

Sparse high-dimensional linear mixed modeling with a partitioned empirical Bayes ECM algorithm

SUPPLEMENTARY MATERIAL

Anja Zgodic¹, Ray Bai², Jiajia Zhang¹, Peter Olejua¹,
Alexander C. McLain^{1*}

^{1*}Department of Epidemiology and Biostatistics, University of South
Carolina, 915 Greene Street, Columbia, South Carolina, 29208, United
States.

²Department of Statistics, University of South Carolina, 1523 Greene
Street, Columbia, South Carolina, 29208, United States.

*Corresponding author(s). E-mail(s): mclaina@mailbox.sc.edu;

A Proofs of Propositions 1 and 2

We reiterate Propositions 1–2 from the main text and provide proofs for them below.

Proposition 1. *Under a standard EM algorithm, the parameters in the M-step for LMM-PROBE are not always estimable. That is, the maximizer of $Q_{EM}(\boldsymbol{\eta} \mid \Theta^{(t-1)})$, where $\boldsymbol{\eta} = (\boldsymbol{\beta} \ \boldsymbol{\omega} \ \boldsymbol{\tau})'$, is not always unique.*

Proposition 2. *Let $Q_{CM}^{M1}(\boldsymbol{\eta} \mid \Theta^{(t-1)})$ and $Q_{CM}^{M2}(\boldsymbol{\eta} \mid \Theta^{(t-1)})$ denote the two M-step quantities in the PX-ECM algorithm for LMM-PROBE. Assuming no perfect collinearity between $\mathbf{X}_k$ and $\mathbf{V}$ for any k , the maximizers of $Q_{CM}^{M1}(\boldsymbol{\eta} \mid \Theta^{(t-1)})$ and $Q_{CM}^{M2}(\boldsymbol{\eta} \mid \Theta^{(t-1)})$ always exist and are unique.*

Proof of Proposition 1. Consider a standard M-step for the regression parameters in (1) without any additional non-sparse predictors. We add the expanded parameter $\boldsymbol{\tau}$ for r -dimensional random effects. For the small p non-sparse situation, Liu, Rubin, and Wu (1998) considered this model (see Section 4.1 therein) and were able to demonstrate superior convergence to the non-expanded parameter version. In this situation,

the M-step maximizes

$$Q_{\text{EM}}(\boldsymbol{\eta} \mid \Theta^{(t-1)}) = -E_{\boldsymbol{\gamma}\boldsymbol{\beta}} \left[\{Y - U(\mathbf{b})'\boldsymbol{\eta}(\boldsymbol{\gamma})\}' \{Y - U(\mathbf{b})'\boldsymbol{\eta}(\boldsymbol{\gamma})\} \mid \Theta^{(t-1)} \right]$$

where $U(\mathbf{b}) = (\mathbf{X} \ \mathbf{V} \ \boldsymbol{\nu}'\mathbf{b})'$ and $\boldsymbol{\eta}(\boldsymbol{\gamma}) = (\boldsymbol{\gamma}\boldsymbol{\beta} \ \boldsymbol{\omega} \ \boldsymbol{\tau})$ with $\boldsymbol{\eta} \equiv \boldsymbol{\eta}(\mathbf{1})$. For brevity, we drop the $(t-1)$ superscript on all expectations and, without loss of generality, assume $p_k > 0$ (predictor $\mathbf{X}_k$ can be taken out of the optimization when $p_k = 0$). Note that

$$E_{\boldsymbol{\gamma}\mathbf{b}} \{U(\mathbf{b})'\boldsymbol{\eta}(\boldsymbol{\gamma})\} = U(\tilde{\mathbf{b}})'\boldsymbol{\eta}(\mathbf{p}) \equiv \begin{bmatrix} (\mathbf{X} \odot \mathbf{P})' \\ \mathbf{V}' \\ \boldsymbol{\nu}'\tilde{\mathbf{b}} \end{bmatrix}' \boldsymbol{\eta} = \tilde{\mathbf{U}}'\boldsymbol{\eta},$$

