arXiv:2112.12719v3 [stat.AP] 27 Feb 2023

A general framework for penalized mixed-effects multitask learning with applications on DNA methylation surrogate biomarkers creation

Andrea Cappozzo * Francesca Ieva*† Giovanni Fiorito ‡

Abstract

Recent evidence highlights the usefulness of DNA methylation (DNAm) biomarkers as surrogates for exposure to risk factors for non-communicable diseases in epidemiological studies and randomized trials. DNAm variability has been demonstrated to be tightly related to lifestyle behavior and exposure to environmental risk factors, ultimately providing an unbiased proxy of an individual state of health. At present, the creation of DNAm surrogates relies on univariate penalized regression models, with elastic-net regularizer being the gold standard when accomplishing the task. Nonetheless, more advanced modeling procedures are required in the presence of multivariate outcomes with a structured dependence pattern among the study samples. In this work we propose a general framework for mixed-effects multitask learning in presence of high-dimensional predictors to develop a multivariate DNAm biomarker from a multi-center study. A penalized estimation scheme based on an expectation-maximization algorithm is devised, in which any penalty criteria for fixed-effects models can be conveniently incorporated in the fitting process. We apply the proposed methodology to create novel DNAm surrogate biomarkers for multiple correlated risk factors for cardiovascular diseases and comorbidities. We show that the proposed approach, modeling multiple outcomes together, outperforms state-of-the-art alternatives, both in predictive power and bio-molecular interpretation of the results.

1 Introduction

DNA methylation (DNAm) is an epigenetic process that regulates gene expression, typically occurring in cytosine within CpG sites (CpGs) in the DNA sequence (Singal and Ginder, 1999). DNAm regulates gene expression in different manners. Specifically, high DNAm has been observed in bodies of highly transcribed genes, whereas DNAm in gene promoters and first introns typically have

*MOX Lab, Department of Mathematics, Politecnico di Milano, andrea.cappozzo@polimi.it, francesca.ieva@polimi.it

†Health Data Science Center, Human Technopole

‡Laboratory of Biostatistics, Department of Biomedical Sciences, University of Sassari gfiorito@uniss.it

an inverse correlation with gene expression (Anastasiadi et al., 2018; Rauluseviciute et al., 2020). Also, recent studies suggest that the relationship between genetic variation, DNAm and gene expression is complex and tissue-specific, highlighting that DNAm in non-CpG island regions regulates the transcription of distal genes (van Eijk et al., 2012). Advanced technology allows measuring whole-genome DNAm for many samples at the same time. The most common ways for DNAm measurements consist of whole-genome bisulphite sequencing and DNAm microarray. The first commercial high-density microarray measuring genome-wide methylation was the HumanMethylation27 (27K CpGs) released by Illumina in 2009, followed by the HumanMethylation450 (450K CpGs), and more recently, by the IlluminaMethylation850 (850K CpGs, Campagna et al., 2021). Since then, a tremendous amount of associations between DNAm at individual CpG sites and different exposures, traits, and diseases have been identified in the so-called epigenome-wide association studies (EWAS, Battram et al., 2022). Concurrently, the development of surrogate scores based on blood DNA methylation has also received thriving attention in recent years: impressive epidemiological evidence has been established between DNAm and individual history of exposure to lifestyle and environmental risk factors (Zhong et al., 2016; Guida et al., 2015; Fiorito et al., 2018). To this extent, multi-CpG DNAm biomarkers have been devised to predict patient-specific state of health indicators; and relevant examples include epigenetic clocks to measure "biological age" (Lu et al., 2019), smoking habits (Guida et al., 2015) and proxies for inflammatory proteins (Stevenson et al., 2020). Remarkably, DNAm based scores have been demonstrated to outperform surveyed exposure measurements when predicting diseases (Zhang et al., 2016; Conole et al., 2020). A possible explanation for this somewhat counter-intuitive behavior being that DNA methylation intrinsically accounts for biases in self-reported exposure (e.g., underestimation of smoked cigarettes) as well as individual responses to risk factors (e.g., the same amount of tobacco may produce different effects in dissimilar patients).

From a modeling perspective, state-of-the-art methods for DNAm biomarkers creation generally rely on standard univariate penalized regression models, with elastic-net (Zou and Hastie, 2005) being the routinely employed technique when accomplishing the task. Indeed, the associated learning problem entirely falls within the "p bigger than N" framework: DNA methylation levels are measured at approximately half million CpG sites for each sample, with the dimension of the latter generally not exceeding the order of thousands in most studies. The afore-described procedure is shown to be widely effective in building DNAm biomarkers, with very recent contributions including surrogate scores for short-term risk of cardiovascular events (Cappozzo et al., 2022), cumulative lead exposure (Colicino et al., 2021), DNAm surrogate for alcohol consumption, obesity indexes, and blood measured inflammatory proteins (Hillary and Marioni, 2021) and the identification of CpG sites associated with clinical severity of COVID-19 disease (Castro de Moura et al., 2021). Nonetheless, elastic-net penalties may be too restrictive when dealing with complex learning problems involving multivariate responses and distinctive dependence patterns across statistical units.

The aforesaid first layer of complexity is encountered when a multi-dimensional DNAm biomarker needs to be created, to jointly model multiple risk factors and to coherently account for the correlation structure among the response variables. Such a multivariate problem, also known as multi-task regression in the machine learning literature (Caruana, 1997), can be fruitfully untangled only if dedicated care is devoted in choosing the most appropriate penalty required for the analysis. For instance, one may opt for the incorporation of $\ell 1/\ell 2$ type of regularizers (Obozinski et al., 2009, 2010; Li et al., 2015), that extend the lasso (Tibshirani, 1996), group-lasso (Yuan and Lin, 2006) and sparse group-lasso (Simon et al., 2013; Laria et al., 2019) to the multiple response framework. Another option could contemplate the inclusion, within the estimation procedure, of prior information related to the association structure among CpG sites: this is effectively achieved by means of graph-based penalties (Li and Li, 2010; Kim et al., 2013; Cheng et al., 2014; Dirmeier et al., 2018). Furthermore, tree-based regularization methods have also been recently introduced in the literature, to account for hierarchical structure over the responses in a single study (Kim and Xing, 2012) as well as when multiple data sources are at our disposal (Zhao and Zucknick, 2020; Zhao et al., 2022). For a thorough and up-to-date survey on the analysis of high-dimensional omics data via structured regularization we refer the interested reader to Vinga (2021), while the monograph of Hastie et al. (2015) provides a general introduction to statistical learning with sparsity.

A second layer of complexity is introduced when DNA samples and related blood measured biomarkers are collected in a study comprising multiple cohorts. In such a situation, an unknown degree of heterogeneity may be included in the data, with patients coming from the same cohort sharing some degree of commonality. Observations in the dataset are thus no longer independent and the cohort-wise covariance structure needs to be properly estimated. Linear Mixed-Effects Models (LMM) provide a convenient solution to this problem by adding a random component to the model specification (see, e.g., Pinheiro and Bates, 2006; Gałecki and Burzykowski, 2013; Demidenko, 2013, for an introduction on the topic). Whilst being able to capture unobserved heterogeneity, standard mixed models, very much like their fixed counterpart, cannot directly handle situations in which the number of predictors exceeds the sample size. In order to overcome this issue Schelldorfer et al. (2011) introduced a procedure for estimating high-dimensional LMM via an ℓ_1-penalization . More recently, Rohart et al. (2014) devised a general-purpose ECM algorithm (Meng and Rubin, 1993) for solving the same issue, but achieving greater flexibility as the proposed framework can be combined with any penalty structure previously developed for linear fixed-effects models.

A Multivariate Mixed-Effects Model (MLMM) is an LMM in which multiple characteristics (response variables) are measured for the statistical units comprising the study. Despite being quite a long-established methodology (Reinsel, 1984; Shah et al., 1997), its further development has not received much attention in the recent literature. Relevant exceptions include the computational strategies for handling missing values proposed in Schafer and Yucel (2002), and the estimation theory based on hierarchical likelihood developed in Chipperfield and

Steel (2012). On this account, to the best of our knowledge, a unified approach for penalized MLMM estimation is still missing in the literature and it could thus be a relevant contribution to the statistics and machine learning fields.

Motivated by the problem of creating a DNAm biomarker for hypertension and hyperlipidemia from a multi-center study, we propose in this article a general framework for high-dimensional multitask learning with random-effects. Leveraging from the algorithm introduced in Rohart et al. (2014) for the univariate response case, the estimation mechanism is effectively constructed to accommodate custom penalty types, building upon existing routines developed for regression with fixed-effects only.

The remainder of the paper is structured as follows. Section 2 describes the EPIC Italy dataset, which gave the motivation for the development of the methodology proposed in this manuscript. In Section 3 we introduce the penalized mixed-effects model for multitask learning, covering its formulation, inference and model selection. Section 4 presents a simulation study on synthetic data for three different scenarios. Section 5 outlines the results of the novel method applied to the EPIC Italy data for creating DNAm surrogates for cardiovascular risk factors and comorbidities, comparing it with state-of-the-art alternatives. Section 6 concludes the paper with a discussion and directions for future research. The `R` package `emlmm` implementing the proposed method accompanies the article and it is freely available at `https://github.com/AndreaCappozzo/emlmm`.

2 EPIC Italy data and study design

The considered dataset belongs to the Italian branch of the European Prospective Investigation into Cancer and Nutrition (EPIC) study, one of the largest cohort study in the world, with participants recruited across 10 European countries and followed for almost 15 years (Riboli et al., 2002). For each participant, lifestyle and personal history questionnaires were recorded, together with anthropomorphic measures and blood samples for DNA extraction. The EPIC Italy dataset is comprised of geographical sub-cohorts identified by the center of recruitment; particularly, we will consider the provinces of Ragusa and Varese and the cities of Turin and Naples. The latter center became associated with EPIC in later times through the Progetto ATENA study (Panico et al., 1992). DNAm was measured with the HumanMethylation450 array following standard laboratory procedures (see Fiorito et al., 2022, for a detailed description), while the pre-processing included removing CpG sites and samples with a call rate lower than 95%, BMIQ method for reducing technical variability and bias introduced by type II probes, and ComBat technique for batch effect adjustment (Marabita et al., 2013).

By profiting from the information recorded in the aforementioned sub-cohorts, we aim at creating a multi-dimensional DNAm biomarker for cardiovascular risk factors and comorbidities. To this extent, we consider a multivariate response comprised of $r = 5$ measures, namely Diastolic Blood Pressure (DBP), Systolic Blood Pressure (SBP), High Density Lipoprotein (HDL), Low Density Lipopro-

tein (LDL) and Triglycerides (TG). These characteristics are chosen as they represent the major risk factors for cardiovascular diseases (Wu et al., 2015). In building a DNAm biomarker, the response variables are regressed on DNA methylation values for each CpG site, adjusted for sex and age. A total of $N = 574$ individuals in the $J = 4$ cohorts showcase non-missing values for every response variable: they comprise the sample onto which all subsequent analyses will be performed. To reconstruct the process of DNAm surrogates creation and validation, the EPIC Italy data is randomly split into two sets: 70% ($N_{tr} = 401$) of it is employed for pre-processing and model fitting, while the remaining 30% ($N_{te} = 173$) acts as test set for assessing prediction accuracy. In addition, we will consider samples from the EXPOsOMICS project (Fiorito et al., 2018) as an external validation dataset to assess out of groups predictive performance. In details, EXPOsOMICS is a case-control study on cardiovascular diseases (CVDs) nested in the EPIC Italy cohort composed by 276 volunteers (not overlapping with the main dataset) whose center of recruitment is unknown or different from the $J = 4$ observed in the learning phase.

Coming back to the data analysis pipeline, an epigenome-wide association study (EWAS, Campagna et al., 2021) is performed on the training set as a pre-screening procedure. In details, log-transformed DBP, SBP, HDL, LDL and TG are separately regressed on each available CpG site, adjusting for sex and age. P-values are then collected and arranged in increasing order. We then screen the set of predictors retaining, for each dimension of the multivariate response, the CpG sites whose p-values are smaller than the fifth percentile of the resulting empirical distributions. The final set of covariates for the multitask learning problem is achieved by taking the union of the resulting CpG sites separately preserved for DBP, SBP, HDL, LDL and TG. In so doing, out of the whole initial set of 295614 CpG sites, 62128 DNA methylation features are retained for subsequent modeling. Together with sex and age, this amounts to a total of $p = 62130$ predictors and a 5-dimensional response for a training sample size of $N_{tr} = 401$. Whilst variable screening in ultra-high feature space is itself an ongoing research field (see, e.g., Fan and Lv, 2008; Fan et al., 2009; Zhong et al., 2021, and references therein), we decided to rely on the EWAS technique as it is the standard approach employed in epigenomics (Fazzari and Greally, 2010).

