Bayesian Shrinkage Estimation of High Dimensional Causal Mediation Effects in Omics Studies

Yanyi Song¹, Xiang Zhou^{1,*}, Min Zhang¹, Wei Zhao², Yongmei Liu³, Sharon L. R. Kardia²,
Ana V. Diez Roux⁴, Belinda L. Needham², Jennifer A. Smith², and Bhramar Mukherjee^{1,**}

¹Department of Biostatistics, University of Michigan, Ann Arbor, MI, U.S.A.

²Department of Epidemiology, University of Michigan, Ann Arbor, MI, U.S.A.

³Department of Epidemiology and Prevention, Wake Forest School of Medicine, Winston-Salem, NC, U.S.A.

⁴Department of Epidemiology and Biostatistics, Drexel University, Philadelphia, PA, U.S.A.

*email: xzhousph@umich.edu

**email: bhramar@umich.edu

Abstract

Causal mediation analysis aims to examine the role of a mediator or a group of mediators that lie in the pathway between an exposure and an outcome. Recent biomedical studies often involve a large number of potential mediators based on high-throughput technologies. Most of the current analytic methods focus on settings with one or a moderate number of potential mediators. With the expanding growth of omics data, joint analysis of molecular-level genomics data with epidemiological data through mediation analysis is becoming more common. However, such joint analysis requires methods that can simultaneously accommodate high-dimensional mediators and that are currently lacking. To address this problem, we develop a Bayesian inference method using continuous shrinkage priors to extend previous causal mediation analysis techniques to a high-dimensional setting. Simulations demonstrate that our method improves the power of global mediation analysis compared to simpler alternatives and has decent performance to identify true non-null mediators. We also construct tests for natural indirect effects using a permutation procedure. The Bayesian method helps us to understand the structure of the composite null hypotheses. We applied our method to Multi-Ethnic Study of Atherosclerosis (MESA) and identified DNA methylation regions that may actively mediate the effect of socioeconomic status

(SES) on cardiometabolic outcome.

1 Introduction

Causal mediation analysis has been of great interest across many disciplines [1, 2]. It investigates how an intermediate variable, referred to as mediator, explains the mechanism through which the exposure variable affects the outcome. Under certain regularity conditions, mediation analysis allows us to disentangle the exposure’s effect into two parts: effect that acts through the mediator of interest (indirect/mediation effect) and effect that is unexplained by the mediator (direct effect). The state-of-the-art causal mediation analysis [3, 4], which is built upon the counterfactual framework [5, 6], establishes rigorous assumptions regarding the exposure-outcome, exposure-mediator and mediator-outcome relationships to justify appropriate use of the classical formulas from Baron and Kenny in the linear regression setting [7, 8] and creates a framework for other general extensions. Many of the existing methods focus on univariate mediator analysis that analyzes one mediator at a time in the causal inference framework, and are applicable to both continuous [9] and binary outcomes [10], and also can account for exposure-mediator interactions [11]. These methods have been widely applied in areas of social, economic, epidemiological and genetic studies [4, 8], including recent extensions to multiple exposure variables that lead to more powerful single nucleotide polymorphism (SNP) set tests in presence of gene expression data [12]. Several studies have recently extended mediation analysis models to jointly account for multiple mediators. However, most of the literature considered settings with two or three mediators, where each mediator is ordered along a priori known mediation pathways and the path-specific effects are estimated [13, 14]. In the presence of multiple unordered mediators, one often has to rely on an *ad hoc* approach to fit a series of mediation models with one mediator and one exposure [15, 16]/outcome [17] at a time and then summarize the mediation effects across all the mediators. Such approach ignores correlation among mediators and the estimated mediation effect does not necessarily have an intuitive interpretation, particularly when the dimension of the potential mediators is truly large.

In this paper, building on the potential outcome framework for causal inference, we de-

velop a Bayesian mediation analysis method in the presence of high-dimensional mediators. Bayesian methods for mediation have primarily been proposed in a principal stratification framework [18], in which the exposure effects on outcome are defined conditional on a single mediator. For estimating natural direct and indirect effects, recent work applied Bayesian non-parametric models, especially Dirichlet process mixture models [19, 20] in multiple mediators analysis. In contrast, here, we rely on Bayesian variable selection models to simultaneously analyze a relatively large number of mediators with potentially a small number being truly active. With sparsity inducing priors on mediator effects, we assume that only a small proportion of mediators may mediate the exposure effect on the outcome. This sparsity assumption allows us to extend previous univariate mediator analysis to a high-dimensional setting by casting the identification of active mediators as a variable selection problem and applying Bayesian methods with continuous shrinkage priors on the effects. Unlike previous methods developed for multiple mediators analysis, ours can simultaneously analyze much larger number of mediators without making any path-specific or causal ordering assumptions on mediators. Our method enables us to identify both the indirect effect of a specific mediator and the joint indirect effects of all the mediators, and propagates uncertainty in inference in a principled way.

While our method is generally applicable to many settings, we examine the performance of our method in the setting of genomics studies. Due to fast advances in high-throughput biological technologies, genomics studies can nowadays measure a large number of molecular-level traits such as gene expression and DNA methylation (DNAm) levels. Recent studies have proposed these molecular traits may act as a mechanism through which various aspects of socioeconomic status (SES) and neighborhood disadvantages affect physical health. For example, childhood SES, adult SES, social mobility, and neighborhood crime rates have recently been shown to influence DNAm in several genes related to stress and inflammation [21, 22]. DNAm of inflammatory markers have also been associated with the status of cardiovascular disease (CVD) [23] and type 2 diabetes (T2D) [24]. Here, we show through simulations and data analysis that our high-dimensional mediation analysis framework can increase power of a joint analysis and facilitate the identification of individual mediators.

2 Notation, Definitions and Assumptions

In this paper, we focus on causal mediation analysis for the setting where there is a single exposure of interest but a high-dimensional set of candidate mediators that may mediate the effect of exposure on an outcome. Suppose our analysis is based on a study of n subjects and for subject i , $i = 1, \dots, n$, we collect data on exposure A_i , p candidate mediators $\mathbf{M}_i = (M_i^{(1)}, M_i^{(2)}, \dots, M_i^{(p)})^T$, outcome Y_i , and q covariates $\mathbf{C}_i = (C_i^{(1)}, \dots, C_i^{(q)})^T$. In particular, we focus on the case where Y_i and $\mathbf{M}_i$ are all quantitative variables.

