



# Adaptive Bi-Level Variable Selection of Conditional Main Effects for Generalized Linear Models

Kexin Xie and Xinwei Deng

Department of Statistics, Virginia Tech, Blacksburg, VA

## ABSTRACT

Understanding interaction effects among variables is important for regression modeling in various applications. The conventional approach of quantifying interactions as the product of variables often lacks clear interpretability, especially in complex systems. The concept of conditional main effects (CME) provides a more intuitive and interpretable framework for capturing interaction effects by quantifying the effect of one variable conditional on the level of another. A recent method called *cmenet* further considered the bi-level selection of CMEs by leveraging their natural grouping structure (e.g., sibling and cousin groups) through penalization. However, there are several limitations in the *cmenet* method, including the coupling ability of penalties for within-group CMEs, lack of adaptiveness for between-group penalties, and restriction to linear models with continuous responses. To overcome these limitations, we propose an adaptive *cmenet* method for CME selection under the generalized linear model (GLM) framework. The proposed method considers a penalized likelihood approach with adaptive weights to enable effective bi-level variable selection, improving both between-group and within-group selection. An efficient algorithm for parameter estimation is also developed by employing an iteratively reweighted least squares procedure. The performance of the proposed method is evaluated by both simulation studies and real-data studies in gene association analysis.

## ARTICLE HISTORY

Received April 2025  
Accepted February 2026

## KEYWORDS

Adaptive weight; Coordinate descent; Conditional effects; Gene association; High-dimensional data; Interaction analysis

## 1. Introduction

In statistical analyses, the interactions between variables are of great interest for modeling and uncovering complex relationships among variables. Traditionally, the interaction between two variables  $x_1$  and  $x_2$  is modeled as their product term  $x_1x_2$  in regression models. However, such representations often lack intuitive interpretability, especially for non-specialists, making them challenging to apply in practice. The conditional main effects (CMEs), introduced by Wu (2015), offer a novel alternative by focusing on the effect of one variable while conditioning on the fixed level of another, that is,  $x_1|x_2$ . The work of Su and Wu (2017) expanded on this idea by proposing a systematic strategy to identify an appropriate set of CMEs for the analysis of fractional factorial designs. The concept of CME offers an interpretable framework for understanding how certain factors may influence outcomes in the presence of specific conditions. For example, in gene association studies, identifying epistatic (gene-gene interaction) effects associated with multiple SNPs is crucial in these oligogenic/polygenic diseases (Shibata et al. 2024). Instead of quantifying the interaction  $AB$  between gene  $A$  and gene  $B$ , the CME  $A|B+$  provides a variable to quantify the significance of gene  $A$  only when gene  $B$  is present, thus, enhancing biological interpretability.

However, when there are  $p$  variables, even each with two levels, the total number of CMEs is  $2p(p - 1)$ , which can be large even for a moderate  $p$ . Effective variable selection is

therefore critical. Moreover, CMEs naturally exhibit structured groupings. Based on Mak and Wu (2019), there are two general group structures: *sibling* groups and *cousin* groups. Specifically, the sibling group of a ME  $A$  consists of CMEs sharing the same ME, such as  $\{A|B+, A|B-, A|C+, A|C-, \dots\}$ . In contrast, the cousin group of  $A$  comprises CMEs that share the same conditional effect, including  $\{B|A+, B|A-, C|A+, C|A-, \dots\}$ . Accommodating such group structures, it will be more appropriate to conduct the bi-level variable selection for identifying both active CME groups and the significant effects within these groups. A pioneering work from Mak and Wu (2019) proposed a so-called *cmenet* method to consider the “exponential-MC+” hierarchical penalization for bi-level variable selection in CMEs under the context of linear models. Their method facilitates bi-level variable selection by controlling both *between-group* and *within-group* selection by adopting composition penalty in the penalized regression (Breheny and Huang 2009).

Although the *cmenet* method has demonstrated good performance in selecting CMEs, several limitations remain. One issue is that the extent of coupling, where the penalty applied to a coefficient is reduced when grouped with other important predictors, is static and cannot adapt to the signal strength within the group, even with the inclusion of the coupling parameter  $\tau$ . Moreover, the tuning parameters  $\lambda_S$  (for sibling groups) and  $\lambda_C$  (for cousin groups) do not vary across all groups, limiting the flexibility in capturing group-specific variations. We also note

that the *cmenet* method primarily works for linear models with continuous responses, making it unsuitable for situations where the response is binary, count-based, or follows a generalized linear model (GLM) structure.

To address these issues, we propose an *adaptive cmenet* approach for bi-level CME selection under the GLM framework. Through systematic investigation of the penalty term used in the *cmenet*, we develop an adaptive penalization strategy to address the static coupling limitations in variable selection and extend the CME selection to generalized linear models. The key contributions of this work are as follows: First, inspired by adaptive Lasso (Zou 2006), the proposed method incorporates both group adaptive weights and individual adaptive weights to enable more accurate variable selection of CMEs. The group adaptive weights mitigate the static coupling issue by assigning higher weights to effects within the same groups that are likely to be coupled. The individual adaptive weights can regulate the shrinkage rate, enabling more accurate coefficient estimation. A theoretical interpretation of the adaptive weights is provided to justify the choice of weights. Furthermore, we extend the CME selection to GLMs and develop an efficient iteratively reweighted least squares (IRLS) algorithm (Green 1984) for parameter estimation. The proposed IRLS algorithm incorporates the coordinate descent methods (Fu 1998; Friedman et al. 2007; Mazumder, Friedman, and Hastie 2011) and a majorization-minimization technique (Hunter and Lange 2000) to have a closed-form solution for the optimization in each iteration. The implementation of the proposed method is deployed through the R package, `glmcmenet`, and will be publicly available for use in practical applications.

The rest of this article is organized as follows. We describe the concept of the CMEs briefly and review the *cmenet* method in Section 2.1. We detail the proposed adaptive *cmenet* approach for GLMs in Section 3 and present the parameter estimation algorithm in Section 4. The performance of the adaptive *cmenet* is illustrated in Section 5 via several numerical simulations. The application to the gene-association studies is analyzed in Section 6. Section 7 concludes this work with some discussions on future work.

## 2. Background and Motivation

### 2.1. Review of Conditional Main Effects

The definition of conditional main effects (CME) is first introduced by Wu (2015). Following the work of Mak and Wu (2019), we use the definition of CME derived from main effects (ME). Suppose that there are  $n$  observations with  $p$  factors of 2-level under consideration, represented by binary covariate vector  $\mathbf{x}_j = (x_{1,j}, x_{2,j}, \dots, x_{n,j})^T \in \{-1, 1\}^n$ ,  $j = 1, \dots, p$  in the model. Briefly, for the covariate  $x_{i,j}$  conditional on  $x_{i,k}$  at level  $+1$  for object  $i$ , the CME covariate is defined as  $x_{i,j|k+} = x_{i,j}$  when  $x_{i,k} = +1$ , and  $x_{i,j|k+} = 0$  when  $x_{i,k} = -1$ . The CME covariate for  $x_{i,j|k-}$  is defined analogously (see Mak and Wu (2019) for further details). For convenience, we adopt the notation  $A, B, C, D, \dots$  to represent the MEs that quantify the effect of covariates  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_p$ , respectively. Similarly, the CMEs that

**Table 1.** Model Matrix for the two MEs A and B, with its interaction effect and four CMEs.

A	B	AB	A B+	A B-	B A+	B A-
+1	+1	+1	+1	0	+1	0
+1	-1	-1	0	+1	-1	0
-1	+1	-1	-1	0	0	+1
-1	-1	+1	0	-1	0	-1

quantify the effect of covariates  $\mathbf{x}_{j|k+}$  and  $\mathbf{x}_{j|k-}$  are denoted as  $J|K+$  and  $J|K-$ . In the CMEs  $J|K+$  and  $J|K-$ , the effects  $J$  and  $K$  are referred to as the *parent effect* and the *child effect*, respectively. This notation will be used consistently throughout the remainder of the article.

Note that the conventional interaction terms are often not straightforward for interpretation. Consider a simple example involving two covariates,  $A$  and  $B$ . Traditionally, their interaction is represented as the product term  $AB$ . However, this term  $AB$  alone does not provide explicit information about the individual values or levels of  $A$  and  $B$ , making it difficult to disentangle their interaction in practical terms. For instance,  $AB = 1$  could result from either  $A = +1, B = +1$  or  $A = -1, B = -1$ . CMEs address this ambiguity by capturing conditional relationships more explicitly. Table 1 presents the four possible CMEs derived from these MEs.

It is evident that there are complex hierarchical structures among the ME covariates and CME covariates, either as parent terms or as child terms. Following the grouping structures in Mak and Wu (2019), there are three general categories: (a) *Sibling* CMEs, which share the same parent effects, such as  $\{A|B+, A|B-, A|C+, A|C-, \dots\}$ ; (b) *Cousin* CMEs, which share the same child effects, such as  $\{B|A+, B|A-, C|A+, C|A-, \dots\}$ ; and (c) *Parent-Child* pairs, such as  $\{A, A|B+\}$ ,  $\{A, A|B-\}$ ,  $\dots$ . Mak and Wu (2019) examined the collinearity and pairwise correlations within these groups and observed that the correlation magnitudes create a natural hierarchy among effect groups, leading to inconsistent selection when using standard variable selection methods. To address this, they proposed a bi-level variable selection method for CMEs, implemented in the *cmenet* framework. This idea builds upon earlier developments in bi-level variable selection, such as the group bridge penalty by Huang et al. (2009), which applies a bridge penalty to the  $\ell_1$ -norm of predefined groups to encourage group-level sparsity. Breheny and Huang (2009) later introduced a general framework for composite penalties, which employs an *outer* penalty to regulate a sum of *inner* penalties. It facilitates bi-level variable selection by controlling both between-group and within-group selection.

Let  $(\mathbf{y}, \mathbf{X})$  represent the complete set of observations, where  $\mathbf{y} \in \mathbb{R}^{n \times 1}$  is the response variable corresponding to the full model matrix  $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_p, \dots, \mathbf{x}_{p'}) \in \mathbb{R}^{n \times p'}$ . Here,  $\{\mathbf{x}_j\}_{j=1}^p$  and  $\{\mathbf{x}_j\}_{j=p+1}^{p'}$  are the normalized vectors for the  $p$  MEs and the  $4\binom{p}{2}$  CMEs, respectively. Define  $\mathcal{S}(j) = \{J, J|A+, J|A-, J|B+, J|B-, \dots\}$  as the sibling group for the parent effect  $j$ , and  $\mathcal{C}(j) = \{J, A|J+, A|J-, B|J+, B|J-, \dots\}$  as the cousin group for the child effect  $j$ ,  $j = 1, \dots, p$ . The parameter estimation is then conducted as follows:

$$\begin{aligned} \min_{\boldsymbol{\beta}} Q(\boldsymbol{\beta}) &= \min_{\boldsymbol{\beta}} \left\{ \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + P_S(\boldsymbol{\beta}) + P_C(\boldsymbol{\beta}) \right\}, \\ P_S(\boldsymbol{\beta}) &= \sum_{j=1}^p f_{out,S} \left\{ \sum_{k \in \mathcal{S}(j)} f_{in,S}(\beta_k) \right\}, \\ P_C(\boldsymbol{\beta}) &= \sum_{j=1}^p f_{out,C} \left\{ \sum_{k \in \mathcal{C}(j)} f_{in,C}(\beta_k) \right\}, \end{aligned} \quad (1)$$

where  $f_{out,S}$  and  $f_{in,S}$  ( $f_{out,C}$  and  $f_{in,C}$ ) represent the outer and inner penalties, which govern the between-group and within-group selection for sibling (or cousin) groups, respectively. Using different combinations of outer and inner penalties, Mak and Wu (2019) adopted the exponential penalty  $\eta_{\lambda_s, \tau}(\theta)$  (or  $\eta_{\lambda_c, \tau}(\theta)$ ) for the outer penalty (Breheny 2015) and the scaled minimax concave-plus (MC+) penalty  $g_{\lambda_s, \gamma}(\theta)$  (or  $g_{\lambda_c, \gamma}(\theta)$ ) for the inner penalty (Zhang 2010). These penalties are defined as follows:

$$\begin{aligned} \eta_{\lambda, \tau}(\theta) &= \frac{\lambda^2}{\tau} \left\{ 1 - \exp\left(-\frac{\tau\theta}{\lambda}\right) \right\}, \\ \text{and } g_{\lambda, \gamma}(\beta) &= \int_0^{|\beta|} \left( 1 - \frac{x}{\lambda\gamma} \right)_+ dx, \end{aligned} \quad (2)$$

where  $\lambda$ ,  $\gamma$ , and  $\tau$  are tuning parameters that regulate the degree of regularization, nonconvexity, and coupling, respectively. Specifically, two penalization parameters  $\lambda_s > 0$  and  $\lambda_c > 0$  are the penalization parameters applied to sibling and cousin groups, remaining constant across all groups. The MC+ penalty incorporates the parameter  $\gamma > 1$ , which acts as a transition between the  $\ell_0$ -norm penalty (approximated as  $\gamma \rightarrow 1$ ) and the  $\ell_1$ -norm penalty in LASSO (approximated as  $\gamma \rightarrow \infty$ ). The coupling parameter  $\tau$  within the exponential penalty governs the strength of the relationship among variables within the same group. These penalty functions can effectively control group-wise coupling and sparsity, addressing challenges in bi-level variable selection for CMEs.