As previously mentioned, the considered training samples belong to four different centers distributed across Italy, with data for 91, 234, 44 and 32 volunteers respectively collected in Turin, Varese, Ragusa and Naples provinces. The boxplots in Figure 1 emphasize the differences in the five response variables by center. To capture the center-wise variability and to maintain generalizability of the devised DNAm biomarker outside the Italy EPIC cohorts, a partial pooling random-intercept model shall be adopted. That is, a $q = 1$ random-effect component is included in the model specification. Furthermore, the biomarkers comprising the response vector showcase some degree of relations, as displayed by the sample correlation matrix of Figure 2, so much so that it is sensible to regress them jointly to take advantage of their association structure in the model formulation. This challenging learning task requires an ad-hoc specification for a multivariate mixed-effects framework applicable to high-dimensional

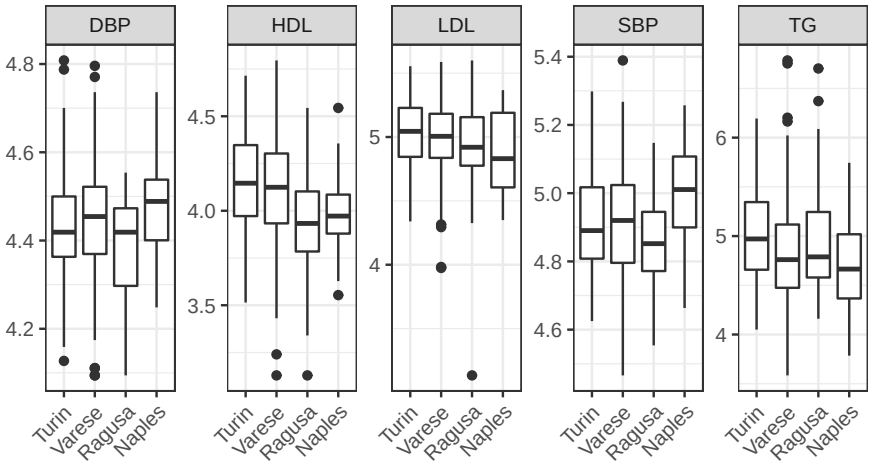

Figure 1: Boxplots of log-transformed Diastolic Blood Pressure (DBP), High Density Lipoprotein (HDL), Low Density Lipoprotein (LDL), Systolic Blood Pressure (SBP) and Triglycerides (TG) for different Center, Italy EPIC training dataset.

predictors.

3 Penalized mixed-effects model for multitask learning

In this section, a novel approach for multivariate mixed-effects modeling based on penalized estimation is proposed.

3.1 Model definition

The multivariate linear mixed-effects model (Shah et al., 1997) expresses the $n_j \times r$ response matrix $\boldsymbol{Y}_j$ for the j-th group as:

$$\boldsymbol{Y}_j = \boldsymbol{X}_j \boldsymbol{B} + \boldsymbol{Z}_j \boldsymbol{\Lambda}_j + \boldsymbol{E}_j, \tag{1}$$

where, for each of the n_j samples in group j and $\sum_{j=1}^{J} n_j = N$, r response variables have been measured. The remainder terms define the following quantities:

- $\boldsymbol{B}$ is the $p \times r$ matrix of fixed-effects (including the intercept)
- $\boldsymbol{\Lambda}_j$ is the $q \times r$ matrix of random-effects
- $\boldsymbol{X}_j$ is the $n_j \times p$ fixed-effects design matrix

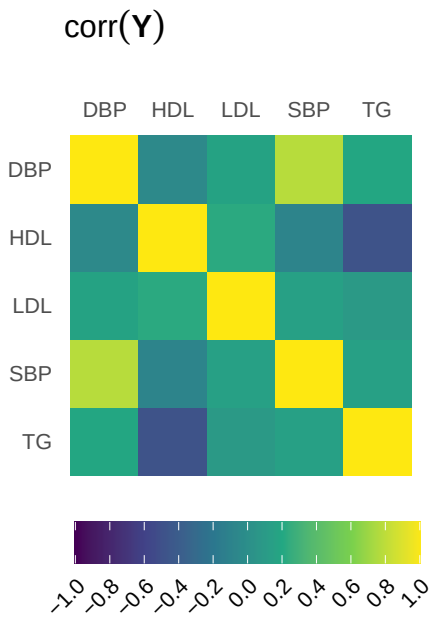

Figure 2: Sample correlation matrix of log-transformed Diastolic Blood Pressure (DBP), High Density Lipoprotein (HDL), Low Density Lipoprotein (LDL), Systolic Blood Pressure (SBP) and Triglycerides (TG), Italy EPIC training dataset.

- $\boldsymbol{Z}_j$ is the $n_j \times q$ random-effects design matrix
- $\boldsymbol{E}_j$ is the $n_j \times r$ within-group error matrix
- $j = 1, \ldots, J$, with J total number of groups.

By employing the vec operator, we assume that:

$$\text{vec}\,(\boldsymbol{\Lambda}_j) \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Psi}),$$

where $\boldsymbol{\Psi}$ is a $qr \times qr$ positive semidefinite matrix, incorporating variations and covariations between the r responses and the q random-effects. We further assume that the error term is distributed as follows:

$$\text{vec}\,(\boldsymbol{E}_j) \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma} \otimes \boldsymbol{I}_{n_j}), \tag{2}$$

where $\boldsymbol{\Sigma}$ is a $r \times r$ covariance matrix, capturing dependence among responses, and $\boldsymbol{I}_{n_j}$ is the identity matrix of dimension $n_j \times n_j$. Formulation in (2) explicitly induces independence between the row vectors of $\boldsymbol{E}_j$. Therefore, the entire model can be rewritten in vec form:

$$\text{vec}\,(\boldsymbol{Y}_j) \sim N\left((\boldsymbol{I}_r \otimes \boldsymbol{X}_j)\,\text{vec}\,(\boldsymbol{B}), (\boldsymbol{I}_r \otimes \boldsymbol{Z}_j)\,\boldsymbol{\Psi}\,(\boldsymbol{I}_r \otimes \boldsymbol{Z}_j)^{'} + \boldsymbol{\Sigma} \otimes \boldsymbol{I}_{n_j}\right).$$

Given a sample of $N = \sum_{j=1}^{J} n_j$, the log-likelihood of model (1) reads:

$$\begin{aligned}\ell(\boldsymbol{\theta}) = \sum_{j=1}^{J} -\frac{n_j}{2}\log 2\pi - \frac{1}{2}\log|\,(\boldsymbol{I}_r \otimes \boldsymbol{Z}_j)\,\boldsymbol{\Psi}\,(\boldsymbol{I}_r \otimes \boldsymbol{Z}_j)^{'} + \boldsymbol{\Sigma} \otimes \boldsymbol{I}_{n_j}| + \\ -\frac{1}{2}\left(\text{vec}\,(\boldsymbol{Y}_j) - (\boldsymbol{I}_r \otimes \boldsymbol{X}_j)\,\text{vec}\,(\boldsymbol{B})\right)^{'}\left((\boldsymbol{I}_r \otimes \boldsymbol{Z}_j)\,\boldsymbol{\Psi}\,(\boldsymbol{I}_r \otimes \boldsymbol{Z}_j)^{'} + \boldsymbol{\Sigma} \otimes \boldsymbol{I}_{n_j}\right)^{-1}\left(\text{vec}\,(\boldsymbol{Y}_j) - (\boldsymbol{I}_r \otimes \boldsymbol{X}_j)\,\text{vec}\,(\boldsymbol{B})\right),\end{aligned} \tag{3}$$

where $\boldsymbol{\theta} = \{\boldsymbol{B}, \boldsymbol{\Sigma}, \boldsymbol{\Psi}\}$ is the set of parameters to be estimated. When the framework outlined in (1) is employed for DNAm biomarker creation, the number of regressors p is most certainly much larger than the sample size N. We are thus not directly interested in maximizing (3), but rather a penalized version of it, generically defined as follows:

$$\ell_{pen}(\boldsymbol{\theta}) = \ell(\boldsymbol{\theta}) - p(\boldsymbol{B}; \lambda), \tag{4}$$

with $p(\boldsymbol{B}; \lambda)$ being a penalty term employed to regularize the fixed-effects $\boldsymbol{B}$ as a function of the complexity parameter $\lambda \geq 0$. Notice that, depending on the chosen penalty, more than one complexity parameter could be involved in the definition of $p(\boldsymbol{B}; \lambda)$ (see Section 3.3 for further details).

A general-purpose algorithm for maximizing (4) can be devised, as described in the next subsection.

3.2 Model estimation

Direct maximization of (4) is unfeasible, as the terms $\text{vec}\,(\boldsymbol{\Lambda}_j)$, $j = 1, \ldots, J$, are unknown. We therefore devise an EM algorithm (Dempster et al., 1977) in which the E-step computes the conditional expectations for the unobserved quantities, while a *complete penalized log-likelihood* is maximized in the M-step.

3.2.1 E-step

The E-step requires the computation of $\mathbb{E}(\text{vec}\,(\boldsymbol{\Lambda}_j)\,|\boldsymbol{Y}_j;\boldsymbol{\theta})$ and $\mathbb{E}(\text{vec}\,(\boldsymbol{\Lambda}_j)\,\text{vec}\,(\boldsymbol{\Lambda}_j)^{'}\,|\boldsymbol{Y}_j;\boldsymbol{\theta})$. This is achieved by noticing that the conditional density $p(\text{vec}\,(\boldsymbol{\Lambda}_j)\,|\boldsymbol{Y}_j;\boldsymbol{\theta})$ is Normal. Updating formulae for the quantities of interest are thus derived as follows:

$$\hat{\boldsymbol{\Gamma}}_j = \mathbb{V}(\text{vec}\,(\boldsymbol{\Lambda}_i)\,|\boldsymbol{Y}_j;\boldsymbol{\theta}) = \left[(\boldsymbol{I}_r \otimes \boldsymbol{Z}_j)^{'}\left(\boldsymbol{\Sigma}\otimes\boldsymbol{I}_{n_j}\right)^{-1}(\boldsymbol{I}_r\otimes\boldsymbol{Z}_j) + \boldsymbol{\Psi}^{-1}\right]^{-1}, \quad (5)$$

$$\widehat{\text{vec}\,(\boldsymbol{\Lambda}_j)} = \mathbb{E}(\text{vec}\,(\boldsymbol{\Lambda}_j)\,|\boldsymbol{Y}_j;\boldsymbol{\theta}) = \hat{\boldsymbol{\Gamma}}_j\,(\boldsymbol{I}_r\otimes\boldsymbol{Z}_j)^{'}\left(\boldsymbol{\Sigma}\otimes\boldsymbol{I}_{n_j}\right)^{-1}(\text{vec}\,(\boldsymbol{Y}_j) - (\boldsymbol{I}_r\otimes\boldsymbol{X}_j)\,\text{vec}\,(\boldsymbol{B}))\,. \quad (6)$$

Consequently, the second moment $\hat{\boldsymbol{R}}_j = \mathbb{E}(\text{vec}\,(\boldsymbol{\Lambda}_j)\,\text{vec}\,(\boldsymbol{\Lambda}_j)^{'}\,|\boldsymbol{Y}_j;\boldsymbol{\theta})$ reads:

$$\hat{\boldsymbol{R}}_j = \hat{\boldsymbol{\Gamma}}_j + \widehat{\text{vec}\,(\boldsymbol{\Lambda}_j)}\widehat{\text{vec}\,(\boldsymbol{\Lambda}_j)}^{'}. \quad (7)$$

At the t-th iteration of the EM algorithm, the E-step requires the computation of (5)-(7) conditioning on the parameter values estimated at iteration $t-1$. Notice that we can directly define the conditional density of $\boldsymbol{Y}_j|\boldsymbol{\Lambda}_j$ by means of the matrix normal distribution

$$\boldsymbol{Y}_j|\boldsymbol{\Lambda}_j \sim m\mathcal{N}\left(\boldsymbol{X}_j\boldsymbol{B} + \boldsymbol{Z}_j\boldsymbol{\Lambda}_j, \boldsymbol{I}_{n_j}, \boldsymbol{\Sigma}\right), \quad (8)$$

where $\boldsymbol{X}_j\boldsymbol{B} + \boldsymbol{Z}_j\boldsymbol{\Lambda}_j$ is the $n_j \times r$ mean matrix, and $\boldsymbol{I}_{n_j}$, $\boldsymbol{\Sigma}$ respectively identify the row and column covariance matrices (Dawid, 1981). Such a representation will be useful in specifying the update for $\boldsymbol{B}$ in the devised M-step: details are provided in the next subsection.