We adopt the counterfactual (or potential outcomes) framework to formally define mediators and their causal effects. Let the vector $\mathbf{M}_i(a) = (M_i^{(1)}(a), M_i^{(2)}(a), \dots, M_i^{(p)}(a))$ denote the i th subject's potential (or counterfactual) value of the p mediators if, possibly contrary to fact, he/she received exposure a . Let $Y_i(a, \mathbf{m})$ denote the i th subject's potential outcome if the subject's exposure were set to a and mediators were set to $\mathbf{m}$. These potential (or counterfactual) variables are hypothetical variables and may not be observed in real data. To connect potential variables to observed data, we make the Stable Unit Treatment Value Assumption (SUTVA) [25, 26], which is a commonly made assumption for performing causal inference. Specifically, the SUTVA assumes there is no interference between subjects and the consistency assumption, which states that the observed variables are the same as the corresponding potential variables when their determinants are set to the observed, i.e., $\mathbf{M}_i = \mathbf{M}_i(a = A_i)$, and $Y_i = Y_i(a = A_i, \mathbf{m} = \mathbf{M}_i)$. For simplicity in notation, we define $Y_i(a) = Y_i(a, \mathbf{M}_i(a))$, i.e., the potential outcome when exposure were set to a and mediators were set to the value that would have been observed had the exposure were set to a . Although potential or counterfactual variables are useful concepts in order to formally define causal effects, they are hypothetical and actually most of them are not observed in real data. For example, if $A_i \neq a$, then $Y_i(a)$ or $M_i(a)$ are not observed. Also $Y_i(a)$ and $Y_i(a^*)$ are never simultaneously observed for a subject.

We may decompose the effect of an exposure into its direct effect and effect mediated through mediators. The controlled direct effect (CDE) of the exposure on the outcome is defined as $Y_i(a, \mathbf{m}) - Y_i(a^*, \mathbf{m})$, which is the effect of changing exposure from level a^* (the reference

level) to a while hypothetically controlling mediators at level $\mathbf{m}$. The natural direct effect (NDE) is defined as $Y_i(a, \mathbf{M}_i(a^*)) - Y_i(a^*, \mathbf{M}_i(a^*))$, which is the CDE when mediators are controlled at the level that would have naturally been had the exposure been a^* . The natural indirect effect (NIE) is defined by $Y_i(a, \mathbf{M}_i(a)) - Y_i(a, \mathbf{M}_i(a^*))$, capturing the part of exposure effect mediated through mediators, i.e., the change in potential outcomes when mediators change from $\mathbf{M}_i(a^*)$ to $\mathbf{M}_i(a)$ while fixing exposure at a . The total effect (TE), $Y_i(a) - Y_i(a^*)$, can then be decomposed into natural direct effect and natural indirect effect, written as $Y_i(a) - Y_i(a^*) = Y_i(a, \mathbf{M}_i(a)) - Y_i(a^*, \mathbf{M}_i(a^*)) = Y_i(a, \mathbf{M}_i(a)) - Y_i(a, \mathbf{M}_i(a^*)) + Y_i(a, \mathbf{M}_i(a^*)) - Y_i(a^*, \mathbf{M}_i(a^*)) = \text{NIE} + \text{NDE}$.

Causal effects are formally defined in terms of potential variables which are not necessarily observed, but the identification of causal effects must be based on observed data. Therefore, similar to missing data problems, further assumptions regarding the confounders are required for the identification of causal effects in mediation analysis [15]. We will use $A \perp\!\!\!\perp B | C$ to denote that A is independent of B conditional on C . For estimating the average CDE, two assumptions on confounding are needed: (1) $Y_i(a, \mathbf{m}) \perp\!\!\!\perp A_i | \mathbf{C}_i$, namely, there is no unmeasured confounding for the exposure effect on the outcome; (2) $Y_i(a, \mathbf{m}) \perp\!\!\!\perp \mathbf{M}_i | \{\mathbf{C}_i, A_i\}$, namely, there is no unmeasured confounding for any of mediator-outcome relationship after controlling for the exposure. The two assumptions are illustrated in the left panel of Figure 1, and controlling for exposure-outcome and mediator-outcome confounding corresponds to controlling for C_1, C_2 in the figure. In practice, both sets of covariates C_1 and C_2 need not to be distinguished from one another and can simply be included in the overall set of C that we adjust for. The identification of the average NDE and NIE requires assumption (1) and (2), along with two additional assumptions: (3) $\mathbf{M}_i(a) \perp\!\!\!\perp A_i | \mathbf{C}_i$, namely, there is no unmeasured confounding for the exposure effect on all the mediators; (4) $Y_i(a, \mathbf{m}) \perp\!\!\!\perp \mathbf{M}_i(a^*) | \mathbf{C}_i$, which can be interpreted as there is no downstream effect of the exposure that confounds the mediator-outcome relationship for any of the mediators. Graphically, assumption (4) implies that there should be no arrow going from exposure A to mediator-outcome confounder C_2 in Figure 1(a). It is thus violated in Figure 1(b) since the mediator-outcome confounder L is itself affected by the exposure. The four assumptions are required to hold

with respect to the whole set of mediators $\mathbf{M}_i(a) = (M_i^{(1)}(a), M_i^{(2)}(a), \dots, M_i^{(p)}(a))$. Finally, as in all mediation analysis, in order for associations to represent causal effects, the temporal ordering assumption also needs to be satisfied, i.e., the exposure precedes the mediators and the mediators precede the outcome.

Figure 1: Left (a): High-dimensional mediators ($(M^{(1)}, M^{(2)}, \dots, M^{(p)})$) between exposure (A) and outcome (Y) with exposure-outcome confounders C_1 and mediator-outcome confounders C_2 ; Right (b): An example of mediator-outcome confounder L that is affected by the exposure A .

Now we show that if the above assumptions hold, then the average natural direct and indirect effects can be identified from the observed data. We first notice that $E[Y_i(a, \mathbf{M}_i(a^*)) | \mathbf{C}_i]$ can be expressed as below (see Supplementary Materials for details),

$$E[Y_i(a, \mathbf{M}_i(a^*)) | \mathbf{C}_i] = \int_{\mathbf{m}} E(Y_i | a, \mathbf{m}, \mathbf{C}_i) P(\mathbf{M}_i = \mathbf{m} | \mathbf{C}_i, a^*) d\mathbf{m} \quad (1)$$

If we replace a with a^* in $E[Y_i(a, \mathbf{M}_i(a^*)) | \mathbf{C}_i]$, then we get $E[Y_i(a^*, \mathbf{M}_i(a^*)) | \mathbf{C}_i] = \int_{\mathbf{m}} E(Y_i | a^*, \mathbf{m}, \mathbf{C}_i) * P(\mathbf{M}_i = \mathbf{m} | \mathbf{C}_i, a^*) d\mathbf{m}$. Therefore, we can express the average natural direct effect conditional on C as,

$$\begin{aligned} & E[Y_i(a, \mathbf{M}_i(a^*)) - Y_i(a^*, \mathbf{M}_i(a^*)) | \mathbf{C}_i] \\ &= \int_{\mathbf{m}} \{E(Y_i | a, \mathbf{m}, \mathbf{C}_i) - E(Y_i | a^*, \mathbf{m}, \mathbf{C}_i)\} P(\mathbf{M}_i = \mathbf{m} | \mathbf{C}_i, a^*) d\mathbf{m}. \end{aligned} \quad (2)$$

If we replace a^* with a in $E[Y_i(a, \mathbf{M}_i(a^*)) | \mathbf{C}_i]$, then we get $E[Y_i(a, \mathbf{M}_i(a)) | \mathbf{C}_i] = \int_{\mathbf{m}} E(Y_i | a, \mathbf{m}, \mathbf{C}_i) *$

$P(\mathbf{M}_i = \mathbf{m} | \mathbf{C}_i, a) d\mathbf{m}$, and thus the average indirect effect conditional on C is given by,

$$\begin{aligned} & E[Y_i(a, \mathbf{M}_i(a)) - Y_i(a, \mathbf{M}_i(a^*)) | \mathbf{C}_i] \\ &= \int_{\mathbf{m}} E(Y_i | a, \mathbf{m}, \mathbf{C}_i) \{P(\mathbf{M}_i = \mathbf{m} | \mathbf{C}_i, a) - P(\mathbf{M}_i = \mathbf{m} | \mathbf{C}_i, a^*)\} d\mathbf{m}. \end{aligned} \quad (3)$$

Finally, one can get the average NDE and NIE by taking expectation over C of the two conditional effects defined in (2) and (3). Importantly, Equations (2), (3) show that, under the assumptions we made, the average NDE and the average NIE can be identified by modeling $Y_i | A_i, \mathbf{M}_i, \mathbf{C}_i$ and $\mathbf{M}_i | A_i, \mathbf{C}_i$ using observed data.