## 2.2. Real-World Motivation

Biologists often observe epistasis, where the effect of a gene (or genetic variant) on a phenotype depends on the presence of another variant or on whether a pathway is active. A main-effects screen can miss such patterns, while a full search over all interactions can be very large and hard to interpret. Conditional main effects (CMEs) can address this gap. A statement such as “SNP  $J$  given SNP  $K = 1$ ” expresses the biological context directly and groups related effects in a way that scientists use. In the biological context, the SNP whose effect is being evaluated ( $J$ ) is the focal locus, and the SNP on which it is conditioned ( $K$ ) is the modifier locus. This representation clarifies when and under what genetic context a locus influences the phenotype, while keeping the model concise and interpretable. To illustrate how CMEs capture such context-dependent effects in practice, we consider a genome-wide association study in *Arabidopsis thaliana*.

We study the visual chlorosis phenotype from the GWAS of Atwell et al. (2010) (details in Section 6). *Arabidopsis* is

a predominantly self-fertilizing model plant with substantial genetic variation across many traits. CMEs can help connect SNPs to specific biological pathways. For example, a CME might show that a SNP near a chlorophyll gene only has an effect if a variant in an iron-related gene is also present. Another CME might show that a redox-enzyme SNP has a stronger effect when its antioxidant partner gene is absent. This type of “conditional” statement facilitates comparison with existing functional annotations, and is associated with experimentally testable hypotheses. Meanwhile, the *cmenet* method also shows its limitations in this data scenario. First, the chlorosis phenotype is a binary outcome, and the *cmenet* method, built on a linear Gaussian model, cannot model binary responses on an appropriate scale. Second, genetic markers are often correlated. When *cmenet* finds one strong CME, it often selects a group of correlated, nearby CMEs as well. This creates redundancy in the results, populating the output with long lists of CMEs that all point to the same underlying genetic source and reducing interpretability.

Our framework is designed to address these challenges. We keep the CME representation and consider a logistic model for the analysis of a binary trait. We then adopt the penalization at both the group level and the individual level. The key advantage is that the relative influence of the two CME groups shifts with the evidence rather than being fixed in advance. It can obtain an interpretable set of context-specific associations that can highlight candidate gene–gene interaction patterns for follow-up study. The details of the proposed method will be introduced in the following Section 3.1.

## 3. The Proposed Method: Adaptive *cmenet*

In this section, we first investigate the coupling issue in the *cmenet* method and identify its limitations in Section 3.1. We then propose a new penalization approach in Section 3.2 to address these limitations. Finally, we examine the effectiveness of the proposed approach in comparison with the *cmenet* method in Section 3.3.

### 3.1. Coupling Ability in the *cmenet*

Note that the choice of the “exponential-MC+” penalty framework is motivated by two key principles: (a) *CME coupling*, which states that a CME is more likely to be included in the model if other effects within the same sibling or cousin group have already been selected; and (b) *CME reduction*, which implies that a ME is more easily selected when effects in its sibling or cousin groups have already entered the model. To elucidate these principles, we adopt the concept of the “selection threshold” introduced by Breheny (2015), which refers to the minimum level of association between a covariate and the outcome that is necessary for the covariate to be selected in the model. Let  $z$  represent the unpenalized estimate for  $\beta_j$ , that is,  $z = \frac{1}{n} \mathbf{x}_j^T (\mathbf{y} - \mathbf{X}_{-j} \boldsymbol{\beta}_{-j})$ . Here,  $\mathbf{X}_{-j} \in \mathbb{R}^{n \times (p'-1)}$  represents the design matrix excluding the  $j$ th covariate, and  $\boldsymbol{\beta}_{-j} \in \mathbb{R}^{(p'-1)}$  is the corresponding coefficient vector. The estimate  $z$  corresponds to the value that minimizes the loss function while keeping all other coefficients fixed. The “selection threshold” is formally defined as  $t_{\beta_j} := \inf\{|z| : \hat{\beta}_j \neq 0\}$  for  $j = 1, \dots, p'$ . In other

words, it represents the infimum of all values of  $|z|$  for which the corresponding coefficient  $\hat{\beta}_j$  is nonzero. Derived from (1), the selection threshold  $t_{\beta_{j|k+}}$  for  $\beta_{j|k+}$  is expressed as

$$t_{\beta_{j|k+}} = \lambda_s \exp\left\{-\frac{\tau \|\beta_{S(j)}\|_{\lambda_s, \gamma}}{\lambda_s}\right\} + \lambda_c \exp\left\{-\frac{\tau \|\beta_{C(k)}\|_{\lambda_c, \gamma}}{\lambda_c}\right\}. \quad (3)$$

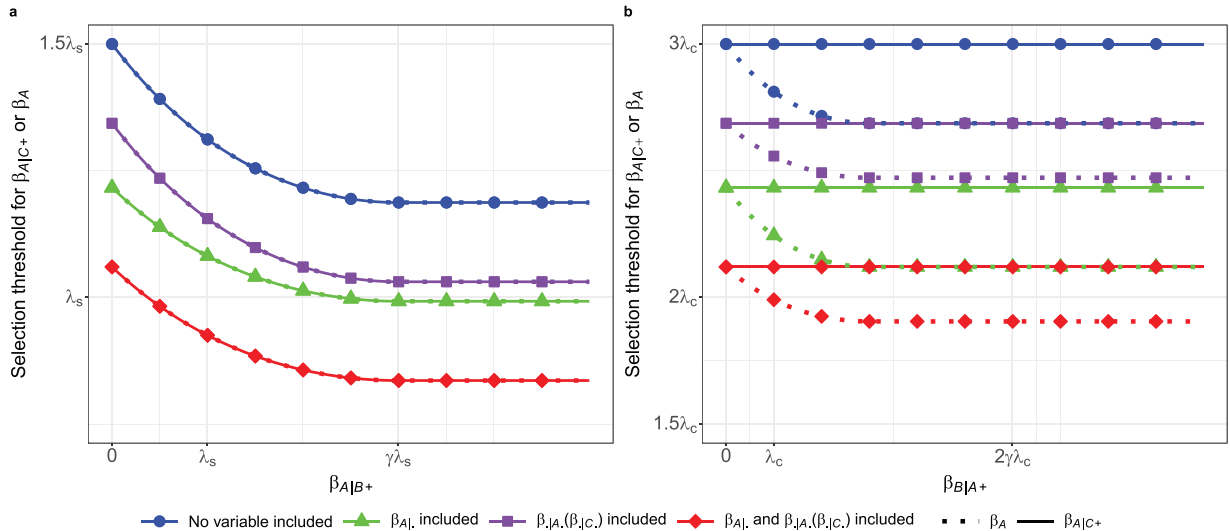
Here, the coefficient vector of a group  $\mathcal{G}$  is denoted by  $\beta_{\mathcal{G}} \in \mathbb{R}^{|\mathcal{G}|}$ , and the penalty term is defined as  $\|\beta_{\mathcal{G}}\|_{\lambda, \gamma} := \sum_{l \in \mathcal{G}} g_{\lambda, \gamma}(\beta_l)$ . The selection threshold  $t_{\beta_j}$  for the coefficient  $\beta_j$  is defined in an analogous manner. For simplicity, we introduce the notations  $\Delta_{S(j)} := \lambda_s \exp\left\{-\frac{\tau \|\beta_{S(j)}\|_{\lambda_s, \gamma}}{\lambda_s}\right\}$  and  $\Delta_{C(k)} := \lambda_c \exp\left\{-\frac{\tau \|\beta_{C(k)}\|_{\lambda_c, \gamma}}{\lambda_c}\right\}$ , which will be used throughout the remainder of the article.

Let us consider a model in which  $A|B+$  has been selected. Two candidate effects are under consideration:  $A|C+$  and  $A$ . Figure 1(a) illustrates how changes in  $\beta_{A|B+}$  will influence the selection thresholds for  $\beta_{A|C+}$  and  $\beta_A$ , respectively. When no variables are initially included in the model, the selection threshold remains high regardless of the coefficient magnitude. However, as a sibling variable ( $\mathbf{x}_{A|}$ ) or a cousin variable ( $\mathbf{x}_{|C}$  for  $A|C+$ , or  $\mathbf{x}_{|A}$  for  $A$ ) is included into the model, the selection thresholds for both  $\beta_{A|C+}$  and  $\beta_A$  decrease. This reduction is particularly pronounced when both sibling and cousin effects are included, facilitating the inclusion of additional effects within the same group. This phenomenon underpins the first principle: *CME Coupling*. To contrast with the behavior of the sibling effect, Figure 1(b) considers a model that includes  $B|A+$  and varies its coefficient  $\beta_{B|A+}$ . The selection threshold for  $\beta_A$  (represented by dotted lines) declines with  $\beta_{B|A+}$ , whereas the threshold for  $\beta_{A|C+}$  (represented by solid lines) remains essentially unchanged because cousins of  $C$  are fixed. The lower threshold value for the  $\beta_A$  indicates that the ME  $A$  is more likely to be included in the model than the CME  $A|C+$ . That is, when effects involving  $A$  have already been selected from  $S(A)$  or from  $C(A)$ , the ME  $A$  enters the model more easily. When many such effects are selected, the set of conditional effects tends to

contract to the underlying main effect. This behavior exemplifies the second principle: *CME Reduction*.

An important observation is that as the coefficient magnitude increases, the selection threshold stabilizes at a fixed level. That is, the extent of coupling is inherently controlled by group size and the coupling parameter  $\tau$ . This phenomenon stems from the characteristics of the inner penalty function  $g_{\lambda, \gamma}(\beta)$  in (2), which is monotonic increasing and gradually reduces penalization as  $|\beta|$  increases, up to the critical point  $|\beta| = \gamma\lambda$ . Beyond this critical value,  $g_{\lambda, \gamma}(\beta)$  becomes constant at its maximum value of  $\gamma\lambda/2$ . As a result, the selection threshold reaches a plateau at a fixed value of  $\eta'_{\lambda, \tau}(\theta)|_{\theta=\gamma\lambda/2}$ , remaining unaffected by any further increases in the coefficient magnitude. This plateau prevents excessive coupling because each term in the exponent in  $\eta_{\lambda, \tau}(\theta)$  has a finite contribution; once terms have saturated, enlarging coefficients cannot push the threshold downward without bound.

Although this plateau is desirable at very large signals, the method suffers from a static coupling limitation, which manifests in two practical ways. First, it has a fixed between-group relative influence determined by the non-adaptive parameters  $\lambda_s$  and  $\lambda_c$ . This balance is locked in from the start and cannot adapt to the data. For example, in Figure 1, setting  $\lambda_s = 2\lambda_c$  causes the green curve to lie below the purple one, regardless of how much signal is present. This fixed setting locks in a predefined ratio between the groups, making the threshold insensitive to changes in the actual strength of evidence. Second, and more importantly, it has a constant marginal effect of adding a new variable within a group, regardless of the current strength of that group. It is clearly shown in Figure 1(a): the gap between threshold curves, such as between the blue and green lines (or between purple and red), remains unchanged as  $\beta_{A|B+}$  increases. The lines run parallel. Moreover, due to the uniform bounds of the inner penalty, the benefit of adding a second sibling is exactly the same as adding the first. Consequently, the threshold keeps dropping as you add more siblings, even if the group is already ‘‘saturated.’’ This lack of diminishing marginal returns



**Figure 1.** Effect of varying CME coefficients on selection thresholds of *cmenet*: (a) Vary  $\beta_{A|B+}$ ; (b) Vary  $\beta_{B|A+}$ . Curves show the threshold for  $\beta_{A|C+}$  (dotted line) and  $\beta_A$  (solid line) under four contexts: when no other variable included (blue), one sibling included (green), one cousin included (purple), and one sibling and one cousin included (red) in the model. Baseline *cmenet* setting:  $(\lambda_s, \lambda_c, \gamma, \tau) = (1, 0.5, 3, 0.25)$ . The adding variable has the same coefficient as  $\beta = 2$ .

can lead to overemphasis on already dominant groups and insufficient responsiveness to emerging signals in competing groups.

These considerations suggest that it is important to allow the penalization to be adaptive to the current level of signal of each group. First, to address the fixed between-group relative influence, it will be more desirable to let the main penalization parameters ( $\lambda_s$  and  $\lambda_c$ ) adjust based on the total signal strength of their respective groups. Thus, the procedure can dynamically “rebalance” the influence of the sibling and cousin groups. Second, to address the constant within-group marginal effect, we need to adjust the inner penalty for each coefficient adaptively. This means the inner penalty applied to a coefficient (e.g.,  $\beta_{A|C+}$ ) should decrease as the signal from a related coefficient ( $\beta_{A|B+}$ ) increases. This makes the selection threshold drop earlier in response to initial weak signals within the group, and mitigates the marginal impact once the group already contains large signals. Thus, we develop an adaptive penalization framework, which is described in detail in the following section.