3.2.2 M-step

In the M-step we maximize the *complete penalized log-likelihood*:

$$\ell_{C\,pen}(\boldsymbol{\theta}) = \sum_{j=1}^{J}\log(p(\text{vec}\,(\boldsymbol{Y}_j)\,|\,\text{vec}\,(\boldsymbol{\Lambda}_j)\,;\boldsymbol{B},\boldsymbol{\Sigma})) + \log(p(\text{vec}\,(\boldsymbol{\Lambda}_j)\,;\boldsymbol{\Psi})) - p(\boldsymbol{B};\lambda) =$$

$$= \sum_{j=1}^{J} -\frac{n_j}{2}\log(2\pi) - \frac{1}{2}\log|\boldsymbol{\Sigma}\otimes\boldsymbol{I}_{n_j}| - \frac{1}{2}\mathbb{E}(\boldsymbol{e}_j^{'}\left(\boldsymbol{\Sigma}\otimes\boldsymbol{I}_{n_j}\right)^{-1}\boldsymbol{e}_j|\boldsymbol{Y}_j,\boldsymbol{\theta}) +$$

$$-\frac{n_j}{2}\log(2\pi) - \frac{1}{2}\log|\boldsymbol{\Psi}| - \frac{1}{2}\mathbb{E}\left(\text{vec}\,(\boldsymbol{\Lambda}_j)^{'}\boldsymbol{\Psi}^{-1}\,\text{vec}\,(\boldsymbol{\Lambda}_j)\,|\boldsymbol{Y}_j,\boldsymbol{\theta}\right) - p(\boldsymbol{B};\lambda), \quad (9)$$

where $\boldsymbol{e}_j = \text{vec}\,(\boldsymbol{Y}_j) - (\boldsymbol{I}_r \otimes \boldsymbol{X}_j)\,\text{vec}\,(\boldsymbol{B}) - (\boldsymbol{I}_r \otimes \boldsymbol{Z}_j)\,\text{vec}\,(\boldsymbol{\Lambda}_j)$ and the maximization is performed with respect to $\boldsymbol{\theta} = \{\boldsymbol{B}, \boldsymbol{\Sigma}, \boldsymbol{\Psi}\}$.

The updating formula for $\boldsymbol{B}$ clearly depends on the considered $p(\boldsymbol{B}; \lambda)$ penalty. All the same, it is convenient to work with the matrix-variate representation defined in (8). In so doing, the objective function to be maximized wrt $\boldsymbol{B}$ reads:

$$Q_{\boldsymbol{B}}(\boldsymbol{B}) = -\frac{1}{2}\sum_{j=1}^{J} \text{tr}\left(\boldsymbol{\Sigma}^{-1}\left(\tilde{\boldsymbol{Y}}_j - \boldsymbol{X}_j\boldsymbol{B}\right)'\left(\tilde{\boldsymbol{Y}}_j - \boldsymbol{X}_j\boldsymbol{B}\right)\right) - p(\boldsymbol{B}; \lambda), \qquad (10)$$

where $\tilde{\boldsymbol{Y}}_j = \boldsymbol{Y}_j - \boldsymbol{Z}_j\hat{\boldsymbol{\Lambda}}_j$. $\hat{\boldsymbol{\Lambda}}_j$ is recovered by applying the inverse of the vectorization operator to $\widehat{\text{vec}\,(\boldsymbol{\Lambda}_j)}$, previously computed in the E-step. Simply put, the $\widehat{\text{vec}\,(\boldsymbol{\Lambda}_j)}$ vector of length qr is rearranged in a $q \times r$ matrix, obtaining $\hat{\boldsymbol{\Lambda}}_j$. Start by noticing that, when no penalty is considered, maximization of (10) agrees with the generalized least squares (GLS) estimator assuming $\boldsymbol{\Sigma}$ and $\boldsymbol{\Psi}$ known (Shah et al., 1997). By exploiting properties of the trace operator, we can rewrite (10) defining the following minimization problem:

$$\text{minimize}_{\boldsymbol{B}\in\mathbb{R}^{p\times r}}\ \frac{1}{2}\sum_{j=1}^{J}\left|\left|\boldsymbol{\Sigma}^{-1/2}\left(\tilde{\boldsymbol{Y}}_j - \boldsymbol{X}_j\boldsymbol{B}\right)'\right|\right|_F^2 + p(\boldsymbol{B}; \lambda), \qquad (11)$$

where $||\cdot||_F^2$ denotes the squared Frobenius norm and $\boldsymbol{\Sigma}^{-1/2}$ is the symmetric positive definite square root of $\boldsymbol{\Sigma}^{-1}$, such that $\boldsymbol{\Sigma}^{-1} = \boldsymbol{\Sigma}^{-1/2}\boldsymbol{\Sigma}^{-1/2}$. The representation in (11) allows to employ standard routines for multivariate penalized fixed-effects models for estimating $\boldsymbol{B}$. In details, for solving (11) a two-step updating scheme is devised. Firstly, we compute

$$\tilde{\boldsymbol{B}} = \arg\min_{\boldsymbol{B}}\ \frac{1}{2}\sum_{j=1}^{J}\left|\left|\boldsymbol{\Sigma}^{-1/2}\tilde{\boldsymbol{Y}}_j - \boldsymbol{X}_j\boldsymbol{B}\right|\right|_F^2 + p(\boldsymbol{B}; \lambda), \qquad (12)$$

that is a fixed-effects penalized regression problem in which the response variable is $\boldsymbol{\Sigma}^{-1/2}\tilde{\boldsymbol{Y}}_j$, $j = 1, \ldots, J$; $\tilde{\boldsymbol{B}}$ is thus easily retrieved via fixed-effects routines for penalized estimation. Secondly, the solution to (11) is obtained post multiplying $\tilde{\boldsymbol{B}}$ by $\boldsymbol{\Sigma}^{1/2}$. Therefore, at each iteration of the EM-algorithm, we firstly compute $\tilde{\boldsymbol{B}}$ and then we set:

$$\hat{\boldsymbol{B}} = \tilde{\boldsymbol{B}}\boldsymbol{\Sigma}^{1/2}, \qquad (13)$$

where $\hat{\boldsymbol{B}}$ maximizes (10). This procedure stems from the rationale outlined in Rohart et al. (2014), where, contrarily to their original solution, in our context the updating steps are made more complex by the multidimensional nature of $\boldsymbol{Y}$. The devised updating scheme allows to easily incorporate any $p(\boldsymbol{B}; \lambda)$ that has been previously defined for the fixed-effects framework, and whose estimating routines are available. A list of possible penalties is proposed in Section 3.3.

Updating formulae for the covariance matrices $\boldsymbol{\Psi}$ and $\boldsymbol{\Sigma}$ agree with those of the unpenalized framework, namely

$$\hat{\boldsymbol{\Psi}} = \frac{1}{J}\sum_{j=1}^{J}\hat{\boldsymbol{R}}_j, \tag{14}$$

and for the (h,k)-th element of matrix $\boldsymbol{\Sigma}$

$$\hat{\boldsymbol{\Sigma}}_{(h,k)} = \frac{1}{N}\sum_{j=1}^{J}\left[\mathbb{E}\left(\boldsymbol{E}_{jh}|\boldsymbol{Y}_j\right)'\mathbb{E}\left(\boldsymbol{E}_{jk}|\boldsymbol{Y}_j\right)\right]+\mathrm{tr}\left[\mathrm{cov}(\boldsymbol{E}_{jh},\boldsymbol{E}_{jk}|\boldsymbol{Y}_j)\right], \quad h,k=1,\ldots,r, \tag{15}$$

where $\boldsymbol{E}_{jh}$ denotes the h-th column of matrix $\boldsymbol{E}_j = \boldsymbol{Y}_j - \boldsymbol{Z}_j\hat{\boldsymbol{\Lambda}}_j - \boldsymbol{X}_j\boldsymbol{B}$, $h = 1,\ldots,r$.

3.3 Definition of $p(\boldsymbol{B};\lambda)$

The EM algorithm devised in the previous section defines a general-purpose optimization strategy for penalized mixed-effects multitask learning. While any penalty type can in principle be defined, three notable examples, commonly used in this context, are the elastic net penalty (Zou and Hastie, 2005), the multivariate group-lasso penalty (Obozinski et al., 2011) and the netReg routines for Network-regularized linear models (Dirmeier et al., 2018). Each of them is briefly described in the next subsections.

3.3.1 Elastic-net penalty

The first penalty type we consider is the renowned convex combination of lasso and ridge regularizers, whose magnitude of the former over the latter is controlled by the mixing parameter α, $0 \leq \alpha \leq 1$. In details, the penalty expression reads:

$$p(\boldsymbol{B};\lambda,\alpha) = \lambda\left[(1-\alpha)\sum_{c=1}^{r}\sum_{l=2}^{p}b_{lc}^2 + \alpha\sum_{c=1}^{r}\sum_{l=2}^{p}|b_{lc}|\right], \tag{16}$$

where b_{lc} denotes the element in the l-th row and c-th column of matrix $\boldsymbol{B}$. Notice that the first row of $\boldsymbol{B}$ contains the r intercepts and it is thus not penalized. Algorithmically, penalty (16) can be enforced employing standard and widely available routines for univariate penalized estimation, like the `glmnet` software (Tay et al., 2021). The only computational detail that shall be examined is how to prevent the default shrinkage of the r intercepts: the `penalty.factor` argument of the `glmnet` function effectively serves the purpose. The latter can also be employed in our framework to force coefficients that need not be penalized to enter the model specification.

3.3.2 Multivariate group-lasso penalty

This type of penalty imposes a group structure on the coefficients, forcing the same subset of predictors to be preserved across all r components of the response

matrix. This feature is particularly desirable when building multivariate DNAm biomarkers, since it automatically identifies the CpG sites that are *jointly* related to the considered risk factors. Such a penalty is defined as follows:

$$p(\boldsymbol{B};\lambda,\alpha) = \lambda \left[(1-\alpha) \sum_{c=1}^{r} \sum_{l=2}^{p} b_{lc}^2 + \alpha \sum_{l=2}^{p} ||\boldsymbol{b}_{l.}||_2 \right], \tag{17}$$

where $\boldsymbol{b}_{l.}$ identifies the l-th row of the matrix $\boldsymbol{B}$, such that each $\boldsymbol{b}_{l.}$, $l = 2, \ldots, p$ is an r-dimensional vector. Likewise Section 3.3.1, summations over rows in (17) start at 2 since we do not penalize the vector of intercepts. This penalty behaves like the lasso, but on the whole group of predictors for each of the r variables: they are either all zero, or else none are zero, but are shrunk by an amount depending on λ. Similarly to (16), the mixing parameter α controls the weight associated to ridge and group-lasso regularizers. The `glmnet` software, with `family = "mgaussian"` is again at our disposal for efficiently incorporating (17) in the framework outlined in the present paper.

3.3.3 Network-Regularized penalty

The last penalty we consider allows for the inclusion of biological graph-prior knowledge in the estimation by accounting for the contribution of two non-negative adjacency matrices $\boldsymbol{G}_X \in \mathbb{R}_+^{(p-1)\times(p-1)}$ and $\boldsymbol{G}_Y \in \mathbb{R}_+^{r\times r}$, respectively related to $\boldsymbol{X}$ and $\boldsymbol{Y}$. In this case, $p(\boldsymbol{B};\lambda)$ assumes the following functional form:

$$p(\boldsymbol{B};\lambda,\lambda_X,\lambda_Y) = \lambda||\boldsymbol{B}_0||_1 + \lambda_X \operatorname{tr}\left(\boldsymbol{B}_0^{'}(\boldsymbol{D}_{G_X} - \boldsymbol{G}_X)\boldsymbol{B}_0\right) + \lambda_Y \operatorname{tr}\left(\boldsymbol{B}_0(\boldsymbol{D}_{G_Y} - \boldsymbol{G}_Y)\boldsymbol{B}_0^{'}\right), \tag{18}$$

where $\boldsymbol{B}_0$ is the $(p-1)\times r$ matrix of coefficients without the intercepts and $\boldsymbol{D}_{G_X}$, $\boldsymbol{D}_{G_Y}$ indicate the degree matrices of $\boldsymbol{G}_X$ and $\boldsymbol{G}_Y$, respectively (Chung and Graham, 1997). $\boldsymbol{G}_X$ and $\boldsymbol{G}_Y$ encode a biological similarity, forcing rows and columns of $\boldsymbol{B}_0$ to be similar. The `netReg` `R` package (Dirmeier et al., 2018) provides a convenient implementation of (18).

3.3.4 On the choice of $p(\boldsymbol{B};\lambda)$

Leaving the flexibility attained by the methodology proposed in Section 3.1 aside; in practice, a functional form for $p(\boldsymbol{B};\lambda)$ must be chosen when performing the analysis. We hereafter highlight pros and cons of the proposed approaches with respect to a mixed-effects multitask learning setting.

The elastic-net penalty in (16) does not take into account the multivariate nature of the problem in (4), as the shrinkage is applied directly to $\text{vec}(\boldsymbol{B})$. This behavior allows for capturing a wide variety of sparsity patterns that may be present in $\boldsymbol{B}$, but does not impose any specific structure that could be desirable in a multivariate context. Differently, the multivariate group-lasso of Section 3.3.2 defines a shrinkage term that forces the same subset of predictors to be preserved across all r components of the response $\boldsymbol{Y}$. This can be seen as the

generalization of the variable selection problem to the multivariate response setting, which is also known as *support union problem* or *row selection problem* in the literature (Obozinski et al., 2011). Lastly, the network-Regularized penalty in 3.3.3 is particularly useful when the interaction among features and/or responses is, at least partially, known, such that it can be profited from within the learning mechanism (Cheng et al., 2014).