3 Models and Estimands

As discussed in Section 2, effects of mediators (average NDE and NIE) defined in terms of potential outcomes can be deduced from two conditional models for $Y_i | A_i, \mathbf{M}_i, \mathbf{C}_i$ and $\mathbf{M}_i | A_i, \mathbf{C}_i$ using observed data. Therefore, we propose two regression models for the two conditional relationships and subsequently deduce the causal effects of mediators. For modeling $Y_i | A_i, \mathbf{M}_i, \mathbf{C}_i$, we assume for subject i ($i = 1, \dots, n$), a continuous outcome of interest Y_i is associated with exposure A_i , p potential mediators $\mathbf{M}_i = (M_i^{(1)}, M_i^{(2)}, \dots, M_i^{(p)})^T$ that may be on the pathway from A_i to Y_i , and q covariates $\mathbf{C}_i$ with the first element being the scalar 1 for the intercept:

$$Y_i = \mathbf{M}_i^T \boldsymbol{\beta}_m + A_i \beta_a + \mathbf{C}_i^T \boldsymbol{\beta}_c + \epsilon_{Y_i} \quad (4)$$

where $\boldsymbol{\beta}_m = (\beta_{m1}, \dots, \beta_{mp})^T$, $\boldsymbol{\beta}_c = (\beta_{c1}, \dots, \beta_{cq})^T$, $\epsilon_{Y_i} \sim N(0, \sigma_e^2)$. Here we assume there is no interaction between A_i and $\mathbf{M}_i$. Next for modeling $\mathbf{M}_i | A_i, \mathbf{C}_i$ we consider a multivariate regression model that jointly analyzes the p potential mediators:

$$\mathbf{M}_i = A_i \boldsymbol{\alpha}_a + \boldsymbol{\alpha}_c \mathbf{C}_i + \boldsymbol{\epsilon}_{M_i} \quad (5)$$

where $\boldsymbol{\alpha}_a = (\alpha_{a1}, \dots, \alpha_{ap})^T$, $\boldsymbol{\alpha}_c = (\boldsymbol{\alpha}_{c1}^T, \dots, \boldsymbol{\alpha}_{cp}^T)^T$, $\boldsymbol{\alpha}_{c1}, \dots, \boldsymbol{\alpha}_{cp}$ are q -by-1 vectors, $\boldsymbol{\epsilon}_{M_i} \sim MVN(\mathbf{0}, \boldsymbol{\Sigma})$, $\boldsymbol{\Sigma}$ captures the correlation among the mediators. ϵ_{Y_i} and $\boldsymbol{\epsilon}_{M_i}$ are assumed independent of A_i , $\mathbf{C}_i$ and each other.

With assumptions made in Section 2 and under the regression models specified for the outcome $E(Y_i|A_i, \mathbf{M}_i, \mathbf{C}_i)$ and for the mediators $P(\mathbf{M}_i|A_i, \mathbf{C}_i)$, we can analytically calculate the right-hand side of Equations (2), (3). We show in Supplementary Materials that the average NDE, NIE and TE can then be computed as below, and in the rest of the paper, we refer to NDE as direct effect and NIE as indirect/mediation effect.

$$\text{NDE: } E[Y_i(a, \mathbf{M}_i(a^*)) - Y_i(a^*, \mathbf{M}_i(a^*))|\mathbf{C}_i] = \beta_a(a - a^*). \quad (6)$$

$$\text{NIE: } E[Y_i(a, \mathbf{M}_i(a)) - Y_i(a, \mathbf{M}_i(a^*))|\mathbf{C}_i] = (a - a^*) \sum_{j=1}^p (\boldsymbol{\alpha}_a)_j (\boldsymbol{\beta}_m)_j. \quad (7)$$

$$\text{TE: } E[Y_i(a) - Y_i(a^*)|\mathbf{C}_i] = (\beta_a + \boldsymbol{\alpha}_a^T \boldsymbol{\beta}_m)(a - a^*). \quad (8)$$

According to Equation (7), to select active mediators among the potential high-dimensional set of mediators is equivalent to identify the ones with marginal indirect effect $(\boldsymbol{\alpha}_a)_j (\boldsymbol{\beta}_m)_j$ being non-zero. On the other hand, any inactive (un-selected) mediator will naturally fall into one of the following three categories: $(\boldsymbol{\beta}_m)_j$ is non-zero while $(\boldsymbol{\alpha}_a)_j$ is zero; $(\boldsymbol{\alpha}_a)_j$ is non-zero while $(\boldsymbol{\beta}_m)_j$ is zero; both are zero. Such a refined partition for the high-dimensional set of mediators provides useful and insightful interpretations on the way in which a mediator links or does not link exposure to outcome, and furthermore facilitates understanding the composite of non-mediating cases.

Regarding a global measure of the indirect effects, we note that the quantity in Equation (7), summation of each mediator's marginal mediation effect, is a good summary of the global mediation effects when the marginal mediation effect for each mediator is of the same direction. However, when marginal mediation effects have opposite directions, their effects may cancel out and result in a small or zero indirect effect. Considering this, we propose to use the L_2 norm of the vector of marginal mediation effects [17] as a global measure of

mediation effects, i.e.,

$$\begin{aligned}\tau &= ||((\alpha_a)_1(\beta_m)_1, (\alpha_a)_1(\beta_m)_2, \dots, (\alpha_a)_p(\beta_m)_p)||^2 \\ &= \sum_{j=1}^p \{(\alpha_a)_j(\beta_m)_j\}^2.\end{aligned}\tag{9}$$

Such quadratic quantity, which involves a sum of squared terms, has been widely studied for testing associations between rare genetic variants and complex human traits, and shown to have reasonable power across a wide range of alternatives with correlated predictor set [27].