where $\odot$ denotes the Hadamard product and $\mathbf{P}$ is an $n \times p$ matrix in which each row is the vector $\mathbf{p}$, and

$$E_{\boldsymbol{\gamma}\mathbf{b}} \left[\{U(\mathbf{b})'\boldsymbol{\eta}(\boldsymbol{\gamma})\}' \{U(\mathbf{b})'\boldsymbol{\eta}(\boldsymbol{\gamma})\} \right] = (\tilde{\mathbf{U}}'\boldsymbol{\eta})'\tilde{\mathbf{U}}'\boldsymbol{\eta} + \text{Var} \{U(\mathbf{b})'\boldsymbol{\eta}(\boldsymbol{\gamma})\}.$$

The variance term reduces to

$$\mathbf{1}' \begin{bmatrix} (\mathbf{X}^2) \odot \mathbf{P} \odot (\mathbf{J} - \mathbf{P}) \\ \mathbf{0}_r \\ \boldsymbol{\nu}'(\sigma^{-2}\boldsymbol{\Psi})^{-1}\boldsymbol{\nu} \end{bmatrix} \boldsymbol{\eta}^2 = \boldsymbol{\lambda}'\boldsymbol{\eta}^2$$

where $\mathbf{X}^2 = \mathbf{X} \odot \mathbf{X}$, $\mathbf{J}$ is a $(p \times p)$ matrix of ones, $\mathbf{0}_r$ is an r -dimensional vector of zeros, and $\boldsymbol{\lambda} = \{\lambda_1, \dots, \lambda_p, \mathbf{0}_r', \boldsymbol{\nu}'(\sigma^{-2}\boldsymbol{\Psi})^{-1}\boldsymbol{\nu}\}'$ where $\lambda_k = p_k(1 - p_k) \sum_i X_{ik}^2$ for $k \leq p$. This yields

$$Q_{\text{EM}}(\boldsymbol{\eta} \mid \Theta^{(t-1)}) = - \left\{ (\mathbf{Y} - \tilde{\mathbf{U}}'\boldsymbol{\eta})' (\mathbf{Y} - \tilde{\mathbf{U}}'\boldsymbol{\eta}) \right\} - [\boldsymbol{\lambda}'\boldsymbol{\eta}^2],$$

resulting in maximizer

$$\hat{\boldsymbol{\eta}} = \left[\tilde{\mathbf{U}}'\tilde{\mathbf{U}} + \boldsymbol{\lambda}\mathbf{I}_{p+2r} \right]^{-1} \tilde{\mathbf{U}}'\mathbf{Y}.$$

As a result, $\hat{\boldsymbol{\eta}}$ is only estimable when $\sum_k I(p_k = 1) < M+1-r$; adding additional non-sparse predictors would lower this limit. Therefore, $Q_{\text{EM}}(\boldsymbol{\eta} \mid \Theta^{(t-1)})$ cannot always be maximized, as the M-step is undefined when $\sum_k I(p_k = 1) \geq M+1-r$. $\square$

Remark 1: A standard M-step for calculating $\hat{\boldsymbol{\eta}}$ requires the inverse of $\tilde{\mathbf{U}}'\tilde{\mathbf{U}} + \boldsymbol{\lambda}\mathbf{I}_{p+2r}$ a $(p+2r) \times (p+2r)$ matrix, which is lower-bounded by $\Omega(p^3)$ computational complexity. Note that this matrix is a function of the moments of $\boldsymbol{\gamma}$ and $\mathbf{b}$, which change at every iteration. Consequently, the inversion is required at every iteration, which compounds the computational complexity of the standard M-step. As demonstrated in the simulation studies in the main text, methods that require this type of inversion

(e.g., PGEE) scale poorly with p and require computation times that are not feasible for our data analysis.