### 3.2. The Proposed Adaptive cmenet and glmcmenet

As discussed in Section 3.1, the issue of static coupling limitation primarily arises from the uniform bounds of the inner penalty and the non-adaptive penalization parameters applied across groups and individuals. Drawing inspiration from the adaptive Lasso in Zou (2006) and the adaptive sparse group Lasso developed by Fang et al. (2015), we consider incorporating adaptive weights for different groups and individual effects. Accordingly, the proposed adaptive cmenet penalization is defined as follows:

$$\tilde{P}_{\mathcal{G}}(\boldsymbol{\beta}) = \sum_{j=1}^p f_{out,\mathcal{G}}^{(j)} \left\{ \sum_{k \in \mathcal{G}(j)} \omega_k f_{in,\mathcal{G}}^{(j)}(\beta_k^{(j)}) \right\}, \quad (4)$$

$\mathcal{G} \in \{\mathcal{S}, \mathcal{C}\}$  and  $k \in \mathcal{G}(j)$

where  $f_{out,\mathcal{G}}^{(j)}(\theta) = \eta_{\lambda_{\mathcal{G}(j)},\tau}(\theta)$  and  $f_{in,\mathcal{G}}^{(j)}(\beta) = g_{\lambda_{\mathcal{G}(j)},\gamma}(\beta)$ . Here  $\lambda_{\mathcal{G}(j)}$  is the adaptive weight for group  $\mathcal{G}(j)$ , and  $\omega_k$  is the individual weight for ME or CME within group  $\mathcal{G}(j)$ . The rationale of introducing adaptive weights is:  $\lambda_{\mathcal{G}(j)}$  resolves the issue of non-adaptive penalization parameters across groups, while  $\omega_k$  relaxes the uniform boundary of the inner penalty. To illustrate the distinction between the cmenet penalty and the proposed adaptive cmenet penalty, readers can consider a simplified example as illustrated in Example 1. Consequently, the parameter estimation can be expressed as

$$\min_{\boldsymbol{\beta}} Q(\boldsymbol{\beta}) = \min_{\boldsymbol{\beta}} \left\{ \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \tilde{P}_{\mathcal{S}}(\boldsymbol{\beta}) + \tilde{P}_{\mathcal{C}}(\boldsymbol{\beta}) \right\}. \quad (5)$$

**Example 1.** Consider a simple model with only two MEs  $A$  and  $B$ , then the corresponding sibling groups are  $\mathcal{S}(A) = \{A, A|B+, A|B-\}$ ,  $\mathcal{S}(B) = \{B, B|A-, B|A+\}$ . For the cmenet penalty in (1), the penalization for sibling groups is defined as

$$P_{\mathcal{S}}(\boldsymbol{\beta}) = \frac{\lambda^2}{\tau} \left\{ 1 - \exp \left[ -\frac{\tau}{\lambda} \left\{ g_{\lambda,\gamma}(\beta_A) + g_{\lambda,\gamma}(\beta_{A|B+}) + g_{\lambda,\gamma}(\beta_{A|B-}) \right\} \right] \right\}$$

$$+ \frac{\lambda^2}{\tau} \left\{ 1 - \exp \left[ -\frac{\tau}{\lambda} \left\{ g_{\lambda,\gamma}(\beta_B) + g_{\lambda,\gamma}(\beta_{B|A+}) + g_{\lambda,\gamma}(\beta_{B|A-}) \right\} \right] \right\}.$$

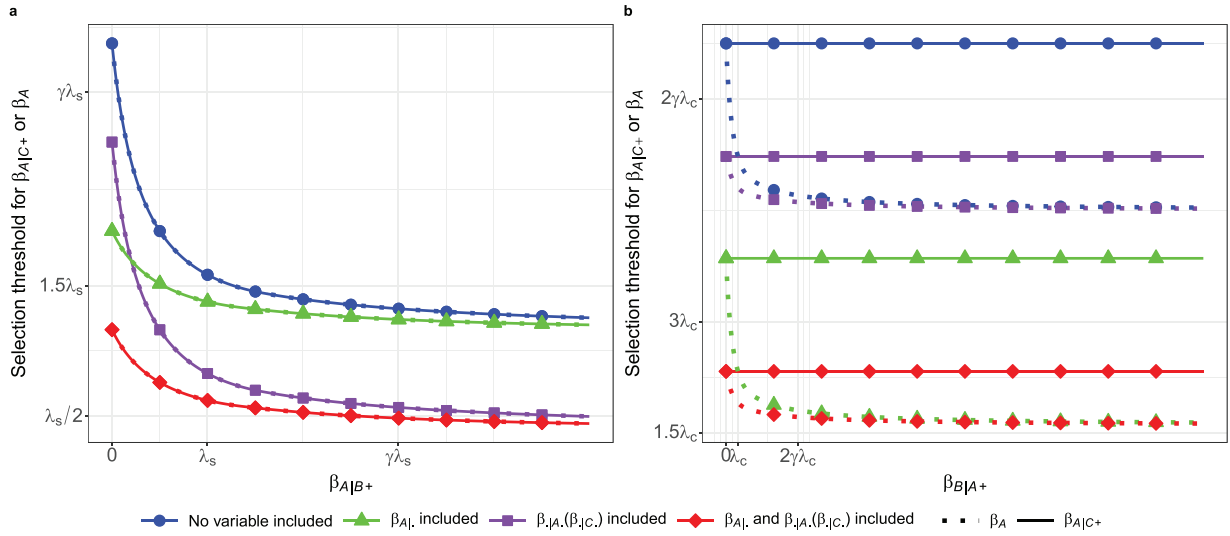
In contrast, for the proposed adaptive cmenet penalty in (4), sibling group weights  $\mathcal{S}(j)$  and individual weights  $\omega_k$  are introduced:

$$\begin{aligned} \tilde{P}_{\mathcal{S}}(\boldsymbol{\beta}) = & \frac{\lambda_{\mathcal{S}(A)}^2}{\tau} \left\{ 1 - \exp \left[ -\tau \left\{ \frac{\omega_A}{\lambda_{\mathcal{S}(A)}} \cdot g_{\lambda_{\mathcal{S}(A)},\gamma}(\beta_A) \right. \right. \right. \\ & + \frac{\omega_{A|B+}}{\lambda_{\mathcal{S}(A)}} \cdot g_{\lambda_{\mathcal{S}(A)},\gamma}(\beta_{A|B+}) \\ & \left. \left. \left. + \frac{\omega_{A|B-}}{\lambda_{\mathcal{S}(A)}} \cdot g_{\lambda_{\mathcal{S}(A)},\gamma}(\beta_{A|B-}) \right\} \right] \right\} \\ & + \frac{\lambda_{\mathcal{S}(B)}^2}{\tau} \left\{ 1 - \exp \left[ -\tau \left\{ \frac{\omega_B}{\lambda_{\mathcal{S}(B)}} \cdot g_{\lambda_{\mathcal{S}(B)},\gamma}(\beta_B) \right. \right. \right. \\ & + \frac{\omega_{B|A+}}{\lambda_{\mathcal{S}(B)}} \cdot g_{\lambda_{\mathcal{S}(B)},\gamma}(\beta_{B|A+}) \\ & \left. \left. \left. + \frac{\omega_{B|A-}}{\lambda_{\mathcal{S}(B)}} \cdot g_{\lambda_{\mathcal{S}(B)},\gamma}(\beta_{B|A-}) \right\} \right] \right\}. \end{aligned}$$

A similar formula can be written for cousin group penalties.

For each group  $\mathcal{G}(j)$ , we consider the group weight to be  $\Omega_{\mathcal{G}(j)} = (\|\hat{\boldsymbol{\beta}}_{\mathcal{G}(j)}^{ridge}\|_1 + \frac{1}{n})^{-1}$ , where  $\|\hat{\boldsymbol{\beta}}_{\mathcal{G}(j)}^{ridge}\|_1$  denotes the  $l_1$ -norm of the ridge estimators within group  $\mathcal{G}(j)$ . Thus, the adaptive penalization parameter of this group is  $\lambda_{\mathcal{G}(j)} = \lambda_{\mathcal{G}} \Omega_{\mathcal{G}(j)}$ . The individual weight within this group is defined as  $\omega_k = (\|\hat{\beta}_k^{ridge}\| + \frac{1}{n})^{-1}$ ,  $k \in \mathcal{G}(j)$ . The use of ridge estimators for weight construction aligns with Zou (2006). Given the presence of implicit multicollinearity in CME group structures, ridge estimation is preferred over  $\hat{\boldsymbol{\beta}}^{MLE}$  due to its superior consistency and robustness. To prevent division by zero, we follow Fang et al. (2015) and add  $1/n$  to the denominator. The adaptive cmenet method preserves the principles of CME coupling and CME reduction. By incorporating adaptive weights, this approach effectively alleviates the issue of static coupling, as further elaborated in Section 4.

Note that the cmenet method is primarily designed for linear models. To broaden the scope of CMEs, we extend the adaptive cmenet framework to generalized linear models (GLM), hereafter named glmcmenet. For GLM, the loss function is typically the negative log-likelihood derived from the exponential family. GLMs extend linear models to accommodate various types of response variables, by linking the mean response to the predictors through a link function. Let  $\eta_i = \mathbf{x}_i^T \boldsymbol{\beta}$  denote the linear predictor, where  $\mathbf{x}_i$  is the vector of predictors for the  $i$ th observation. The relationship between the response mean  $\mu_i$  and the linear predictor  $\eta_i$  is specified by a canonical link function  $g(\mu_i) = \eta_i$ . The loss function (i.e., the negative log-likelihood function) for GLMs is  $-\frac{1}{n} \sum_{i=1}^n [y_i \cdot \eta_i - b(\eta_i) + c(y_i)]$ , where  $b(\eta_i)$  and  $c(y_i)$  are functions defined by the exponential family distribution of the response variable. The optimization for the proposed method



**Figure 2.** Effect of varying CME coefficients on selection thresholds of *adaptive cmenet*: (a) Vary  $\beta_{A|B+}$ ; (b) Vary  $\beta_{B|A+}$ . Curves show the threshold for  $\beta_{A|C+}$  (dotted line) and  $\beta_A$  (solid line) under four contexts: when no other variable included (blue), one sibling included (green), one cousin included (purple), and one sibling and one cousin included (red) in the model. Baseline adaptive cmenet setting:  $(\lambda_s, \lambda_c, \gamma, \tau) = (1, 0.5, 3, 0.25)$ . The adding variable has the same coefficient as  $\beta = 0.5$ .

can be written as

$$\min_{\beta} Q(\beta) = \min_{\beta} \left\{ -\frac{1}{n} \sum_{i=1}^n [y_i \cdot \eta_i - b(\eta_i) + c(y_i)] + \tilde{P}_S(\beta) + \tilde{P}_C(\beta) \right\}, \quad (6)$$

where  $\tilde{P}_S(\beta)$  and  $\tilde{P}_C(\beta)$  are defined in (4). The details on parameter estimation will be described in Section 4.

### 3.3. Coupling Ability of the Adaptive cmenet

In this section, we further investigate the coupling ability outlined in Section 3.1, focusing on the selection threshold  $\tilde{t}_{\beta_{j|k+}}$  for  $J|K+$  after incorporating adaptive weights, derived from (5):

$$\tilde{t}_{\beta_{j|k+}} = \lambda_{S(j)} \exp \left\{ -\frac{\tau \|\beta_{S(j)}\|_{\lambda_{S(j)}, \omega, \gamma}}{\lambda_{S(j)}} \right\} + \lambda_{C(k)} \exp \left\{ -\frac{\tau \|\beta_{C(k)}\|_{\lambda_{C(k)}, \omega, \gamma}}{\lambda_{C(k)}} \right\}. \quad (7)$$

In contrast to the fixed plateaus in Figure 1, Figure 2 shows that the adaptive weights create lower and signal-dependent plateaus. This pattern comes from the adaptive weights: small coefficients in  $S(A)$  receive larger weights, so the threshold reacts strongly when the signal first appears. As coefficients in  $S(A)$  become larger, its weight shrinks and the penalty term approaches its plateau. This makes the threshold drop quickly at the start and then flatten earlier. Moreover, Figure 2 demonstrates a clearer form of *CME coupling* principle: once a sibling or a cousin has been selected, the threshold shifts downward. A similar downward trend is observed for the dotted curves for  $A$  in Figure 2(b), while the solid curves for  $A|C+$  remain flat as  $\beta_{B|A+}$  increases, aligning with the *CME reduction* principle. The proposed method addresses both of the previous limitations: First, it dynamically “rebalances” the between-group influence, which is proven in Figure 2(a) by the crossover of the green and purple lines. At first, when the sibling signal is weak (near

$\beta_{A|B+} = 0$ ), the green line is lower. However, as the sibling signal strengthens, the purple line drops more steeply and crosses the green line, proving the relative group influence is no longer fixed. Second, the proposed method successfully diminishes the within-group marginal effect. In Figure 2(a), as  $\beta_{A|B+}$  increases, the green line approaches the blue line at the higher plateau, and the red line approaches the purple line at the lower plateau. This indicates that, once the sibling group becomes strong, the extra sibling has a negligible marginal effect as its contribution is downweighted and the inner penalty is near saturation. Signals coming from the cousin side still matter, since the adjustments for cousins are not driven by how much activity we have in  $S(A)$ . A similar observation can be seen for the dotted line in Figure 2(b): the marginal effect of adding a new cousin diminishes as the cousin effect  $\beta_{B|A+}$  increases, while adding a new sibling still matters.

### 4. Parameter Estimation

For parameter estimation in GLM, the iteratively reweighted least squares (IRLS) algorithm is commonly used. The IRLS uses a second-order Taylor expansion to approximate the loss function quadratically around the current estimate  $\beta^{(m)}$ . The coefficients are then iteratively updated under the format of a weighted least squares as  $\frac{1}{2n}(\tilde{\mathbf{y}} - \mathbf{X}\beta)^T \mathbf{W}(\tilde{\mathbf{y}} - \mathbf{X}\beta)$ , where  $\tilde{\mathbf{y}} = \mathbf{X}\beta^{(m)} + \mathbf{W}^{-1}(\mathbf{y} - \mathbf{v})$  represents the working response with  $\mathbf{v}$  and  $\mathbf{W}$  to be the first and second derivatives of  $L(\eta)$  with respect to  $\eta$ , evaluated at  $\eta^{(m)} = \mathbf{X}\beta^{(m)}$ . Notably,  $\mathbf{W}$  is a diagonal matrix, where each diagonal entry corresponds to the weight of the  $i$ th observation under the least squares loss. At each iteration, the objective function with adaptive cmenet penalization can be expressed as

$$Q(\beta) = \frac{1}{2n}(\tilde{\mathbf{y}} - \mathbf{X}\beta)^T \mathbf{W}(\tilde{\mathbf{y}} - \mathbf{X}\beta) + \tilde{P}_S(\beta) + \tilde{P}_C(\beta). \quad (8)$$

The  $Q(\cdot)$  will exclusively refer to the formulation in (8) for the remainder of the article.