4 Bayesian Method for Estimation

4.1 Prior Specification

In order to conduct high-dimensional mediation analysis, we need to make certain model assumptions on the effect sizes. In genome-wide association studies, Bayesian sparse regression models, such as Bayesian variable selection regression models (BVS), have been proven to yield better power in detecting relevant covariates [28]. For high-dimensional mediation analysis, we also make the reasonable sparsity assumption, which implies that only a small proportion of mediators mediate the exposure effects on the outcome. In practice, the exposure effects on the mediators and the mediator effects on the outcome may be small but not exactly zero. Linear mixed models (LMM), on the other hand, assume that every mediator transmits certain effects from exposure to outcome, with the effect sizes normally distributed. Therefore, in this paper, we propose Bayesian Sparse Linear Mixed Model (BSLMM), a hybrid between LMM and BVS [29] that imposes continuous shrinkage on the effects, for high-dimensional mediation analysis. The BSLMM is capable of learning the underlying mediation architecture from the data, producing good performances across a wide range of scenarios. Specifically, we assume a mixture of two normal components a

priori for the j th mediator, $j = 1, 2, \dots, p$,

$$\begin{aligned}(\boldsymbol{\beta}_m)_j &\sim \pi_m N(0, \sigma_{m1}^2) + (1 - \pi_m) N(0, \sigma_{m0}^2) \\ (\boldsymbol{\alpha}_a)_j &\sim \pi_a N(0, \sigma_{ma1}^2) + (1 - \pi_a) N(0, \sigma_{ma0}^2)\end{aligned}$$

where $\sigma_{m1}^2 > \sigma_{m0}^2$, $\sigma_{ma1}^2 > \sigma_{ma0}^2$, and π_m, π_a denote the proportion of coefficients that belong to the normal distribution with a larger variance.

For the other coefficients, we assume,

$$\beta_a \sim N(0, \sigma_a^2) \quad \text{and} \quad \beta_c, \alpha_c \sim MVN(\mathbf{0}, \sigma_c^2 \mathbf{I}), \sigma_c^2 \rightarrow \infty$$

Here we use a limiting normal prior for β_c, α_c with its variance going to infinity, since we often have sufficient information from the data to overwhelm any prior assumptions. For the convenience of modeling, we set the correlation structure among mediators $\boldsymbol{\Sigma}$ as $\sigma_g^2 \mathbf{I}$.

For the hyper-parameters of variances in the model, we use the standard conjugate priors,

$$\begin{aligned}\sigma_{ms}^2 &\sim \text{inverse-gamma}(k_{ms}, l_{ms}), s = 0, 1 \\ \sigma_a^2 &\sim \text{inverse-gamma}(k_a, l_a) \\ \sigma_{mas}^2 &\sim \text{inverse-gamma}(k_{mas}, l_{mas}), s = 0, 1 \\ \sigma_e^2 &\sim \text{inverse-gamma}(k_e, l_e) \\ \sigma_g^2 &\sim \text{inverse-gamma}(k_g, l_g)\end{aligned}$$

We set $k_{m0} = k_{m1} = k_a = k_{ma0} = k_{ma1} = k_e = k_g = 2.0$, and $l_{m0} = l_{ma0} = 10^{-4}, l_a = l_{m1} = l_{ma1} = l_e = l_g = 1.0$. Following [29], we place a uniform prior on $\log(\pi_m), \log(\pi_a), :$

$$\log(\pi_m), \log(\pi_a) \sim U(\log(10/p), \log(1))$$

where p is the number of mediators. The priors were chosen so that π_m and π_a range from $10/p$ to 1, and the lower and upper bounds correspond to an expectation of 10 and p

covariates in each model. A uniform prior on $\log(\pi_m)$ and $\log(\pi_a)$ reflects the fact that the uncertainty in π_m, π_a spans orders of magnitude.

4.2 Posterior Sampling Algorithm

We develop a Markov chain Monte Carlo (MCMC) sampling algorithm to obtain the posterior samples from our Bayesian method. To facilitate MCMC, we introduce indicator variables $\mathbf{r}_m, \mathbf{r}_a \in \{0, 1\}^p$ to indicate which normal component $(\beta_m)_j$ and $(\alpha_a)_j$ come from, and for the j th mediator, $r_{mj} = I((\beta_m)_j \sim N(0, \sigma_{m1}^2)), r_{aj} = I((\alpha_a)_j \sim N(0, \sigma_{ma1}^2))$. Let $\theta_1 = (\beta_m, \beta_a, \pi_m, \mathbf{r}_m, \sigma_{m1}^2, \sigma_{m0}^2, \sigma_a^2, \sigma_e^2)$ denote all the unknown parameters in model (4), and $\theta_2 = (\alpha_a, \pi_a, \mathbf{r}_a, \sigma_{ma1}^2, \sigma_{ma0}^2, \sigma_g^2)$ for model (5). The joint log posterior distribution is,

$$\begin{aligned} & \log P(\theta_1, \theta_2 | (Y_i, \mathbf{M}_i, A_i)_{i=1}^n) \\ &= \sum_{i=1}^n \log P(Y_i | \theta_1, A_i, \mathbf{M}_i) + \sum_{i=1}^n \log P(\mathbf{M}_i | \theta_2, A_i) + \log P(\theta_1) + \log P(\theta_2) \end{aligned}$$

We use a Hastings-within-Gibbs algorithm to obtain posterior samples, and full details of the sampling algorithm appear in Supplementary Materials.

For the j th mediator, we can estimate the posterior probability of both $(\beta_m)_j$ and $(\alpha_a)_j$ being in the normal components with larger variances as the posterior inclusion probability (PIP), defined as $P(r_{mj} = 1, r_{aj} = 1 | \text{Data})$ in our model. The PIP estimated in this way measures the association strength between exposure and mediators in model (5) and between mediators and outcome in model (4). Therefore, we select mediators with the highest PIP as the potentially active mediators.

4.3 Mediator Categorization

Under the above Bayesian mediation framework, active mediators are the ones whose $(\beta_m)_j$ and $(\alpha_a)_j$ both come from larger normal components. The three categories for the inactive mediators are: $(\beta_m)_j$ from larger normal component while $(\alpha_a)_j$ from smaller normal component; $(\alpha_a)_j$ from larger normal component while $(\beta_m)_j$ from smaller normal component; both from smaller components. In addition to identifying true mediators, our method

	$(\beta_m)_j$	Large component	Small component
$(\alpha_a)_j$			
Large component		$r_{mj} * r_{aj} = 1$ (Group 1)	$r_{mj} = 0, r_{aj} = 1$ (Group 2)
Small component		$r_{mj} = 1, r_{aj} = 0$ (Group 3)	$r_{mj} = r_{aj} = 0$ (Group 4)

Table 1: Mediators are categorized into four groups based on their relationships with exposure and outcome; Group 1: Both $(\beta_m)_j$ and $(\alpha_a)_j$ come from larger normal components; Group 2: $(\alpha_a)_j$ from larger normal component while $(\beta_m)_j$ from smaller normal component; Group 3: $(\beta_m)_j$ from larger normal component while $(\alpha_a)_j$ from smaller normal component; Group 4: Both $(\beta_m)_j$ and $(\alpha_a)_j$ come from smaller normal components.

automatically classifies all the mediators into four groups based on their relationship with exposure and outcome. In practice, we have the indicator variables r_{mj} and r_{aj} to denote which component the coefficients $(\beta_m)_j$, $(\alpha_a)_j$ belong to and can easily obtain the posterior probabilities for each group. The four groups are illustrated in Table 1,

4.4 Global Test of Mediation Effects

Typically the majority of mediators are not actively mediating the exposure effect on the outcome, so it is natural to focus on the global null hypothesis and test for $H_0 : \tau = 0$,

$$\tau = \|((\alpha_a)_1(\beta_m)_1, (\alpha_a)_1(\beta_m)_2, \dots, (\alpha_a)_p(\beta_m)_p)\|^2 = \sum_{j=1}^p \{(\alpha_a)_j(\beta_m)_j\}^2 \quad (10)$$

One estimate for τ from our method is the sampling mean calculated from the posterior samples, $\hat{\tau} = \frac{1}{L} \sum_{i=1}^L \{\sum_{j=1}^p \{(\hat{\alpha}_a^{(i)})_j(\hat{\beta}_m^{(i)})_j\}^2\}$, where $\{(\hat{\alpha}_a^{(i)}, \hat{\beta}_m^{(i)}) : i = 1, \dots, L\}$ are the L samples generated from posterior distributions.