Remark 2: The benefit of the estimate of, say, β_k from the standard M-step is that it would adjust for the remaining predictors in the model. This points to the motivation for including α , the expanded parameters for β , in the proposed LMM-PROBE method. These expanded parameters adjust for the remaining signal, denoted by $\mathbf{W}_k$ in the main text, while estimating β_k . Without α_k , β_k would be estimated using $(\mathbf{Y} - \mathbf{W}_k)$ which is *corrected for* $\mathbf{W}_k$ instead of *adjusted for* $\mathbf{W}_k$. The differences in the estimate of β_k with correction versus adjustment can be large when $\mathbf{X}_k$ and $\mathbf{W}_k$ are related.

Proof of Proposition 2. In the PX-ECM algorithm for LMM-PROBE, the ℓ th PX-CM partition maximizes the expected posterior of ξ_ℓ (where η is a subset of ξ , $\eta \in \xi$), which is a function of $\mathbf{X}_\ell$, $\mathbf{V}$, and the latent terms $\mathbf{U}_\ell$,

$$\hat{\xi}_\ell^{(M1)} = \operatorname{argmax}_{\xi_\ell} E_{\mathbf{U}_\ell} \left\{ l(\xi_\ell | \mathbf{Y}, \mathbf{U}_\ell, \mathbf{\Gamma}_\ell) | \Theta_{/\ell}^{(t-1)} \right\} \text{ for } \ell = 0, 1, \dots, p.$$

With the parameter expansion, this reduces the first M-step (M1) to a straightforward calculation for each partition ℓ , akin to a $2(r+1)$ -dimensional linear regression with maximizer

$$\hat{\xi}_\ell^{(M1)} = \left\{ (\mathbf{Z}_\ell' \mathbf{Z}_\ell)^{(t-1)} \right\}^{-1} \mathbf{Z}_\ell^{(t-1)'} \mathbf{Y},$$

where $\mathbf{Z}_\ell$ is defined in Section 2.2 of the main text. This maximizer is fully identifiable under the PX-ECM framework as long as there is not perfect collinearity between $\mathbf{X}_k$ and $\mathbf{V}$ for all k . As a result, $Q_{\text{CM}}^{M1}(\eta | \Theta^{(t-1)})$ can be maximized. The second M-step M2 updates $\hat{\xi}_0^{(M2)}$, a subset of η , conditional on the first M-step M1 and E-steps. This yields the same proof that the maximizer is identifiable and exists. Therefore, the maximizers of $Q_{\text{CM}}^{M1}(\eta | \Theta^{(t-1)})$ and $Q_{\text{CM}}^{M2}(\eta | \Theta^{(t-1)})$ both exist and are unique. $\square$

B Convergence Assessment

This Section provides additional results and discussion with regards to convergence. Our implementation of LMM-PROBE is an ‘all-at-once’ optimization, where the parameter value updates at iteration t are performed accounting for updates from the previous iteration $t-1$ (analogous to the Jacobi least-squares optimization approach, Mascarenhas, 1995). In previous work, we have found that the ‘all-at-once’ optimization was less sensitive to the updating order and more efficiently maximized the likelihood than the ‘one-at-a-time’ approach (McLain, Zgodic, & Bondell, 2022). In the ‘one-at-a-time’ approach, parameter values are sequentially updated, with each update accounting for all previous updates within current iteration t (analogously to Gauss-Seidel least-squares optimization, Ma, Needell, & Ramdas, 2015). The ‘one-at-a-time’ approach fulfills the monotonicity and space-filling properties of EM and ECM algorithms (Meng & Rubin, 1993). Our ‘all-at-once’ implementation of LMM-PROBE is less prone to getting stuck in local maxima (in practice) compared to a ‘one-at-a-time’ approach but does not provide convergence guarantees.

