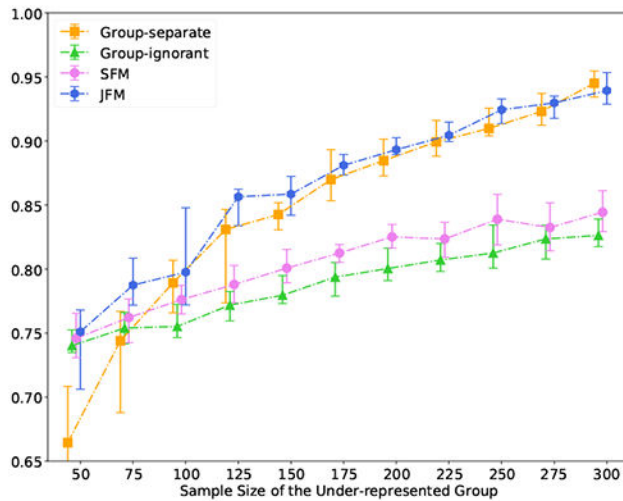


(c) Overall AUC

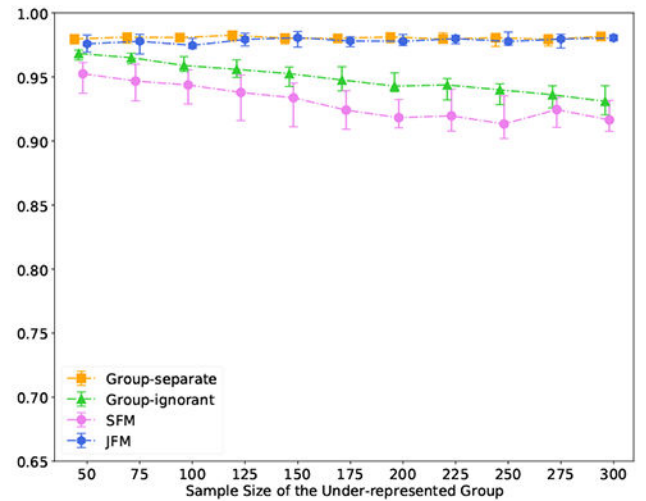


(d) Disparity of AUC

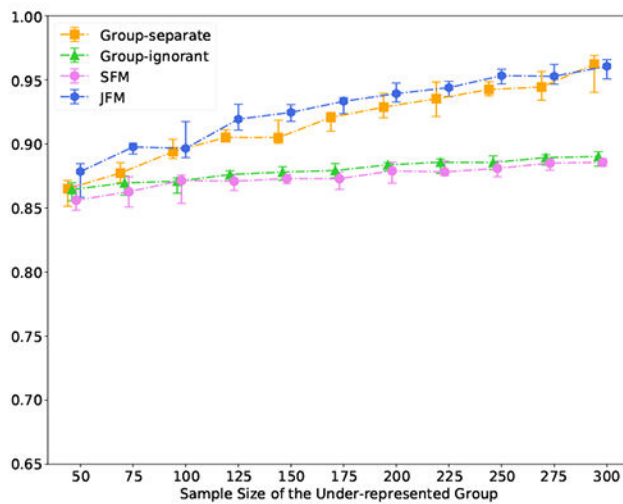
**Figure 1:**  
Experimental Results for Scenario 1