It is difficult to analytically derive the composite null distribution of τ because the null distribution depends on the proportion of mediators belonging to each of the four categories. Instead, we resort to the permutation method. The global null indicates that none of the mediators are active, and to construct the three null components for each mediator, the following permutation procedures are brought up: (a) permute $\{Y_i\}_{i=1}^n$ in Equation (4) to dissolve the relationship between outcome and mediators and obtain the null estimates of $(\beta_m)_j^{(null)}$, (b) permute $\{M_i\}_{i=1}^n$ in Equation (5) to dissolve the relationship between mediators and exposure and obtain the null estimates of $(\alpha_a)_j^{(null)}$. We then propose the following three quantities: (i) $\{(\beta_m)_j^{(null)}(\alpha_a)_j^{(original)}\}^2$, (ii) $\{(\beta_m)_j^{(original)}(\alpha_a)_j^{(null)}\}^2$, (iii)

$\{(\beta_m)_j^{(null)}(\alpha_a)_j^{(null)}\}^2$, where $(\beta_m)_j^{(original)}$ and $(\alpha_a)_j^{(original)}$ are estimated from the original data. For the j th mediator, (i), (ii) and (iii) correspond to the three components of the null distribution: $r_{aj} = 1, r_{mj} = 0$ (Group 2), $r_{mj} = 1, r_{aj} = 0$ (Group 3), $r_{mj} = 0, r_{aj} = 0$ (Group 4), respectively. Each quantity is then weighted by the posterior probability for that group estimated from original data, and we finally sum the results across all the mediators to obtain draws from the empirical null distribution of τ .

5 Simulations

We evaluate the performance of the proposed Bayesian mediation method and compare it with two other existing mediation methods in our simulations. The two existing methods include single mediation analysis and multivariate mediation analysis. Single mediation analysis tests one mediator at a time for its mediation effect. We use the R package **mediation** to run single mediation analysis with the nonparametric bootstrap option for standard error estimation. Multivariate mediation analysis [15], on the other hand, jointly analyzes all the mediators in both model (4) and (5) and tests the product term $(\beta_m)_j(\alpha_a)_j$ for each j at a time while controlling for all other variables. This method can only be fit when a multivariate ordinary least squares regression model can be fit for the outcome model (4). We implement the multivariate mediation analysis and compute the standard error based on first and second order Taylor series approximation of the product [30]: $se((\hat{\beta}_m)_j(\hat{\alpha}_a)_j) = \sqrt{(\hat{\beta}_m)_j^2 Var((\hat{\alpha}_a)_j) + (\hat{\alpha}_a)_j^2 Var((\hat{\beta}_m)_j) + Var((\hat{\alpha}_a)_j)Var((\hat{\beta}_m)_j)}$. Afterwards, we obtain a z -statistics for the j th mediator by dividing $(\hat{\beta}_m)_j(\hat{\alpha}_a)_j$ with its standard error and compute the corresponding p -value based on asymptotic normality. We use p -values for the two existing frequentist methods (univariate and multivariate) and PIP for our Bayesian method as measures of the evidence for mediation. We compare the power to identify active mediators based on either 5% or 10% false discovery rate (FDR).

We consider various simulation settings with $n = 1,000$ samples and p mediators ($p = 100$ or $2,000$). Since the multivariate mediation analysis can only be applied to settings where the number of mediators is smaller than the number of observations (i.e. $p < n$), we first examine the settings where $p = 100$ in order to include the multivariate mediation

analysis for comparison. We will later consider the high-dimensional setting where $p = 2,000$. For each simulation setting, we first simulate a set of continuous exposure variables $\{A_i, i = 1, \dots, 1000\}$ independently from a standard normal distribution. We then generate a p -vector of mediators for the i th individual from $\mathbf{M}_i = A_i \boldsymbol{\alpha}_a + \boldsymbol{\epsilon}_{Mi}$. Here, each element of $\boldsymbol{\alpha}_a$, $(\boldsymbol{\alpha}_a)_j$ ($j = 1, \dots, p$), is simulated from a point-normal prior: $\pi_a N(0, 1) + (1 - \pi_a) \delta_0$, where δ_0 is a point mass at zero. The residual errors $\boldsymbol{\epsilon}_{Mi}$ are simulated from a multivariate normal distribution with mean zero and a covariance $\boldsymbol{\Sigma}$. $\boldsymbol{\Sigma}$ accounts for the correlation among mediators commonly seen in real data, and we use the sample covariance estimated from the Multi-Ethnic Study of Atherosclerosis (MESA) data to serve as $\boldsymbol{\Sigma}$. Because our Bayesian mediation model does not explicitly account for the correlation structure of mediators in the model between mediators and exposure, the simulations with correlated mediators allow us to examine the robustness of our modeling assumption regarding independence. After simulating $A_i \boldsymbol{\alpha}_a$ and $\boldsymbol{\epsilon}_{Mi}$, we scale these two terms further so that the first term explains a fixed proportion of variance: $PVE_A = \text{Var}(A_i \boldsymbol{\alpha}_a) / \text{Var}(\mathbf{M}_i)$, where Var denotes the sample variance.

Given the exposure and mediators, we then generate the outcome Y_i from the linear model: $Y_i = \mathbf{M}_i^T \boldsymbol{\beta}_m + A_i \beta_a + \epsilon_{Yi}$. Here, each element of $\boldsymbol{\beta}_m$, $(\boldsymbol{\beta}_m)_j$ ($j = 1, \dots, p$), is simulated from $\pi_m N(0, 1) + (1 - \pi_m) \delta_0$, where δ_0 is a point mass at zero, and β_a from a standard normal distribution. The residual error ϵ_{Yi} is simulated independently from a standard normal distribution. We assume that only 10% of the mediators are truly mediating the exposure effects on the outcome (i.e. active mediators), whose $(\boldsymbol{\beta}_m)_j$ and $(\boldsymbol{\alpha}_a)_j$ are both sampled from the large variance normal distribution. After simulating $\mathbf{M}_i^T \boldsymbol{\beta}_m$, $A_i \beta_a$ and ϵ_{Yi} , we scale these three terms further to achieve two desirable PVE s: $PVE_{IE} = \text{Var}(\boldsymbol{\alpha}_a^T \boldsymbol{\beta}_m A_i) / \text{Var}(Y_i)$ and $PVE_{DE} = \text{Var}(A_i \beta_a) / \text{Var}(Y_i)$.