Various studies have demonstrated that the PX-EM algorithm and some parameterizations of ECM have better global rates of convergence than a standard EM algorithm (Liu et al., 1998; Meng, 1994; Sexton & Swensen, 2000). However, for LMM-PROBE, the parameters in the complete data model of the standard EM algorithm are nonidentifiable when $p \gg n$ (Proposition 1), and thus the M-step is often undefined. Therefore, we cannot compare the theoretical rate of convergence of a PX multi-cycle ECM as in LMM-PROBE to that of a standard EM. Instead, we assess empirical evidence of convergence. For eight simulation settings chosen to illustrate all simulation factors, Figure B.1 shows that the log-likelihood sharply increased in early iterations and flattened quickly in a small number of iterations, indicating a rapid rate of convergence. The empirical convergence assessment was similar for other simulation settings (omitted from figures).




Fig. B.1 Log-Likelihood for LMM-PROBE over iterations, across various simulation settings, including r , σ^2 , $\mathbf{G}$, and β values, when $p = (15^2, 75^2)$ and $\pi = 0.1$. Within each simulation scenario, lines represent simulation repetitions (10 repetitions).

To empirically assess the maximization of the conditional likelihood, we examined $-\frac{n}{2} \log(\hat{\sigma}^2) - \frac{1}{2\hat{\sigma}^2} \sum_i \left[\mathbf{Y}_i - \left\{ \mathbf{X}_i(\tilde{\alpha}_0 \tilde{\mathbf{p}} \tilde{\beta}) + \mathbf{V}_i \tilde{\omega}_0 + \mathbf{V}_i \tilde{\mathbf{b}}_i \right\} \right]^2$ for all methods. Figure B.2 shows that LMM-PROBE maximized the conditional likelihood better than LASSO and PROBE in the settings displayed. Additionally, LMM-PROBE maximized the conditional likelihood better than LASSO+ in all settings except two, and better than LMM-LASSO when $\beta = 0.5$ and variances were higher. When $\pi = 0.05$, LMM-PROBE maximized the conditional likelihood most effectively in all settings (figure omitted). Generally, methods for LMMs maximized the conditional likelihood better than LASSO and PROBE, which are not designed for LMMs.




Fig. B.2 Log-Likelihood for LMM-PROBE and four comparison methods (excluding PGEE) across various simulation settings, including p , σ^2 , $\mathbf{G}$, and β values, when $r = 2$ and $\pi = 0.1$. Vertical lines display the interquartile range of the Log-Likelihood values across simulation iterations. Comparison methods LMM-LASSO and LASSO+ are methods for linear mixed models or repeated measures, while LASSO and PROBE are methods for linear models.

C Additional Simulation Results

This Section provides additional simulation results. Figure C.3 corresponds to Figure 1 in Section 4 of the main text, with the additional method of penalized generalized estimating equations (PGEE) for high-dimensional longitudinal data. The Mean Squared Predictive Error (MSPE) was the highest for this additional method. Figure C.4 also mirrors Figure 1 from Section 4 in the main text, but shows MSPEs for the simulation settings with one random effect ($r = 1$). Figure C.5 shows Median Absolute Deviations (MADs) for simulation settings when $r = 1$ and $\pi = 0.1$. In all settings except one, MADs were lowest for LMM-PROBE. Figures C.6 and C.7 show sensitivity, specificity, and the Matthews Correlation Coefficient (MCC, Matthews, 1975) when $p = 25^2$ and $p = 75^2$, respectively. Figure C.6 shows that when the effect size of signals was higher ($\beta = 0.75$), LMM-PROBE had a similar sensitivity as PROBE, and outperformed other methods. LMM-PROBE had the highest specificity for all settings. In Figure C.7, LMM-PROBE had higher sensitivity when $\beta = 0.75$ as well as high sensitivity in all settings except one.