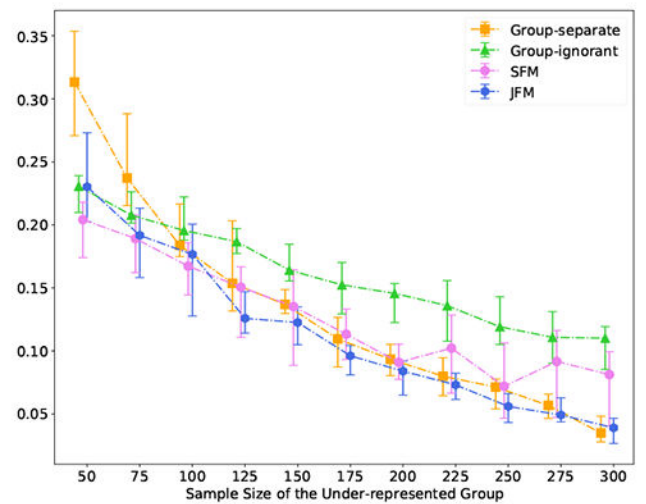
(a) AUC of the Under-represented Group



(b) AUC of the Over-represented Group

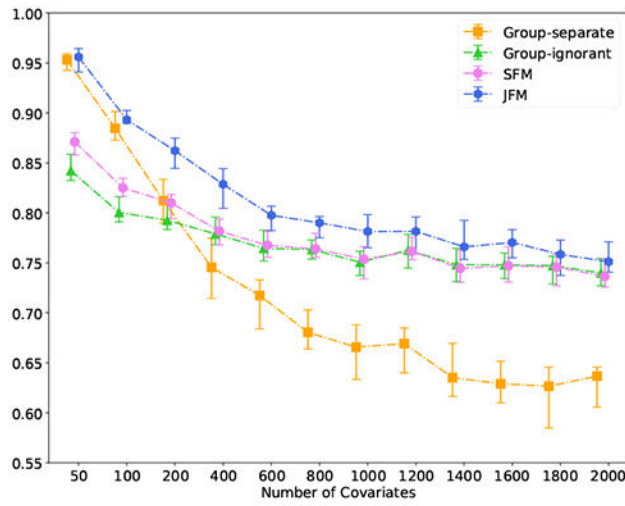


(c) Overall AUC

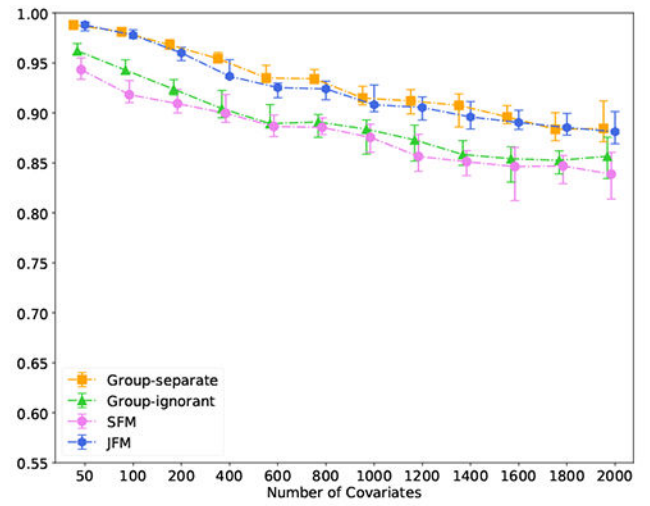


(d) Disparity of AUC

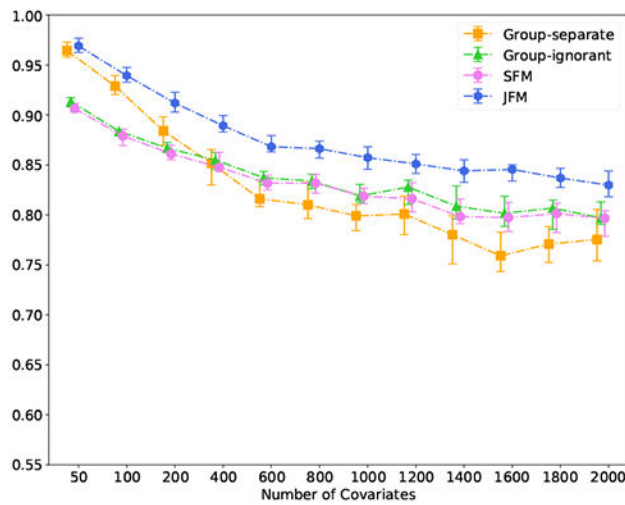
**Figure 2:**  
Experimental Results for Scenario 2