To explore a variety of simulation scenarios, we first examine a baseline scenario where we set $PVE_A = 0.5$, $PVE_{IE} = 0.4$, $PVE_{DE} = 0.1$, $\pi_a = 0.3$, $\pi_m = 0.2$. We then vary each of the four parameters ($PVE_A, PVE_{IE}, \pi_a, \pi_m$) one at a time in different scenarios to investigate their individual influences on the results. We perform 200 replicates for each simulation scenario to do the power comparison.

Figure 2: Power comparison among our Bayesian mediation method (yellow), multivariate mediation method (red) and single mediation method (orange) when the number of mediators is 100 and sample size 1,000. The x-axis marks the one parameter we change at a time from the baseline setting. The average TPR at FDR = 0.05/0.1 and its error bar are calculated across 200 replicates.

We first examine the settings for $p = 100$ and display the comparative results in Figure 2. The results show that our Bayesian multivariate mediation method outperforms both the univariate and multivariate mediation analysis methods in all scenarios. For example, in the baseline setting, at 10% FDR, Bayesian mediation method achieves a power of 0.725, while the univariate and multivariate methods achieve a power of 0.527 and 0.676, respectively. The power of the three approaches increases with increasing PVE_{IE} , which increases the effect sizes of β_m . In addition, the power of various approaches reduces with increased π_a or π_m , which reduces the effect sizes of either α_a or β_m , respectively. As one would expect, the advantage of our Bayesian method over the other two methods is more apparent in sparse settings with smaller values of π_a and π_m . In terms of PVE_A , which determines the effect size of α_a , we found that the power of different methods first increases slightly when PVE_A changes from 0.3 to 0.5 and then decreases slightly as PVE_A changes further to 0.8. The later decrease in power in the setting of $PVE_A = 0.8$ is presumably due to the increased correlation between the exposure and mediators, which makes it difficult for all the three methods to distinguish between direct effects and indirect effects in model (4). Between the two competing methods, the multivariate mediation analysis method yields better power than the single mediation analysis method in all scenarios, as the multivariate mediation analysis properly controls for the correlation among mediators.

Next, we examine the settings for $p = 2,000$. Now we select 1% of the mediators to be active and set $\pi_m = 2\%$, $\pi_a = 3\%$ as the baseline setting with all other configurations being same as in the baseline setting of $p = 100$. Since the multivariate mediation analysis is unfeasible when $p > n$, we compare our method with single mediation analysis alone. We use a threshold of 1% false positive rate (FPR) instead of false discovery rate due to low power in the $p = 2,000$ settings. The comparisons are shown in Figure 3. The Bayesian mediation method yields more power than the single mediator analysis in all scenarios. For example, in the baseline setting, at 1% FPR, Bayesian mediation method achieves a power of 0.470, while the univariate method has a power of 0.357. The power of the two approaches again increases with increasing PVE_{IE} and decreases with increasing π_a or π_m . The power of the Bayesian method decreases with increasing PVE_A , while the power of the univariate

Figure 3: Power comparison between our Bayesian mediation method (yellow) and single mediation method (orange) when the number of mediators is 2,000 and sample size 1,000. The x-axis marks the one parameter we change at a time from the baseline setting. The average TPR at FPR = 0.01 and its error bar are calculated across 200 replicates.

analysis method is relatively stable. In addition, our Bayesian method with continuous shrinkage is more powerful than the univariate method especially when the model is most sparse.

Finally, we examine the ability of our method to detect the overall mediation effect and estimate the proportion of mediators in the four different categories. The four categories are characterized by the effect of the mediator on the outcome and the effect of the exposure on the mediator as shown in Table 1. We use π_{g1} , π_{g2} , π_{g3} , π_{g4} to represent the proportion of mediators in Group 1, Group 2, Group 3 and Group 4, respectively. We examine eight different simulation scenarios based on different combinations of π_{g1} , π_{g2} , π_{g3} and π_{g4} , which include four null scenarios with $\pi_{g1} = 0$ and four alternative scenarios with $\pi_{g1} \neq 0$. In these simulations, we set PVE s to be the same as in the baseline setting ($PVE_A = 0.5$, $PVE_{IE} = 0.4$, $PVE_{DE} = 0.1$; except when $\pi_{g4} = 1$ where PVE_A and PVE_{IE} are zero).

To detect the overall mediation effect, we obtain the posterior mean of τ . For each simulated data, we apply the permutation procedure described in Section 4.4 with 100 permutations and compare the estimated $\hat{\tau}$ with the empirical null distribution. We find that the global test controls well the Type I error under four different null scenarios (Table 2). In addition, the global test yields reasonable power under the four different alternative scenarios (Table 3). In most scenarios, the power is close to or above 0.8. Finally, we estimate the proportion of mediators in each of the four different categories using posterior samples and find that our method provides decent estimates for π_{g1} , π_{g2} , π_{g3} and π_{g4} across different scenarios (Tables 2 and 3). Note that our estimates for π_{g1} , π_{g2} , π_{g3} are slightly conservative due to the fact that our model does not have full power to detect all the mediators.

6 Data Analysis

We applied the proposed Bayesian method to investigate the mediation mechanism of DNAm in the pathway from adult socioeconomic status (SES) to glycated hemoglobin (HbA1c) in the Multi-Ethnic Study of Atherosclerosis (MESA) [31]. The exposure, adult SES, is indicated by adult educational attainment and is an important risk factor for cardiovascular diseases. The outcome, HbA1c, is a long-term measurement of average blood glucose levels

p	Composition	π_{g1}	π_{g2}	π_{g3}	π_{g4}	Size	$\hat{\pi}_{g1}$	$\hat{\pi}_{g2}$	$\hat{\pi}_{g3}$	$\hat{\pi}_{g4}$	$\widehat{PVE}_{IE}$
100	A	0	0.2	0.1	0.7	0.005	0.003	0.195	0.074	0.728	0.397
	B	0	0.1	0.2	0.7	0.003	0.003	0.097	0.115	0.786	0.420
	C	0	0.1	0.1	0.8	0.005	0.001	0.098	0.076	0.824	0.399
	D	0	0	0	1	0.008	0.001	0.004	0.021	0.974	0.006
2,000	A	0	0.03	0.02	0.95	0.005	0.000	0.028	0.005	0.967	0.417
	B	0	0.1	0.02	0.88	0.006	0.000	0.098	0.007	0.895	0.429
	C	0	0.1	0.1	0.8	0.006	0.000	0.098	0.001	0.901	0.450
	D	0	0	0	1	0.010	0.000	0.000	0.000	1.000	0.001

Table 2: Empirical size of the proposed global test at level of 0.01 based on 1,000 simulations when $p = 100/2,000$. We denote π_{g1} , π_{g2} , π_{g3} and π_{g4} to represent the proportion of mediators in Group 1, Group 2, Group 3 and Group 4 as defined in Table 1, and $\hat{\pi}_{g1}$, $\hat{\pi}_{g2}$, $\hat{\pi}_{g3}$, $\hat{\pi}_{g4}$ are the estimated proportions from our Bayesian method. The true $PVE_{IE} = 0.4$, and $\widehat{PVE}_{IE}$ is the estimated value.