Fig. C.3 Mean Squared Predictive Errors (MSPE) for LMM-PROBE and five comparison methods (including PGEE) across various simulation settings, including p , σ^2 , $\mathbf{G}$, and β values, when $r = 2$ and $\pi = 0.1$. The MSPEs are based on both fixed and random effects. Vertical lines display the interquartile range of the MSPEs. Comparison methods LMM-LASSO, LASSO+, and PGEE are methods for linear mixed models or repeated measures, while LASSO and PROBE are methods for linear models.

D Additional Data Analysis Results

In this Section, we examine additional results from the data analysis on pediatric systemic lupus erythematosus (SLE) and provide a second data analysis on the popular riboflavin dataset (Buhlmann, Kalisch, & Meier, 2014). For the SLE analysis, Figure D.8 shows that LMM-PROBE had the lowest estimates of $\tilde{\sigma}^2$ across Cross-Validation (CV) folds, compared to PROBE and LASSO. LASSO showed the largest variation in its $\tilde{\sigma}^2$ estimates across CV folds. The average random effect variance $\tilde{\mathbf{G}}$ estimate was approximately 0.02 giving an average ICC of 0.15. Figure D.9 shows the average number of selected predictors for each method, across CV folds. Computation time for the three methods varied slightly. The methods for traditional high-dimensional linear regression (LASSO, PROBE) required around nine and 48 seconds per fold, respectively, while LMM-PROBE required 69 seconds on average.

We examined residuals for each CV fold. Figure D.10 shows conditional residuals (Nobre & Singer, 2007) plotted against observations for each CV fold. These residuals show that the homogeneous variance assumption is fulfilled. Quantile-Quantile (Q-Q) plots in Figure D.11 show that the normality assumption is not violated. Figure D.12 shows the predicted random intercepts plotted against subjects, where no subject had




Fig. C.4 Mean Squared Predictive Errors (MSPE) for LMM-PROBE and four comparison methods across various simulation settings, including p , σ^2 , $\mathbf{G}$, and β values, when $r = 1$ and $\pi = 0.1$. The MSPEs are based on both fixed and random effects. Vertical lines display the interquartile range of the MSPEs. Comparison methods LMM-LASSO and LASSO+ are methods for linear mixed models, while LASSO and PROBE are methods for linear models.

an outlying predicted random intercept. Finally, the Q-Q plots in Figure D.13 show that the random effect normality assumption is not violated either.

We now showcase the performance and characteristics of the LMM-PROBE method on the riboflavin (vitamin B₂) dataset (Buhlmann et al., 2014), a popular high-throughput genomic dataset about riboflavin production using a genetically modified bacterium. Riboflavin is grown in recombinant *Bacillus subtilis*, with the goal of maintaining the production rate of the riboflavin over time. Riboflavin grows across multiple generations of the bacterium for a different duration, leading to a longitudinal design with repeated observations. The dataset contains a log-transformed riboflavin production rate, which we use as the outcome, as well as 4088 predictor variables, each consisting of gene expression levels on the log scale. There are 28 clusters in the dataset, measured over time, with the number of observations per cluster ranging between 2 and 6. In total, there are 111 observations. The original aim of the research that generated this dataset was to identify the genes that impacted riboflavin production. To analyze this data, we used a linear mixed-effects model. We considered all 4088 gene expression predictors as sparse fixed effects and considered an intercept as well as a time predictor as both non-sparse fixed and random effects.




Fig. C.5 Median Absolute Deviations (MAD) for LMM-PROBE and four comparison methods across various simulation settings, including p , σ^2 , $\mathbf{G}$, and β values, when $r = 1$ and $\pi = 0.1$. The MADs are based on both fixed and random effects. Vertical lines display the interquartile range of the MADs. Comparison methods LMM-LASSO and LASSO+ are methods for linear mixed models, while LASSO and PROBE are methods for linear models.