In relation to the DNAm surrogate creation task motivating our methodological proposal, the multivariate group-lasso is definitely the most appropriate penalty, as it not only showcases better prediction performances but it is also supported by biological reasons: a thorough analysis for the EPIC dataset is reported in Section 5.

3.4 Further aspects

Hereafter, we discuss some practical considerations related to the presented methodology.

- **Initialization:** we start the algorithm with an M-step, setting $\hat{\theta}^{(0)} = \{\hat{\boldsymbol{B}}^{(0)}, \hat{\boldsymbol{\Sigma}}^{(0)}, \hat{\boldsymbol{\Psi}}^{(0)}\}$. In details, both $\hat{\boldsymbol{\Sigma}}^{(0)}$ and $\hat{\boldsymbol{\Psi}}^{(0)}$ are initialized with identity matrices of dimension $r \times r$ and $qr \times qr$ respectively, while $\hat{\boldsymbol{B}}^{(0)}$ is estimated from a penalized linear model (without the random-effects) employing the chosen penalty function with the associated hyper-parameters.

- **Convergence:** the EM algorithm is considered to have converged once the relative difference in the objective function for two subsequent iterations is smaller than ε, for a given $\varepsilon > 0$:

$$\frac{|\ell_{pen}(\hat{\boldsymbol{\theta}}^{(t+1)}) - \ell_{pen}(\hat{\boldsymbol{\theta}}^{(t)})|}{|\ell_{pen}(\hat{\boldsymbol{\theta}}^{(t)})|} < \varepsilon,$$

  where $\hat{\boldsymbol{\theta}}^{(t)} = \{\hat{\boldsymbol{B}}^{(t)}, \hat{\boldsymbol{\Sigma}}^{(t)}, \hat{\boldsymbol{\Psi}}^{(t)}\}$ is the set of estimated values at the end of the t-th iteration. In our analyses, ε is set equal to 10^{-6}. The procedure described in Section 3.2 falls within the class of Expectation Conditional Maximization (ECM) algorithms, whose convergence properties have been proved in Meng and Rubin (1993) and in Section 5.2.3 of McLachlan and Krishnan (2008).

- **Model selection:** a standard 10-fold cross validation (CV) strategy is employed for selecting the tuning factors. Alternatively, as suggested in Rohart et al. (2014), one could employ a modified version of the Bayesian Information Criterion (BIC, Schwarz, 1978):

$$BIC = 2\ell(\hat{\boldsymbol{\theta}}) - d_0 \log(N), \tag{19}$$

  where $\ell(\hat{\boldsymbol{\theta}})$ is the log-likelihood evaluated at $\hat{\boldsymbol{\theta}}$, obtained maximizing (4), and d_0 is the number of non-zero parameters resulting from the penalized estimation. Another option would be to rely on an interval search

algorithm, like the efficient parameter selection via global optimization (Frohlich and Zell, 2005): an implementation is available in the `c060 R` package (Sill et al., 2014).

- **Scalability:** the devised methodology provides a framework for incorporating any penalty in a high-dimensional mixed-effects multitask learning framework. To this extent, the data dimensionality our procedure can cope with, as well as the overall computing time, very much depends on the scalability and efficiency associated to the chosen shrinkage term. Typically nevertheless, penalized likelihood approaches fail to be directly applied to ultrahigh-dimensional problems (Fan et al., 2009), and pre-processing procedures such as variable screening are thus required prior to modeling. The epigenetic application that motivated our work naturally called for an EWAS pre-screening strategy (see Section 2), but clearly other dimensionality reduction techniques could be considered when dealing with massive datasets. The interested reader is referred to Jordan (2013) for a thought-provoking investigation on the topic.

- **Implementation:** routines for fitting the penalized mixed-effects multitask learning method have been implemented in `R` (R Core Team, 2022), and the source code is freely available in the Supplementary Material and at `https://github.com/AndreaCappozzo/emlmm` in the form of an `R` package. The three penalties described in Section 3.3 are included in the software, and can be selected via the `penalty_type` argument of the `ecm_mlmm_penalized` function. As described in Section 3.3, the M-step heavily relies on previously developed fast and stable subroutines, while the E-step and the objective function evaluation have been implemented in `c++` to reduce the overall computing time.

- **Response-specific random-effects:** model in (1) assumes that each and every response requires a random-effects component. Whilst in principle reasonable, it may happen in specific applications that only a subset of the r characteristics in $\boldsymbol{Y}$ enjoys group-dependent heterogeneity. The occurrence of such a scenario can be unveiled by looking at the r diagonal elements of dimension q in $\hat{\boldsymbol{\Psi}}$: a response may be considered group-independent when the magnitude of the associated elements in $\text{diag}(\hat{\boldsymbol{\Psi}})$ is significantly lower than the remaining ones. Doing this way, the impact a given random-effect has on the different characteristics is retrieved as a by-product of the modeling procedure.

4 Simulation study

In this section, we evaluate the model introduced in Section 3 on synthetic data. The aim of the analyses reported hereafter is twofold. On the one hand, we would like to validate the predictive power of the proposed procedure against its fixed-effects counterpart when the random-effects vary across dimensions in

the multivariate response. On the other hand, we assess the estimated model parameters and the recovery of the underlying sparsity structure for different values of the shrinkage factor λ.

4.1 Experimental setup

We generate $N = 600$ data points according to model (1) with the following parameters:

$$\boldsymbol{\Psi} = \begin{bmatrix} 50.00 & -1.59 & -0.60 & -0.22 & 2.38 \\ -1.59 & 40.00 & -0.96 & -0.91 & 0.37 \\ -0.60 & -0.96 & 30.00 & -0.43 & 0.50 \\ -0.22 & -0.91 & -0.43 & 20.00 & 0.80 \\ 2.38 & 0.37 & 0.50 & 0.80 & 0.16 \end{bmatrix}, \quad \boldsymbol{\Sigma} = \begin{bmatrix} 2.16 & 0.09 & -0.80 & 0.91 & -0.26 \\ 0.09 & 2.16 & -0.33 & 0.55 & -0.10 \\ -0.80 & -0.33 & 2.16 & -0.03 & -0.13 \\ 0.91 & 0.55 & -0.03 & 2.16 & 0.02 \\ -0.26 & -0.10 & -0.13 & 0.02 & 2.16 \end{bmatrix},$$

implying that $r = 5$ and $q = 1$. Notice that $\boldsymbol{\Psi}$ is purposely constructed for the random-effects to differently affect the five dimensional response: while the first component showcases high variance (first entry in the main diagonal) the last one is very small and close to 0. Further, the error variances (diagonal elements of $\boldsymbol{\Sigma}$) are held constant across dimensions to better highlight the impact the variability of the random-effects has on the models performance. The data generating process assumes ten equally-sized subpopulations, resulting in $J = 10$. The matrix of fixed-effects $\boldsymbol{B}$ is of dimension 10001×5, with distinct sparsity pattern according to three scenarios:

- $\boldsymbol{B}$ *row-wise sparse*: $\boldsymbol{B}$ has entries equal to 0.5 for the first 100 rows, while all the other entries are equal to 0,
- $\boldsymbol{B}$ *sparse at random*: $\boldsymbol{B}$ is equal to 0.5 for approximately 70% of its entries, while all the others are equal to 0,
- $\boldsymbol{B}$ *with dependence structure*: $\boldsymbol{B}$ has entries whose magnitude agrees with the correlation structure between the covariates, inducing coefficients to be similar when the absolute correlation between two predictors is high.

Lastly, $\boldsymbol{Z}_j$ is an all-one column vector $\forall j = 1, \ldots, 10$, while $\boldsymbol{X}_j$ has the first column equal to 1, meaning that the intercept is included in $\boldsymbol{X}_j$ in our model specification. The remaining 10000 dimensions are generated according to a normal random vector with independent marginals for the $\boldsymbol{B}$ *row-wise sparse* and $\boldsymbol{B}$ *sparse at random* scenarios, while the `cormat_from_triangle` function from the `faux` package (DeBruine, 2021) has been used to simulate correlated predictors in the $\boldsymbol{B}$ *with dependence structure* experiment.

Taking a cue from the Monte Carlo simulations of Li and Li (2010), for each replication of our experiment the learning framework is structured as follows: we equally divide the $N = 600$ units in a training set, an independent validation set and an independent test set, retrieving a sample size of 200 for each. Notice that, as to mimic the process of DNAm surrogates creation, the total number of variables ($p = 10001$) is much larger than the sample size. Seven different models, varying λ within a grid, are fitted on the training data:

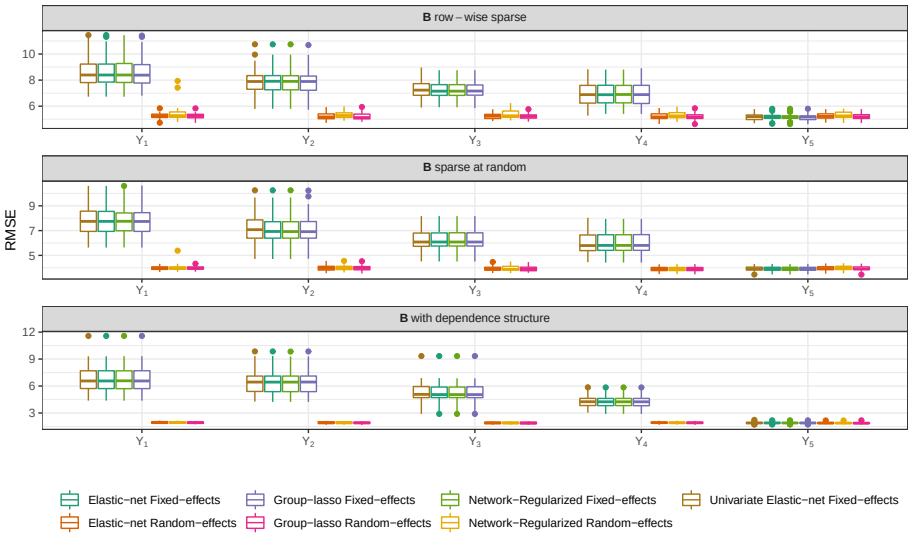

Figure 3: Boxplots of the Root Mean Squared Error (RMSE) for $MC = 100$ repetitions of the simulated experiment. RMSE is computed on 200 test points for different methods and three scenarios varying sparsity pattern for $\boldsymbol{B}$.

- *Univariate Elastic-net Fixed-effects:* univariate elastic-net regression, obtained fitting independent models to each dimension of the multivariate response,
- *Elastic-net Fixed-effects:* a penalized multitask learning model with elastic-net regularization. The considered penalty is described in Section 3.3.1,
- *Group-lasso Fixed-effects:* a penalized multitask learning model with multivariate group-lasso regularization. The considered penalty is described in Section 3.3.2,
- *Network-Regularized Fixed-effects:* graph-regularized multitask learning model with edge-based regularization. The considered penalty is described in Section 3.3.3,
- *Elastic-net Random-effects:* the penalized MLMM methodology introduced in the paper with elastic-net regularization (Section 3.3.1),
- *Group-lasso Random-effects:* the penalized MLMM methodology introduced in the paper, with group-lasso regularization (Section 3.3.2),
- *Network-Regularized Random-effects:* the penalized MLMM methodology introduced in the paper, with edge-based regularization (Section 3.3.3).

Such an extensive comparison can be regarded as performing a within-scenario ablation study, in which we start from a complex method and we subsequently

remove the random-effects component and, finally, the borrow strength property of multivariate regression, to be left with Univariate Elastic-net Fixed-effects models. In this way, we investigate the contribution of our proposal to the overall system. The mixing parameter α was set equal to 0.5 for methods with elastic-net and group-lasso regularizers, while for the Network-Regularized penalty we employ 5-fold CV to tune λ_X and λ_Y on the training set. For the latter penalty, the adjacency matrices $\boldsymbol{G}_X$ and $\boldsymbol{G}_Y$ are computed via a thresholding procedure on the correlation matrices of $\boldsymbol{X}$ and $\boldsymbol{Y}$, respectively, with a threshold equal to 0.1 (Langfelder and Horvath, 2008). Subsequently, the validation dataset is used to select the best shrinkage parameter λ minimizing the RMSE for every model. The predictive performance is then evaluated on the test set. Lastly, to assess out of groups prediction, models are further validated on 100 external samples, generated according to (1), coming from five extra subpopulations not observed in the training set. The devised simulated experiment is replicated $MC = 100$ times: results are reported in the next subsection.