p	Composition	π_{g1}	π_{g2}	π_{g3}	π_{g4}	Power	$\hat{\pi}_{g1}$	$\hat{\pi}_{g2}$	$\hat{\pi}_{g3}$	$\hat{\pi}_{g4}$	$\widehat{PVE}_{IE}$
100	A	0.1	0.2	0.1	0.6	0.788	0.064	0.234	0.103	0.600	0.341
	B	0.1	0.1	0.2	0.6	0.780	0.058	0.140	0.145	0.657	0.375
	C	0.1	0.1	0.1	0.7	0.806	0.078	0.120	0.139	0.663	0.382
	D	0.1	0	0	0.9	0.823	0.041	0.066	0.068	0.823	0.335
2,000	A	0.01	0.02	0.01	0.96	0.773	0.001	0.030	0.001	0.968	0.404
	B	0.01	0.04	0.01	0.94	0.728	0.001	0.049	0.004	0.946	0.397
	C	0.01	0.09	0.09	0.81	0.604	0.001	0.097	0.010	0.892	0.403
	D	0.01	0	0	0.99	0.804	0.010	0.001	0.001	0.988	0.363

Table 3: Empirical power of the proposed global test at level of 0.01 based on 1,000 simulations when $p = 100/2,000$. We denote π_{g1} , π_{g2} , π_{g3} and π_{g4} to represent the proportion of mediators in Group 1, Group 2, Group 3 and Group 4 as defined in Table 1, and $\hat{\pi}_{g1}$, $\hat{\pi}_{g2}$, $\hat{\pi}_{g3}$, $\hat{\pi}_{g4}$ are the estimated proportions from our Bayesian method. The true $PVE_{IE} = 0.4$, and $\widehat{PVE}_{IE}$ is the estimated value.

and a critical variable for various diseases including T2D and CVD [32]. Thus, understanding how methylation at different CpG sites mediates the effects of adult SES on HbA1c can shed light on the molecular mechanisms of CVD. We provide our detailed processing steps for MESA data in the Supplementary Materials. Briefly, we selected 1,231 individuals with both adult SES and HbA1c measurements as well as DNA methylation profiles measured from purified monocytes. Due to computational reasons, we focused on a final set of 2,000 CpG sites that have the strongest marginal associations with adult SES for the following mediation analysis.

We applied both univariate mediation analysis and our Bayesian multivariate mediation analysis to analyze the selected 2,000 CpG sites. For the multivariate analysis, we consider

$$Y_i = \mathbf{M}_i^T \boldsymbol{\beta}_m + A_i \beta_a + \mathbf{C}_{2i}^T \boldsymbol{\beta}_c + \epsilon_{Yi} \quad (11)$$

$$\mathbf{M}_i = A_i \boldsymbol{\alpha}_a + \boldsymbol{\alpha}_c \mathbf{C}_{1i} + \boldsymbol{\epsilon}_{Mi} \quad (12)$$

where Y_i represents HbA1c levels; A_i represents adult SES values; and $\mathbf{M}_i$ represents methylation level for 2,000 CpG sites. In Equation (11), the model controls for age, gender and race/ethnicity, and in Equation (12), we adjust for age, gender, race/ethnicity and enrichment scores for 4 major blood cell types (neutrophils, B cells, T cells and natural killer cells). All the continuous variables are standardized to have zero mean and unit variance. The univariate analysis is applied in a similar fashion except that it is used to analyze one site at a time.

To compare the power between the univariate analysis and our Bayesian mediation analysis, we perform 100 permutations to obtain an empirical null distribution of PIP values following the procedure in Section 4.4, with which we obtain empirical estimates of FDRs for both univariate and Bayesian multivariate analysis. Consistent with simulations, our Bayesian multivariate mediation method is more powerful than the univariate mediation method. For example, at an FDR of 0.05, our Bayesian mediation method is able to identify 406 mediators while the univariate method is only able to identify 137. At an FDR of 0.10, our Bayesian mediation method can identify 612 mediators while the univariate mediation

Figure 4: Consider the trio: Adult SES $\rightarrow$ DNAm $\rightarrow$ HbA1c. The black dots are the estimated posterior inclusion probability (PIP) for each CpG site from the Bayesian mediation method and the red dots are the estimated PIPs when we permute the outcome once and fit the Bayesian mediation method.

method can only identify 356.

We display PIP values for each of the 2,000 CpG sites from the Bayesian multivariate analysis in Figure 4. Two CpG sites were identified with strong evidence ($PIP > 0.5$) for mediating the adult SES effects on HbA1c. These two CpG sites are also among the top ten sites with the smallest p -values obtained from univariate mediation analysis. In addition, these two CpG sites are close to genes *CCDC54* and *CCND2*, both of which are known candidates associated with HbA1c. Specifically, the expression of *CCND2* has been shown to be associated with risk of T2D and the related glycemic traits of glucose, HbA1c, and insulin [33]. The gene *CCDC54* interacts with valproic acid and acrylamide, both of which are associated with diabetes and blood insulin [34, 35]. Therefore, strong evidence suggests that adult SES may act through these two genes to affect HbA1c.

We also estimate the global mediation effects $\hat{\tau}$ as 0.0084 and its 95% credible interval from the posterior samples as (0.0063, 0.0115). The p -value of the global test based on the empirical null distribution is 0.02. And the $\widehat{PVE}_{IE}$ is 0.096, indicating that approximately 10% of the outcome variance is indirectly explained by DNAm after controlling for covariates. In addition, we estimate the proportion of CpG sites in each of the four categories as defined in Section 4.3: $\hat{\pi}_{g1} = 0.002$, $\hat{\pi}_{g2} = 0.031$, $\hat{\pi}_{g3} = 0.001$, $\hat{\pi}_{g4} = 0.966$. We find that a small proportion of DNAm has large effects on the HbA1c level, and a small proportion of DNAm is notably associated with adult SES. The results also suggest that adult SES acts through certain important DNAm sites to influence HbA1c.

7 Discussion

In this paper, we develop a Bayesian sparse linear mixed model for high-dimensional mediation analysis. The advantage of a Bayesian method is to propagate uncertainty for functions of parameters in a natural way instead of resorting to Delta methods or two-step approaches. Our method can simultaneously analyze a large number of unordered mediators without making any causally ordering assumptions. By imposing continuous shrinkage priors on the key regression coefficients for the mediation analysis, our method achieves up to 30% power gain in identifying true non-null mediators compared with univariate mediation method and approximately 10% power gain compared with multivariate method based on simulations. The Bayesian method also provides better interpretations of the way in which a mediator links or does not link exposure to outcome, and we construct tests for global indirect effects based on the structure of the composite null hypothesis. Our global test is slightly conservative under the null and yields decent power under the alternatives. Implementing our method to MESA, our Bayesian mediation method can detect more active mediators than univariate mediation method at fixed FDR levels. We also identified two genes, *CCDC54* and *CCND2*, with strong evidence for actively mediating the adult SES effects on HbA1c. Both of them are candidate genes associated with diabetes and blood insulin.

Recent literature proposes a convex penalty on the product term of indirect effect [36], which improves power of pathway selection and reduces estimation bias in the indirect

effects. Under the Bayesian framework, direct shrinkage on the product term may be a more appropriate choice, as it takes into account the correlation between the two models in the mediation analysis and is more straightforward when the goal is to identify non-null mediators. Directly incorporating the correlation between the mediators will be another avenue to pursue. We leave that for future work.

References

- [1] David P MacKinnon, Amanda J Fairchild, and Matthew S Fritz. “Mediation analysis”. In: *Annu. Rev. Psychol.* 58 (2007), pp. 593–614.
- [2] Kristopher J Preacher and Andrew F Hayes. “Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models”. In: *Behavior research methods* 40.3 (2008), pp. 879–891.
- [3] Thomas R Ten Have and Marshall M Joffe. “A review of causal estimation of effects in mediation analyses”. In: *Statistical Methods in Medical Research* 21.1 (2012), pp. 77–107.
- [4] Tyler J VanderWeele. “Mediation analysis: a practitioner’s guide”. In: *Annual review of public health* 37 (2016), pp. 17–32.
- [5] James M Robins and Sander Greenland. “Identifiability and exchangeability for direct and indirect effects”. In: *Epidemiology* (1992), pp. 143–155.
- [6] Kosuke Imai, Luke Keele, and Dustin Tingley. “A general approach to causal mediation analysis.” In: *Psychological methods* 15.4 (2010), p. 309.
- [7] Reuben M Baron and David A Kenny. “The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations.” In: *Journal of personality and social psychology* 51.6 (1986), p. 1173.
- [8] David Peter MacKinnon. *Introduction to statistical mediation analysis*. Routledge, 2008.
- [9] Kosuke Imai, Luke Keele, Teppei Yamamoto, et al. “Identification, inference and sensitivity analysis for causal mediation effects”. In: *Statistical science* 25.1 (2010), pp. 51–71.
- [10] Tyler J VanderWeele and Stijn Vansteelandt. “Odds ratios for mediation analysis for a dichotomous outcome”. In: *American journal of epidemiology* 172.12 (2010), pp. 1339–1348.
- [11] Yin-Hsiu Chen et al. “Mediation Formula for a Binary Outcome and a Time-Varying Exposure and Mediator, Accounting for Possible Exposure-Mediator Interaction”. In: *American journal of epidemiology* 184.2 (2016), pp. 157–159.

- [12] Yen-Tsung Huang, Tyler J VanderWeele, and Xihong Lin. “Joint analysis of SNP and gene expression data in genetic association studies of complex diseases”. In: *The annals of applied statistics* 8.1 (2014), p. 352.
- [13] Yen-Tsung Huang. “Integrative modeling of multi-platform genomic data under the framework of mediation analysis”. In: *Statistics in medicine* 34.1 (2015), pp. 162–178.
- [14] RM Daniel et al. “Causal mediation analysis with multiple mediators”. In: *Biometrics* 71.1 (2015), pp. 1–14.
- [15] Tyler VanderWeele and Stijn Vansteelandt. “Mediation analysis with multiple mediators”. In: *Epidemiologic methods* 2.1 (2014), pp. 95–115.
- [16] Masataka Taguri, John Featherstone, and Jing Cheng. “Causal mediation analysis with multiple causally non-ordered mediators”. In: *Statistical methods in medical research* (2015), p. 0962280215615899.
- [17] Yen-Tsung Huang and Wen-Chi Pan. “Hypothesis test of mediation effect in causal mediation model with high-dimensional continuous mediators”. In: *Biometrics* 72.2 (2016), pp. 402–413.
- [18] Michael R Elliott, Trivellore E Raghunathan, and Yun Li. “Bayesian inference for causal mediation effects using principal stratification with dichotomous mediators and outcomes”. In: *Biostatistics* 11.2 (2010), pp. 353–372.
- [19] Michael J Daniels et al. “Bayesian inference for the causal effect of mediation”. In: *Biometrics* 68.4 (2012), pp. 1028–1036.
- [20] Chanmin Kim et al. “A framework for Bayesian nonparametric inference for causal effects of mediation”. In: *Biometrics* 73.2 (2017), pp. 401–409.
- [21] Belinda L Needham et al. “Life course socioeconomic status and DNA methylation in genes related to stress reactivity and inflammation: the multi-ethnic study of atherosclerosis”. In: *Epigenetics* 10.10 (2015), pp. 958–969.
- [22] Jennifer A Smith et al. “Neighborhood characteristics influence DNA methylation of genes involved in stress response and inflammation: The Multi-Ethnic Study of Atherosclerosis”. In: *Epigenetics* 12.8 (2017), pp. 662–673.

- [23] Jia Zhong, Golareh Agha, and Andrea A Baccarelli. “The role of DNA methylation in cardiovascular risk and disease”. In: *Circulation research* 118.1 (2016), pp. 119–131.
- [24] Taulant Muka et al. “The role of global and regional DNA methylation and histone modifications in glycemic traits and type 2 diabetes: A systematic review”. In: *Nutrition, Metabolism and Cardiovascular Diseases* 26.7 (2016), pp. 553–566.
- [25] Donald B Rubin. “Randomization analysis of experimental data: The Fisher randomization test comment”. In: *Journal of the American Statistical Association* 75.371 (1980), pp. 591–593.
- [26] Donald B Rubin. “Comment: Which ifs have causal answers”. In: *Journal of the American Statistical Association* 81.396 (1986), pp. 961–962.
- [27] Andriy Derkach, Jerry F Lawless, and Lei Sun. “Pooled association tests for rare genetic variants: a review and some new results”. In: *Statistical Science* (2014), pp. 302–321.
- [28] Yongtao Guan and Matthew Stephens. “Bayesian variable selection regression for genome-wide association studies and other large-scale problems”. In: *The Annals of Applied Statistics* (2011), pp. 1780–1815.
- [29] Xiang Zhou, Peter Carbonetto, and Matthew Stephens. “Polygenic modeling with Bayesian sparse linear mixed models”. In: *PLoS genetics* 9.2 (2013), e1003264.
- [30] Leo A Aroian. “The probability function of the product of two normally distributed variables”. In: *The Annals of Mathematical Statistics* (1947), pp. 265–271.
- [31] Diane E Bild et al. “Multi-ethnic study of atherosclerosis: objectives and design”. In: *American journal of epidemiology* 156.9 (2002), pp. 871–881.
- [32] Elizabeth Selvin et al. “Glycated hemoglobin, diabetes, and cardiovascular risk in nondiabetic adults”. In: *New England Journal of Medicine* 362.9 (2010), pp. 800–811.
- [33] Hanieh Yaghootkar et al. “Association analysis of 29,956 individuals confirms that a low frequency variant at CCND2 halves the risk of type 2 diabetes by enhancing insulin secretion.” In: *Diabetes* (2015), db141456.

- [34] Silu Zuo et al. “Valproic Acid as a potentiator of metabolic syndrome in institutionalized residents on concomitant antipsychotics: fat chance, or slim to none?” In: *Pharmacy and Therapeutics* 40.2 (2015), p. 126.
- [35] Chien-Yu Lin et al. “Association Between Acrylamide, Blood Insulin and Insulin Resistance in Adults”. In: *Diabetes care* (2009).
- [36] Yi Zhao and Xi Luo. “Pathway lasso: estimate and select sparse mediation pathways with high dimensional mediators”. In: *arXiv preprint arXiv:1603.07749* (2016).