To minimize the objective function in (8), we employ the local coordinate descent (LCD) algorithm (Fu 1998; Friedman et al. 2007; Mazumder, Friedman, and Hastie 2011) to efficiently derive a closed-form threshold operator of the coordinate-wise objective  $Q_j(\cdot)$  with respect to coefficient  $\beta_j$  while keeping the remaining coefficients fixed:

$$Q_j(\boldsymbol{\beta}) = \frac{1}{2n} (\tilde{\mathbf{y}} - \mathbf{X}_{-j}\boldsymbol{\beta}_{-j} - \mathbf{x}_j^T \boldsymbol{\beta}_j)^T \mathbf{W} (\tilde{\mathbf{y}} - \mathbf{X}_{-j}\boldsymbol{\beta}_{-j} - \mathbf{x}_j^T \boldsymbol{\beta}_j) + \eta_{\lambda_{S(j)}, \tau} \{ \|\boldsymbol{\beta}_{S(j)}\|_{\lambda_{S(j)}, \omega, \gamma} \} + \eta_{\lambda_{C(j)}, \tau} \{ \|\boldsymbol{\beta}_{C(j)}\|_{\lambda_{C(j)}, \omega, \gamma} \}. \quad (9)$$

The coordinate-wise objective  $Q_{j|k+}(\cdot)$  with respect to coefficient  $\beta_{j|k+}$  is defined similarly:

$$Q_{j|k+}(\boldsymbol{\beta}) = \frac{1}{2n} (\tilde{\mathbf{y}} - \mathbf{X}_{-j|k+}\boldsymbol{\beta}_{-j|k+} - \mathbf{x}_{j|k+}^T \boldsymbol{\beta}_{j|k+})^T \mathbf{W} (\tilde{\mathbf{y}} - \mathbf{X}_{-j|k+}\boldsymbol{\beta}_{-j|k+} - \mathbf{x}_{j|k+}^T \boldsymbol{\beta}_{j|k+}) + \eta_{\lambda_{S(j)}, \tau} \{ \|\boldsymbol{\beta}_{S(j)}\|_{\lambda_{S(j)}, \omega, \gamma} \} + \eta_{\lambda_{C(k)}, \tau} \{ \|\boldsymbol{\beta}_{C(k)}\|_{\lambda_{C(k)}, \omega, \gamma} \}. \quad (10)$$

The LCD algorithm iteratively optimizes one coordinate at a time while holding the others fixed, offering computational efficiency and stability, particularly for high-dimensional problems. However, due to the presence of penalty terms, the objective function does not have a closed-form one-dimensional solution. To address this limitation, the majorization-minimization (MM) approach (Hunter and Lange 2000) is applied to construct a surrogate function for  $Q_j(\cdot)$  or  $Q_{j|k+}(\cdot)$ , enabling the derivation of a closed-form threshold operator.

#### 4.1. Convexity and Stability

For parameter estimation, it is useful to examine the convexity property of the objective function in optimization. For simplicity and clarity, we assume throughout the remainder of this article that each column  $\mathbf{x}_j$  of  $\mathbf{X}$  is normalized, such that  $\mathbf{x}_j^T \mathbf{I}_n = 0$  and  $\frac{1}{n} \|\mathbf{x}_j\|_2^2 = 1$  for all  $j = 1, \dots, p'$ . Unless otherwise specified, we denote all coordinate-wise objective functions in (9) and (10) as  $Q_j(\cdot)$ , for  $j = 1, \dots, p'$  without differentiating between MEs and CMEs.

**Proposition 1.** Define  $\omega_{\max} = \max_{1 \leq j \leq p'} \omega_j$  as the maximum individual weight. The objective function  $Q(\boldsymbol{\beta})$  in (8) is strictly convex whenever  $\tau + \frac{1}{\gamma \omega_{\max}} < \frac{\zeta_{\min}(\mathbf{X}^T \mathbf{W} \mathbf{X})}{2n\omega_{\max}^2}$ , where  $\mathbf{W}$  is evaluated at  $\boldsymbol{\beta}$  and  $\zeta_{\min}$  is the minimum eigenvalue.

**Proposition 1** implies that a small value of  $\tau + \frac{1}{\gamma \omega_{\max}}$  ensures the convexity of the objective function, leading to a global minimum. However, this condition is quite restrictive and typically applies only in low-dimensional settings where  $n \leq p'$ . In high-dimensional scenarios, the coordinate descent algorithm remains effective, as it only requires local convexity for each coordinate. For linear regression with  $\mathbf{W} = \mathbf{I}$ , the local convexity of the coordinate-wise objective function  $Q_j(\beta_j)$  is guaranteed if  $\tau + \frac{1}{\gamma \omega_{\max}} \leq \frac{1}{2\omega_{\max}^2}$ .

For generalized linear models, the situation becomes more complex due to the introduction of a nonidentity weight matrix  $\mathbf{W}$ , where  $\frac{1}{n} \mathbf{X}^T \mathbf{W} \mathbf{X} \neq \mathbf{I}$ . Despite this, local convexity can still be

ensured by imposing an upper bound on  $\mathbf{W}$ , which corresponds to the Hessian matrix of the loss function with respect to the linear predictor  $\boldsymbol{\eta}$ . Let  $\bar{w} = \max_i \sup_{\boldsymbol{\eta}} \{ \nabla^2 L_i(\boldsymbol{\eta}) \}$ , where  $L_i(\boldsymbol{\eta})$  is the loss function for the  $i$ th observation. For logistic regression, this bound is easily derived as  $\bar{w} = 1/4$ , since  $w_i = \pi_i(1 - \pi_i)$ . The following corollary provides some insights into maintaining local convexity and ensuring the effectiveness of the coordinate descent algorithm in high-dimensional settings for the logistic regression.

**Corollary 1.1.** For a coordinate-wise objective function  $Q_j(\beta_j)$  of logistic regression, the  $Q_j(\beta_j)$  is strictly convex for any  $j = 1, \dots, p'$ , whenever  $\tau + \frac{1}{\gamma \omega_{\max}} \leq \frac{1}{8\omega_{\max}^2}$ .

**Remark 1.** It is important to note that **Corollary 1.1** does not necessarily extend to other generalized linear models, as the Hessian matrices of loss functions derived from different exponential family distributions may lack a finite upper bound. However, this limitation is not prohibitive, as suitable values for  $\tau$  and  $\gamma$  can still be identified to ensure the convexity of the objective function. Alternatively, a pseudo-upper bound can be introduced by initializing  $\bar{w} \leftarrow \max_i \{ \nabla^2 L_i(\boldsymbol{\eta}_i^*) \}$  at the start of each iteration.

#### 4.2. Threshold Operators

To derive the closed-form threshold operator, we first define the majorizing function of  $Q_j(\beta_j)$  and  $Q_{j|k+}(\beta_{j|k+})$  through the first-order Taylor expansion of the adaptive penalties, as established in **Lemma 1**.

**Lemma 1.** With the adaptive cmenet penalization defined in (4), let  $\tilde{\eta}_{\lambda_{\mathcal{G}}, \tau}(\boldsymbol{\beta}_{\mathcal{G}} | \boldsymbol{\beta}_{\mathcal{G}}^*)$  be the first-order Taylor expansion of  $\eta_{\lambda_{\mathcal{G}}, \tau}(\boldsymbol{\beta}_{\mathcal{G}})$  at fixed  $\boldsymbol{\beta}_{\mathcal{G}}^* \in \mathbb{R}^{|\mathcal{G}|}$ , where  $\mathcal{G}$  represents the sibling or cousin group of one pre-specified effect. Then,

$$\tilde{\eta}_{\lambda_{\mathcal{G}}, \tau}(\boldsymbol{\beta}_{\mathcal{G}} | \boldsymbol{\beta}_{\mathcal{G}}^*) = \eta_{\lambda_{\mathcal{G}}, \tau} \{ \|\boldsymbol{\beta}_{\mathcal{G}}^*\|_{\lambda_{\mathcal{G}}, \omega, \gamma} \} + \Delta_{\mathcal{G}}^* \times \{ \|\boldsymbol{\beta}_{\mathcal{G}}\|_{\lambda_{\mathcal{G}}, \omega, \gamma} - \|\boldsymbol{\beta}_{\mathcal{G}}^*\|_{\lambda_{\mathcal{G}}, \omega, \gamma} \} \quad (11)$$

is a majorizing function of  $\eta_{\lambda_{\mathcal{G}}, \tau}(\boldsymbol{\beta}_{\mathcal{G}})$ , where  $\|\boldsymbol{\beta}_{\mathcal{G}}\|_{\lambda_{\mathcal{G}}, \omega, \gamma} := \sum_{l \in \mathcal{G}} \omega_l g_{\lambda_{\mathcal{G}}, \gamma}(\beta_l)$ . Furthermore, we can derive  $\tilde{Q}_{j|k+}(\beta_{j|k+} | \boldsymbol{\beta}^*)$  as the majorizing function of  $Q_{j|k+}(\beta_{j|k+})$  at fixed  $\boldsymbol{\beta}^* \in \mathbb{R}^{p'}$ ,

$$\begin{aligned} \tilde{Q}_{j|k+}(\beta_{j|k+} | \boldsymbol{\beta}^*) &= L(\beta_{j|k+} | \boldsymbol{\beta}^*) + \eta_{\lambda_{S(j)}, \tau} (\|\boldsymbol{\beta}_{S(j)}^*\|_{\lambda_{S(j)}, \omega, \gamma}) \\ &\quad + \eta_{\lambda_{C(k)}, \tau} (\|\boldsymbol{\beta}_{C(k)}^*\|_{\lambda_{C(k)}, \omega, \gamma}) \\ &\quad + \Delta_{S(j)}^* \omega_{j|k+} \{ g_{\lambda_{S(j)}, \gamma}(\beta_{j|k+}) - g_{\lambda_{S(j)}, \gamma}(\boldsymbol{\beta}_{j|k+}^*) \} \\ &\quad + \Delta_{C(k)}^* \omega_{j|k+} \{ g_{\lambda_{C(k)}, \gamma}(\beta_{j|k+}) - g_{\lambda_{C(k)}, \gamma}(\boldsymbol{\beta}_{j|k+}^*) \}, \end{aligned} \quad (12)$$

where  $*$  indicates the quantity is computed with the current (i.e., most recently updated) value of  $\boldsymbol{\beta}^*$ . The majorizing function  $\tilde{Q}_j(\beta_j | \boldsymbol{\beta}^*)$  for  $Q_j(\beta_j)$  is defined in the similar manner.

For a given coordinate, minimizing the objective function  $Q_j(\beta_j)$  is equivalent to regressing the partial residuals  $\mathbf{y} - \mathbf{X}_{-j}\boldsymbol{\beta}_{-j}^*$  on  $\mathbf{x}_j$ . Expressing the gradient of  $L(\boldsymbol{\beta})$  with respect to  $\beta_j$ , while holding other coefficients fixed, yields  $\nabla L(\beta_j | \boldsymbol{\beta}^*) = -z_j + v_j \beta_j$ , where  $z_j = \frac{1}{n} \mathbf{x}_j^T \mathbf{W} (\tilde{\mathbf{y}} - \mathbf{X}_{-j}\boldsymbol{\beta}_{-j}^*)$  represents the unpenalized solution

for  $\beta_j$ , and  $v_j = \frac{1}{n} \mathbf{x}_j^T \mathbf{W} \mathbf{x}_j$ . Therefore, the coordinate-wise update for  $\beta_j$  (or  $\beta_{j|k+}$ ) in generalized linear regression with adaptive cmenet penalization relies on a closed-form threshold function with respect to  $z_j$  (or  $z_{j|k+}$ ), expressed as

$$\begin{aligned} \beta_j &\leftarrow S_{\lambda_{S(j)}, \lambda_{C(j)}}(z_j; \Delta_{S(j)}, \Delta_{C(j)}), \\ \text{or } \beta_{j|k+} &\leftarrow S_{\lambda_{S(j)}, \lambda_{C(k)}}(z_{j|k+}; \Delta_{S(j)}, \Delta_{C(k)}). \end{aligned}$$

$$S_{\lambda_1, \lambda_2}(z_j; \Delta_1, \Delta_2) = \begin{cases} \begin{cases} z_j, & \text{if } |z_j| \in [v_j \gamma \lambda_{(1)}, \infty), \\ v_j, & \end{cases} \\ \begin{cases} \frac{\text{sgn}(z_j)(|z_j| - \Delta_{(1)} \omega_j)}{v_j - \frac{\Delta_{(1)} \omega_j}{\lambda_{(1)} \gamma}}, & \text{if } |z_j| \in [v_j \lambda_{(2)} \gamma + \Delta_{(1)} \omega_j (1 - \frac{\lambda_{(2)}}{\lambda_{(1)}}), v_j \lambda_{(1)} \gamma), \\ \frac{\text{sgn}(z_j)(|z_j| - \Delta_{(1)} \omega_j - \Delta_{(2)} \omega_j)}{v_j - \frac{\Delta_{(1)} \omega_j}{\lambda_{(1)} \gamma} - \frac{\Delta_{(2)} \omega_j}{\lambda_{(2)} \gamma}}, & \text{if } |z_j| \in [\Delta_{(1)} \omega_j + \Delta_{(2)} \omega_j, v_j \lambda_{(2)} \gamma + \Delta_{(1)} \omega_j (1 - \frac{\lambda_{(2)}}{\lambda_{(1)}})), \\ 0, & \text{if } |z_j| \in [0, \Delta_{(1)} \omega_j + \Delta_{(2)} \omega_j). \end{cases} \end{cases} \quad (13)$$