To evaluate the performance of LMM-PROBE, we used MSPEs resulting from five-fold CV. In the CV, we balanced clusters across the folds, meaning that a cluster's observations were all in the same fold. At a given iteration of the CV, 80% of the clusters were in training folds, and 20% were in the validation-test fold. The validation-test fold was further split into validation (with earlier *time* values for each cluster in the fold, i.e., 1–2) and testing (with subsequent *time* values, i.e., 3–6) subfolds. The testing subfold had two observations from each cluster with four or more measurements and one observation otherwise. We used the validation subfold to obtain the predicted random effects and the testing subfold to calculate MSPEs. To generate random effect predictions for the validation subfold, we used the MAP estimate of LMM-PROBE. To tune the penalty parameter in LASSO, we performed an additional five-fold CV using the training set.

We calculated two types of MSPEs. The first type of MSPE is the difference between the true outcome value and prediction obtained based on both the fixed and random effect estimates. Specifically, using estimates $\hat{\mathbf{G}}$ and $\hat{\sigma}^2$ from the training data, we obtained predicted random effects $\hat{\mathbf{b}}_i$ using the validation data and finally obtained complete predictions for future time points using the testing data. The second type of MSPE is the difference between the true outcome value and predictions obtained only




Fig. C.6 Sensitivity, Specificity, and the Matthews Correlation Coefficient (MCC) for LMM-PROBE and four comparison methods across various simulation settings, including π , r , and β values, when $\sigma^2 = \sigma_1^2$, $\mathbf{G} = \mathbf{G}_1$, and $p = 25^2$. Vertical lines display the interquartile range of the sensitivity, specificity, and MCC. Comparison methods LMM-LASSO and LASSO+ are methods for linear mixed models, while LASSO and PROBE are methods for linear models.

based on the fixed effect estimates. Figure D.14 shows the two types of MSPEs as well as MADs. For MSPEs based on both fixed and random effect estimates, LMM-PROBE had the lowest MSPEs compared to PROBE and LASSO. When examining MSPEs based on fixed effect estimates only, LMM-PROBE performed best as well. Generally, the range of MSPE values across the CV folds was wider for LASSO and PROBE, and narrower for LMM-PROBE. Trends in MADs results were overwhelmingly similar to those of MSPEs.

Figure D.15 shows estimates for within-sample variation, $\tilde{\sigma}^2$. For PROBE and LASSO, $\tilde{\sigma}^2$ is the model residual error estimates since these methods do not delineate between within-unit and between-unit variation. LMM-PROBE had a mid-range estimate for $\tilde{\sigma}^2$ with the lowest variability across CV fold. LASSO had the lowest average $\tilde{\sigma}^2$ estimate but varied markedly across CV folds. Figure D.16 shows the average number of predictors selected by LMM-PROBE (1), PROBE (1), and LASSO (33), with LMM-PROBE selecting the fewest predictors across CV folds. Prior to choosing the final model described above with the LMM-PROBE, LASSO, and PROBE approaches, we examined additional methods (LASSO+, LASSO+EN, and BRMS) but did not continue with them due to computation time. Figure D.17 shows the average computation time in seconds for one CV fold of an intermediate model in our data analysis. The




Fig. C.7 Sensitivity, Specificity, and the Matthews Correlation Coefficient (MCC) for LMM-PROBE and two comparison methods across various simulation settings, including π , r , and β values, when $\sigma^2 = \sigma_1^2$, $\mathbf{G} = \mathbf{G}_1$, and $p = 75^2$. Vertical lines display the interquartile range of the sensitivity, specificity, and MCC. Comparison methods LASSO and PROBE are methods for linear models.

methods for traditional high-dimensional linear regression (LASSO, PROBE) required around one second per fold. The remaining methods required between 200 and 6,000 seconds, on average, per fold. Overall, the total computation time for the intermediate model was 1,026 minutes for BRMS, 202 minutes for LASSO+EN, 51 minutes for LASSO+, 2 minutes for LMM-PROBE, 31 seconds for PROBE, and 18 seconds for LASSO.




Fig. D.8 Average residual ($\hat{\sigma}^2$) and random effect (G) variance estimates for LMM-PROBE and two comparison methods. Vertical lines represent $\pm$ the standard error of $\hat{\sigma}^2$ or G , divided by $\sqrt{5}$, based on the number of Cross-Validation folds.




Fig. D.9 Average number of predictors selected for LMM-PROBE and two comparison methods. Vertical lines display the range across the Cross-Validation folds.




Fig. D.10 Conditional residuals plotted against observations for each Cross-Validation fold.




Fig. D.11 Quantile-Quantile plots of conditional residuals for each Cross-Validation fold.




Fig. D.12 Predicted random intercepts plotted against subjects for each Cross-Validation fold.




Fig. D.13 Quantile-Quantile plots of predicted random intercepts for each Cross-Validation fold.




Fig. D.14 Mean Squared Predictive Errors (MSPE) and Median Absolute Deviations (MAD) for LMM-PROBE and two comparison methods, based on both random and fixed effects and fixed effects only. Vertical lines represent $\pm$ the standard error of MSPE or MAD, divided by $\sqrt{5}$, based on the number of Cross-Validation folds.




Fig. D.15 Average residual ($\tilde{\sigma}^2$) variance estimates for LMM-PROBE and two comparison methods. Vertical lines represent $\pm$ the standard error of $\tilde{\sigma}^2$ divided by $\sqrt{5}$, based on the number of Cross-Validation folds.




Fig. D.16 Average number of predictors selected for LMM-PROBE and two comparison methods. Vertical lines display the range across the Cross-Validation folds.




Fig. D.17 Average computation time per Cross-Validation fold (in seconds) for LMM-PROBE and five comparison methods, on the $\log_{10}$ scale.

References

- Buhlmann, P., Kalisch, M., Meier, L. (2014). High-dimensional statistics with a view towards applications in biology. *Annual Review of Statistics and its Applications*, 1, 255–278,
- Liu, C., Rubin, D.B., Wu, Y.N. (1998, 12). Parameter expansion to accelerate EM: The PX-EM algorithm. *Biometrika*, 85(4), 755–770, <https://doi.org/10.1093/biomet/85.4.755> Retrieved from <https://doi.org/10.1093/biomet/85.4.755>
- Ma, A., Needell, D., Ramdas, A. (2015). Convergence properties of the randomized extended gauss–seidel and kaczmarz methods. *SIAM Journal on Matrix Analysis and Applications*, 36(4), 1590–1604,
- Mascarenhas, W.F. (1995). On the convergence of the jacobi method for arbitrary orderings. *SIAM journal on matrix analysis and applications*, 16(4), 1197–1209,
- Matthews, B. (1975). Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA) - Protein Structure*, 405(2), 442–451,
- McLain, A.C., Zgodic, A., Bondell, H. (2022). *Sparse high-dimensional linear regression with a partitioned empirical bayes ecm algorithm*. <https://arxiv.org/abs/2209.08139>. arXiv. Retrieved from <https://arxiv.org/abs/2209.08139>
- Meng, X.-L. (1994). On the rate of convergence of the ECM algorithm. *Ann. Statist.*, 22(1), 326–339, <https://doi.org/10.1214/aos/1176325371> Retrieved from <https://doi.org/10.1214/aos/1176325371>
- Meng, X.L., & Rubin, D.B. (1993). Maximum likelihood estimation via the ECM algorithm: A general framework. *Biometrika*, 80(2), 267–278, <https://doi.org/10.1093/biomet/80.2.267>
- Nobre, J., & Singer, J. (2007). Residual analysis for linear mixed models. *Biometrical Journal*, 49(6), 863–875,
- Sexton, J., & Swensen, A.R. (2000). Ecm algorithms that converge at the rate of em. *Biometrika*, 87(3), 651–662,