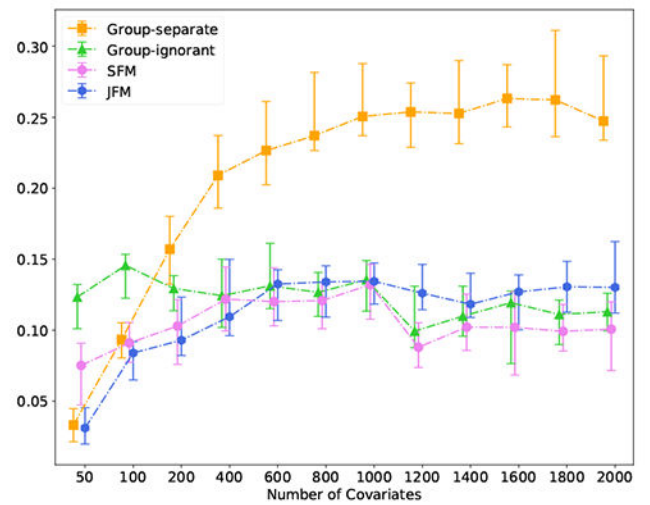
(a) AUC of the Under-represented Group



(b) AUC of the Over-represented Group

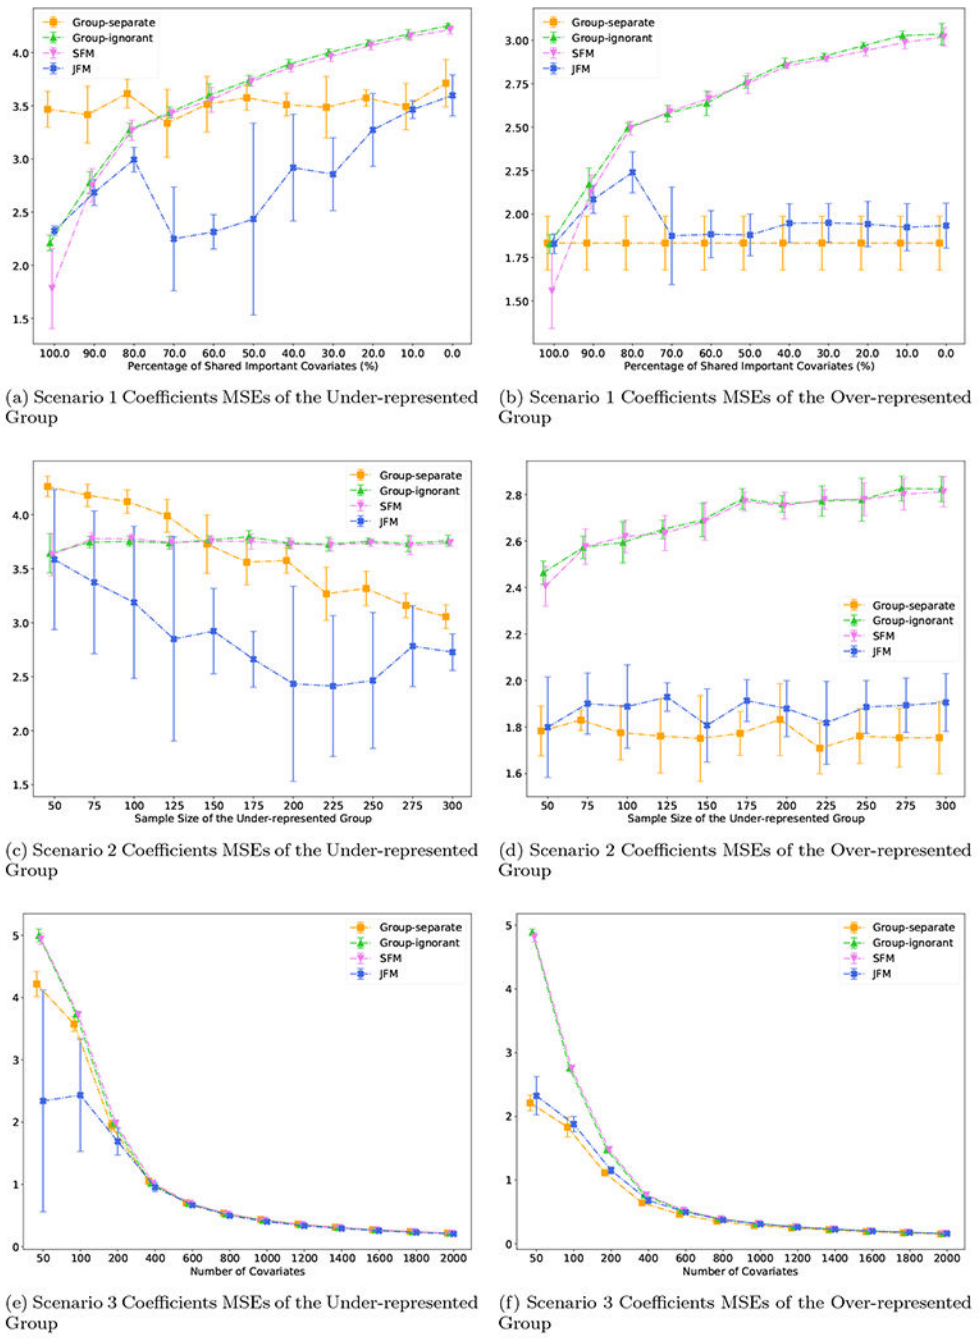


(c) Overall AUC

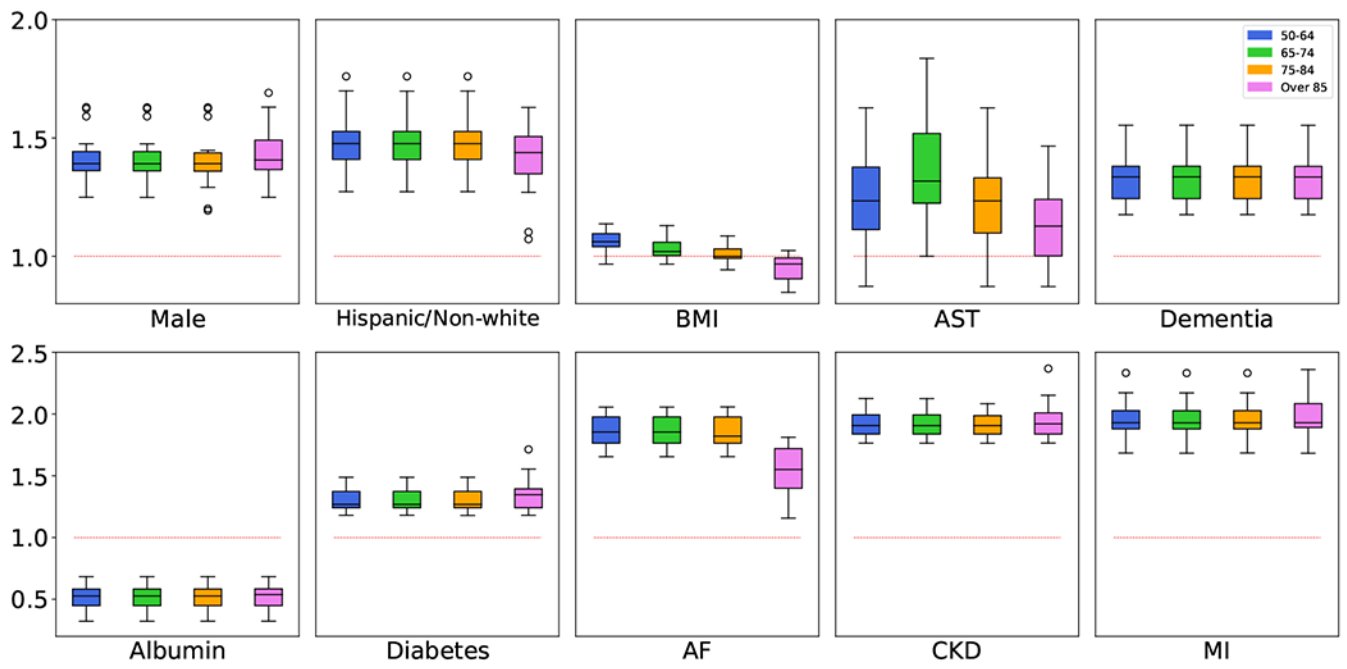


(d) Disparity of AUC

**Figure 3:**  
Experimental Results for Scenario 3



**Figure 4:**  
Coefficients MSEs for Scenario 1 – 3



**Figure 5:**  
Estimated Odds Ratios for COVID-19 Dataset

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**Algorithm 1**

Accelerated Smoothing Proximal Gradient (ASPG) Algorithm for the JFM

- 
- 1: **Input:** Data  $\mathbf{X}^k, \mathbf{y}^k$  for  $k = 1 \dots K$ , hyperparameters  $\lambda_F, \lambda_{\text{Sim}}, \lambda_{\text{Sp}}, \epsilon, \mu$
  - 2: **Output:**  $\hat{\boldsymbol{\beta}} = (\hat{\boldsymbol{\beta}}^1, \dots, \hat{\boldsymbol{\beta}}^K)$  solving the joint fairness objective function (2).