4.2 Simulation results

Figure 3 displays boxplots of the Root Mean Squared Error, computed for each component of the 5-dimensional response on the test set. For all scenarios we observe that the component-wise predictive performance is heavily affected by the magnitude of the related diagonal entry in the $\boldsymbol{\Psi}$ matrix. When the grouping effect is negligible (fifth dimension Y_5), all methods showcase comparable predictive performance under both scenarios. Contrarily, the RMSE deteriorates for fixed-effects models in those response components for which the grouping impact is more relevant. The same does not happen for the mixed-effects counterparts, as the random intercept effectively captures baseline differences across groups. Interestingly, the penalty type does not seem to influence the RMSE metric, with our proposal displaying excellent results irrespective of the chosen shrinkage functional for all scenarios. On the other hand, when it comes to perform out of groups prediction the gain achieved by including random-effects decreases and the outcome of models with fixed and mixed-effects are fairly similar. In details, for the latter class of methods the unconditional (population level) intercepts are employed when making predictions for unobserved groups. Notwithstanding, we recognize that results are no worse than those obtained with fixed-effects procedures, corroborating the generalizability of our proposals in external cohorts.

Figure 4 displays the analogue of the Percentage of Variation due to Random Effects (PVRE) metric, computed taking the ratio between the diagonal elements of $\hat{\boldsymbol{\Psi}}$ and the sum of the diagonals of $\hat{\boldsymbol{\Psi}}$ and $\hat{\boldsymbol{\Sigma}}$. From the plot, it clearly emerges how the grouping impact differently affects the variability in the five components of the response.

Figure 5 reports boxplots of the Frobenius distance between true and estimated matrices of fixed-effects $\boldsymbol{B}$. When looking at $||\boldsymbol{B} - \hat{\boldsymbol{B}}||_F$ under the three scenarios we observe some interesting facts. First off, it is immediately noticed that the *Univariate Elastic-net Fixed-effects* model showcases the poorest per-

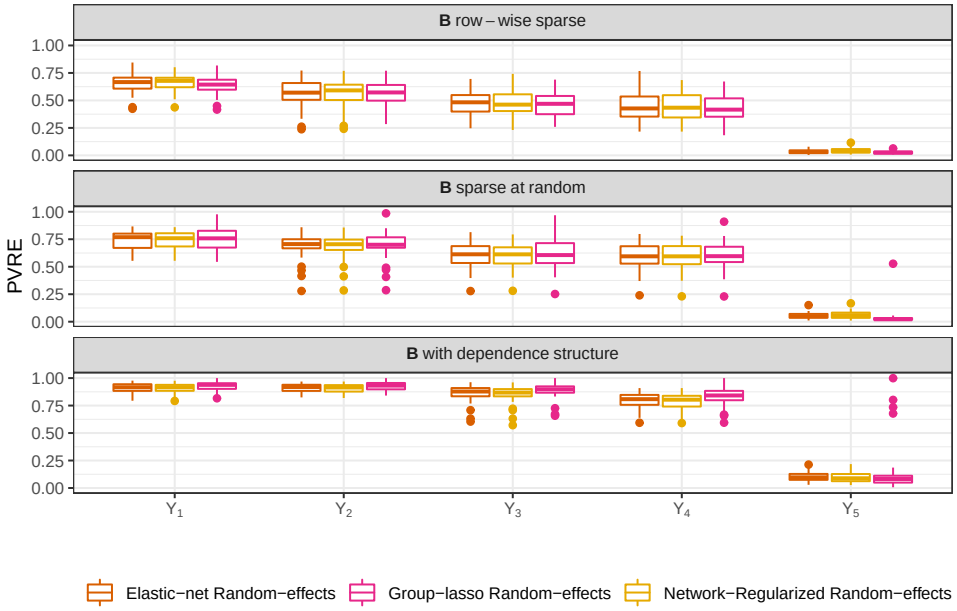

Figure 4: Boxplots of the Percentage of Variation due to Random Effects (PVRE) for $MC = 100$ repetitions of the simulated experiment.

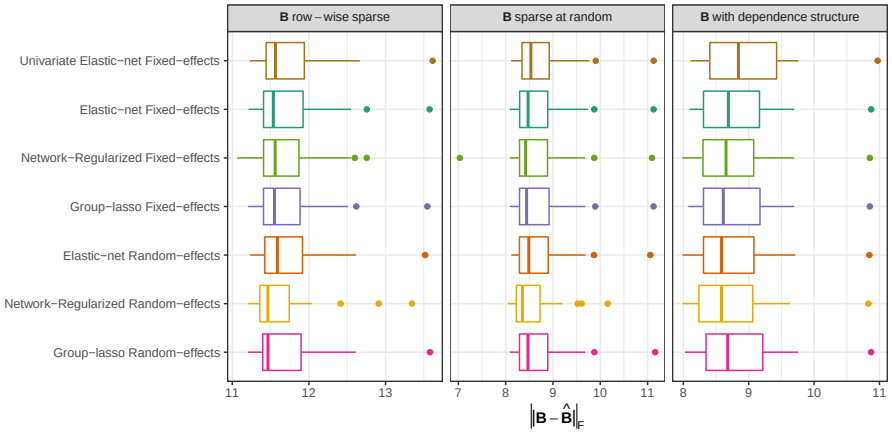

Figure 5: Boxplots of the Frobenius distance between true and estimated matrices of fixed-effects $\boldsymbol{B}$ for $MC = 100$ repetitions of the simulated experiment.

formance, in particular for the $\boldsymbol{B}$ *with dependence structure* experiment. This is due to the fact that the different components of the response vector are related in our simulated specification, and therefore fitting separate regression models results in a loss of quality for the estimator. Secondly, we observe that for the $\boldsymbol{B}$ *row-wise sparse* scenario the *Group-lasso Random-effects* model is the best performing one among all the competitors, displaying the lowest median distance to $\boldsymbol{B}$. This may be expected, as such method is precisely constructed to identify a matrix of fixed-effects with row-wise sparsity patterns. Furthermore, notice that the performance of the *Group-lasso Random-effects* is slightly better than its Fixed-effects counterpart. For the remaining scenarios the superiority of the mixed-effects procedures is not so apparent, and both fixed and random-effects models demonstrate a comparable performance. An only modest gain is showcased by the *Network-Regularized Random-effects* method, for which the inclusion of the adjacency matrices $\boldsymbol{G}_X$ and $\boldsymbol{G}_Y$ in the penalty specification helps in better recovering the $\boldsymbol{B}$ structure.

We now look at the ability of the competing procedures in recovering the true underlying sparsity patterns in the matrix of fixed-effects $\boldsymbol{B}$ under the different scenarios. In so doing we compute, for each replication of the simulated experiment, the F_1 score defined as follows:

$$F_1 = \frac{\texttt{tp}}{\texttt{tp} + 0.5(\texttt{fp} + \texttt{fn})}, \tag{20}$$

where with `tp` we denote the number of zero entries in $\boldsymbol{B}$ correctly estimated as such, while `fp` and `fn` represent the number of non-zero entries wrongly shrunk to 0 and the number of zero entries not shrunk to 0, respectively. Figure 6 displays boxplots of such metric for different methods and scenarios. We notice that $\boldsymbol{B}$ *row-wise sparse* structure displays much higher F_1 score, irrespective of the considered method, than the other two cases. Intuitively, the former scenario is less challenging since, while all penalty types can potentially accommodate a row-wise sparse $\boldsymbol{B}$, group-lasso regularizers only force entire rows of $\boldsymbol{B}$ to be shrunk to 0. We further observe that the F_1 score is higher for the *Group-lasso Random-effects* than for its fixed-effects counterpart, highlighting that a penalized mixed-effects modeling strategy, in presence of grouped data and a row-wise sparse $\boldsymbol{B}$, not only increases the predictive accuracy but also improves the recovery of the sparsity pattern in the fixed-effects matrix. The same does not happen in the remaining two scenarios, for which all methods display a comparable empirical distribution of the F_1 metric across simulations. The same behavior was already observed in high-dimensional linear mixed-effects modeling for univariate responses (Schelldorfer et al., 2011).

As a last worthy note we acknowledge that, as rightly underlined by an anonymous reviewer, the present simulation study does not consider any violation in the distributional assumptions of the involved quantities. In this regard, we replicated the experiment using both a Multivariate skew-normal and Multivariate skew-t distributions (Azzalini and Capitanio, 2013) as generative models for the error term, but we did not report the results in the paper since no dramatic changes were observed in model performances. While clearly

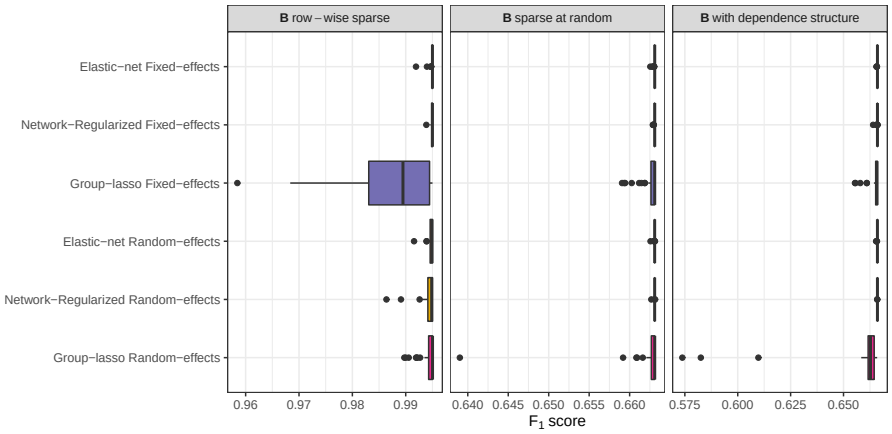

Figure 6: Boxplots of the F_1 score for $MC = 100$ repetitions of the simulated experiment for different methods and three scenarios varying sparsity pattern for $\boldsymbol{B}$.

more extreme scenarios could be considered, results in the literature have previously validated the robustness of linear mixed-effects models to violations of distributional assumptions (McCulloch and Neuhaus, 2011). For further details about the simulation study, the Supplementary Material (Cappozzo et al., 2023) provides additional figures and a note on the overall computing times.

All in all, the good performances displayed by our proposal, particularly when coupled with a multivariate group-lasso penalty, encourage its usage in multivariate DNAm surrogates creation: promising results are reported in the next Section.

5 DNAm biomarkers analysis for EPIC and EXPOsOMICS datasets

The methodology described in Section 3 is employed to build a 5-dimensional DNAm biomarker of hypertension and hyperlipidemia. As mentioned in the introduction, DNAm surrogates possess extensive advantages over their blood-measured counterparts since:

1. DNAm biomarkers directly account for genetic susceptibility and subject specific response to risk factors;

2. DNAm biomarkers can immediately be computed whenever DNAm values are accessible. This is particularly useful when the risk factors of interest have not been directly measured;

3. further understanding of the biomolecular mechanisms associated with complex phenotypes can be acquired through a pathway enrichment anal-

ysis (Reimand et al., 2019), allowing to identify molecular pathways over-represented among the regressors involved in the surrogate construction (i.e., the CpG sites whose associated parameters are not shrunk to 0).

We subsequently assess how well the so-devised surrogates perform, for both internal and external cohorts, in reconstructing the blood measured biomarkers (Section 5.1) and in predicting the clinical endpoint of interest, namely the future presence/absence of CVD events (Section 5.2). Lastly, we study from a biological perspective the CpG sites selection operated by the multivariate group-lasso penalty, comparing it with previous findings available in the literature (Section 5.3).

5.1 DNAm surrogates creation and validation

To construct multivariate DNAm surrogates, several penalized models are fitted to the EPIC Italy training set, varying shrinkage factors and considering both fixed and random-effects components. As mentioned in Section 2, the design matrix comprises of $p = 62130$ variables. Thus, redundancies are likely to occur as the feature space is constituted by the union of CpG sites pre-screened by univariate epigenome-wide analyses. After having standardized the covariates, for each model the penalty term λ is tuned via 10-fold CV, while the mixing parameter α is kept fixed and equal to 0.5. Results on the internal cohort are summarized in Table 1, where the Root Mean Squared Error (RMSE) computed on the EPIC Italy test set, the number of active CpG sites and the overall elapsed time are reported. The first two rows are related to the novel penalized MLMM methodology with a random-effects design matrix that includes a $q = 1$ random intercept, coupled with multivariate group-lasso (Section 3.3.2) and elastic-net (Section 3.3.1) penalties, respectively. The corresponding fixed-effects counterparts are reported in the third and fourth rows, while univariate elastic-net metrics, obtained fitting $r = 5$ separate models, one for each response, are detailed in the last row of Table 1. Notice that our proposal outperforms the state-of-the-art approach (univariate elastic-net) for 4 out of 5 dimensions of the response variable. The reason being that our method takes advantage of the borrowing information asset typical of multivariate models (the correlation between SBP and DBP is equal to 0.77 in the training set), whilst allowing for center-wise difference to be captured by the random intercept. Furthermore, thanks to the multivariate group-lasso penalty, our penalized MLMM approach directly identifies the CpG sites that are jointly related to hypertension and hyperlipidemia, with a total number of features that is lower with respect to univariate elastic-nets. By taking the ratio between the diagonal elements of $\hat{\boldsymbol{\Psi}}$ and the sum of the diagonals of $\hat{\boldsymbol{\Psi}}$ and $\hat{\boldsymbol{\Sigma}}$ it is possible to compute, for each component of the response matrix $\boldsymbol{Y}$, the analogue of the Percentage of Variation due to Random Effects (PVRE) index. For the EPIC Italy dataset, the estimated PVRE amounts to 7.97%, 16.05%, 5.56%, 14.26% and 6.01% for DBP, HDL, LDL, SBP and TG respectively; giving reason for the performance improvement showcased by the random-effects models.

Table 1: Root Mean Squared Error (RMSE), active number of CpG sites and overall computing times for different penalized regression models, EPIC Italy test set. Bold numbers indicate lowest RMSE for each of the $r = 5$ dimension of the response matrix.

	Framework		Root Mean Squared Error					Active #	Elapsed time
Model	Penalty type	Response	DBP	HDL	LDL	SBP	TG	CpG sites	(secs)
Random-effects	Group-lasso	Multivariate	**0.1167**	**0.2442**	**0.3236**	0.1374	**0.4745**	417	304
Random-effects	Elastic-net	Multivariate	0.1178	0.2480	0.3318	0.1379	0.4853	773	549
Fixed-effects	Group-lasso	Multivariate	0.1317	0.2602	0.3321	0.1364	0.4853	382	54
Fixed-effects	Elastic-net	Multivariate	0.1176	0.2513	0.3324	0.1362	0.4841	115	311
Fixed-effects	Elastic-net	Univariate	0.1179	0.2596	0.3383	**0.1359**	0.4996	1712	115

The employment of the multivariate group-lasso penalty within a mixed-effects multitask learning framework is also supported by biological reasons. In fact, it is more likely that multiple correlated phenotypes affect (or are affected by, depending on the causal relationship between DNAm and the exposure variable) the same set of CpG sites. This mechanism is known as pleiotropic effect (Tyler et al., 2013; Richard et al., 2017).

Internal validation results obtained for the EPIC Italy test set highlight the benefits of mixed-effects modeling while constructing DNAm surrogates for data possessing a grouping-structure. The random intercept allows to account for center-wise variability that is induced by geographic genetical variation, as well as by samples collection and storage that are likely to differ across centers. Notice that in general, if not properly modeled, the unexplained heterogeneity in multi-center studies could be taken care of with dedicated batch-effect removal procedures (see, e.g., Johnson et al., 2007). Nonetheless, when developing DNAm biomarkers it is of interest to devise study-invariant surrogates, to be readily computed also for samples not belonging to the learning cohort. To this aim, we validate the performance of the models estimated on the EPIC training set in constructing surrogates for the external EXPOsOMICS cohort (see Section 2). In this context the grouping information (i.e., the center of recruitment) cannot be considered when performing predictions with mixed-effects models, and the unconditional (population-level) intercepts are thus utilized. RMSE between estimated and blood-measured biomarkers for the EXPOsOMICS validation cohort are reported in Table 2. Likewise for the EPIC Italy test set, the lowest RMSEs for all but SBP biomarker are retained employing a penalized random-intercept model with multivariate group-lasso penalty. Interestingly, the predictive outcomes obtained in the EXPOsOMICS validation dataset are comparable with the ones reported in Table 1, with slightly worse performances, as it may be expected, for those dimensions of the response displaying higher PVRE indexes. All in all, the proposed approach exhibits promising results when it comes to multivariate DNAm biomarker creation, outperforming the current employed procedure in both internal and external validation cohorts.

Table 2: Root Mean Squared Error (RMSE) for different penalized regression models, EXPOsOMICS validation set. Bold numbers indicate lowest RMSE for each of the $r = 5$ dimension of the response matrix.

Model	Penalty type	Response	DBP	HDL	LDL	SBP	TG
Random-effects	Group-lasso	Multivariate	**0.1314**	**0.2384**	**0.2707**	0.1412	**0.4735**
Random-effects	Elastic-net	Multivariate	0.1335	0.2504	0.2821	0.1450	0.4890
Fixed-effects	Group-lasso	Multivariate	0.1409	0.2750	0.2969	0.1574	0.5142
Fixed-effects	Elastic-net	Multivariate	0.1286	0.2479	0.2859	0.1359	0.5002
Fixed-effects	Elastic-net	Univariate	0.1331	0.2733	0.3136	**0.1368**	0.5251

5.2 Association of DNAm surrogates with CVD risk

DNAm surrogates creation is not a standalone regression problem, as its primary aim is to provide reliable covariates for diseases prediction models (Fernández-Sanlés et al., 2021; Odintsova et al., 2021; Hidalgo et al., 2021). We therefore validate whether employing the estimated DNAm surrogates acts as a superior proxy of blood measured biomarkers in association analyses. In details, within the cohort of patients in the EPIC Italy test set we build logistic regression models to predict the probability of cardiovascular risk using as regressors either the blood measured biomarkers or the two best performing DNAm surrogates devised in the previous section, adjusting for sex and age. The Receiver Operating Characteristic (ROC) curves and the associated Area Under the Curve (AUC) metrics for the considered methods are displayed in Figure 7. As expected, in light of the RMSE results reported in Table 1, we notice that classification performances are similar among the competing models. Nonetheless, the logistic curves regressed on the DNAm based surrogates seem to outperform the blood measured counterparts. Interestingly, all surrogates in Table 1 define logistic regression models whose AUCs are higher than those retrieved by the blood measured biomarkers, with the best performance attained by the penalized random-intercept model with multivariate group-lasso penalty.

We further assess the association of DNAm surrogates with CVD risk in the external EXPOsOMICS study. For this dataset values of the blood based biomarkers are available only for a subset of volunteers, we thus construct the logistic regression models by means of the surrogates only. Such a situation is quite common in validation data and in line with the principle DNAm surrogates were devised in the first place. Also in this context the predictive performance of our novel proposal, coupled with a multivariate group lasso penalty, is higher with respect to state-of-the-art surrogates created via Elastic-net Fixed-effects models. The associated Receiver Operating Characteristic curves, as well as additional figures related to the DNAm biomarkers analysis, are reported in the Supplementary Material (Cappozzo et al., 2023).

The association analyses reported in this section cast light on the applicability of the devised DNAm surrogates as an enriched and patient-specific proxy of their blood measured counterparts, in both internal and external cohorts. These

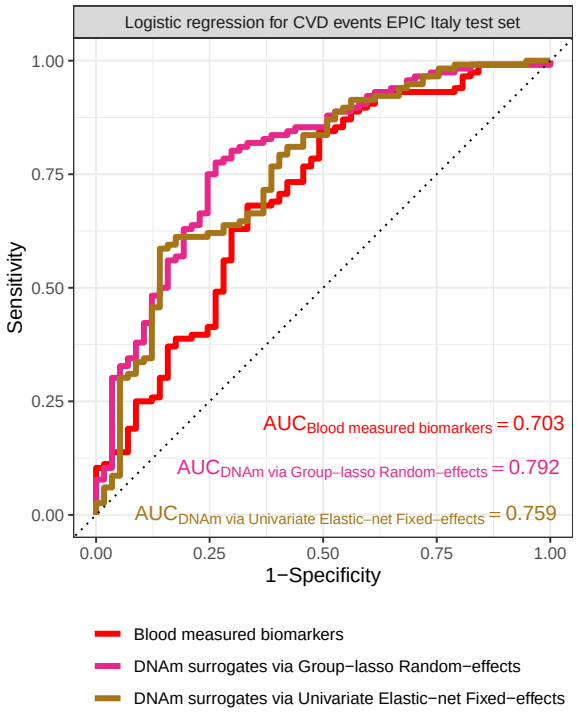

Figure 7: Receiver Operating Characteristic curves and Area Under the Curve metrics for the association analyses of DNAm surrogates with CVD risk, Italy EPIC test set.

favorable outcomes indicate that using models based on DNAm surrogates could be more appropriate for prediction tasks, such as CVD prevention, since they can possibly incorporate individual characteristics not directly recorded in the blood measured biomarkers.

5.3 CpG sites selection and gene set enrichment analyses of inflammatory pathways

In the previous sections the newly devised random-intercept model for multi-task learning has demonstrated superior predictive performance when it comes to DNAm surrogates creation and CVD prediction. We hereafter examine the epidemiological rationale of the multivariate group-lasso penalty compared to univariate elastic-nets, investigating the biological reliability of the selected features (CpG sites). The univariate elastic-nets extract 178, 504, 518, 79, 497 CpG sites for diastolic blood pressure, HDL cholesterol, LDL cholesterol, systolic blood pressure and triglycerides, respectively. As reported in Table 1, the total number of unique CpGs is 1712. However, despite the high degree of correlation among the multivariate outcomes, no CpGs were in common in the five sets, and only a minor percentage of CpGs was shared among two or more responses. Instead, as previously described , our MLMM procedure regularized with a multivariate group-lasso penalty extracts 417 features that are associated with the five outcomes at the same time, a biological mechanism known as pleiotropy (Atchley and Hall, 1991a,b).

We are further interested in assessing whether the selected CpG sites are associated with specific biological pathways. To do so, gene set enrichment analyses (Subramanian et al., 2005) are performed on the features retained by penalized MLMM and univariate elastic-net models. Specifically, given that the number of CpGs extracted for each method is small (hundreds of CpGs selected from an initial set of 295614 features), an analysis based on all the biomolecular pathways described in the canonical datasets, i.e., KEGG (Yi et al., 2020), GO (Gene Ontology Consortium, 2004) and Reactome (Fabregat et al., 2018), would be under-powered. Therefore, to overcome this limitation and with inflammation being the main mechanism involved in the onset of the majority of chronic diseases, we focus our analyses on the 17 inflammatory pathways described in Loza et al. (2007). Enrichment analyses results are summarized in Table 3. For each list of CpGs we test for over-representation of features in inflammatory pathways using the method implemented in the `missMethyl R` package (Phipson et al., 2016). Considering the CpGs extracted by our multivariate approach, we find significant enrichment for CpGs in four inflammatory pathways. These results agree with previous literature suggesting that hypertension and hyperlipidaemia are associated with the dysregulation of molecular pathways regulating apoptosis, oxidative stress, and the immune system (Senoner and Dichtl, 2019; Dong et al., 2020). Instead, modeling the outcomes one by one using univariate models leads to a less consistent pattern of associations. In fact, we find only one (and always different) significant pathway per analysis.

All in all, the results reported in this section support the advantages of mod-

Table 3: Empirical p-values of the enrichment analyses computed using a permutation procedure via the `gometh` function in the `missMethyl R` package. Empirical p-values lower than 0.05, highlighted in bold, indicate significant over-representation.

Inflammatory pathway	**Group-lasso Random-effects**	**Univariate Elastic-net Fixed-effects**				
		DBP	**HDL**	**LDL**	**SBP**	**TG**
Leukocyte signaling	**0.006**	0.14	0.35	0.01	1	1
ROS/Glutathione/Cytotoxic granules	**0.009**	0.06	0.15	1	**0.03**	0.15
Apoptosis Signaling	**0.01**	1	1	0.17	1	0.07
Natural Killer Cell Signaling	**0.01**	0.36	0.10	1	1	1
PI3K/AKT Signaling	0.18	1	0.26	0.03	1	0.22
Innate pathogen detection	0.26	0.12	0.30	0.31	1	1
Cytokine signaling	0.36	1	1	**0.0003**	1	0.10
Adhesion-Extravasation-Migration	1	0.18	**0.003**	0.02	0.09	0.10
Calcium Signaling	1	1	1	0.08	1	1
Complement Cascase	1	1	1	1	1	0.16
Glucocorticoid/PPAR signaling	1	1	0.10	1	0.17	0.10
G-Protein Coupled Receptor Signaling	1	1	1	1	0.05	0.27
MAPK signaling	1	1	0.14	0.49	1	**0.005**
NF-kB signaling	1	1	0.16	0.16	1	0.15
Phagocytosis-Ag presentation	1	1	1	0.26	1	1
Eicosanoid Signaling	1	1	1	1	1	1
TNF Superfamily Signaling	1	1	0.01	0.14	1	1

eling multiple correlated outcomes not only from a prediction perspective, for both blood measured biomarkers and endpoint of interest, but also considering the biological reliability of the extracted features.

6 Discussion and further work

In the present paper we have proposed a novel framework for mixed-effects multitask learning suitable for high-dimensional data. The ubiquitous presence in modern applications of "p bigger than N" problems asks for the development of ad-hoc statistical tools able to cope with such scenarios. By resorting to penalized likelihood estimation, we have devised a general purpose EM algorithm capable of accommodating any penalty type that has been previously defined for fixed-effects models. We have examined three functional forms for the penalty term, discussing pros and cons of each and providing convenient routines for model fitting. The proposal has been accompanied by some considerations on distinguishing features, like how to quantify response specific random-effects, and other more general issues concerning initialization, convergence and model selection.

The work has been motivated by the problem of developing a multivariate

DNAm biomarker of cardiovascular and high blood pressure comorbidities from a multi-center study. The EPIC Italy dataset has been analyzed using Diastolic Blood Pressure, Systolic Blood Pressure, High Density Lipoprotein, Low Density Lipoprotein and Triglycerides as response variables, regressing them on 62128 CpG sites and accounting for between-center heterogeneity. Our modeling framework, coupled with a multivariate group-lasso penalty, has demonstrated to outperform the state-of-the-art alternative, both in terms of predictive power and biomedical interpretation. Remarkably, the number of CpG sites deemed as relevant in the multi-dimensional surrogate creation was found to be lower than those identified by separately fitting penalized models for each risk factor. Decreasing the amount of relevant CpG sites is crucial to reduce sequencing costs for future studies, with the final aim of querying only a limited number of targeted genomic regions. Such a result may thereupon favor the adoption of our methodological approach for building DNAm surrogates.

The devised pipeline possesses also some limitations. The EWAS results are adjusted for clinical covariates external to the analysis, and this may thus affect the pre-processing outcome. Moreover, the level of strictness in the screening process is influenced by the chosen threshold on the p-values. On this wise two different, yet both sensible, strategies can be adopted. On the one hand, one may rely on the "Occam's razor" principle, preferring to use a stricter threshold being it the simplest and fastest option. On the other hand, one can positively include many redundant variables in the design matrix, relying on the model ability to shrink coefficients of irrelevant features to zero. Concurrently, further insights about the associations between DNA methylation and blood-measured biomarkers may be unraveled by means of sparse multiple canonical correlation analysis (Witten et al., 2009; Witten and Tibshirani, 2009; Rodosthenous et al., 2020); while other modeling approaches such as deep-learning (Nguyen et al., 2022; Yuan et al., 2022), Bayesian methods (Zhao et al., 2021a,b) and boosting machines (Sigrist, 2022) could be profitably adapted to build DNAm multidimensional surrogates.

A direction for future research concerns promoting the application of the proposed procedure in creating additional multi-dimensional DNAm biomarkers, conveniently embedding mixed-effects and customized penalty types. In this regard, of particular interest may be the definition of a shrinkage term for which the grouping in $\boldsymbol{B}$ is introduced from both responses and predictors: the former is induced by the multivariate nature of $\boldsymbol{Y}$ (i.e., the r responses), while the latter can stem from any structure present in $\boldsymbol{X}$ (e.g., CpG islands). Such a problem could be solved by extending the Multivariate Sparse Group Lasso, proposed by Li et al. (2015), to the mixed-effects framework.

In addition, having assumed random intercepts for each and every component in a low-dimensional response framework was only motivated by the application at hand, and it may not be valid in general. Thus, a two-fold methodological development naturally arises: a first one concerning the definition of response-specific random-effects in multitask learning and another accounting for the inclusion of custom penalties when dealing with high-dimensional response variables. Furthermore, the latter may also possess a mixed-type struc-

ture, with components simultaneously be nominal, ordinal, discrete and/or continuous. Some proposals are currently under study and they will be the object of future work.

References

Anastasiadi D, Esteve-Codina A, Piferrer F (2018) Consistent inverse correlation between DNA methylation of the first intron and gene expression across tissues and species. Epigenetics & Chromatin 11(1):37

Atchley WR, Hall BK (1991a) A model for development and evolution of complex morphological structures. Biological Reviews of the Cambridge Philosophical Society 66(2):101–157

Atchley WR, Hall BK (1991b) A model for development and evolution of complex morphological structures. Biological Reviews 66(2):101–157

Azzalini A, Capitanio A (2013) The Skew-Normal and Related Families, vol 3. Cambridge University Press

Battram T, Yousefi P, Crawford G, Prince C, Babaei MS, Sharp G, Hatcher C, Vega-Salas MJ, Khodabakhsh S, Whitehurst O, Others (2022) The EWAS Catalog: a database of epigenome-wide association studies. Wellcome open research 7

Campagna MP, Xavier A, Lechner-Scott J, Maltby V, Scott RJ, Butzkueven H, Jokubaitis VG, Lea RA (2021) Epigenome-wide association studies: current knowledge, strategies and recommendations. Clinical Epigenetics 13(1):214

Cappozzo A, McCrory C, Robinson O, Freni Sterrantino A, Sacerdote C, Krogh V, Panico S, Tumino R, Iacoviello L, Ricceri F, Sieri S, Chiodini P, McKay GJ, McKnight AJ, Kee F, Young IS, McGuinness B, Crimmins EM, Arpawong TE, Kenny RA, O'Halloran A, Polidoro S, Solinas G, Vineis P, Ieva F, Fiorito G (2022) A blood DNA methylation biomarker for predicting short-term risk of cardiovascular events. Clinical Epigenetics 14(1):121

Cappozzo A, Ieva F, Fiorito G (2023) Supplement to "A general framework for penalized mixed-effects multitask learning with applications on DNA methylation surrogate biomarkers creation"

Caruana R (1997) Multitask learning. Machine learning 28(1):41–75

Castro de Moura M, Davalos V, Planas-Serra L, Alvarez-Errico D, Arribas C, Ruiz M, Aguilera-Albesa S, Troya J, Valencia-Ramos J, Vélez-Santamaria V, Rodríguez-Palmero A, Villar-Garcia J, Horcajada JP, Albu S, Casasnovas C, Rull A, Reverte L, Dietl B, Dalmau D, Arranz MJ, Llucià-Carol L, Planas AM, Pérez-Tur J, Fernandez-Cadenas I, Villares P, Tenorio J, Colobran R, Martin-Nalda A, Soler-Palacin P, Vidal F, Pujol A, Esteller M (2021)

Epigenome-wide association study of COVID-19 severity with respiratory failure. EBioMedicine 66:103339

Cheng W, Zhang X, Guo Z, Shi Y, Wang W (2014) Graph-regularized dual Lasso for robust eQTL mapping. Bioinformatics 30(12):139–148

Chipperfield JO, Steel DG (2012) Multivariate random effect models with complete and incomplete data. Journal of Multivariate Analysis 109:146–155

Chung FRK, Graham FC (1997) Spectral graph theory. 92, American Mathematical Soc.

Colicino E, Just A, Kioumourtzoglou MA, Vokonas P, Cardenas A, Sparrow D, Weisskopf M, Nie LH, Hu H, Schwartz JD, Wright RO, Baccarelli AA (2021) Blood DNA methylation biomarkers of cumulative lead exposure in adults. Journal of Exposure Science & Environmental Epidemiology 31(1):108–116

Conole ELS, Stevenson AJ, Green C, Harris SE, Maniega SM, Valdés-Hernández MdC, Harris MA, Bastin ME, Wardlaw JM, Deary IJ, Miron VE, Whalley HC, Marioni RE, Cox SR (2020) An epigenetic proxy of chronic inflammation outperforms serum levels as a biomarker of brain ageing. medRxiv p 2020.10.08.20205245

Dawid AP (1981) Some Matrix-Variate Distribution Theory: Notational Considerations and a Bayesian Application. Biometrika 68(1):265

DeBruine L (2021) faux: Simulation for Factorial Designs

Demidenko E (2013) Mixed models: theory and applications with R. John Wiley & Sons

Dempster AP, Laird NM, Rubin DB (1977) Maximum Likelihood from Incomplete Data Via the EM Algorithm. Journal of the Royal Statistical Society: Series B (Methodological) 39(1):1–22

Dirmeier S, Fuchs C, Mueller NS, Theis FJ (2018) NetReg: Network-regularized linear models for biological association studies. Bioinformatics 34(5):896–898

Dong W, Chen H, Wang L, Cao X, Bu X, Peng Y, Dong A, Ying M, Chen X, Zhang X, Yao L (2020) Exploring the shared genes of hypertension, diabetes and hyperlipidemia based on microarray. Brazilian Journal of Pharmaceutical Sciences 56:1–12

van Eijk KR, de Jong S, Boks MP, Langeveld T, Colas F, Veldink JH, de Kovel CG, Janson E, Strengman E, Langfelder P, Kahn RS, van den Berg LH, Horvath S, Ophoff RA (2012) Genetic analysis of DNA methylation and gene expression levels in whole blood of healthy human subjects. BMC Genomics 13(1):636

Fabregat A, Jupe S, Matthews L, Sidiropoulos K, Gillespie M, Garapati P, Haw R, Jassal B, Korninger F, May B, Milacic M, Roca CD, Rothfels K, Sevilla C, Shamovsky V, Shorser S, Varusai T, Viteri G, Weiser J, Wu G, Stein L, Hermjakob H, D'Eustachio P (2018) The Reactome Pathway Knowledgebase. Nucleic Acids Research 46(D1):D649–D655

Fan J, Lv J (2008) Sure independence screening for ultrahigh dimensional feature space. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 70(5):849–911

Fan J, Samworth R, Wu Y (2009) Ultrahigh Dimensional Feature Selection: Beyond the Linear Model. Journal of Machine Learning Research 10(2009):2013–2038

Fazzari MJ, Greally JM (2010) Introduction to Epigenomics and Epigenome-Wide Analysis, Humana Press, Totowa, NJ, pp 243–265

Fernández-Sanlés A, Sayols-Baixeras S, Subirana I, Sentí M, Pérez-Fernández S, de Castro Moura M, Esteller M, Marrugat J, Elosua R (2021) DNA methylation biomarkers of myocardial infarction and cardiovascular disease. Clinical Epigenetics 13(1):86

Fiorito G, Vlaanderen J, Polidoro S, Gulliver J, Galassi C, Ranzi A, Krogh V, Grioni S, Agnoli C, Sacerdote C, Panico S, Tsai MY, Probst-Hensch N, Hoek G, Herceg Z, Vermeulen R, Ghantous A, Vineis P, Naccarati A (2018) Oxidative stress and inflammation mediate the effect of air pollution on cardio- and cerebrovascular disease: A prospective study in nonsmokers. Environmental and Molecular Mutagenesis 59(3):234–246

Fiorito G, Pedron S, Ochoa-Rosales C, McCrory C, Polidoro S, Zhang Y, Dugué PA, Ratliff S, Zhao WN, McKay GJ, Costa G, Solinas MG, Harris KM, Tumino R, Grioni S, Ricceri F, Panico S, Brenner H, Schwettmann L, Waldenberger M, Matias-Garcia PR, Peters A, Hodge A, Giles GG, Schmitz LL, Levine M, Smith JA, Liu Y, Kee F, Young IS, McGuinness B, McKnight AJ, van Meurs J, Voortman T, Kenny RA, Vineis P, Carmeli C (2022) The Role of Epigenetic Clocks in Explaining Educational Inequalities in Mortality: A Multicohort Study and Meta-analysis. The Journals of Gerontology: Series A 77(9):1750–1759

Frohlich H, Zell A (2005) Efficient parameter selection for support vector machines in classification and regression via model-based global optimization. In: Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005., IEEE, vol 3, pp 1431–1436

Gałecki A, Burzykowski T (2013) Linear Mixed-Effects Models Using R. Springer Texts in Statistics, Springer New York, New York, NY

Gene Ontology Consortium (2004) The Gene Ontology (GO) database and informatics resource. Nucleic Acids Research 32(90001):258D–261

Guida F, Sandanger TM, Castagné R, Campanella G, Polidoro S, Palli D, Krogh V, Tumino R, Sacerdote C, Panico S, Severi G, Kyrtopoulos SA, Georgiadis P, Vermeulen RC, Lund E, Vineis P, Chadeau-Hyam M (2015) Dynamics of smoking-induced genome-wide methylation changes with time since smoking cessation. Human Molecular Genetics 24(8):2349–2359

Hastie T, Tibshirani R, Wainwright M (2015) Statistical Learning with Sparsity. Chapman and Hall/CRC

Hidalgo BA, Minniefield B, Patki A, Tanner R, Bagheri M, Tiwari HK, Arnett DK, Irvin MR (2021) A 6-CpG validated methylation risk score model for metabolic syndrome: The HyperGEN and GOLDN studies. PLOS ONE 16(11):e0259836

Hillary RF, Marioni RE (2021) MethylDetectR: a software for methylation-based health profiling. Wellcome Open Research 5:283

Johnson WE, Li C, Rabinovic A (2007) Adjusting batch effects in microarray expression data using empirical Bayes methods. Biostatistics 8(1):118–127

Jordan MI (2013) On statistics, computation and scalability. Bernoulli 19(4):1378–1390

Kim S, Xing EP (2012) Tree-guided group lasso for multi-response regression with structured sparsity, with an application to eQTL mapping. The Annals of Applied Statistics 6(3):1095–1117

Kim S, Pan W, Shen X (2013) Network-Based Penalized Regression With Application to Genomic Data. Biometrics 69(3):582–593

Langfelder P, Horvath S (2008) WGCNA: an R package for weighted correlation network analysis. BMC Bioinformatics 9(1):559

Laria JC, Carmen Aguilera-Morillo M, Lillo RE (2019) An Iterative Sparse-Group Lasso. Journal of Computational and Graphical Statistics 28(3):722–731

Li C, Li H (2010) Variable selection and regression analysis for graph-structured covariates with an application to genomics. The Annals of Applied Statistics 4(3):1498–1516

Li Y, Nan B, Zhu J (2015) Multivariate sparse group lasso for the multivariate multiple linear regression with an arbitrary group structure. Biometrics 71(2):354–363

Loza MJ, McCall CE, Li L, Isaacs WB, Xu J, Chang BL (2007) Assembly of Inflammation-Related Genes for Pathway-Focused Genetic Analysis. PLoS ONE 2(10):e1035

Lu AT, Quach A, Wilson JG, Reiner AP, Aviv A, Raj K, Hou L, Baccarelli AA, Li Y, Stewart JD, Whitsel EA, Assimes TL, Ferrucci L, Horvath S (2019) DNA methylation GrimAge strongly predicts lifespan and healthspan. Aging 11(2):303–327

Marabita F, Almgren M, Lindholm ME, Ruhrmann S, Fagerström-Billai F, Jagodic M, Sundberg CJ, Ekström TJ, Teschendorff AE, Tegnér J, Gomez-Cabrero D (2013) An evaluation of analysis pipelines for DNA methylation profiling using the Illumina HumanMethylation450 BeadChip platform. Epigenetics 8(3):333–346

McCulloch CE, Neuhaus JM (2011) Misspecifying the Shape of a Random Effects Distribution: Why Getting It Wrong May Not Matter. Statistical Science 26(3):388–402

McLachlan GJ, Krishnan T (2008) The EM Algorithm and Extensions, 2E, Wiley Series in Probability and Statistics, vol 54. John Wiley & Sons, Inc., Hoboken, NJ, USA

Meng XL, Rubin DB (1993) Maximum Likelihood Estimation via the ECM Algorithm: A General Framework. Biometrika 80(2):267

Nguyen TM, Le HL, Hwang KB, Hong YC, Kim JH (2022) Predicting High Blood Pressure Using DNA Methylome-Based Machine Learning Models. Biomedicines 10(6):1406

Obozinski G, Wainwright MJ, Jordan MI (2009) High-dimensional support union recovery in multivariate regression. Advances in Neural Information Processing Systems 21 - Proceedings of the 2008 Conference pp 1217–1224

Obozinski G, Taskar B, Jordan MI (2010) Joint covariate selection and joint subspace selection for multiple classification problems. Statistics and Computing 20(2):231–252

Obozinski G, Wainwright MJ, Jordan MI (2011) Support union recovery in high-dimensional multivariate regression. Annals of Statistics 39(1):1–47

Odintsova VV, Rebattu V, Hagenbeek FA, Pool R, Beck JJ, Ehli EA, van Beijsterveldt CEM, Ligthart L, Willemsen G, de Geus EJC, Hottenga JJ, Boomsma DI, van Dongen J (2021) Predicting Complex Traits and Exposures From Polygenic Scores and Blood and Buccal DNA Methylation Profiles. Frontiers in Psychiatry 12(July):1–17

Panico S, Dello Iacovo R, Celentano E, Galasso R, Muti P, Salvatore M, Mancini M (1992) Progetto ATENA, A study on the etiology of major chronic diseases in women: Design, rationale and objectives. European Journal of Epidemiology 8(4):601–608

Phipson B, Maksimovic J, Oshlack A (2016) missMethyl: an R package for analyzing data from Illumina's HumanMethylation450 platform. Bioinformatics 32(2):286–288

Pinheiro J, Bates D (2006) Mixed-effects models in S and S-PLUS. Springer science & business media

R Core Team (2022) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria

Rauluseviciute I, Drabløs F, Rye MB (2020) DNA hypermethylation associated with upregulated gene expression in prostate cancer demonstrates the diversity of epigenetic regulation. BMC Medical Genomics 13(1):6

Reimand J, Isserlin R, Voisin V, Kucera M, Tannus-Lopes C, Rostamianfar A, Wadi L, Meyer M, Wong J, Xu C, Merico D, Bader GD (2019) Pathway enrichment analysis and visualization of omics data using g:Profiler, GSEA, Cytoscape and EnrichmentMap. Nature Protocols 14(2):482–517

Reinsel G (1984) Estimation and Prediction in a Multivariate Random Effects Generalized Linear Model. Journal of the American Statistical Association 79(386):406–414

Riboli E, Hunt K, Slimani N, Ferrari P, Norat T, Fahey M, Charrondière U, Hémon B, Casagrande C, Vignat J, Overvad K, Tjønneland A, Clavel-Chapelon F, Thiébaut A, Wahrendorf J, Boeing H, Trichopoulos D, Trichopoulou A, Vineis P, Palli D, Bueno-de Mesquita H, Peeters P, Lund E, Engeset D, González C, Barricarte A, Berglund G, Hallmans G, Day N, Key T, Kaaks R, Saracci R (2002) European Prospective Investigation into Cancer and Nutrition (EPIC): study populations and data collection. Public Health Nutrition 5(6b):1113–1124

Richard MA, Huan T, Ligthart S, Gondalia R, Jhun MA, Brody JA, Irvin MR, Marioni R, Shen J, Tsai PC, Montasser ME, Jia Y, Syme C, Salfati EL, Boerwinkle E, Guan W, Mosley TH, Bressler J, Morrison AC, Liu C, Mendelson MM, Uitterlinden AG, van Meurs JB, Franco OH, Zhang G, Li Y, Stewart JD, Bis JC, Psaty BM, Chen YDI, Kardia SL, Zhao W, Turner ST, Absher D, Aslibekyan S, Starr JM, McRae AF, Hou L, Just AC, Schwartz JD, Vokonas PS, Menni C, Spector TD, Shuldiner A, Damcott CM, Rotter JI, Palmas W, Liu Y, Paus T, Horvath S, O'Connell JR, Guo X, Pausova Z, Assimes TL, Sotoodehnia N, Smith JA, Arnett DK, Deary IJ, Baccarelli AA, Bell JT, Whitsel E, Dehghan A, Levy D, Fornage M, Heijmans BT, 't Hoen PA, van Meurs J, Isaacs A, Jansen R, Franke L, Boomsma DI, Pool R, van Dongen J, Hottenga JJ, van Greevenbroek MM, Stehouwer CD, van der Kallen CJ, Schalkwijk CG, Wijmenga C, Zhernakova A, Tigchelaar EF, Slagboom PE, Beekman M, Deelen J, van Heemst D, Veldink JH, van den Berg LH, van Duijn CM, Hofman A, Uitterlinden AG, Jhamai PM, Verbiest M, Suchiman HED, Verkerk M, van der Breggen R, van Rooij J, Lakenberg N, Mei H, van

Iterson M, van Galen M, Bot J, van 't Hof P, Deelen P, Nooren I, Moed M, Vermaat M, Zhernakova DV, Luijk R, Bonder MJ, van Dijk F, Arindrarto W, Kielbasa SM, Swertz MA, van Zwet EW (2017) DNA Methylation Analysis Identifies Loci for Blood Pressure Regulation. The American Journal of Human Genetics 101(6):888–902

Rodosthenous T, Shahrezaei V, Evangelou M (2020) Integrating multi-OMICS data through sparse canonical correlation analysis for the prediction of complex traits: a comparison study. Bioinformatics 36(17):4616–4625

Rohart F, San Cristobal M, Laurent B (2014) Selection of fixed effects in high dimensional linear mixed models using a multicycle ECM algorithm. Computational Statistics & Data Analysis 80:209–222

Schafer JL, Yucel RM (2002) Computational Strategies for Multivariate Linear Mixed-Effects Models With Missing Values. Journal of Computational and Graphical Statistics 11(2):437–457

Schelldorfer J, Bühlmann P, De Geer SV (2011) Estimation for High-Dimensional Linear Mixed-Effects Models Using L1-Penalization. Scandinavian Journal of Statistics 38(2):197–214

Schwarz G (1978) Estimating the dimension of a model. The Annals of Statistics 6(2):461–464

Senoner T, Dichtl W (2019) Oxidative stress in cardiovascular diseases: Still a therapeutic target? Nutrients 11(9)

Shah A, Laird N, Schoenfeld D (1997) A Random-Effects Model for Multiple Characteristics with Possibly Missing Data. Journal of the American Statistical Association 92(438):775–779

Sigrist F (2022) Latent Gaussian Model Boosting. IEEE Transactions on Pattern Analysis and Machine Intelligence pp 1–1

Sill M, Hielscher T, Becker N, Zucknick M (2014) c060 : Extended Inference with Lasso and Elastic-Net Regularized Cox and Generalized Linear Models. Journal of Statistical Software 62(5)

Simon N, Friedman J, Hastie T, Tibshirani R (2013) A Sparse-Group Lasso. Journal of Computational and Graphical Statistics 22(2):231–245

Singal R, Ginder GD (1999) DNA Methylation. Blood 93(12):4059–4070

Stevenson AJ, McCartney DL, Hillary RF, Campbell A, Morris SW, Bermingham ML, Walker RM, Evans KL, Boutin TS, Hayward C, McRae AF, McColl BW, Spires-Jones TL, McIntosh AM, Deary IJ, Marioni RE (2020) Characterisation of an inflammation-related epigenetic score and its association with cognitive ability. Clinical Epigenetics 12(1):113

Subramanian A, Tamayo P, Mootha VK, Mukherjee S, Ebert BL, Gillette MA, Paulovich A, Pomeroy SL, Golub TR, Lander ES, Mesirov JP (2005) Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. Proceedings of the National Academy of Sciences 102(43):15545–15550

Tay JK, Narasimhan B, Hastie T (2021) Elastic Net Regularization Paths for All Generalized Linear Models

Tibshirani R (1996) Regression Shrinkage and Selection Via the Lasso. Journal of the Royal Statistical Society: Series B (Methodological) 58(1):267–288

Tyler AL, Crawford DC, Pendergrass SA (2013) Detecting and characterizing pleiotropy: new methods for uncovering the connection between the complexity of genomic architecture and multiple phenotypes. In: Biocomputing 2014, WORLD SCIENTIFIC, pp 183–187

Vinga S (2021) Structured sparsity regularization for analyzing high-dimensional omics data. Briefings in Bioinformatics 22(1):77–87

Witten DM, Tibshirani RJ (2009) Extensions of Sparse Canonical Correlation Analysis with Applications to Genomic Data. Statistical Applications in Genetics and Molecular Biology 8(1):1–27

Witten DM, Tibshirani R, Hastie T (2009) A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. Biostatistics 10(3):515–534

Wu CY, Hu HY, Chou YJ, Huang N, Chou YC, Li CP (2015) High Blood Pressure and All-Cause and Cardiovascular Disease Mortalities in Community-Dwelling Older Adults. Medicine 94(47):e2160

Yi Y, Fang Y, Wu K, Liu Y, Zhang W (2020) Comprehensive gene and pathway analysis of cervical cancer progression. Oncology Letters 19(4):3316–3332

Yuan M, 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 T, Edelmann D, Fan Z, Alwers E, Kather JN, Brenner H, Hoffmeister M (2022) Machine learning in the identification of prognostic DNA methylation biomarkers among patients with cancer: a systematic review of epigenome-wide studies. medRxiv

Zhang Y, Elgizouli M, Schöttker B, Holleczek B, Nieters A, Brenner H (2016) Smoking-associated DNA methylation markers predict lung cancer incidence. Clinical Epigenetics 8(1):1–12

Zhao Z, Zucknick M (2020) Structured penalized regression for drug sensitivity prediction. Journal of the Royal Statistical Society: Series C (Applied Statistics) 69(3):525–545

Zhao Z, Banterle M, Bottolo L, Richardson S, Lewin A, Zucknick M (2021a) BayesSUR : An R Package for High-Dimensional Multivariate Bayesian Variable and Covariance Selection in Linear Regression . Journal of Statistical Software 100(11)

Zhao Z, Banterle M, Lewin A, Zucknick M (2021b) Structured Bayesian variable selection for multiple related response variables and high-dimensional predictors. arXiv preprint arXiv:210105899 pp 1–33

Zhao Z, Wang S, Zucknick M, Aittokallio T (2022) Tissue-specific identification of multi-omics features for pan-cancer drug response prediction. iScience 25(8):104767

Zhong J, Agha G, Baccarelli AA (2016) The Role of DNA Methylation in Cardiovascular Risk and Disease. Circulation Research 118(1):119–131

Zhong W, Wang J, Chen X (2021) Censored mean variance sure independence screening for ultrahigh dimensional survival data. Computational Statistics & Data Analysis 159:107206

Zou H, Hastie T (2005) Regularization and variable selection via the elastic net. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 67(5):768–768