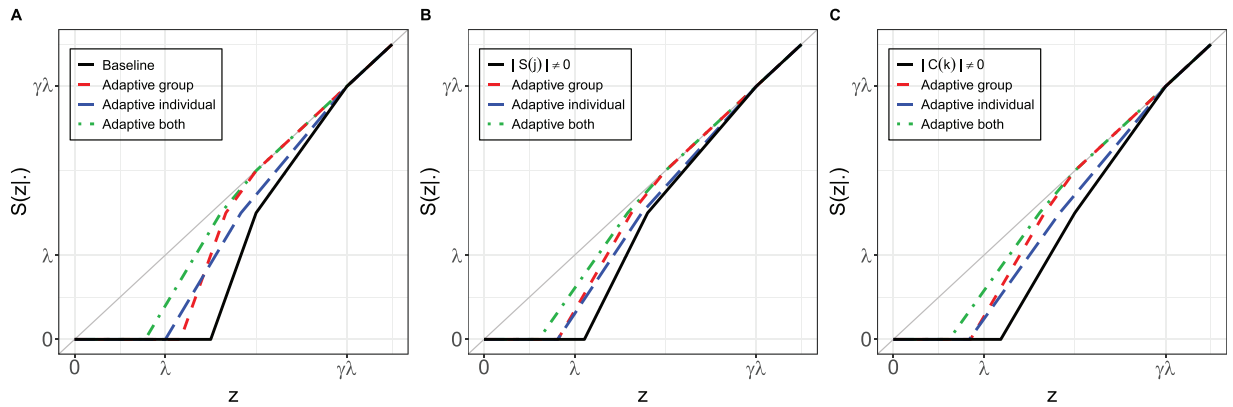
where  $\lambda_{(1)} = \max(\lambda_1, \lambda_2)$  and  $\lambda_{(2)} = \min(\lambda_1, \lambda_2)$ , with  $\Delta_{(1)}$  and  $\Delta_{(2)}$  its corresponding slopes.

It is critical to understand how group and individual weights will influence the threshold function  $S(\cdot; \Delta_1, \Delta_2)$ . As illustrated in Figure 3, the threshold function demonstrates how changes in penalty weights affect both the turning point and the slope of the curve. Introducing group weights into the penalty terms, represented by the dashed line across the three plots, shifts the turning point of the threshold function to the left. This shift indicates that the model becomes more likely to include effects at a lower selection threshold. Additionally, incorporating individual weights, shown by the long dashed line, modifies the slope of the threshold function. This adjustment reduces the shrinkage rate of the coefficients, enabling a more gradual penalization. When both group and individual weights are applied, as depicted by the dot dashed line, the resulting threshold function is the least penalized. This configuration aligns more closely with the diagonal line, reflecting a generalized and less restrictive

approach to penalization. These results highlight the flexibility of the adaptive cmenet penalty structure. Specifically, group weights primarily adjust the selection threshold for effect inclusion, while individual weights regulate the rate of shrinkage. By combining these two weighting schemes, the proposed framework provides a more flexible and adaptive approach to variable selection.

### 4.3. The *glmcmnet* Algorithm and Parameter Tuning

The detailed algorithm of *glmcmnet* using the threshold function is available in the supplemental materials. For the parameter tuning, we use K-fold cross-validation to find an optimal set of parameters  $(\lambda_s, \lambda_c, \gamma, \tau)$ . To efficiently explore the tuning parameter space, we adopt the same strategy as Mak and Wu (2019). Specifically, we first scan a log-spaced grid for the structure parameters  $(\gamma, \tau)$  and pick the best pair by cross-validation (CV); conditional on that choice, we search



**Figure 3.** Comparison of the threshold function  $S(\cdot; \Delta_1, \Delta_2)$  between *cmenet* and *adaptive cmenet* across three scenarios: panel (A) No variable included, panel (B) inclusion of one sibling with coefficient  $\beta = 2$ , and panel (C) inclusion of one cousin with coefficient  $\beta = 2$ . Each plot illustrates four threshold functions: (solid) unweighted penalty, that is, baseline *cmenet* with settings  $(\lambda_s, \lambda_c, \gamma, \tau) = (1, 0.5, 3, 0.25)$ ; (dash) group-weighted penalty with adaptive weights  $(\Omega_{S(j)}, \Omega_{C(j)}, \omega_j) = (2/3, 1, 1)$ ; (long dash) individual-weighted penalty with adaptive weights  $(\Omega_{S(j)}, \Omega_{C(j)}, \omega_j) = (1, 1, 1.5)$ ; and (dot dash) combined group and individual-weighted penalty with adaptive weights  $(\Omega_{S(j)}, \Omega_{C(j)}, \omega_j) = (2/3, 1, 1.5)$ .

a log-spaced grid for the penalty magnitudes  $(\gamma, \tau)$  pair, and then a log-spaced grid for the penalty magnitudes  $(\lambda_s, \lambda_c)$ . The penalties  $\lambda_s$  and  $\lambda_c$  are global, while adaptivity is introduced via sibling and cousin group weights  $\Omega_{S(j)}, \Omega_{C(j)}$ , and an individual weight  $w_j$  for effect  $J$ . To avoid trivial solutions and reduce the search space, we use a weighted start rule: with standardized  $\mathbf{X}$  and centered  $\mathbf{y}$ , define

$$\lambda_{\max}^{\omega}(\rho) = \max_j \frac{|\mathbf{x}_j^T \mathbf{y}|/n}{\rho \Omega_{S(j)} w_j + (1-\rho) \Omega_{C(j)} w_j},$$

where  $\rho = \lambda_s / (\lambda_s + \lambda_c) \in (0, 1)$ . Then the all-zero solution satisfies the KKT conditions (i.e., is a stationary solution) whenever  $\lambda_s + \lambda_c \geq \lambda_{\max}^{\omega}$ . Accordingly, for each fixed  $\rho$  we restrict the grid to  $\lambda_s + \lambda_c \geq \lambda_{\max}^{\omega}(\rho)$ , and set  $(\lambda_s, \lambda_c) = (\rho(\lambda_s + \lambda_c), (1-\rho)(\lambda_s + \lambda_c))$ .

For GLMs with canonical links  $g(\mu) = \eta$ , we use the analogous weighted start rule  $\lambda_{\max}^{\omega} = \max_j \frac{|\mathbf{x}_j^T \mathbf{u}|/n}{\rho \Omega_{S(j)} w_j + (1-\rho) \Omega_{C(j)} w_j}$  with  $\mathbf{u} = \frac{y-\mu}{\text{var}(Y|\eta)} \cdot \frac{\partial \mu}{\partial \eta}$ . Moreover, if  $Q(\boldsymbol{\beta})$  is strictly convex, it is the unique minimizer based on Proposition 1. Therefore, we also enforce the stability condition to guarantee the strict convexity of the coordinate-wise subproblem. For computational efficiency during CV, we leverage the same techniques as those employed in *cmenet*, including warm starts and active set optimization as in Mak and Wu (2019). Thus, the same stability guarantees and practical behavior apply.

One notable advantage of the developed algorithm is the guaranteed descent property due to the MM framework, combined with the cycling procedure of LCD. These techniques ensure that the objective function decreases monotonically at each iteration. Moreover, by leveraging the local convexity property, we can establish that the algorithm converges to a stationary point, providing enhanced stability to the optimization process.

**Proposition 2.** Let  $\{\boldsymbol{\beta}^{(m)}\}$  denote the sequence of coefficient estimates generated at each iteration of the proposed algorithm for logistic regression with adaptive *cmenet* penalization. For all iterations  $m = 0, 1, 2, \dots$ , the following inequality holds:

$$Q(\boldsymbol{\beta}^{(m+1)}) \leq Q(\boldsymbol{\beta}^{(m)}).$$

Furthermore, when  $\tau + \frac{1}{\gamma \omega_{\max}} \leq \frac{1}{8\omega_{\max}^2}$ , the algorithm is guaranteed to converge to a stationary point of  $Q(\boldsymbol{\beta})$ .

## 5. Simulation Study

In this section, two scenarios are considered to examine different aspects of the proposed method. In the first scenario (Section 5.1), we focus on continuous responses, using the exact same simulation settings as those used in the *cmenet* paper (Mak and Wu 2019). Here, the active effects consist entirely of pure siblings, pure cousins, or pure main effects. The objective of this scenario is to assess the effectiveness of adding adaptive weights compared to the original *cmenet* method. This scenario serves as a controlled experiment to demonstrate the improvements because of adaptive weights. In the second scenario (Section 5.2), we consider mixed effects where the active effects include combinations of main effects and their associated siblings or cousins. It reflects more realistic applications. This

scenario evaluates the performance of the proposed method for both continuous and binary responses, providing a broader assessment of its effectiveness across different data types. Each simulation case in both scenarios is replicated 100 times to ensure the robustness of the results.


The model matrix  $\mathbf{X}$  is generated using the equicorrelated latent model described in Mak and Wu (2019). This approach constructs the MEs  $\{\mathbf{x}_j\}_{j=1}^p \in \{-1, +1\}^n$  from a latent matrix  $\mathbf{Z} = (z_{ij})_{i=1, j=1}^{n, p} \in \mathbb{R}^{n \times p}$ , where each row follows a multivariate normal distribution  $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \rho \mathbf{J}_p + (1-\rho) \mathbf{I}_p)$  ( $\mathbf{J}_p$  is the  $p \times p$  matrix of one,  $\mathbf{I}_p$  is the  $p \times p$  identity matrix, and  $\rho \in [0, 1]$  is the correlation parameter). The binary MEs are subsequently derived from this latent structure:

$$x_{ij} = -1\{z_{ij} \leq 0\} + 1\{z_{ij} > 0\}, \quad i = 1, \dots, n, j = 1, \dots, p. \quad (14)$$

Using the equicorrelated latent model, it is straightforward to construct  $4 \binom{p}{2}$  CMEs based on the definition provided in Section 2.1. As a result, the full model matrix for each simulation is represented as  $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n \in \mathbb{R}^{n \times p'}$ , where each row vector is defined as  $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,p}, x_{i,p+1}, \dots, x_{i,p'})' \in \mathbb{R}^{p' \times 1}$ . The correlation structure of this design matrix for the case where  $p = 3$  is illustrated in Figure 4 (see Mak and Wu (2019), sec. 2.2, for additional details).

Two different types of response variables are considered: continuous and binary. Here we focus on logistic regression as a representative case, but the methodology applies similarly to other GLMs. The continuous response variable  $y_i$  in this study is generated by  $y_i = \beta_0 + \sum_{j=1}^{p'} \beta_j x_{i,j} + \epsilon_i$ , where  $\epsilon_i \sim N(0, 1)$ . The binary response is generated based on a logistic regression model. The linear predictor,  $\eta_i$ , is calculated as  $\eta_i = \beta_0 + \sum_{j=1}^{p'} \beta_j x_{i,j}$ . Then the binary response variable  $y_i$  is simulated from a Bernoulli distribution  $y_i \sim \text{Bernoulli}\left(\frac{1}{1+\exp(-\eta_i)}\right)$ .

The *cmenet* (R package *cmenet* V0.1.2; Mak 2022) is compared to the proposed approach to assess the effectiveness of adaptive weights. Several benchmark methods of variable selection are also considered, including the Lasso (R package *glmnet* V4.1.8; Friedman, Tibshirani, and Hastie 2010; Tay, Narasimhan, and Hastie 2023), the adaptive Lasso (R package *glmnet* V4.1.8; Zou 2006), the non-concave penalty method MCP (R package *ncvreg* V3.14.3; Breheny and Huang 2011; Zhang 2010), the bi-level selection method GEL (R package *grpreg* V3.5.0; Breheny 2015), and the popular two-way interaction selection method *hierNet* (R package *hierNet* V1.9; Bien and Tibshirani 2020). Three criteria are used for measuring the selection accuracy:  $F_1$  score, Precision, and True Positive Rate (TPR) or Recall. Denote the number of true positives, false positives, true negatives, and false negatives as TP, FP, TN, and FN, respectively. We can calculate the Precision as  $\text{Precision} = \text{TP}/(\text{TP} + \text{FP})$ , the Recall as  $\text{Recall} = \text{TP}/(\text{TP} + \text{FN})$ , and the  $F_1$  score as  $F_1 = (2\text{Precision} \times \text{Recall})/(\text{Precision} + \text{Recall})$ . In addition to evaluating selection performance, we also measure the prediction performance. The metrics for prediction evaluation depend on the type of response. For continuous responses, we evaluate the performance using the mean squared prediction error (MSPE). For binary responses, we assess performance based on the misclassification rate (MCR). Here,  $\text{MSPE} =$



**Figure 4.** Correlation matrix of the simulated full model matrix ( $n = 100, p = 3$ ) generated from the equicorrelation latent model with (A)  $\rho = 0$  and (B)  $\rho = 1/\sqrt{2}$ .

$\frac{1}{n} \sum_{i=1}^n (y_{i,new} - \hat{y}_{i,new})^2$ , and  $MCR = \frac{1}{n} \sum_{i=1}^n I(y_{i,new} \neq \hat{y}_{i,new})$ . Since `hierNet` does not encode conditional relationships within interactions, element-wise selection metrics are not directly comparable. Therefore, for `hierNet` we report model size (number of selected effects) and out-of-sample prediction error as indicators of parsimony and predictive performance. According to the principle of parsimony (Stoica and Söderström 1982; Wu and Hamada 2011), when two models achieve comparable prediction performance, the model with fewer parameters is preferable to one with a larger number of parameters.

### 5.1. Evaluation on Adaptive Weights

In this section, we assess the performance of adaptive `cmenet` under the same simulation settings as those used for `cmenet` in Mak and Wu (2019), with focus exclusively on continuous responses. Following their framework, we consider three distinct types of active effects in the underlying true model: (a) all active effects are main effects, such as  $\{A, B, C, D\}$ ; (b) all active effects are sibling effects, such as  $\{A|B+, A|D+, B|C-, B|E+\}$ ; and (c) all active effects are cousin effects, such as  $\{C|A+, B|A-, D|B+, E|B+\}$ . For the cases involving sibling and cousin effects, we ensure that no other types of group structures are included among the active effects. In this study, the sample size is  $n = 100$ , and the number of MEs is  $p = 40$ , resulting in a total of  $p' = p + 4\binom{p}{2} = 3160$  effects (including both MEs and CMEs). The coefficients for active effects are set to  $\beta_j = 5$  and the intercept is  $\beta_0 = 12$ . To evaluate the impact of correlations among MEs, this study considers two correlation settings: uncorrelated ( $\rho = 0$ ) and moderately correlated ( $\rho = 1/\sqrt{2}$ ). The number of active groups is varied ( $\{4, 6, 8, 10, 12\}$ ), with two active effects within each group.

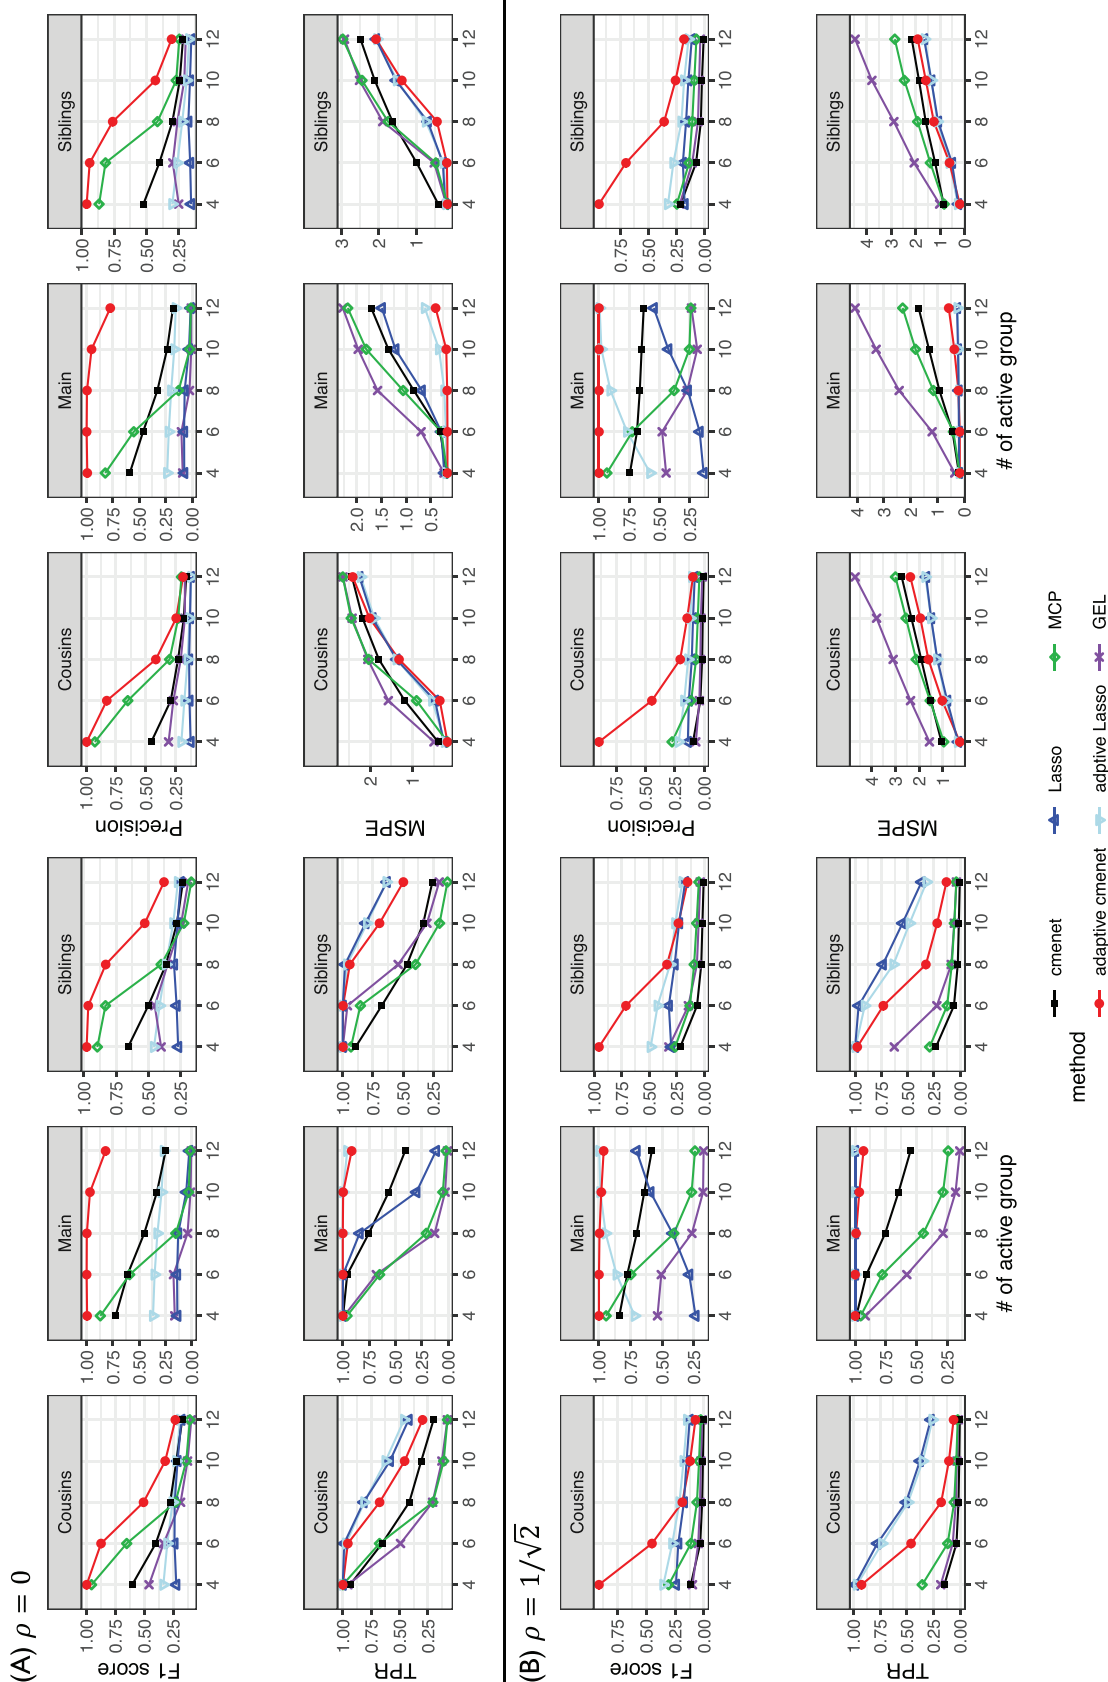
Figure 5 shows the simulation results for scenario 1. Across nearly all simulation settings of the uncorrelated case, the proposed method performs significantly better than not only the `cmenet` but also other well-established methods. By incorporat-

ing adaptive weights, it achieves the highest  $F_1$  scores, striking a better balance between Precision and TPR. The adaptive `cmenet` always demonstrates much higher Precision by selecting fewer irrelevant variables, while maintaining competitive TPR. It also produces lower or comparable MSPE to other methods, highlighting its capability to deliver strong prediction performance alongside improved selection accuracy.

Comparing the adaptive `cmenet` to the Lasso and the adaptive Lasso, a tradeoff becomes apparent. Sometimes, the Lasso-based methods demonstrate higher TPR than adaptive `cmenet`, which means they identify more true active effects. However, the adaptive `cmenet` achieves higher Precision, highlighting its ability to select fewer irrelevant variables. This pattern is persistent across cousins, main effects, and siblings, emphasizing the relative strengths of adaptive `cmenet` in maximizing true positive detection while producing parsimonious models.

When comparing the relative performance across the three types of active effects, there are notable differences. Main effects are the easiest to identify, as indicated by consistently high selection criteria values, with many approaching or reaching 1. This suggests a near-perfect selection of all active effects without introducing false positives. In contrast, cousin effects are more difficult to detect, as evidenced by lower selection criteria values, particularly as the number of active groups increases. Sibling effects fall between these two extremes, exhibiting moderate performance in selection accuracy. The superior identification of main effects compared to cousin and sibling effects highlights the effectiveness of the CME reduction principle within the penalization framework.

Now considering the MEs are correlated ( $\rho = 1/\sqrt{2}$ ) as shown in Figure 5(B), the proposed adaptive `cmenet` demonstrates a clear advantage over other methods in terms of selection accuracy, particularly when the number of active groups is small. The improvement in performance under correlated settings is more substantial than in the uncorrelated case. For instance, in the uncorrelated scenario, certain methods, such as the MCP,



**Figure 5.** Example 1. Simulation results for continuous responses with  $n = 100, \rho = 40$  for (A) uncorrelated MEs  $\rho = 0$  and (B) correlated MEs  $\rho = 1/\sqrt{2}$ . The x-axis shows the number of active groups, while the y-axis displays the F1 score, Precision, True Positive Rate (TPR), and MSPE. Each column corresponds to a type of active effect: cousins (left), main effects (middle), and siblings (right).

achieve comparable  $F_1$  scores to adaptive cmenet for cousin and sibling effects. However, in the correlated case, the gap of  $F_1$  scores between the adaptive cmenet and the MCP becomes more pronounced compared to the uncorrelated setting. This superiority only diminishes when the number of active groups is large. The result in the correlated case underscores the effectiveness of the CME coupling principle in capturing group structures within CMEs, as the CME group structure becomes more evident with higher correlations (Figure 4).

## 5.2. Evaluation on Complex Effect Structures

In the second simulation scenario, we consider two distinct types of active effects in the underlying true model for both continuous and binary responses: (a) active effects consist of main effects and their siblings, such as  $\{A, B, A|B+, A|D+, B|C-, B|E+\}$ ; and (b) active effects consist of main effects and their cousins, such as  $\{A, B, C|A+, B|A-, D|B+, E|B+\}$ . The inclusion of main effects introduces overlapping and ambiguous groupings. Unlike the first example, we do not ensure the absence of other types of group structures among the active effects. This setup, which closely resembles real-world scenarios, provides additional challenges for selection methods. Here, the sample size is  $n = 50$ , and the number of MEs is  $p = 20$ , leading to a total of  $p' = p + 4\binom{p}{2} = 780$  effects. For this example, the coefficients of MEs are set to  $\beta_{ME} = 5$ , while the coefficients of CMEs are  $\beta_{CME} = 1$ . The correlation settings and the number of active groups follow the same configurations as in Scenario 1. As a sensitivity analysis, the supplemental materials report an additional experiment with a fixed total of four groups and exactly two active effects, where the effect size  $\beta_{CME}$  varies from 0.25 to 16 (with fixed  $\beta_{ME}=5$ ) to evaluate the impact of signal-to-noise ratio and effect-size ratio  $\beta_{ME}/\beta_{CME}$  on selection and prediction.

Considering continuous responses in Figure 6(A) and (B), the adaptive cmenet method outperforms other methods across both uncorrelated and correlated settings for main+cousin effects. The improvement is particularly clear in the correlated case, where the stronger group dependencies enhance adaptive cmenet's ability to balance true positive and false positive rates. For main+sibling effects, the adaptive cmenet performs better than other methods in the uncorrelated case, reflecting the effectiveness of detecting sibling effects when group structures are weakly defined. In the correlated case, the adaptive cmenet demonstrates improved performance compared to the uncorrelated case, particularly as the number of active groups increases. However, the  $F_1$  score curves of adaptive cmenet and adaptive Lasso largely overlap, indicating comparable overall performance, although the adaptive cmenet achieves higher Precision and lower TPR. A possible explanation is that both methods rely on the same ridge estimator for generating initial weights, which influences the penalty structure in both approaches. Across all scenarios, the adaptive cmenet consistently delivers the lowest MSPE, underscoring its ability to integrate accurate selection with strong predictive performance.

For the results of binary responses shown in Figure 6, the adaptive cmenet exhibits competitive performance across uncorrelated and correlated settings, effectively balancing selection accuracy and prediction. For main+cousin effects,

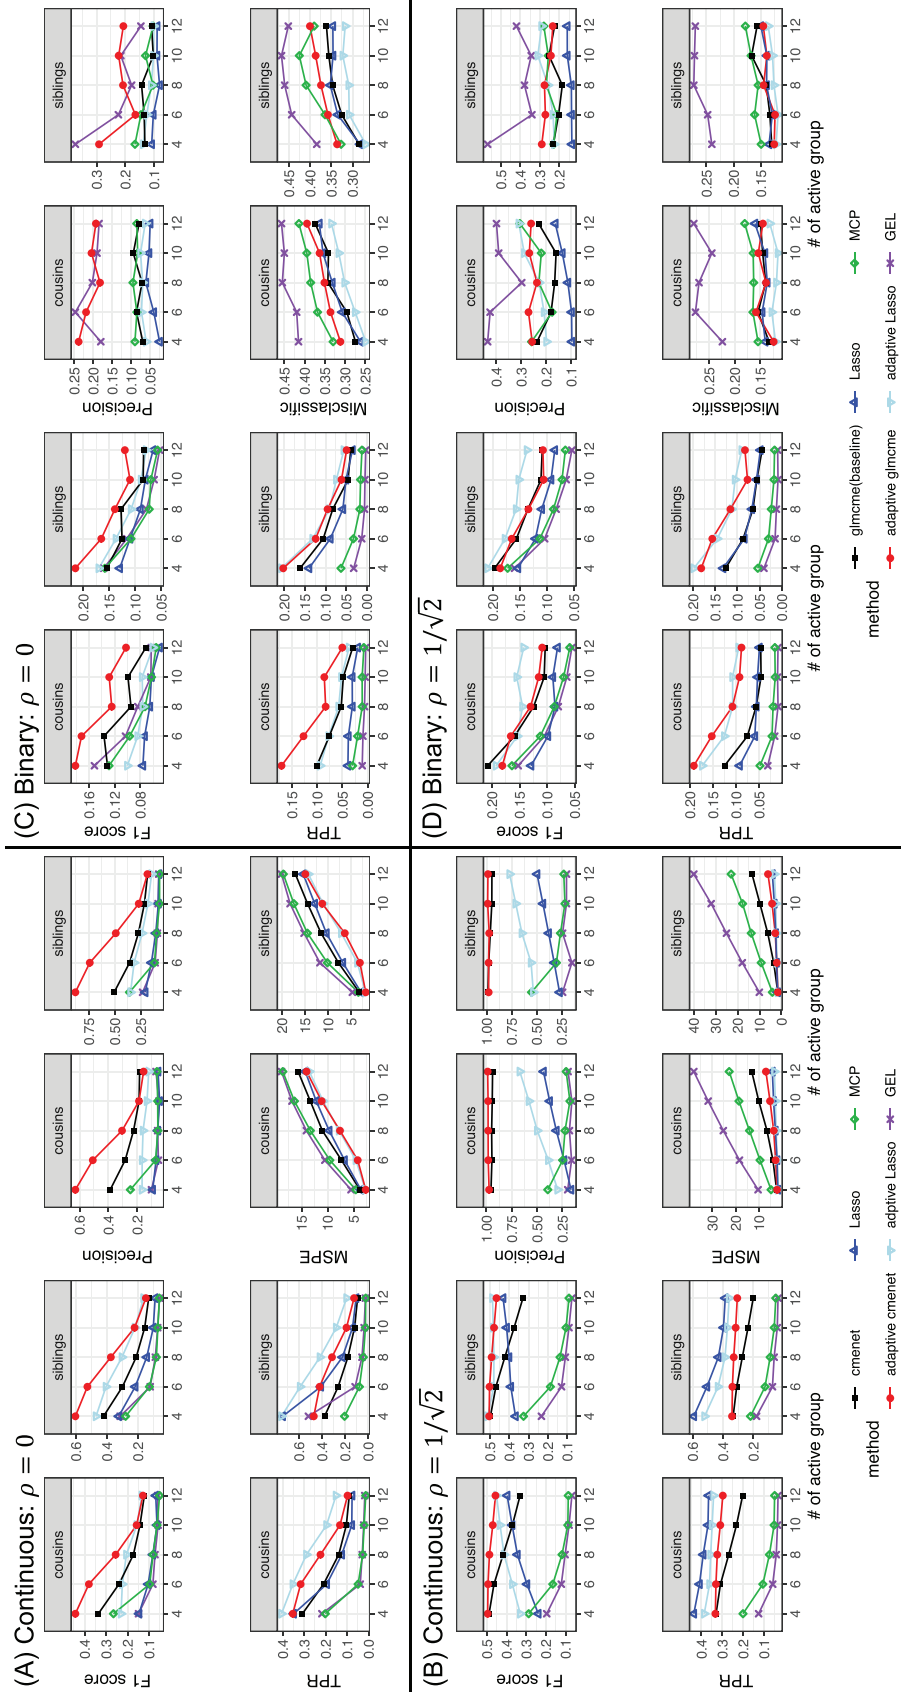
the adaptive cmenet achieves the highest  $F_1$  scores and TPR in the uncorrelated setting, highlighting its ability to capture true positive relationships while maintaining strong precision. For main+sibling effects, the adaptive cmenet performs competitively with the adaptive Lasso in both uncorrelated and correlated cases, offering a strong tradeoff between the precision and TPR. While it does not surpass the adaptive Lasso in  $F_1$  scores in some scenarios, it still maintains lower false discovery rates. Although the MCP and GEL methods exhibit slightly higher Precision when  $\rho = 1/\sqrt{2}$ , the adaptive cmenet strikes a better balance overall by effectively managing both false discoveries and true positives.

In Figure 7, hierNet achieves relatively better predictive accuracy than adaptive cmenet with the advantage widening under correlation ( $\rho = 1/\sqrt{2}$ ) and as the number of active groups increases. However, the tradeoff is obvious: hierNet selects noticeably larger models, suggesting it leverages additional pairwise interactions to capture dependence structure and boost prediction. By contrast, the adaptive cmenet produces more parsimonious models with an interpretable CME structure, making it preferable when interpretability and parsimony are the primary objectives.

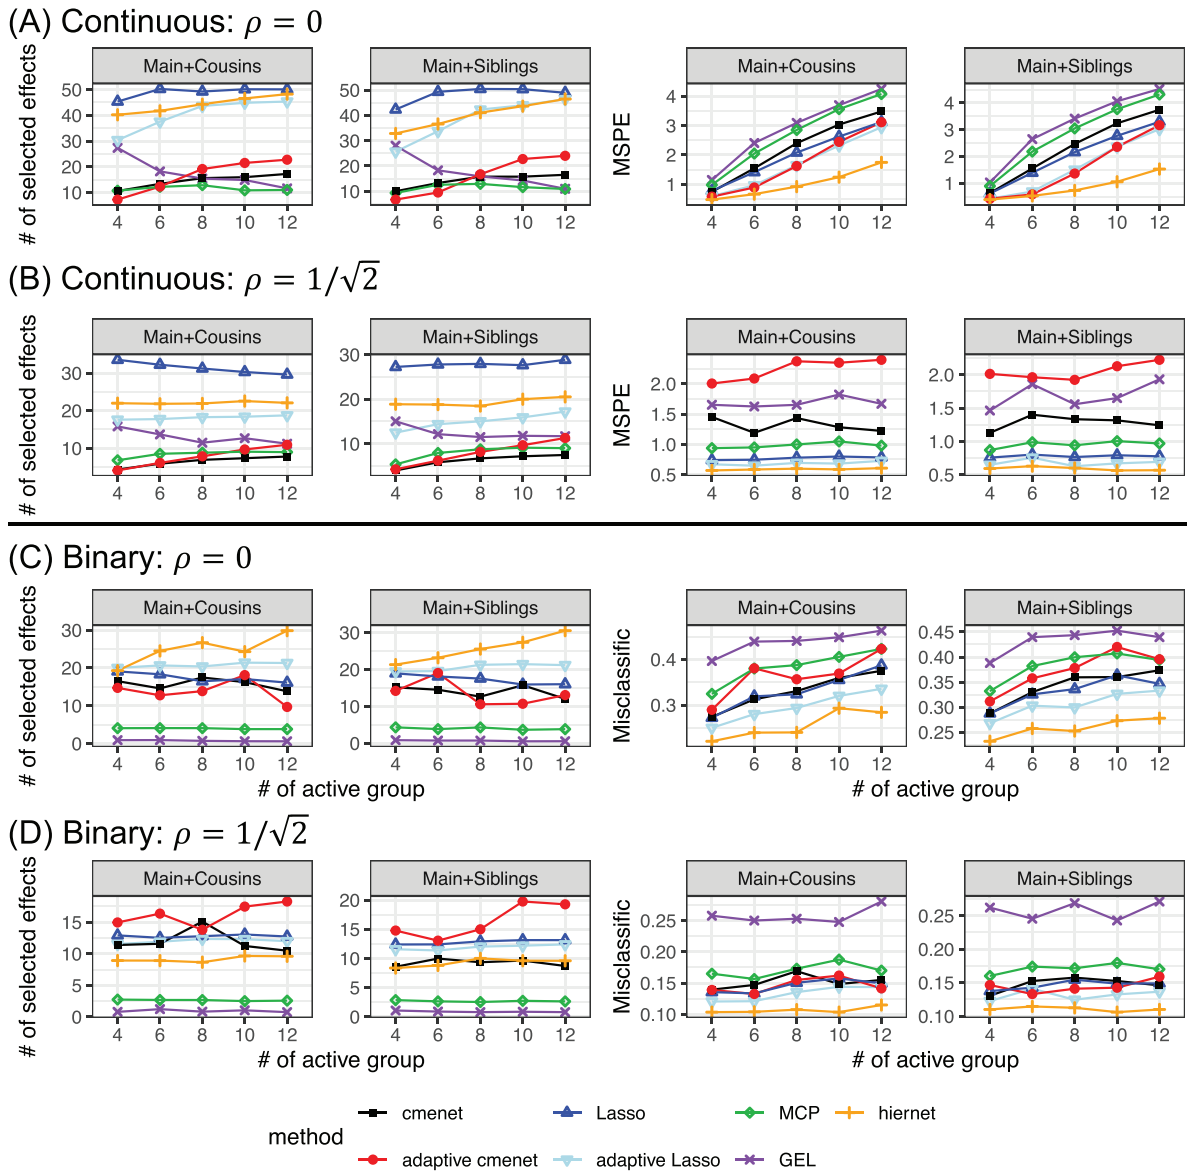
## 6. Case Study of Gene Association

In this section, we apply the proposed method to two gene association studies to evaluate its performance in real-world datasets: one with a continuous response and the other with a binary response. In both studies, the main effects are binary single-nucleotide polymorphism (SNP) markers. Each SNP represents a specific genetic variation at a particular location in the genome and can exist in three genotypic forms: homozygous reference (AA), heterozygous (Aa), or homozygous variant (aa). Traditionally, SNPs are encoded using an additive  $\{0, 1, 2\}$  coding scheme to represent AA, Aa, and aa, respectively (Lewis 2002). In both studies, we analyze datasets from inbred lines, where the inbreeding process ensures that all single nucleotide polymorphisms (SNPs) are fully homozygous (either AA or aa), eliminating heterozygous genotypes. Therefore, the SNPs in these datasets are inherently binary-coded, without the heterozygous category. To align with the framework of our proposed method, we relabel  $\{0, 1\}$  to  $\{-1, 1\}$ . All tuning and performance reporting are conducted out-of-sample using 5-fold cross-validation. The goal of these analyses is to identify significant gene CMEs, which quantify the additive or dominant effect of one genetic variant conditional on the presence or absence of another genetic variant.

In both studies, we compare the proposed approach to the cmenet, Lasso, MCP, GEL, and hierNet (R package hierNet V1.9; Bien and Tibshirani 2020). The hierNet method (Bien, Taylor, and Tibshirani 2013) is a widely used method for selecting two-way interactions (2FI) in high-dimensional data. By comparing the adaptive cmenet to the hierNet, we aim to evaluate its practical advantage in identifying interaction relationships between genes. Note that the true underlying active effects are typically unknown in these real-world studies. Here we focus on assessing whether the estimated model can select a smaller number of effects while maintaining high prediction accuracy (Stoica and Söderström 1982). Such a criteria is crucial



**Figure 6.** Example 2. Simulation results with  $n = 50, p = 20$  for (A) continuous response with uncorrelated MEs  $\rho = 0$ , (B) continuous response with correlated MEs  $\rho = 1/\sqrt{2}$ , (C) binary response with uncorrelated MEs  $\rho = 0$ , and (D) binary response with correlated MEs  $\rho = 1/\sqrt{2}$ . The x-axis shows the number of active groups, while the y-axis displays the  $F_1$  score, Precision, True Positive Rate (TPR), and MSPE (for continuous) or Misclassification Rate (for binary). Each column corresponds to a type of active effect: main+cousins (left), and main+siblings (right).



**Figure 7.** Example 2. Comparison to hierNet with  $n = 50, p = 20$  for (A) continuous response with uncorrelated MEs  $\rho = 0$ , (B) continuous response with correlated MEs  $\rho = 1/\sqrt{2}$ , (C) binary response with uncorrelated MEs  $\rho = 0$ , and (D) binary response with correlated MEs  $\rho = 1/\sqrt{2}$ . The x-axis shows the number of active groups, while the y-axis displays the total # of selected effects, and MSPE (for continuous) or Misclassification Rate (for binary). Each column corresponds to a type of active effect: main+cousins (left), and main+siblings (right).

in genomic studies, where overly complex models with excessive effects may lead to reduced interpretability and overfitting. A model that selects fewer effects with comparable or better prediction accuracy is more likely to capture meaningful biological relationships while avoiding spurious findings. Additionally, we report the  $p$ -values of the selected effects using regular (generalized) linear regression to validate the relevance of the selected effects. This step can provide further evidence for the method's ability to identify biologically meaningful associations.

The first study focuses on maize flowering time, a continuous variable measuring the number of days to flowering. We analyze a subset of the dataset from Buckler et al. (2009) containing  $n = 150$  observations, and  $p = 40$  main effects represented by binary SNP markers. Our proposed method demonstrates superior performance compared to several established methods. As

shown in Table 2(a), the adaptive cmenet identifies 14 significant effects out of the 19 selected, maintaining a lower MSPE than the hierNet. While the hierNet selects more significant effects, it also includes more insignificant ones, resulting in a higher MSPE. On the other hand, the MCP method selects fewer total effects (most of which are significant) but still exhibits a higher MSPE than the adaptive cmenet. These results demonstrate that the proposed method can efficiently capture key gene interactions while minimizing prediction error. Additionally, the  $p$ -value analysis of selected effects in Table 3 reveals that adaptive cmenet identifies more significant effects than other methods, with considerable overlap between the significant effects selected by the adaptive cmenet and the hierNet.

The second study investigates chlorosis phenotypes in *Arabidopsis thaliana* using data from a genome-wide association

**Table 2.** Selection performance and predictive performance of six methods in the gene association study.

	Adaptive cmenet	cmenet	Lasso	MCP	GEL	hierNet
<b>(A) Maize flowering time</b>						
# of selected effects	19	24	43	12	21	66
# of significant selected effects ( $p < 0.05$ )	14	5	11	11	10	16
MSPE	3.83	5.47	4.95	6.38	5.81	4.68
<b>(B) Arabidopsis thaliana</b>						
# of selected effects	9	—	13	11	5	93
# of significant selected effects ( $p < 0.05$ )	7	—	4	1	5	29
Misclassification Rate	0.1	—	0.1	0.13	0.3	0.01

**Table 3.**  $P$ -values of some selected effects for each method from a linear regression in Maize flowering time study.

Method	$g1 g17-$	$g1 g39-$	$g9 g10-$	$g11 g38+$	$g12 g38-$	$g14 g38+$	$g15 g35+$	$g18 g33-$	$g27 g26-$
Adaptive cmenet	—	0.0004	0.0183	0.0000	0.0351	0.0042	0.0003	0.0062	0.0143
cmenet	—	0.1214	0.1281	0.0091	0.0124	0.0966	0.0007	0.1629	0.0180
Lasso	0.0392	0.0953	0.1951	0.0043	0.1050	0.0790	0.0372	0.0343	0.0200
MCP	0.2573	0.0544	0.0196	0.0333	0.0646	0.8667	0.0022	0.0133	0.0242
GEL	0.0085	0.0483	0.0049	0.0029	0.1048	0.0001	0.0001	0.0008	0.0081
hierNet	0.0259	0.0510	0.0002	0.0057	0.0000	0.0259	0.0002	0.0001	0.0001

**Table 4.** The  $p$ -values of some selected effects for each method from a logistic regression in Arabidopsis thaliana study.

	Adaptive cmenet	Lasso	MCP	GEL	hierNet	Univariate regression
2 : 5804076 2 : 8750047+	0.0069	0.0076	—	0.0000	0.0214	0.0007
2 : 8750047 5 : 1931167+	0.0337	0.0960	—	0.0004	0.0197	0.0005
2 : 18752840 3 : 4818602+	0.0122	0.0298	—	0.0008	0.0077	0.0066
2 : 18752840 4 : 11580143—	0.0096	0.0133	—	0.0011	0.0276	0.0006
3 : 7392098 3 : 21632608+	0.0028	0.0491	—	—	0.3579	0.0034
4 : 17038400 2 : 9256095+	0.5980	0.0502	—	—	0.0020	0.0012
4 : 17038400 5 : 1931167—	0.0314	0.2270	—	—	0.0197	0.0013
5 : 20409250 2 : 9256095—	0.0063	0.5506	—	—	0.0020	0.0153

study (GWAS) conducted by Atwell et al. (2010). Arabidopsis thaliana is a widely distributed, predominantly self-fertilizing model plant that exhibits substantial genetic variation for numerous adaptively significant traits (Atwell et al. 2010). The dataset includes genotypic information for 250,000 single-nucleotide polymorphisms (SNPs) and phenotypic measurements for 107 traits across 199 strains, with data publicly available at <https://cynin.gmi.oeaw.ac.at/home/resources/atpolydb>. In terms of the visual chlorosis phenotype, assessed at 16°C after five weeks of growth, 84 strains exhibit chlorosis, while 92 do not. For this analysis, a subset of the dataset is selected, consisting of 100 strains (55 with visual chlorosis and 45 without), with  $p = 63$  main effects corresponding to binary SNP markers. The response variable is binary, representing the presence or absence of visual chlorosis. Consider Table 2(B), the adaptive cmenet demonstrates superior performance compared to other methods by identifying a higher number of significant effects while maintaining a competitive misclassification rate. In contrast, the hierNet method selects the highest number of significant effects but also includes a larger number of insignificant effects, resulting in an overfitted model. The MCP method selects a comparable number of effects to adaptive cmenet; however, only one of them is statistically significant, and notably, it does not identify any effects that overlap with those selected by other methods. While the GEL method identifies fewer effects, of which five are significant, it exhibits higher misclassification rates. The adaptive cmenet achieves a better balance by selecting a greater number of significant dominant effects while minimizing the inclusion of insignificant

effects, effectively combining selection accuracy with predictive reliability. The  $p$ -value analysis of selected effects in Table 4 shows that both the adaptive cmenet method and the hierNet method consistently identify overlapping significant dominant effects. Notably, all CMEs identified by the adaptive cmenet are validated as statistically significant in the univariate regression model, further highlighting the superior selection performance of the proposed method.

To investigate biological interpretations for the selected CMEs (Conditional Main Effects), we mapped each associated SNP to its nearest TAIR10 gene (within  $\pm 10$  kb) and functionally annotated the genes using Gene Ontology, KEGG, and AraCyc. Intriguingly, the top CMEs frequently pair loci whose annotated functions suggest a conditional relationship relevant to chlorosis. For instance, we detected a CME (4 : 17038400|5 : 1931167—) between APX5 (AT4G35970) and 2-Cys Peroxiredoxin B (AT5G06290). Both enzymes are critical components of ROS detoxification and redox buffering in photosynthetic tissues. The negative polarity of this CME indicated that the lack of 2-Cys Prx B removes its protective or regulatory function, making the APX5-associated genetic variation a more critical factor in determining the chlorotic phenotype. This finding is consistent with the compensation mechanism between the ascorbate peroxidase and peroxiredoxin redox pathways: loss of 2-Cys Prx activity can be compensated by other chloroplast antioxidant systems, most notably the ascorbate—glutathione cycle acting with superoxide dismutase, so that reliance shifts toward APX-linked routes (Puerto-Galán et al. 2013). This example elaborates the contribution of CMEs beyond traditional

interactions: they capture context-dependent effects, where the influence of one locus emerges only in the presence (or absence) of a second, biologically related locus.

## 7. Discussion

This work has focused on adaptive bi-level variable selection of conditional main effects for generalized linear models. The adaptive penalty structure used in the proposed method can enable a more flexible and adaptable selection of CMEs. The proposed adaptive cmenet enhances the ability to capture important conditional relationships while reducing the coupling constraints.

Here we would like to remark that there are several aspects of improving the proposed method. First, similar to the IRLS algorithm, the algorithm for parameter estimation in Section 4 may not converge to a global optimal solution with respect to the actual objective function. It will be interesting to further investigate the properties of numerical convergence of the parameter estimation in the proposed method. Nevertheless, the proposed approach demonstrates robust performance in numerical studies. Second, while the adaptive cmenet incorporates additional penalization to account for group structures such as cousins and siblings, the choice of initial weights could limit its ability to fully exploit these structures. Thus, alternative weighting schemes tailored to the CME framework may further improve the performance of adaptive cmenet. Third, while the similar theoretical framework for adaptive cmenet can be extended to canonical-link Poisson and Gamma regressions, the computational details are case-specific: variance functions (Poisson  $\text{var} = \mu$ ; Gamma  $\text{var} = \phi\mu^2$ ) and working weights alter curvature, step sizes, and numerical safeguarding. Cross-validation proceeds with the appropriate deviance loss, and warm starts/active-set updates remain effective, but require these case-specific investigations. Fourth, it is important to further improve the computational efficiency of the algorithm for scalability. Last, one can consider extending the proposed methodology to handle CMEs for multi-level or continuous variables, broadening its applicability to diverse datasets and experimental settings. Finally, given the origins of CMEs in experimental design, one can explore the use of CME in optimal experimental design and sequential design.

## Supplementary Materials

The Supplement.pdf contains additional methodology, proofs, algorithm details, and extended results. The Simulation.zip contains R scripts used to reproduce simulations and figures reported in the manuscript.

## Disclosure Statement

No potential conflict of interest was reported by the author(s).

## References

Atwell, S., Huang, Y. S., Vilhjálmsson, B. J., Willems, G., Horton, M., Li, Y., Meng, D., Platt, A., Tarone, A. M., Hu, T. T., et al. (2010), "Genome-Wide Association Study of 107 Phenotypes in Arabidopsis Thaliana Inbred Lines," *Nature*, 465, 627–631. [3,15]

Bien, J., Taylor, J., and Tibshirani, R. (2013), "A Lasso for Hierarchical Interactions," *Annals of Statistics*, 41, 1111. [12]

Bien, J., and Tibshirani, R. (2020), *hierNet: A Lasso for Hierarchical Interactions*. R package version 1.9. [9,12]

Breheeny, P. (2015), "The Group Exponential Lasso for Bi-level Variable Selection," *Biometrics*, 71, 731–740. [3,9]

Breheeny, P., and Huang, J. (2009), "Penalized Methods for Bi-level Variable Selection," *Statistics and Its Interface*, 2, 369. [1,2]

——— (2011), "Coordinate Descent Algorithms for Nonconvex Penalized Regression, with Applications to Biological Feature Selection," *Annals of Applied Statistics*, 5, 232–253. [9]

Buckler, E. S., Holland, J. B., Bradbury, P. J., Acharya, C. B., Brown, P. J., Browne, C., Ersoz, E., Flint-Garcia, S., Garcia, A., Glaubitz, J. C., et al. (2009), "The Genetic Architecture of Maize Flowering Time," *Science*, 325, 714–718. [14]

Fang, K., Wang, X., Zhang, S., Zhu, J., and Ma, S. (2015), "Bi-Level Variable Regression via Adaptive Sparse Group Lasso," *Journal of Statistical Computation and Simulation*, 85, 2750–2760. [5]

Friedman, J., Hastie, T., Höfling, H., and Tibshirani, R. (2007), "Pathwise Coordinate Optimization," *The Annals of Applied Statistics*, 1, 302–332. [2,7]

Friedman, J., Tibshirani, R., and Hastie, T. (2010), "Regularization Paths for Generalized Linear Models via Coordinate Descent," *Journal of Statistical Software*, 33, 1–22. [9]

Fu, W. J. (1998), "Penalized Regressions: The Bridge versus the Lasso," *Journal of Computational and Graphical Statistics*, 7, 397–416. [2,7]

Green, P. J. (1984), "Iteratively Reweighted Least Squares for Maximum Likelihood Estimation, and Some Robust and Resistant Alternatives," *Journal of the Royal Statistical Society, Series B*, 46, 149–170. [2]

Huang, J., Ma, S., Xie, H., and Zhang, C.-H. (2009), "A Group Bridge Approach for Variable Selection," *Biometrika*, 96, 339–355. [2]

Hunter, D. R., and Lange, K. (2000), "Quantile Regression via an MM Algorithm," *Journal of Computational and Graphical Statistics*, 9, 60–77. [2,7]

Lewis, C. M. (2002), "Genetic Association Studies: Design, Analysis and Interpretation," *Briefings in Bioinformatics*, 3, 146–153. [12]

Mak, S. (2022), *cmenet: Bi-Level Selection of Conditional Main Effects*. R package version 0.1.2. [9]

Mak, S., and Wu, C. J. (2019), "cmenet: A New Method for Bi-level Variable Selection of Conditional Main Effects," *Journal of the American Statistical Association*, 114, 844–856. [1,2,3,8,9,10]

Mazumder, R., Friedman, J. H., and Hastie, T. (2011), "Sparsenet: Coordinate Descent with Nonconvex Penalties," *Journal of the American Statistical Association*, 106, 1125–1138. [2,7]

Puerto-Galán, L., Pérez-Ruiz, J. M., Ferrández, J., Cano, B., Naranjo, B., Nájera, V. A., González, M., Lindahl, A. M., and Cejudo, F. J. (2013), "Overoxidation of Chloroplast 2-cys Peroxiredoxins: Balancing Toxic and Signaling Activities of Hydrogen Peroxide," *Frontiers in Plant Science*, 4, 310. [15]

Shibata, M., Terada, A., Kawaguchi, T., Kamatani, Y., Okada, D., Nagashima, K., Ohmura, K., Matsuda, F., Kawaguchi, S., Sese, J., et al. (2024), "Identification of Epistatic SNP Combinations in Rheumatoid Arthritis Using Lamplink and Japanese Cohorts," *Journal of Human Genetics*, 69, 541–547. [1]

Stoica, P., and Söderström, T. (1982), "On the Parsimony Principle," *International Journal of Control*, 36, 409–418. [10,12]

Su, H., and Wu, C. J. (2017), "Cme Analysis: A New Method for Unraveling Aliased Effects in Two-Level Fractional Factorial Experiments," *Journal of Quality Technology*, 49, 1–10. [1]

Tay, J. K., Narasimhan, B., and Hastie, T. (2023), "Elastic Net Regularization Paths for all Generalized Linear Models," *Journal of Statistical Software*, 106, 1–31. [9]

Wu, C. J. (2015), "Post-Fisherian Experimentation: From Physical to Virtual," *Journal of the American Statistical Association*, 110, 612–620. [1,2]

Wu, C. J., and Hamada, M. S. (2011), *Experiments: Planning, Analysis, and Optimization*, Hoboken, NJ: Wiley & Sons [10]

Zhang, C.-H. (2010), "Nearly Unbiased Variable Selection under Minimax Concave Penalty," *The Annals of Statistics*, 38, 894–942. [3,9]

Zou, H. (2006), "The Adaptive Lasso and its Oracle Properties," *Journal of the American Statistical Association*, 101, 1418–1429. [2,5,9]