  - 3: **Initialize:**  $\boldsymbol{\beta}^{(0)} = \mathbf{0}, \boldsymbol{\gamma}^{(0)} = \mathbf{0}, s^{(0)} = 1$
  - 4:  $L = \frac{1}{4} \max \{ \lambda_{\max}(\mathbf{X}^k \mathbf{T} \mathbf{X}^k) : k = 1, \dots, K \} + \mu^{-1} \|\mathbf{D}\lambda_F, \lambda_{\text{Sim}}\|_2^2$
  - 5: **for**  $t \geq 1$  **do**
  - 6:  $\boldsymbol{\alpha}^{(t)} = \boldsymbol{\gamma}^{(t-1)} - L^{-1} (-\nabla \ell(\boldsymbol{\gamma}^{(t-1)})) + \nabla f_{\mu}(\boldsymbol{\gamma}^{(t-1)})$
  - 7:  $\boldsymbol{\beta}^{(t)} = \mathcal{S}(\boldsymbol{\alpha}^{(t)}; L^{-1}\lambda_{\text{Sp}})$
  - 8: **if**  $\|\boldsymbol{\beta}^{(t)} - \boldsymbol{\beta}^{(t-1)}\|_2 \leq \epsilon$
  - 9:  $s^{(t)} = \frac{1 + \sqrt{1 + 4s^{(t-1)}}}{2}$
  - 10:  $\boldsymbol{\gamma}^{(t)} = \boldsymbol{\beta}^{(t)} + \left( \frac{s^{(t-1)} - 1}{s^{(t)}} \right) (\boldsymbol{\beta}^{(t)} - \boldsymbol{\beta}^{(t-1)})$
  - 11:  $t \leftarrow t + 1$
  - 12: **end for**
  - 13:  $\hat{\boldsymbol{\beta}} \leftarrow \boldsymbol{\beta}^{(t)}$ .
-

**Table 1:**

Predictive Performance on COVID-19 Case Study

Models	AUCs				Average of TPR and TNR			
	50-64	65-74	75-84	Over 85	50-64	65-74	75-84	Over 85
Group-separate	0.838	0.773	0.709	0.649	0.780	0.722	0.669	0.632
Group-ignorant	<b>0.855</b>	0.786	0.735	0.659	<b>0.803</b>	<b>0.731</b>	0.687	0.639
SFM	0.847	0.774	0.728	0.660	0.791	0.724	0.688	0.640
JFM	0.852	<b>0.791</b>	<b>0.736</b>	<b>0.672</b>	0.794	<b>0.731</b>	<b>0.690</b>	<b>0.659</b>

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript