

MAGISTERARBEIT | MASTER'S THESIS

Titel | Title

Linear Mixed Models with Ill-posed Features

verfasst von | submitted by
Oleksandr Volotko

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of
Magister der Sozial- und Wirtschaftswissenschaften (Mag.rer.soc.oec.)

Wien | Vienna, 2026

Studienkennzahl lt. Studienblatt | Degree
programme code as it appears on the
student record sheet:

UA 066 951

Studienrichtung lt. Studienblatt | Degree
programme as it appears on the student
record sheet:

Magisterstudium Statistik

Betreut von | Supervisor:

Univ.-Prof. Dr. Tatyana Krivobokova

Contents

List of Tables	3
List of Figures	4
List of Abbreviations	5
1 The expectation–maximization (EM) algorithm and its application to linear mixed Models	6
1.1 Formulation of the EM Algorithm	6
1.2 Monotonicity of $\ell(\theta^{(t)})$ in t	8
1.3 Convergence properties	10
1.4 Linear mixed models as an EM algorithm (adapted from Laird, 1982)	20
1.5 Theoretical convergence in LMM implementation	25
1.6 Summary	29
2 Regularisation of the EM algorithm with Partial least squares	30
2.1 Ill-posed Linear Mixed Models	30
2.2 Existing results on EM regularization	32
2.3 Sample EM–PLS algorithm	34
2.4 Population EM–PLS algorithm	36
2.5 Sample EM–PLS algorithm	38
2.6 Simulation study	41
2.7 Data Description and Exploratory Analysis	46
2.8 Summary	53
3 Discussion	54
Bibliography	55

List of Tables

- | | | |
|---|---|----|
| 1 | Marginal prediction performance of EM and EM-PLS. Correlation and RMSE are reported for the training data (fit) and the test data (prediction), based on the fixed-effects component $\hat{y} = X\hat{\beta}$ | 48 |
|---|---|----|

List of Figures

1	Illustration of Example 1.2. The log-likelihood sequence $l(\theta^{(t)})$ converges monotonically (right), while the parameter sequence $\theta^{(t)}$ fails to converge and instead cycles around the limit circle $r = 1$ (left). This example, adapted from Boyles (1983), demonstrates that convergence of the likelihood does not necessarily imply convergence of the EM iterates.	12
2	Log-likelihood surface $\ell(\theta_1, \theta_2) = -\frac{1}{2} \theta - y ^2$ with the trajectory of the GEM iterates $\theta^{(t)}$. The likelihood converges while the parameter sequence rotates indefinitely around the limit circle.	13
3	Average training-set and test-set prediction correlations of the marginal fixed-effects predictor $\hat{y} = X\hat{\beta}^{(t)}$ and ℓ_2 error of the parameters across EM iterations. The plot of ℓ_2 error of $\hat{\beta}^{(t)}$ is log-scaled in y-axis.	43
4	Train and test marginal correlation of EM-PLS in the simulation study as a function of the number of PLS components K , with latent dimension fixed at $q_{\text{lat}} = 40$	44
5	Final estimation error and predictive performance of EM and EM-PLS as a function of the number of groups m . Shaded areas indicate variability across simulation runs. . . .	45
6	Ion (purple) inside the gramicidin A channel.	46
7	First 100 eigenvalues of the empirical Gram matrix $X^\top X$ on a logarithmic scale. The rapidly decaying spectrum, spanning several orders of magnitude, indicates strong collinearity and a pronounced separation between a low-dimensional signal subspace and a large number of weakly identified directions.	47
8	Bar plot of group sample size n_i of group $i \in 1, \dots, 83$. The minimum number of observations per group is $\min_i n_i = 19$, while the maximum is $\max_i n_i = 320$	48
9	Evolution of marginal correlation over EM iterations for the training data (left) and the test data (right). The EM-PLS algorithm converges in few iterations and attains a worse fit on the training data, but exhibits superior generalization.	49
10	Evolution of marginal RMSE over EM iterations for the training data (left) and the test data (right). While the unconstrained EM algorithm gradually minimizes the training error, this is accompanied by the decrease in the predictive accuracy. The EM-PLS algorithm achieves higher stability and consistently lower prediction error, reflecting the regularizing effect of the Krylov subspace restriction.	49
11	Observed versus fitted values for EM-PLS. The red curve (training data) shows the conditional fit $X_i\beta^{(t)} + Z_i u_i^{(t)}$ using BLUP random effects, while the blue curve (test data) shows the marginal prediction $X_i\beta^{(t)}$	51
12	Contribution of the estimated random effects $Z_i u_i^{(t)}$ to the training BLUP predictions. The model includes a random intercept and random slope. Gray vertical lines depict the first observation in each group.	51
13	Train and test marginal correlation of EM-PLS on the real dataset as a function of the number of PLS components K	52

List of Abbreviations

Abbreviation	Definition
EM	Expectation–Maximization
LMM	Linear Mixed Model
PLS	Partial Least Squares
MM	Mixed Model
REML	Restricted Maximum Likelihood
ML	Maximum Likelihood
BLUP	Best Linear Unbiased Predictor
AR(1)	First-order Autoregressive Structure
AIC	Akaike Information Criterion
BIC	Bayesian Information Criterion

1 The expectation–maximization (EM) algorithm and its application to linear mixed Models

1.1 Formulation of the EM Algorithm

We consider a statistical model in which the data are partly unobserved. Let $(\mathbf{Y}, \mathbf{U})$ be random vectors with joint density $f(\mathbf{Y}, \mathbf{U} \mid \theta^*)$, where $\mathbf{Y} \in \mathcal{Y} \subseteq \mathbb{R}^n$ denotes the observable component and $\mathbf{U} \in \mathcal{U} \subseteq \mathbb{R}^q$ represents latent (unobserved) variables—for instance, random effects in a mixed model. Their realizations in the sample are written (y, u) .

Our goal is to estimate the parameter $\theta^* \in \Theta \subseteq \mathbb{R}^d$ by maximum likelihood, based only on the observed data y . The observed–data likelihood integrates out the latent variables:

$$\ell(\theta) = \log f(y \mid \theta) = \log \int f(y, u \mid \theta) du.$$

Direct maximization of $\ell(\theta)$ is often challenging, as the integral over u is typically intractable. By contrast, the *complete–data* log–likelihood

$$\ell_c(\theta) = \log f(y, u \mid \theta)$$

is typically easier to handle.

Example 1.1 (McLachlan and Krishnan (2008), Ex. 1.2). A *finite mixture model* represents the distribution of an observation Y as a convex combination of K component densities,

$$f(y \mid \theta) = \sum_{k=1}^K \pi_k f_k(y \mid \theta), \quad \pi_k > 0, \quad \sum_{k=1}^K \pi_k = 1,$$

where θ denotes the full parameter vector of the model.

Introduce latent allocation variables $U_i \in \{e_1, \dots, e_K\}$, where $U_{ik} = 1$ indicates that observation Y_i originates from component k . The *complete–data* log–likelihood is

$$\ell_c(\theta \mid y, u) = \sum_{i=1}^n \sum_{k=1}^K u_{ik} \{ \log \pi_k + \log f_k(y_i \mid \theta) \}.$$

For Gaussian components, the complete–data log–likelihood is separable across components and quadratic in the corresponding parameters, yielding closed–form maximizers in the M–step. In contrast, the *observed–data* log–likelihood

$$\ell(\theta) = \sum_{i=1}^n \log \left(\sum_{k=1}^K \pi_k f_k(y_i \mid \theta) \right)$$

is a coupled log–sum–exp expression and admits no closed–form maximizer, which motivates the use of the EM algorithm.

General case. As seen in the example before, directly maximizing $\ell(\theta)$ typically requires repeated integration or numerical approximation of $f(y \mid \theta)$ for each θ , which is computationally prohibitive and analytically opaque, as seen in the example before. This motivates the EM algorithm, which maximizes the conditional expectation of $\ell_c(\theta)$ given the observed data.

The EM algorithm resolves the intractability of the observed-data likelihood by working with its conditional expectation. At iteration t , let the current parameter estimate be $\theta^{(t)}$. Consider the conditional distribution of the latent variables given the observed data, $f(u \mid y, \theta^{(t)})$. The algorithm alternates between:

$$\text{E-step: } Q(\theta \mid \theta^{(t)}) = \mathbb{E}_{\mathbf{U} \mid \mathbf{Y}=y, \theta^{(t)}} [\log f(y, \mathbf{U} \mid \theta)], \quad \text{M-step: } \theta^{(t+1)} = \arg \max_{\theta \in \Theta} Q(\theta \mid \theta^{(t)}),$$

where the iteration index $t = 0, 1, 2, \dots, T$ denotes successive EM updates. The total number of iterations T is determined by a convergence criterion, such as stabilization of the parameter estimates or the log-likelihood value, e.g.

$$\|\theta^{(t+1)} - \theta^{(t)}\| < \varepsilon \quad \text{or} \quad |\ell(\theta^{(t+1)}) - \ell(\theta^{(t)})| < \varepsilon.$$

Thus, in the E-step, the conditional expectation with respect to the current posterior distribution of $\mathbf{U}$ replaces the unobserved component of the likelihood, while in the M-step this expected log-likelihood is maximized to update the parameter estimate.

1.2 Monotonicity of $\ell(\theta^{(t)})$ in t

Developments in this section follow closely Dempster et al. (1977, Sec. 3) and McLachlan and Krishnan (2008, Sec. 3.2-3.3).

Theorem 1.1: Monotonicity of the EM algorithm (Dempster et al., 1977, Theorem 1)

Assume the following hold for all $\theta, \theta^{(t)} \in \Theta$:

(M1) $f(y, u \mid \theta) > 0$ almost everywhere on $\mathcal{Y} \times \mathcal{U}$ and is measurable in (y, u) .

(M2) The Q -function

$$Q(\theta \mid \theta^{(t)}) = \mathbb{E}_{\mathbf{U} \mid \mathbf{Y}=y, \theta^{(t)}} [\log f(y, \mathbf{U} \mid \theta)]$$

exists and is finite.

(M3) The marginal (observed-data) density

$$f(y \mid \theta) = \int f(y, u \mid \theta) \, du$$

satisfies $0 < f(y \mid \theta) < \infty$.

Then, at every EM iteration $t = 0, 1, 2, \dots, T$,

$$\ell(\theta^{(t+1)}) - \ell(\theta^{(t)}) \geq 0, \quad \ell(\theta) = \log f(y \mid \theta).$$

Proof. The joint, conditional, and marginal densities satisfy

$$f(y, u \mid \theta) = f(u \mid y, \theta) f(y \mid \theta).$$

Taking logarithms and rearranging yields

$$\log f(y \mid \theta) = \log f(y, u \mid \theta) - \log f(u \mid y, \theta).$$

Taking the conditional expectation with respect to the current posterior distribution of $\mathbf{U}$ given $\mathbf{Y} = y$ under $\theta^{(t)}$, that is, $\mathbb{E}_{\mathbf{U} \mid \mathbf{Y}=y, \theta^{(t)}}[\cdot]$, gives

$$\boxed{\ell(\theta) = Q(\theta \mid \theta^{(t)}) - H(\theta \mid \theta^{(t)})},$$

where $\ell(\theta) = \log f(y \mid \theta)$ and

$$Q(\theta \mid \theta^{(t)}) = \mathbb{E}_{\mathbf{U} \mid \mathbf{Y}=y, \theta^{(t)}} [\log f(y, \mathbf{U} \mid \theta)], \quad H(\theta \mid \theta^{(t)}) = \mathbb{E}_{\mathbf{U} \mid \mathbf{Y}=y, \theta^{(t)}} [\log f(\mathbf{U} \mid \mathbf{Y} = y, \theta)].$$

The term $H(\theta \mid \theta^{(t)})$ can be written as

$$H(\theta \mid \theta^{(t)}) = \int f(u \mid y, \theta^{(t)}) \log f(u \mid y, \theta) \, du.$$

By adding and subtracting $\log f(u \mid y, \theta^{(t)})$ inside the integral, we obtain

$$\begin{aligned} H(\theta \mid \theta^{(t)}) &= \int f(u \mid y, \theta^{(t)}) \log f(u \mid y, \theta^{(t)}) \, du \\ &\quad - \int f(u \mid y, \theta^{(t)}) \log \left(\frac{f(u \mid y, \theta^{(t)})}{f(u \mid y, \theta)} \right) \, du. \end{aligned}$$

The second term is the Kullback–Leibler divergence between $f(u | y, \theta^{(t)})$ and $f(u | y, \theta)$:

$$D_{\text{KL}}\left(f(u | y, \theta^{(t)}) \parallel f(u | y, \theta)\right) = \int f(u | y, \theta^{(t)}) \log\left(\frac{f(u | y, \theta^{(t)})}{f(u | y, \theta)}\right) du \geq 0.$$

Hence,

$$H(\theta | \theta^{(t)}) = C(\theta^{(t)}) - D_{\text{KL}}\left(f(u | y, \theta^{(t)}) \parallel f(u | y, \theta)\right),$$

where

$$C(\theta^{(t)}) = \int f(u | y, \theta^{(t)}) \log f(u | y, \theta^{(t)}) du$$

depends only on $\theta^{(t)}$ and is therefore constant in θ . Consequently,

$$-H(\theta | \theta^{(t)}) = D_{\text{KL}}\left(f(u | y, \theta^{(t)}) \parallel f(u | y, \theta)\right) - C(\theta^{(t)}),$$

that is, $-H$ equals a nonnegative Kullback–Leibler term up to an additive constant independent of θ .

Substituting the above expression for $H(\theta | \theta^{(t)})$ into the EM decomposition

$$\ell(\theta) = Q(\theta | \theta^{(t)}) - H(\theta | \theta^{(t)}),$$

yields

$$\ell(\theta) = Q(\theta | \theta^{(t)}) + D_{\text{KL}}\left(f(u | y, \theta^{(t)}) \parallel f(u | y, \theta)\right) - C(\theta^{(t)}).$$

The constant $C(\theta^{(t)})$ depends only on the previous iterate $\theta^{(t)}$ and can therefore be ignored when comparing likelihood values across θ .

Let $\theta^{(t+1)}$ denote the maximizer of $Q(\theta | \theta^{(t)})$ in the M-step. Then

$$Q(\theta^{(t+1)} | \theta^{(t)}) \geq Q(\theta^{(t)} | \theta^{(t)}).$$

Since $D_{\text{KL}} \geq 0$, we obtain

$$\begin{aligned} \ell(\theta^{(t+1)}) &\geq Q(\theta^{(t+1)} | \theta^{(t)}) - C(\theta^{(t)}), \\ \ell(\theta^{(t)}) &= Q(\theta^{(t)} | \theta^{(t)}) + \underbrace{D_{\text{KL}}\left(f(u | y, \theta^{(t)}) \parallel f(u | y, \theta^{(t)})\right)}_{=0} - C(\theta^{(t)}). \end{aligned}$$

Combining these two inequalities gives

$$\ell(\theta^{(t+1)}) - \ell(\theta^{(t)}) \geq Q(\theta^{(t+1)} | \theta^{(t)}) - Q(\theta^{(t)} | \theta^{(t)}) \geq 0.$$

Hence, the observed–data log–likelihood is nondecreasing along the EM sequence:

$$\boxed{\ell(\theta^{(t+1)}) \geq \ell(\theta^{(t)}) \quad \text{for all } t.} \quad (*)$$

The property (*) is referred to as the *monotonicity of EM*. It follows directly from the nonnegativity of the Kullback–Leibler divergence and guarantees that each iteration of the EM algorithm moves the parameter vector toward a region of higher likelihood. $\square$

In the standard EM formulation, the M–step requires finding the exact maximizer of the expected complete–data log–likelihood $Q(\theta | \theta^{(t)})$. As it follows from above, it is sufficient to choose $\theta^{(t+1)}$ such that

$$Q(\theta^{(t+1)} | \theta^{(t)}) \geq Q(\theta^{(t)} | \theta^{(t)}),$$

without necessarily achieving the global maximum. Such an iteration is called a **generalized EM (GEM)** step. The monotonicity property of the EM algorithm continues to hold under this weaker condition, since each iteration still guarantees $\ell(\theta^{(t+1)}) \geq \ell(\theta^{(t)})$.

1.3 Convergence properties

Before presenting the detailed proofs, it is useful to summarize the main convergence results for the EM and GEM algorithms that will be discussed in this section (Theorems 1.2–1.4). These results build upon one another, progressing from the stationarity of generalized EM (GEM) likelihood function to the convergence of corresponding iterates and their local rate.

Theorem 1.2 (*Convergence of the likelihood sequence for GEM*) The theorem lists the assumptions under which the scalar sequence $\{\ell(\theta^{(t)})\}_{t \in \mathbb{N}}$ converges to a finite limit $\tilde{\ell}$.

Theorem 1.3 (*Convergence with differentiable Q*) This theorem builds upon Theorem 1.2 and lists additional assumptions under which the sequence of iterates $\{\theta^{(t)}\}_{t \in \mathbb{N}}$ converges to a $\tilde{\theta}$, which is a stationary point of $\ell(\theta)$.

Theorem 1.4 (*Local convergence rate*) Near a convergence point $\tilde{\theta}$, the theorem determines the linear approximation of the convergence rate of iterates sequence $\{\theta^{(t)}\}_{t \in \mathbb{N}}$.

Convergence results. McLachlan and Krishnan (2008) have gathered several results on convergence properties of the EM (GEM) algorithm, starting with the weakest one.

Theorem 1.2: Convergence of the likelihood sequence for GEM (McLachlan and Krishnan, 2008, Section 3.4.2.)

Let $\{\theta^{(t)}\}_{t \in \mathbb{N}}$ be the sequence generated by the GEM algorithm, and let $\ell(\theta) = \log f(y | \theta)$ denote the observed-data log-likelihood. Suppose the assumptions of Theorem 1.1 hold (so that $\ell(\theta^{(t+1)}) \geq \ell(\theta^{(t)})$ for all t), and that:

(A1) The parameter space Θ is a subset of the d -dimensional Euclidean space $\mathbb{R}^d$.

(A2) For any fixed point $\theta_0 \in \Theta$ with $\ell(\theta_0) > -\infty$, the *upper level set*

$$\Theta_{\theta_0} = \{\theta \in \Theta : \ell(\theta) \geq \ell(\theta_0)\}$$

is compact.

(A3) The likelihood function $\ell(\theta)$ is continuous on Θ and continuously differentiable in the interior of Θ .

Then the scalar sequence $\{\ell(\theta^{(t)})\}_{t \in \mathbb{N}}$ converges to a finite limit $\tilde{\ell}$.

Proof. By Theorem 1.1, $\ell(\theta^{(t)})$ is nondecreasing. Pick any $\theta_0 \in \Theta$ with $\ell(\theta_0) > -\infty$. Then $\ell(\theta^{(t)}) \geq \ell(\theta_0)$ for all t , hence $\theta^{(t)} \in \Theta_{\theta_0} := \{\theta : \ell(\theta) \geq \ell(\theta_0)\}$. By (A2), Θ_{θ_0} is compact; by (A3), ℓ is continuous on Θ , so $\sup_{\theta \in \Theta_{\theta_0}} \ell(\theta) < \infty$. Therefore the nondecreasing sequence $\{\ell(\theta^{(t)})\}_{t \in \mathbb{N}}$ is bounded above and thus converges to some finite limit $\tilde{\ell}$. □

The previous theorem established that the sequence of log-likelihood values $\{\ell(\theta^{(t)})\}_{t \in \mathbb{N}}$ generated by the GEM algorithm is monotone and convergent to a finite limit $\tilde{\ell}$. However, this result does not reveal anything about the behavior of the parameter iterates $\{\theta^{(t)}\}_{t \in \mathbb{N}}$ themselves.

Example 1.2 (Boyles, 1983). Let $\Theta = \mathbb{R}^2$ and take the observed-data log-likelihood

$$\ell(\theta) = \log f(y \mid \theta) = \text{const} - \frac{1}{2} \|\theta - y\|^2, \quad \theta = (\theta_1, \theta_2)^\top, \quad y \in \mathbb{R}^2.$$

Write θ in polar coordinates centered at y :

$$\theta = y + r(\cos \alpha, \sin \alpha)^\top, \quad r \geq 0, \quad \alpha \in \mathbb{R}.$$

Define a GEM update mapping $M : \Theta \rightarrow \Theta$ (given in polar form) by

$$(r, \alpha) \mapsto (r^+, \alpha^+) = \begin{cases} (2 - r^{-1}, \alpha + 1 - r^{-1}), & r > 1, \\ (r, \alpha + 1 - r^{-1}), & r \leq 1, \end{cases}$$

and set $\theta^{(t+1)} = M(\theta^{(t)})$ with initial $r_0 = 2, \alpha_0 = 0$.

For the update rule defined above, we first verify that the likelihood sequence $\{\ell(\theta^{(t)})\}$ converges, while the parameters $\{\theta^{(t)}\}$ do not.

Since $\ell(\theta) = \text{const} - \frac{1}{2}r^2$, we obtain

$$\ell(\theta^{(t+1)}) - \ell(\theta^{(t)}) = \frac{1}{2}(r_t^2 - (r_t^+)^2) = \frac{1}{2}(r_t^2 - (2 - r_t^{-1})^2) \geq 0 \quad (r_t > 1),$$

so the likelihood sequence is nondecreasing and bounded above, hence convergent.

Next, we derive explicit formulas for the updated parameters. We prove by induction that

$$r_t = 1 + \frac{1}{t+1}, \quad r_0 = 2.$$

For $t = 0$, the formula gives $r_0 = 2 = 1 + \frac{1}{1}$. Assume $r_t = 1 + \frac{1}{t+1}$. Then

$$r_{t+1} = 2 - \frac{1}{r_t} = 2 - \frac{1}{1 + \frac{1}{t+1}} = 2 - \frac{t+1}{t+2} = 1 + \frac{1}{t+2},$$

confirming the recursive formula. In particular, $r_t > 1$ for all t and $r_t \downarrow 1$.

For the angular coordinate, the update rule gives

$$\alpha_{t+1} = \alpha_t + 1 - \frac{1}{r_t}.$$

Substituting the expression for r_t yields

$$1 - \frac{1}{r_t} = 1 - \frac{1}{1 + \frac{1}{t+1}} = \frac{1}{t+2},$$

hence $\alpha_{t+1} - \alpha_t = \frac{1}{t+2}$ with $\alpha_0 = 0$. Therefore,

$$\alpha_t = \sum_{j=0}^{t-1} \frac{1}{j+2} = \sum_{k=1}^t \frac{1}{k+1}.$$

Thus (r_t, α_t) satisfies the update rule with $r_t \downarrow 1$ while $\alpha_t \rightarrow \infty$ (divergent harmonic series).

Consequently,

$$\theta^{(t)} = y + r_t(\cos \alpha_t, \sin \alpha_t)^\top$$

rotates indefinitely around the circle of radius 1 centered at y and does not converge, even though

$$\ell(\theta^{(t)}) \xrightarrow{t \rightarrow \infty} \text{const} - \frac{1}{2} \cdot 1^2$$

converges and the assumptions of Theorem 1.2 are satisfied (see below).




Figure 1: Illustration of Example 1.2. The log-likelihood sequence $l(\theta^{(t)})$ converges monotonically (right), while the parameter sequence $\theta^{(t)}$ fails to converge and instead cycles around the limit circle $r = 1$ (left). This example, adapted from Boyles (1983), demonstrates that convergence of the likelihood does not necessarily imply convergence of the EM iterates.

Verification of (A1)–(A3).

(A1) Parameter space. $\Theta = \mathbb{R}^2$ is a subset of the d -dimensional Euclidean space with $d = 2$.

(A2) Compact upper level sets. For any fixed $\theta_0 \in \Theta$,

$$\ell(\theta) \geq \ell(\theta_0) \iff -\frac{1}{2}\|\theta - y\|^2 \geq -\frac{1}{2}\|\theta_0 - y\|^2 \iff \|\theta - y\| \leq \|\theta_0 - y\|.$$

Hence

$$\Theta_{\theta_0} = \{\theta \in \Theta : \ell(\theta) \geq \ell(\theta_0)\} = \overline{B}(y, \|\theta_0 - y\|),$$

the closed ball centered at y with radius $\|\theta_0 - y\|$. In $\mathbb{R}^2$, closed and bounded sets are compact (Heine–Borel), so Θ_{θ_0} is compact.

(A3) Continuity and differentiability of ℓ . $\ell(\theta) = C - \frac{1}{2}(\theta - y)^\top(\theta - y)$ is C^∞ on $\mathbb{R}^2$ with

$$\nabla \ell(\theta) = y - \theta, \quad \nabla^2 \ell(\theta) = -I_2,$$

so ℓ is continuous on Θ and continuously differentiable on its interior.


Log-likelihood surface and path of $\theta^{(t)}$ 

Figure 2: Log-likelihood surface $\ell(\theta_1, \theta_2) = -\frac{1}{2}|\theta - y|^2$ with the trajectory of the GEM iterates $\theta^{(t)}$. The likelihood converges while the parameter sequence rotates indefinitely around the limit circle.

The following theorem strengthens the previous result by establishing convergence of the EM iterates under additional regularity conditions.

Theorem 1.3: Convergence of GEM with differentiable Q -function (McLachlan and Krishnan (2008), Th. 3.5)

Assume that the conditions of Theorem 1.2 are satisfied and that the complete-data log-likelihood is continuously differentiable in θ . Suppose further that:

(D1) For each iteration t , the M-step satisfies the first-order condition

$$\left. \frac{\partial Q(\theta; \theta^{(t)})}{\partial \theta} \right|_{\theta=\theta^{(t+1)}} = \mathbf{0}.$$

(D2) The gradient $\partial Q(\theta; \theta^{(t)})/\partial \theta$ is continuous in both arguments $(\theta, \theta^{(t)})$.

(D3) One of the following conditions holds:

- i) The set $\mathcal{S}(\tilde{\ell}) = \{\theta \in \mathcal{S} : \ell(\theta) = \tilde{\ell}\}$ contains a single point $\tilde{\theta}$, where $\mathcal{S}$ is a set of accumulation points of $\{\theta^{(t)}\}_{t \in \mathbb{N}}$,
- ii) $\|\theta^{(t+1)} - \theta^{(t)}\| \rightarrow 0$ and the corresponding limit set $\mathcal{S}(\tilde{\ell}) = \{\theta \in \mathcal{S} : \ell(\theta) = \tilde{\ell}\}$ is discrete,
- iii) $\sum_{t=1}^{\infty} \|\theta^{(t+1)} - \theta^{(t)}\| < \infty$.

Then the sequence of likelihood values $\{\ell(\theta^{(t)})\}_{t \in \mathbb{N}}$ converges monotonically to a finite limit $\tilde{\ell}$, and the parameter sequence $\{\theta^{(t)}\}_{t \in \mathbb{N}}$ converges to a stationary point $\tilde{\theta}$ of ℓ .

Proof. Theorem 1.2 ensures that the sequence of likelihood values $\{\ell(\theta^{(t)})\}_{t \in \mathbb{N}}$ converges to some finite limit $\tilde{\ell}$.

- i) If the level set $\mathcal{S}(\tilde{\ell}) = \{\theta \in \mathcal{S} : \ell(\theta) = \tilde{\ell}\}$ consists of a single point $\tilde{\theta}$, then $\theta^{(t)} \xrightarrow[t \rightarrow \infty]{} \tilde{\theta}$ trivially.
- ii) If $\|\theta^{(t+1)} - \theta^{(t)}\| \rightarrow 0$ and the corresponding limit set

$$\mathcal{S}(\ell^*) = \{\theta \in \mathcal{S} : \ell(\theta) = \ell^*\},$$

is discrete, then every accumulation point of $\{\theta^{(t)}\}_{t \in \mathbb{N}}$ lies in $\mathcal{S}(\ell^*)$ since $\ell(\theta^{(t)}) \xrightarrow[t \rightarrow \infty]{} \ell^*$. Moreover, vanishing steps imply that $\mathcal{S}(\ell^*)$ is nonempty, compact, and connected. Because a connected subset of a discrete set in $\mathbb{R}^d$ is a singleton, we must have $\mathcal{S}(\ell^*) = \{\theta^*\}$, and hence $\theta^{(t)} \xrightarrow[t \rightarrow \infty]{} \theta^*$.

- iii) If $\sum_{t=1}^{\infty} \|\theta^{(t+1)} - \theta^{(t)}\| < \infty$, then for any $m > n$,

$$\|\theta^{(m)} - \theta^{(n)}\| \leq \sum_{t=n}^{m-1} \|\theta^{(t+1)} - \theta^{(t)}\| \rightarrow 0 \quad (n \rightarrow \infty),$$

so $\{\theta^{(t)}\}_{t \in \mathbb{N}}$ is a Cauchy sequence and therefore convergent in $\mathbb{R}^d$.

Now, differentiating both sides of the decomposition

$$\ell(\theta) = Q(\theta; \theta^{(t)}) - H(\theta; \theta^{(t)}),$$

with respect to θ yields

$$\frac{\partial \ell(\theta)}{\partial \theta} = \frac{\partial Q(\theta; \theta^{(t)})}{\partial \theta} - \frac{\partial H(\theta; \theta^{(t)})}{\partial \theta}.$$

From the inequality $H(\theta; \theta^{(t)}) \leq H(\theta^{(t)}; \theta^{(t)})$ and since the gradient of the KL-divergence vanishes at zero, it follows that

$$\left[\frac{\partial H(\theta; \theta^{(t)})}{\partial \theta} \right]_{\theta=\theta^{(t)}} = 0.$$

Let θ_0 be an arbitrary value of θ . Substituting $\theta^{(t)} = \theta_0$ into the previous equation gives

$$\frac{\partial \ell(\theta_0)}{\partial \theta} = \left[\frac{\partial Q(\theta; \theta_0)}{\partial \theta} \right]_{\theta=\theta_0},$$

where the left-hand side denotes the gradient of $\ell(\theta)$ evaluated at $\theta = \theta_0$.

Now suppose that $\theta_0 = \tilde{\theta}$, where $\tilde{\theta}$ is an accumulation point of $\{\theta^{(t)}\}_{t \in \mathbb{N}}$. Then (since the gradient $\partial Q(\theta; \theta^{(t)})/\partial \theta$ is continuous in both arguments $(\theta, \theta^{(t)})$),

$$\frac{\partial \ell(\tilde{\theta})}{\partial \theta} = \left[\frac{\partial Q(\theta; \tilde{\theta})}{\partial \theta} \right]_{\theta=\tilde{\theta}} = \lim_{t \rightarrow \infty} \left[\frac{\partial Q(\theta; \theta^{(t)})}{\partial \theta} \right]_{\theta=\theta^{(t)}} \stackrel{*}{=} \lim_{t \rightarrow \infty} \left[\frac{\partial Q(\theta; \theta^{(t)})}{\partial \theta} \right]_{\theta=\theta^{(t+1)}} = \lim_{t \rightarrow \infty} 0 = 0,$$

where step $*$ is justified because of the assumption (D2) and convergence of $\{\theta^{(t)}\}_{t \in \mathbb{N}}$. Thus, the limit point $\tilde{\theta}$ is a stationary point of $\ell(\theta)$.

Finally, if (D3) i) holds, then $\tilde{\theta}$ is the maximum likelihood estimator θ_{MLE} . $\square$

Remarks.

- Assumption **(D1)** is automatically satisfied by any EM algorithm that exactly maximizes the Q -function at each iteration.
- Note that the sequence $\{\theta^{(t)}\}_{t \in \mathbb{N}}$ of Example 1.2 doesn't satisfy any of the conditions in **(D3)**.
- The main advantage of Theorem 1.3 over its counterparts Wu (1983, Theorem 1) and McLachlan and Krishnan (2008, Theorems 3.3–3.4) is that it does not require the multi-valued GEM update mapping

$$\mathcal{M} : \Theta \rightarrow 2^\Theta, \quad \theta \mapsto \mathcal{M}(\theta),$$

to be *closed outside the stationary set* $\mathcal{S} = \{\theta : \nabla \ell(\theta) = 0\}$. Formally, closedness outside $\mathcal{S}$ means:

$$\begin{aligned} \forall \{\theta^{(t)}\}_{t \in \mathbb{N}}, \{\phi^{(t)}\}_{t \in \mathbb{N}} \subset \Theta, \quad & \left(\theta^{(t)} \xrightarrow[t \rightarrow \infty]{} \theta, \phi^{(t)} \in \mathcal{M}(\theta^{(t)}) \forall t, \phi^{(t)} \xrightarrow[t \rightarrow \infty]{} \phi, \theta \notin \mathcal{S} \right) \\ \implies & \phi \in \mathcal{M}(\theta). \end{aligned}$$

This condition is hard to verify if M-step is implicit.

- Assumption **(D1)** can be weakened to accommodate Expectation-Conditional Maximization (ECM) algorithms, in which the M-step is decomposed into a sequence of exact conditional maximizations over disjoint parameter blocks. In this case, the full first-order condition

$$\left. \frac{\partial Q(\theta; \theta^{(t)})}{\partial \theta} \right|_{\theta=\theta^{(t+1)}} = \mathbf{0}$$

need not hold at finite t .

Nevertheless, if the blockwise first-order conditions are satisfied and the gradient $\partial Q(\theta; \theta')/\partial \theta$ is continuous in both arguments, then vanishing step sizes $\|\theta^{(t+1)} - \theta^{(t)}\| \rightarrow 0$ imply

$$\left\| \frac{\partial Q(\theta; \theta^{(t)})}{\partial \theta} \right\|_{\theta=\theta^{(t+1)}} \xrightarrow{t \rightarrow \infty} 0.$$

As a consequence, the asymptotic identity

$$\lim_{t \rightarrow \infty} \left[\frac{\partial Q(\theta; \theta^{(t)})}{\partial \theta} \right]_{\theta=\theta^{(t)}} = \lim_{t \rightarrow \infty} \left[\frac{\partial Q(\theta; \theta^{(t)})}{\partial \theta} \right]_{\theta=\theta^{(t+1)}} = \mathbf{0}$$

remains valid, and the conclusion that every accumulation point of $\{\theta^{(t)}\}$ is a stationary point of the observed-data log-likelihood $\ell(\theta)$ continues to hold.

So far, the presented Theorems 1.2 and 1.3 rely on the compactness assumption **(A2)** on the upper-level sets

$$\Theta_{\theta_0} = \{\theta \in \Theta : \ell(\theta) \geq \ell(\theta_0)\},$$

which may fail, for instance, in the case of mixture models.

Example 1.3. Consider a univariate two-component Gaussian mixture:

$$f(y | \theta) = \pi \phi(y | \mu_1, \sigma_1^2) + (1 - \pi) \phi(y | \mu_2, \sigma_2^2), \quad \theta = (\pi, \mu_1, \mu_2, \sigma_1^2, \sigma_2^2),$$

where $\phi(y | \mu, \sigma^2) = (2\pi\sigma^2)^{-1/2} \exp[-(y - \mu)^2/(2\sigma^2)]$. Let y_i be one of the observations. Choose

$$\pi = \frac{1}{n}, \quad \mu_1 = y_i, \quad \sigma_1^2 = \varepsilon \rightarrow 0,$$

and denote with θ_ε this parameter choice. Then,

$$\ell(\theta_\varepsilon) = \sum_{j=1}^n \log \left[\pi \phi(y_j | \mu_1, \sigma_{1,\varepsilon}^2) + (1 - \pi) \phi(y_j | \mu_2, \sigma_2^2) \right].$$

Case $j = i$. Using $\log(a + b) \geq \log a$,

$$\log \left[\pi \phi(y_i | y_i, \varepsilon) + (1 - \pi) \phi(y_i | \mu_2, \sigma_2^2) \right] \geq \log [\pi \phi(y_i | y_i, \varepsilon)] = \log \pi - \frac{1}{2} \log(2\pi\varepsilon).$$

Hence this single term is

$$\log \pi - \frac{1}{2} \log(2\pi) - \frac{1}{2} \log \varepsilon \xrightarrow{\varepsilon \downarrow 0} +\infty.$$

Case $j \neq i$. For $j \neq i$,

$$\phi(y_j | y_i, \varepsilon) = \frac{1}{\sqrt{2\pi\varepsilon}} \exp\left(-\frac{(y_j - y_i)^2}{2\varepsilon}\right) \xrightarrow{\varepsilon \downarrow 0} 0.$$

Combining both cases,

$$\ell(\theta_\varepsilon) \geq \underbrace{\left(\log \pi - \frac{1}{2} \log(2\pi) - \frac{1}{2} \log \varepsilon \right)}_{\text{from } j=i} + \sum_{j \neq i} \log \left[(1 - \pi) \phi(y_j | \mu_2, \sigma_2^2) \right] \xrightarrow{\varepsilon \downarrow 0} +\infty.$$

Therefore $\sup_{\theta \in \Theta} \ell(\theta) = \infty$ and, for any θ_0 with $\ell(\theta_0)$ finite, the upper level set $\Theta_{\theta_0} = \{\theta : \ell(\theta) \geq \ell(\theta_0)\}$ is unbounded (noncompact). Although in this setting Theorems 1.2–1.3 no longer apply directly, the EM algorithm still ensures monotone ascent of $\ell(\theta^{(t)})$ because of Theorem 1.1.

For a detailed discussion of this pathology and possible remedies (e.g. constrained variances), see McLachlan and Krishnan (2008, Section 3.5.4).

Convergence speed of the EM mapping. While the preceding results establish global convergence of the EM sequence to a stationary point, they do not describe the *rate* at which $\theta^{(t)}$ approaches its limit. The local behavior can be characterized by a linearization of the EM mapping around a stationary point, yielding a first-order approximation of the convergence speed.

Theorem 1.4: Local convergence rate of the EM algorithm (McLachlan and Krishnan, 2008, Section 3.9.1)

Let $\ell(\theta)$ denote the observed log-likelihood and $M(\theta)$ the EM mapping, $\theta^{(t+1)} = M(\theta^{(t)})$. Assume that the following regularity conditions hold in a neighborhood of a fixed point θ^* satisfying $M(\theta^*) = \theta^*$:

- (R1) The functions $\ell(\theta)$ and $Q(\theta; \theta^{(t)})$ are twice continuously differentiable in θ for all $\theta, \theta^{(t)}$ in a neighborhood of θ^* .
- (R2) The observed-data and complete-data information matrices (expectations taken with respect to the complete-data model $f(y, u | \theta)$) are defined as

$$I_{\text{obs}}(\theta) = -\mathbb{E}_{\theta} \left[\frac{\partial^2 \ell(\theta)}{\partial \theta \partial \theta^{\top}} \right], \quad I_{\text{com}}(\theta) = -\mathbb{E}_{\theta} \left[\frac{\partial^2 \log f(y, u | \theta)}{\partial \theta \partial \theta^{\top}} \right],$$

and are assumed to exist, be finite, and continuous in θ .

- (R3) The matrix $I_{\text{com}}(\theta^*)$ is nonsingular (positive definite).
- (R4) Expectation and differentiation can be interchanged in a neighborhood of θ^* , that is,

$$\frac{\partial}{\partial \theta} \mathbb{E}_{\theta}[g(y, u, \theta)] = \mathbb{E}_{\theta} \left[\frac{\partial g(y, u, \theta)}{\partial \theta} \right].$$

Then, near θ^* , the EM iteration admits the local linear expansion

$$\theta^{(t+1)} - \theta^* = R(\theta^*) (\theta^{(t)} - \theta^*) + o(\|\theta^{(t)} - \theta^*\|),$$

where

$$R(\theta^*) = I_{\text{com}}(\theta^*)^{-1} I_{\text{mis}}(\theta^*), \quad I_{\text{mis}}(\theta) = I_{\text{com}}(\theta) - I_{\text{obs}}(\theta).$$

The matrix $R(\theta^*)$ is the Jacobian of the EM mapping at θ^* , and its *spectral radius* $\rho(R(\theta^*)) := \max_{i=1, \dots, k} |\lambda_i|$, i.e. the largest modulus of its eigenvalues λ_i , determines the local linear convergence rate.

Proof. Let $\partial \ell(\theta) / \partial \theta$ denote the observed-data score function. By a first-order Taylor expansion of $\partial \ell(\theta) / \partial \theta$ around $\theta^{(t)}$,

$$\frac{\partial \ell(\theta)}{\partial \theta} \approx \frac{\partial \ell(\theta^{(t)})}{\partial \theta} - I_{\text{obs}}(\theta^{(t)}) (\theta - \theta^{(t)}).$$

Setting $\theta = \theta^*$ and using $\partial \ell(\theta^*) / \partial \theta = 0$ gives

$$\theta^* - \theta^{(t)} \approx I_{\text{obs}}(\theta^{(t)})^{-1} \frac{\partial \ell(\theta^{(t)})}{\partial \theta}. \quad (1)$$

Next, consider the M-step condition

$$\left. \frac{\partial Q(\theta; \theta^{(t)})}{\partial \theta} \right|_{\theta=\theta^{(t+1)}} = \mathbf{0}_d.$$

Expanding $\partial\ell(\theta^{(t)})/\partial\theta = \partial Q(\theta; \theta^{(t)})/\partial\theta$ around $\theta^{(t)}$ yields

$$\frac{\partial\ell(\theta^{(t)})}{\partial\theta} \approx \mathbf{0}_d + I_{\text{com}}(\theta^{(t)})(\theta^{(t+1)} - \theta^{(t)}). \quad (2)$$

Substituting (2) into (1) gives

$$\theta^* - \theta^{(t)} \approx I_{\text{obs}}(\theta^{(t)})^{-1} I_{\text{com}}(\theta^{(t)})(\theta^{(t+1)} - \theta^{(t)}).$$

Rearranging,

$$\theta^{(t+1)} - \theta^* \approx \left\{ I_d - I_{\text{com}}(\theta^{(t)})^{-1} I_{\text{obs}}(\theta^{(t)}) \right\} (\theta^{(t)} - \theta^*).$$

Evaluating at $\theta = \theta^*$ gives

$$\theta^{(t+1)} - \theta^* \approx R(\theta^*)(\theta^{(t)} - \theta^*), \quad R(\theta^*) = I_{\text{com}}(\theta^*)^{-1} I_{\text{mis}}(\theta^*),$$

where $I_{\text{mis}}(\theta) = I_{\text{com}}(\theta) - I_{\text{obs}}(\theta)$.

One can also rewrite the approximation as

$$\|\theta^{(t+1)} - \theta^*\| \lesssim \rho(R(\theta^*)) \|\theta^{(t)} - \theta^*\|,$$

where the spectral radius $\rho(R(\theta^*)) := \max_{i=1, \dots, k} |\lambda_i| \in [0, 1)$ – the largest modulus of the eigenvalues λ_i of $R(\theta^*)$ – determines the local linear convergence rate. $\square$

If no information is missing, then $I_{\text{mis}} = 0$ and hence $R(\theta^*) = 0$; the EM algorithm converges in a single step, since direct maximization of the complete-data log-likelihood is possible. Conversely, when the fraction of missing information (quantified by $R(\theta^*)$)—and thus its spectral radius $\rho(R(\theta^*))$ —is large, EM makes only a fraction $\rho(R(\theta^*))$ of the remaining distance from $\theta^{(t)}$ to θ^* per iteration instead of reaching the optimum immediately.

Example 1.4. Let $Y_1, \dots, Y_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mu, \sigma^2)$ and suppose each Y_i is observed with probability $1 - p$ and missing with probability p (MCAR).

For simplicity, assume σ^2 is known. The complete-data log-likelihood for μ is

$$\ell_{\text{com}}(\mu) = -\frac{n}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - \mu)^2.$$

The complete-data score and observed information are

$$\frac{\partial\ell_{\text{com}}(\mu)}{\partial\mu} = \frac{1}{\sigma^2} \sum_{i=1}^n (y_i - \mu), \quad I_{\text{com}} = -\mathbb{E} \left[\frac{\partial^2 \ell_{\text{com}}(\mu)}{\partial\mu^2} \right] = \frac{n}{\sigma^2}.$$

Now suppose each Y_i is missing completely at random (MCAR) with probability p , independently of the data. Let $\delta_i \in \{0, 1\}$ be the indicator of observation ($\delta_i = 1$ if Y_i is observed), so that $\delta_i \sim \text{Bernoulli}(1 - p)$ independently of Y_i . Then the observed-data log-likelihood is

$$\ell_{\text{obs}}(\mu) = -\frac{1}{2\sigma^2} \sum_{i:\delta_i=1} (y_i - \mu)^2 + \text{const.}$$

Conditional on $\{\delta_i\}$, only $(1 - p)n$ observations contribute, so

$$I_{\text{obs}} = -\mathbb{E} \left[\frac{\partial^2 \ell_{\text{obs}}(\mu)}{\partial\mu^2} \right] = \frac{(1 - p)n}{\sigma^2}.$$

Hence, the fraction of missing information is

$$I_{\text{mis}} = I_{\text{com}} - I_{\text{obs}} = \frac{p n}{\sigma^2}.$$

and the local EM rate is

$$\rho = \frac{I_{\text{mis}}}{I_{\text{com}}} = p.$$

Thus if 30% of values are missing ($p = 0.3$), the error contracts by ≈ 0.3 per EM step (fast); if 95% are missing, $\rho \approx 0.95$ (slow).

Practical stopping rules. Possible criteria are (i) relative or absolute increases in the observed log-likelihood, e.g. $|\ell(\theta^{(t+1)}) - \ell(\theta^{(t)})| < \varepsilon$ or (ii) parameter stabilization, $\|\theta^{(t+1)} - \theta^{(t)}\| < \varepsilon$.

The stationary-point result holds for both EM and GEM since GEM maintains monotonicity of ℓ while ensuring closedness of the update mapping (Wu, 1983). When $\rho(R(\theta^*))$ is close to 1 (high missing information), EM may converge slowly; ECM/ECME, PX-EM, or quasi-Newton accelerations can reduce $\rho(R(\theta^*))$ (Meng and van Dyk, 1997; McLachlan and Krishnan, 2008).

1.4 Linear mixed models as an EM algorithm (adapted from Laird, 1982)

This section follows the EM formulation of linear mixed models introduced by Laird (1982).

Model definition (general case). Consider m independent groups (clusters) indexed by $i = 1, \dots, m$, with group sizes n_i and total sample size $n = \sum_{i=1}^m n_i$. For each group, let $y_i \in \mathbb{R}^{n_i}$ denote the response vector, $X_i \in \mathbb{R}^{n_i \times p}$ the fixed-effect design matrix with the full column rank, and $Z_i \in \mathbb{R}^{n_i \times q}$ the random-effect design (e.g., $q = 2$ for a random intercept and slope).

The linear mixed model (LMM) for group i is

$$\mathbf{Y}_i = X_i \beta^* + Z_i \mathbf{U}_i + \boldsymbol{\varepsilon}_i, \quad \mathbf{U}_i \sim \mathcal{N}(0, D_i^*), \quad \boldsymbol{\varepsilon}_i \sim \mathcal{N}(0, \Sigma_i^*),$$

with $\mathbf{U}_i$ independent across groups and independent of $\boldsymbol{\varepsilon}_i$; $\text{Cov}(\mathbf{U}_i, \boldsymbol{\varepsilon}_i) = 0 \ \forall i = 1, \dots, m$ and $D_i^* = D_j^* \ \forall i, j = 1, \dots, m$. Here $\Sigma_i^* \in \mathbb{S}_{++}^{n_i}$ denotes a positive-definite covariance matrix and $\mathbb{S}_{++}^{n_i}$ denotes the set of real symmetric positive-definite $n_i \times n_i$ matrices.

Stacking all groups gives

$$\mathbf{Y} = \begin{bmatrix} \mathbf{Y}_1 \\ \vdots \\ \mathbf{Y}_m \end{bmatrix} \in \mathbb{R}^n, \quad X = \begin{bmatrix} X_1 \\ \vdots \\ X_m \end{bmatrix} \in \mathbb{R}^{n \times p}, \quad Z = \text{blockdiag}(Z_1, \dots, Z_m) \in \mathbb{R}^{n \times mq},$$

and

$$\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \vdots \\ \mathbf{U}_m \end{bmatrix} \in \mathbb{R}^{mq}, \quad D^* = \text{blockdiag}(D_1^*, \dots, D_m^*), \quad \boldsymbol{\varepsilon} = \begin{bmatrix} \boldsymbol{\varepsilon}_1 \\ \vdots \\ \boldsymbol{\varepsilon}_m \end{bmatrix} \in \mathbb{R}^n, \quad \Sigma^* = \text{blockdiag}(\Sigma_1^*, \dots, \Sigma_m^*).$$

The stacked model reads

$$\mathbf{Y} = X\beta^* + Z\mathbf{U} + \boldsymbol{\varepsilon}, \quad \mathbf{U} \sim \mathcal{N}(0, D^*), \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(0, \Sigma^*), \quad \text{Cov}(\mathbf{Y}) = V^* = ZD^*Z^\top + \Sigma^*.$$

We denote the full parameter vector by

$$\boldsymbol{\theta} = (\beta, D, \Sigma).$$

Alternatively, one may impose a parametric structure on the covariance components (D, Σ) , assuming that they are fully specified by a parameter $\eta \in \mathbb{R}^{p'}$. In this case, the parameter vector reduces to

$$\boldsymbol{\theta} = (\beta, \eta) \in \mathbb{R}^{p+p'}.$$

Complete-data log-likelihood. Treat the random effects $\mathbf{U} = (\mathbf{U}_1^\top, \dots, \mathbf{U}_m^\top)^\top$ as latent variables. With $\boldsymbol{\theta} = (\beta, D, \Sigma)$,

$$\mathbf{Y} \mid \mathbf{U} = u, \boldsymbol{\theta} \sim \mathcal{N}(X\beta + Zu, \Sigma), \quad \mathbf{U} \mid \boldsymbol{\theta} \sim \mathcal{N}(0, D).$$

The joint density factorizes as

$$f(y, u \mid \boldsymbol{\theta}) = f(y \mid u, \boldsymbol{\theta}) f(u \mid \boldsymbol{\theta}),$$

and both factors are Gaussian:

$$f(y \mid u, \boldsymbol{\theta}) = (2\pi)^{-n/2} |\Sigma|^{-1/2} \exp\left\{-\frac{1}{2}(y - X\beta - Zu)^\top \Sigma^{-1}(y - X\beta - Zu)\right\},$$

$$f(u \mid \theta) = (2\pi)^{-mq/2} |D|^{-1/2} \exp\left\{-\frac{1}{2} u^\top D^{-1} u\right\}.$$

Taking logarithms and summing the two contributions gives the complete-data log-likelihood:

$$\begin{aligned} \ell_c(\theta \mid y, u) &= \log f(y, u \mid \theta) = \log f(y \mid u, \theta) + \log f(u \mid \theta) \\ &= -\frac{1}{2} [\log |2\pi\Sigma| + \log |2\pi D| + (y - X\beta - Zu)^\top \Sigma^{-1} (y - X\beta - Zu) + u^\top D^{-1} u]. \end{aligned}$$

Omitting constants independent of θ (i.e. dropping $\log(2\pi)$ terms) yields

$$\boxed{\ell_c(\theta \mid y, u) = -\frac{1}{2} [\log |\Sigma| + \log |D| + (y - X\beta - Zu)^\top \Sigma^{-1} (y - X\beta - Zu) + u^\top D^{-1} u].}$$

E-step (conditional moments of random effects)

At iteration t , let $\theta^{(t)} = (\beta^{(t)}, D^{(t)}, \Sigma^{(t)})$ denote the current parameter estimates. The latent random effects u_i are updated via their conditional distribution given the observed data y_i and current parameters.

From the joint Gaussian structure

$$\begin{bmatrix} \mathbf{Y}_i \\ \mathbf{U}_i \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} X_i \beta^* \\ 0 \end{bmatrix}, \begin{bmatrix} V_i^* & Z_i D_i^* \\ D_i^* Z_i^\top & D_i^* \end{bmatrix}\right), \quad V_i^* = Z_i D_i^* Z_i^\top + \Sigma_i^*,$$

the conditional distribution of $\mathbf{u}_i$ given y_i and θ is again multivariate normal (a standard result for the conditional of a joint Gaussian):

$$\mathbf{U}_i \mid \mathbf{Y}_i = y_i, \theta \sim \mathcal{N}\left(D_i Z_i^\top V_i^{-1} (y_i - X_i \beta), D_i - D_i Z_i^\top V_i^{-1} Z_i D_i\right).$$

Because the conditional distribution is Gaussian, its **conditional mean equals its mode**. Hence, in the E-step, the conditional mean and covariance (which enter the expected complete-data log-likelihood) are computed as

$$m_i^{(t)} = \mathbb{E}[\mathbf{U}_i \mid \mathbf{Y}_i = y_i, \theta^{(t)}] = D_i^{(t)} Z_i^{(t)\top} (V_i^{(t)})^{-1} (y_i - X_i \beta^{(t)}),$$

$$C_i^{(t)} = \text{Var}(\mathbf{U}_i \mid \mathbf{Y}_i = y_i, \theta^{(t)}) = D_i^{(t)} - D_i^{(t)} Z_i^{(t)\top} (V_i^{(t)})^{-1} Z_i D_i^{(t)},$$

where $V_i^{(t)} = Z_i D_i^{(t)} Z_i^\top + \Sigma_i^{(t)}$. The vector $m_i^{(t)}$ is thus the best linear unbiased predictor (BLUP) of u_i , and $C_i^{(t)}$ is its posterior covariance matrix.

M-step (parameter updates)

In the M-step we maximize the conditional expectation of the complete-data log-likelihood,

$$Q(\theta \mid \theta^{(t)}) = \mathbb{E}_{\mathbf{U} \mid \mathbf{Y}=y, \theta^{(t)}} [\ell_c(\theta) \mid y, \theta^{(t)}], \quad \theta = (\beta, D, \Sigma),$$

where

$$\ell_c(\theta \mid y, u) = -\frac{1}{2} [\log |\Sigma| + \log |D| + (y - X\beta - Zu)^\top \Sigma^{-1} (y - X\beta - Zu) + u^\top D^{-1} u].$$

The expectation is taken with respect to the conditional distribution $\mathbf{U} \mid (\mathbf{Y} = y, \theta^{(t)}) \sim \mathcal{N}(m^{(t)}, C^{(t)})$, where $m^{(t)} = \mathbb{E}[\mathbf{U} \mid \mathbf{Y} = y, \theta^{(t)}]$ and $C^{(t)} = \text{Var}(\mathbf{U} \mid \mathbf{Y} = y, \theta^{(t)})$.

Term 1: $(y - X\beta - Zu)^\top \Sigma^{-1}(y - X\beta - Zu)$.

Let $a = y - X\beta$ and $B = Z$. Then

$$\mathbb{E}[(a - BU)^\top \Sigma^{-1}(a - BU) \mid \mathbf{Y} = y, \theta^{(t)}] = \mathbb{E}[(a - BU)^\top \Sigma^{-1}(a - BU)].$$

Expanding the product and taking expectations gives

$$\begin{aligned} \mathbb{E}[(a - BU)^\top \Sigma^{-1}(a - BU)] &= a^\top \Sigma^{-1}a - 2a^\top \Sigma^{-1}B \mathbb{E}[\mathbf{U}] + \mathbb{E}[\mathbf{U}^\top B^\top \Sigma^{-1}B \mathbf{U}] \\ &= a^\top \Sigma^{-1}a - 2a^\top \Sigma^{-1}Bm^{(t)} + \left\{ \text{tr}(B^\top \Sigma^{-1}B C^{(t)}) + (m^{(t)})^\top B^\top \Sigma^{-1}Bm^{(t)} \right\}. \end{aligned}$$

Collecting terms,

$$\begin{aligned} \mathbb{E}[(a - BU)^\top \Sigma^{-1}(a - BU) \mid \mathbf{Y} = y, \theta^{(t)}] &= (a - Bm^{(t)})^\top \Sigma^{-1}(a - Bm^{(t)}) + \text{tr}(B^\top \Sigma^{-1}B C^{(t)}) \\ &= (y - X\beta - Zm^{(t)})^\top \Sigma^{-1}(y - X\beta - Zm^{(t)}) + \text{tr}(Z^\top \Sigma^{-1}Z C^{(t)}). \end{aligned}$$

Hence

$$\mathbb{E}[(y - X\beta - Z\mathbf{U})^\top \Sigma^{-1}(y - X\beta - Z\mathbf{U}) \mid \mathbf{Y} = y, \theta^{(t)}] = (y - X\beta - Zm^{(t)})^\top \Sigma^{-1}(y - X\beta - Zm^{(t)}) + \text{tr}(Z^\top \Sigma^{-1}Z C^{(t)}).$$

Term 2: $\mathbf{U}^\top D^{-1}\mathbf{U}$.

Applying the same identity with $A = D^{-1}$ gives

$$\begin{aligned} \mathbb{E}[\mathbf{U}^\top D^{-1}\mathbf{U} \mid \mathbf{Y} = y, \theta^{(t)}] &= \text{tr}(D^{-1} \text{Var}(\mathbf{U})) + (\mathbb{E}\mathbf{U})^\top D^{-1}(\mathbb{E}\mathbf{U}) \\ &= \text{tr}(D^{-1}C^{(t)}) + (m^{(t)})^\top D^{-1}m^{(t)}. \end{aligned}$$

Substitution into $Q(\theta \mid \theta^{(t)})$.

Plugging the two expectations into $\ell_c(\theta \mid y, u)$ yields

$$\begin{aligned} Q(\theta \mid \theta^{(t)}) &= -\frac{1}{2} \left[\log |\Sigma| + \log |D| + (y - X\beta - Zm^{(t)})^\top \Sigma^{-1}(y - X\beta - Zm^{(t)}) + \text{tr}(Z^\top \Sigma^{-1}Z C^{(t)}) \right. \\ &\quad \left. + \text{tr}(D^{-1}C^{(t)}) + (m^{(t)})^\top D^{-1}m^{(t)} \right]. \end{aligned} \tag{1}$$

Update for β (GLS / FGLS form). Only the first quadratic term depends on β . Differentiating the Q -function with respect to β and setting the derivative equal to zero yields the normal equation

$$X^\top (\Sigma^{(t)})^{-1} (y - X\beta^{(t+1)} - Zm^{(t)}) = 0,$$

where $m^{(t)} = D^{(t)} Z^\top (V^{(t)})^{-1} (y - X\beta^{(t)})$.

Hence,

$$\boxed{\beta^{(t+1)} = \left(X^\top (\Sigma^{(t)})^{-1} X \right)^{-1} X^\top (\Sigma^{(t)})^{-1} (y - Zm^{(t)})}.$$

Note that this normal equation does *not* coincide with the normal equations of the feasible generalized least-squares (FGLS) estimator based on the estimated marginal covariance matrix

$$V^{(t)} = ZD^{(t)}Z^\top + \Sigma^{(t)}.$$

However, at a β -fixed point of the iteration, that is when $\beta^{(t)} = \beta^{(t+1)}$, we obtain

$$\begin{aligned} y - X\beta^{(t+1)} - Zm^{(t)} &= y - X\beta^{(t+1)} - ZD^{(t)}Z^\top(V^{(t)})^{-1}(y - X\beta^{(t)}) \\ &= y - X\beta^{(t+1)} - ZD^{(t)}Z^\top(V^{(t)})^{-1}(y - X\beta^{(t+1)}) \\ &= \left(I - ZD^{(t)}Z^\top(V^{(t)})^{-1}\right)(y - X\beta^{(t+1)}). \end{aligned}$$

Since $V^{(t)} = \Sigma^{(t)} + ZD^{(t)}Z^\top$, the Woodbury identity yields

$$I - ZD^{(t)}Z^\top(V^{(t)})^{-1} = \Sigma^{(t)}(V^{(t)})^{-1}.$$

Therefore,

$$y - X\beta^{(t+1)} - Zm^{(t)} = \Sigma^{(t)}(V^{(t)})^{-1}(y - X\beta^{(t+1)}).$$

and the EM normal equation can be rewritten as

$$X^\top(\Sigma^{(t)})^{-1}(y - X\beta^{(t+1)} - Zm^{(t)}) = X^\top(V^{(t)})^{-1}(y - X\beta^{(t+1)}) = 0.$$

This equation coincides with the feasible generalized least-squares (FGLS) estimator based on the estimated marginal covariance matrix $V^{(t)}$. In particular, if $\theta^{(t+1)} = \theta^{(t)} = \theta^*$, the EM update for β coincides with the GLS estimator based on the true marginal covariance matrix V^* .

Update for D (via derivative w.r.t. D^{-1}). Recall that the conditional second moment of the stacked random effects is

$$\mathbb{E}(uu^\top \mid y, \theta^{(t)}) = C^{(t)} + m^{(t)}m^{(t)\top},$$

where $m^{(t)} = (m_1^{(t)}, \dots, m_m^{(t)})$ and $C^{(t)} = \text{blockdiag}(C_1^{(t)}, \dots, C_m^{(t)})$.

The D -dependent part of the Q -function can be written in global form as

$$Q_D(D) = -\frac{1}{2} \left[m \log |D| + \text{tr}(D^{-1}(C^{(t)} + m^{(t)}m^{(t)\top})) \right].$$

Using the matrix identities (for symmetric positive-definite D)

$$\frac{\partial \log |D|}{\partial D^{-1}} = -D, \quad \frac{\partial}{\partial D^{-1}} \text{tr}(D^{-1}S) = S,$$

we obtain

$$\frac{\partial Q_D}{\partial D^{-1}} = -\frac{1}{2} \left[-mD + (C^{(t)} + m^{(t)}m^{(t)\top}) \right].$$

Setting the derivative equal to zero yields

$$mD^{(t+1)} = C^{(t)} + m^{(t)}m^{(t)\top},$$

and therefore

$$D^{(t+1)} = \frac{1}{m} (C^{(t)} + m^{(t)}m^{(t)\top}).$$

Update for Σ . The Σ -dependent part of the Q -function can be written in stacked form as

$$Q_{\Sigma}(\Sigma) = -\frac{1}{2} \left[\log |\Sigma| + \text{tr}(\Sigma^{-1} R^{(t)}) \right],$$

where

$$R^{(t)} = (y - X\beta^{(t+1)} - Zm^{(t)})(y - X\beta^{(t+1)} - Zm^{(t)})^{\top} + ZC^{(t)}Z^{\top}.$$

Using the identities

$$\frac{\partial \log |\Sigma|}{\partial \Sigma^{-1}} = -\Sigma, \quad \frac{\partial}{\partial \Sigma^{-1}} \text{tr}(\Sigma^{-1} R) = R,$$

we obtain

$$\frac{\partial Q_{\Sigma}}{\partial \Sigma^{-1}} = -\frac{1}{2} \left[-\Sigma + R^{(t)} \right].$$

Setting this derivative equal to zero yields the M-step update

$$\boxed{\Sigma^{(t+1)} = R^{(t)}}.$$

ECM formulation and GEM property. Consider the parameter vector $\theta = (\beta, D, \Sigma)$ and denote by

$$Q(\theta \mid \theta^{(t)}) = \mathbb{E}_{\theta^{(t)}} [\log f(y, u \mid \theta) \mid y]$$

the expected complete-data log-likelihood at iteration t . In the algorithm described before, we proceed by first performing an E-step under $\theta^{(t)}$ and then updating the parameter blocks sequentially via conditional maximization steps.

This procedure constitutes an *Expectation-Conditional Maximization (ECM)* algorithm Meng and van Dyk (1997), since the M-step is replaced by a sequence of exact conditional maximizations over the parameter blocks β , D , and Σ , each carried out while holding the remaining blocks fixed and with respect to the same Q -function.

More precisely, define the intermediate iterates

$$\begin{aligned} \theta^{(t+0)} &:= (\beta^{(t)}, D^{(t)}, \Sigma^{(t)}), & \theta^{(t+1/3)} &:= (\beta^{(t+1)}, D^{(t)}, \Sigma^{(t)}), \\ \theta^{(t+2/3)} &:= (\beta^{(t+1)}, D^{(t+1)}, \Sigma^{(t)}), & \theta^{(t+1)} &:= (\beta^{(t+1)}, D^{(t+1)}, \Sigma^{(t+1)}). \end{aligned}$$

Each update is defined as an exact maximization of $Q(\cdot \mid \theta^{(t)})$ over the corresponding parameter block while keeping the remaining blocks fixed. Consequently, the following monotonicity chain holds:

$$Q(\theta^{(t+1)} \mid \theta^{(t)}) \geq Q(\theta^{(t+2/3)} \mid \theta^{(t)}) \geq Q(\theta^{(t+1/3)} \mid \theta^{(t)}) \geq Q(\theta^{(t)} \mid \theta^{(t)}).$$

In particular,

$$Q(\theta^{(t+1)} \mid \theta^{(t)}) \geq Q(\theta^{(t)} \mid \theta^{(t)}),$$

which shows that the proposed ECM algorithm satisfies the defining inequality of a generalized EM (GEM) algorithm. As a consequence, the observed-data log-likelihood is nondecreasing along the iterates, i.e.,

$$\ell(\theta^{(t+1)}) \geq \ell(\theta^{(t)}),$$

and it is therefore possible to apply the convergence theory from Section 1.3.

1.5 Theoretical convergence in LMM implementation

We apply Theorem 1.3 to establish:

- **Monotone ascent:** $\ell(\theta^{(t+1)}) \geq \ell(\theta^{(t)})$ for all t .
- **Convergence of likelihood and stationarity:** $\ell(\theta^{(t)}) \xrightarrow[t \rightarrow \infty]{} \tilde{\ell}$ and every accumulation point $\tilde{\theta}$ of $\{\theta^{(t)}\}$ satisfies $\partial \ell(\tilde{\theta}) = 0$ (hence $\tilde{\ell} = \ell(\tilde{\theta})$).
- **Convergence of the parameters:** under any of the alternatives in **(D3)** (e.g., singleton top level set, discrete limit set with vanishing steps, or summable steps), the EM iterates converge: $\theta^{(t)} \xrightarrow[t \rightarrow \infty]{} \tilde{\theta}$.

Model and parameter space. Let $\mathbf{Y} = X\beta^* + Z\mathbf{U} + \varepsilon$ with $\mathbf{U} \sim \mathcal{N}(\mathbf{0}, D^*)$, $\varepsilon \sim \mathcal{N}(\mathbf{0}, \Sigma^*)$ be independent, and the parameter space be defined as

$$\Theta := \left\{ (\beta, D, \Sigma) : \beta \in \mathbb{R}^p, D = \text{blockdiag}(D_i)_{i=1}^m, D_i \in \mathbb{S}_{++}^q, \right. \\ \left. \Sigma = \text{blockdiag}(\Sigma_i)_{i=1}^m, \Sigma_i \in \mathbb{S}_{++}^{n_i}, \lambda_{\min} I \preceq D_i, \Sigma_i \preceq \lambda_{\max} I \right\},$$

for fixed constants $0 < \lambda_{\min} < \lambda_{\max} < \infty$ and all $i \in \mathbb{N}$. We assume throughout that the true parameter $\theta^* = (\beta^*, D^*, \Sigma^*)$ lies in Θ . The joint and observed densities evaluated at the realized data (y, u) are

$$f(y, u \mid \theta) \propto |D|^{-1/2} |\Sigma|^{-1/2} \exp \left\{ -\frac{1}{2} [(y - X\beta - Zu)^\top \Sigma^{-1} (y - X\beta - Zu) + u^\top D^{-1} u] \right\}, \\ f(y \mid \theta) = (2\pi)^{-n/2} |V|^{-1/2} \exp \left\{ -\frac{1}{2} (y - X\beta)^\top V^{-1} (y - X\beta) \right\}, \quad V = ZDZ^\top + \Sigma.$$

Verification of (M1)–(M3).

(M1) Strict positivity holds on Θ : $f(y, u \mid \theta) > 0$ for all (y, u) if $D \succ 0$ and $\Sigma \succ 0$.

(M2) Finiteness of $Q(\theta \mid \theta^{(t)})$. The complete-data log-likelihood is

$$\ell_c(\theta \mid y, u) = -\frac{1}{2} [\log |\Sigma| + \log |D| + (y - X\beta - Zu)^\top \Sigma^{-1} (y - X\beta - Zu) + u^\top D^{-1} u].$$

From equation (1), we have

$$Q(\theta \mid \theta^{(t)}) = \mathbb{E}_{\mathbf{U} \mid \mathbf{Y}=y, \theta^{(t)}} [\ell_c(\theta \mid y, \mathbf{U})] \\ = -\frac{1}{2} \left[\log |\Sigma| + \log |D| + (y - X\beta - Zm^{(t)})^\top \Sigma^{-1} (y - X\beta - Zm^{(t)}) + \text{tr}(Z^\top \Sigma^{-1} Z C^{(t)}) \right. \\ \left. + \text{tr}(D^{-1} C^{(t)}) + m^{(t)\top} D^{-1} m^{(t)} \right].$$

All terms are finite for $\theta, \theta^{(t)} \in \Theta = \{\beta \in \mathbb{R}^p, \Sigma \succ 0, D \succ 0\}$: $\log |\Sigma|$ and $\log |D|$ are finite on the interior, and the remaining quantities are quadratic or trace forms in finite matrices. Since $V \succ 0$ implies $m^{(t)}$ and $C^{(t)}$ are well defined with finite moments, $Q(\theta \mid \theta^{(t)})$ is finite for all $\theta, \theta^{(t)} \in \Theta$.

(M3) The observed density $f(y \mid \theta)$ is a proper Gaussian, so $0 < f(y \mid \theta) < \infty$ for all $\theta \in \Theta$.

Verification of (A1)–(A3). The linear mixed model is

$$\mathbf{Y} = X\beta^* + Z\mathbf{U} + \varepsilon, \quad \mathbf{U} \sim \mathcal{N}(\mathbf{0}, D^*), \quad \varepsilon \sim \mathcal{N}(\mathbf{0}, \Sigma^*),$$

with $\mathbf{U}$ and ε independent.

(A1) Parameter space. Identifying D with its $q(q+1)/2$ distinct entries and Σ with its $n(n+1)/2$ distinct entries, Θ is a subset of a d -dimensional Euclidean space, hence (A1) holds.

To verify (A2)–(A3), we derive the closed form of the observed log-likelihood $\ell(\theta)$ (*note that in many models $\ell(\theta)$ has no tractable closed form!*).

Conditional on $\theta = (\beta, \Sigma, D)$, the marginal distribution of $\mathbf{Y}$ is Gaussian:

$$\mathbf{Y} \mid \theta \sim \mathcal{N}(X\beta, V), \quad V = ZDZ^\top + \Sigma.$$

The marginal density of the observed sample y is obtained by integrating out the random effects:

$$f(y \mid \theta) = \int_{\mathbb{R}^{mq}} f(y, u \mid \theta) du = \int_{\mathbb{R}^{mq}} f(y \mid u, \theta) f(u \mid \theta) du.$$

Both terms are Gaussian:

$$f(y \mid u, \theta) = (2\pi)^{-n/2} |\Sigma|^{-1/2} \exp\left\{-\frac{1}{2} (y - X\beta - Zu)^\top \Sigma^{-1} (y - X\beta - Zu)\right\},$$

$$f(u \mid \theta) = (2\pi)^{-qm/2} |D|^{-1/2} \exp\left(-\frac{1}{2} u^\top D^{-1} u\right).$$

Multiplying gives

$$f(y, u \mid \theta) = (2\pi)^{-(n+qm)/2} |\Sigma|^{-1/2} |D|^{-1/2} \exp\left[-\frac{1}{2} \left((y - X\beta - Zu)^\top \Sigma^{-1} (y - X\beta - Zu) + u^\top D^{-1} u\right)\right].$$

Completing the square in u . Let $r = y - X\beta$. Expanding the first quadratic form yields

$$(r - Zu)^\top \Sigma^{-1} (r - Zu) = r^\top \Sigma^{-1} r - 2r^\top \Sigma^{-1} Zu + u^\top Z^\top \Sigma^{-1} Zu,$$

so that

$$(r - Zu)^\top \Sigma^{-1} (r - Zu) + u^\top D^{-1} u = u^\top \left(D^{-1} + Z^\top \Sigma^{-1} Z\right) u - 2r^\top \Sigma^{-1} Zu + r^\top \Sigma^{-1} r.$$

Define

$$A = D^{-1} + Z^\top \Sigma^{-1} Z, \quad b = Z^\top \Sigma^{-1} r.$$

Then (note that $A \succ 0$ and hence invertible)

$$u^\top Au - 2b^\top u = (u - m)^\top A(u - m) - b^\top A^{-1}b, \quad m = A^{-1}b.$$

Hence,

$$(r - Zu)^\top \Sigma^{-1} (r - Zu) + u^\top D^{-1} u = (u - m)^\top A(u - m) + r^\top \Sigma^{-1} r - b^\top A^{-1}b.$$

Substitute this back into the joint density:

$$f(y, u \mid \theta) = (2\pi)^{-(n+mq)/2} |\Sigma|^{-1/2} |D|^{-1/2} \exp\left[-\frac{1}{2} (u - m)^\top A(u - m)\right] \exp\left[-\frac{1}{2} \left(r^\top \Sigma^{-1} r - b^\top A^{-1}b\right)\right].$$

Integration over u . The first exponential is the kernel of a Gaussian in u with mean m and precision matrix A . Therefore,

$$\int_{\mathbb{R}^{mq}} \exp\left[-\frac{1}{2}(u-m)^\top A(u-m)\right] du = (2\pi)^{mq/2} |A|^{-1/2}.$$

Integrating over u gives

$$\begin{aligned} f(y \mid \theta) &= (2\pi)^{-(n+mq)/2} |\Sigma|^{-1/2} |D|^{-1/2} (2\pi)^{mq/2} |A|^{-1/2} \exp\left[-\frac{1}{2}\left(r^\top \Sigma^{-1} r - b^\top A^{-1} b\right)\right] \\ &= (2\pi)^{-n/2} |\Sigma|^{-1/2} |D|^{-1/2} |A|^{-1/2} \exp\left[-\frac{1}{2}\left(r^\top \Sigma^{-1} r - b^\top A^{-1} b\right)\right]. \end{aligned}$$

Expressing in LMM form. Using the matrix identities

$$A^{-1} = D - DZ^\top (ZDZ^\top + \Sigma)^{-1} ZD, \quad m = A^{-1}b = DZ^\top (ZDZ^\top + \Sigma)^{-1} r,$$

define

$$V = ZDZ^\top + \Sigma.$$

At the same time,

$$|A| = |D|^{-1} |V| |\Sigma|^{-1}, \quad b^\top A^{-1} b = r^\top \Sigma^{-1} r - r^\top V^{-1} r.$$

Substituting these expressions,

$$f(y \mid \theta) = (2\pi)^{-n/2} |V|^{-1/2} \exp\left(-\frac{1}{2} r^\top V^{-1} r\right), \quad V = ZDZ^\top + \Sigma.$$

Log-likelihood. Taking logarithms gives the observed-data log-likelihood:

$$\ell(\theta) = \log f(y \mid \theta) = -\frac{1}{2} \{\log |V| + r^\top V^{-1} r\} + \text{const}, \quad r = y - X\beta.$$

Now we check assumptions (A2)–(A3) using the closed form of $\ell(\theta)$.

(A2) Compact upper-level sets. Since X has full column rank, $\sigma_{\min}(X) > 0$. For θ_0 with $\ell(\theta_0) > -\infty$ define $\Theta_{\theta_0} = \{\theta \in \Theta : \ell(\theta) \geq \ell(\theta_0)\}$. Continuity of $(\beta, \Sigma, D) \mapsto (V, r)$ and of $V \mapsto \log |V|$, $V \mapsto r^\top V^{-1} r$ on $\mathbb{S}_{++}^n$ implies that ℓ is continuous on Θ , so Θ_{θ_0} is closed. If $\|\beta\| \rightarrow \infty$ while (Σ, D) remain fixed, then $aI \preceq V \preceq bI$ for some $0 < a \leq b < \infty$ and

$$r^\top V^{-1} r \geq \frac{1}{b} \|r\|^2 \geq \frac{1}{b} (\sigma_{\min}(X) \|\beta\| - \|y\|)^2 \rightarrow \infty,$$

hence $\ell(\theta) \rightarrow -\infty$; thus β is bounded on Θ_{θ_0} . Approaching the variance boundary,

$$\lambda_{\max}(\Sigma) \uparrow \infty \text{ or } \lambda_{\max}(D) \uparrow \infty \Rightarrow \log |V| \uparrow \infty.$$

Hence $\ell(\theta) \rightarrow -\infty$ and the corresponding parameters are bounded on Θ_{θ_0} .

Unfortunately, there is no equivalent implication for $\lambda_{\min}(\Sigma) \downarrow 0$ or $\lambda_{\min}(D) \downarrow 0$, since either degeneracy can be compensated by the other matrix remaining uniformly positive definite.

Indeed, take a unit vector $v \in \mathbb{R}^n$ such that $v \perp r$ and $v \perp \text{col}(Z)$ (exists if $q+1 < n$). Let $\Sigma_\varepsilon = \varepsilon vv^\top + I_n - vv^\top$ and fix $D = D_0 \succ 0$. Then $V_\varepsilon = ZD_0Z^\top + \Sigma_\varepsilon$ satisfies

$$\lambda_{\min}(V_\varepsilon) = \varepsilon \downarrow 0, \quad \log |V_\varepsilon| = \log |ZD_0Z^\top + \Sigma_\varepsilon| \rightarrow -\infty,$$

while the quadratic term remains constant, $r^\top V_\varepsilon^{-1} r = r^\top r$, because $r \perp v$. Consequently,

$$\ell(\theta_\varepsilon) = -\frac{1}{2} \{ \log |V_\varepsilon| + r^\top V_\varepsilon^{-1} r \} \xrightarrow{\varepsilon \downarrow 0} +\infty.$$

Thus, the upper-level set Θ_{θ_0} is unbounded in this direction, and compactness fails unless we restrict the parameter space away from degeneracy. This pathology is similar to Example 1.3.

To prevent such compensating singularities, we have assumed that there exist a uniform constant $\lambda_{\min} > 0$ such that

$$D \succeq \lambda_{\min} I_m q, \quad \Sigma \succeq \lambda_{\min} I_n. \quad (2)$$

In practice, $\lambda_{\min}$ may be chosen in accordance with the numerical precision of the computational environment.

This assumption constrains the parameter space Θ to matrices with uniformly bounded condition numbers, ensuring that $V = Z D Z^\top + \Sigma$ remains well-conditioned.

With the additional assumption (2), the eigenvalues of Σ and D are confined to closed intervals on Θ_{θ_0} . Therefore Θ_{θ_0} is closed and bounded, hence compact (Heine–Borel).

(A3) Regularity of ℓ . The mappings $(\beta, \Sigma, D) \mapsto V$ and $(\beta, \Sigma, D) \mapsto r$ are smooth on Θ , and on $\mathbb{S}_{++}^n$ the functions $V \mapsto \log |V|$ and $V \mapsto r^\top V^{-1} r$ are smooth as well. Hence, ℓ is continuous on Θ and continuously differentiable (C^1) in the interior of Θ . In particular,

$$\frac{\partial \ell}{\partial \beta} = -X^\top V^{-1} r, \quad \frac{\partial \ell}{\partial V} = -\frac{1}{2} (V^{-1} - V^{-1} r r^\top V^{-1}),$$

and since $dV = d\Sigma$, we can equivalently write

$$\boxed{\frac{\partial \ell}{\partial \Sigma} = -\frac{1}{2} (V^{-1} - V^{-1} r r^\top V^{-1}).}$$

Similarly,

$$\frac{\partial \ell}{\partial D} = -\frac{1}{2} (Z^\top V^{-1} Z - Z^\top V^{-1} r r^\top V^{-1} Z)$$

are continuous on the interior of Θ . Therefore, condition (A3) holds.

Verification of (D1)–(D2).

- (D1) The M–step consists of exact conditional maximizations of $Q(\cdot \mid \theta^{(t)})$ over the parameter blocks (β, D, Σ) , yielding blockwise first–order conditions at $\theta^{(t+1)}$, which satisfies the weakened, blockwise version of condition (D1) of the Theorem 1.3.
- (D2) Since $(m^{(t)}, C^{(t)})$ depend smoothly on $\theta^{(t)}$, and $\ell_c(\theta \mid y, u)$ is C^1 in θ with partial derivatives at most quadratic in u , for any compact $K \subset \Theta$ there exists a constant $K_1 < \infty$ such that

$$\left\| \frac{\partial \ell_c(\theta \mid y, u)}{\partial \theta} \right\| \leq K_1 (1 + \|u\|^2) \quad (\theta \in K).$$

Taking expectations under the Gaussian distribution $u \mid (y, \theta^{(t)})$ yields a finite bound uniformly for $\theta^{(t)} \in K$. By the dominated convergence theorem, differentiation and conditional expectation interchange:

$$\frac{\partial Q(\theta \mid \theta^{(t)})}{\partial \theta} = \int \frac{\partial \ell_c(\theta \mid y, u)}{\partial \theta} f(u \mid y, \theta^{(t)}) du,$$

and the right–hand side is continuous in $(\theta, \theta^{(t)})$ by the smooth dependence of both the integrand and the Gaussian kernel.

Conclusion (monotonicity and convergence). On the constrained space

$$\Theta_{c1,c2} = \{(\beta, \Sigma, D) \in \Theta : D \succeq c_1 I_q, \Sigma \succeq c_2 I_n\}, \quad c_1, c_2 > 0$$

the assumptions of Theorem 1.3 up to **(D3)** hold, and therefore the likelihood values $\{\ell(\theta^{(t)})\}_{t \in \mathbb{N}}$ increase monotonically and converge to some finite $\tilde{\ell}$. To obtain convergence of the parameter sequence, it is enough that one of the alternatives in **(D3)** holds. In practice, **(D3)**(i)–(ii) are typically hard to verify a priori, as they depend on the geometry of the likelihood on the top level set at $\tilde{\ell}$. A practical sufficient assumption is **(D3)**(iii), namely summable steps:

$$\sum_{t=1}^{\infty} \|\theta^{(t+1)} - \theta^{(t)}\| < \infty,$$

which, for example, is implied if the steps eventually decay faster than harmonically, $\|\theta^{(t+1)} - \theta^{(t)}\| = O(t^{-1-\varepsilon})$ for some $\varepsilon > 0$.

1.6 Summary

Section 1 introduces the statistical modeling framework and algorithmic foundation that motivate the developments in the remainder of the thesis. The scientific problem is formulated as explaining the relationship between protein channel configurations and ion permeation, where repeated measurements give rise to grouped observations with latent sources of variability. To capture this structure, a linear mixed model is adopted, allowing for both fixed effects describing the systematic signal and random effects accounting for group-specific deviations.

Since the marginal covariance structure of the observed response vector $\mathbf{Y}$ is unknown and depends on latent random effects, direct maximization of the likelihood is not feasible. An expectation–maximization (EM) approach is therefore proposed, in which the latent random effects are treated as missing data and inference proceeds via conditional expectations. Section 1 formally derives the EM algorithm for this setting and places it within the broader class of generalized EM (GEM) algorithms.

The section then establishes basic monotonicity and convergence results for a generic GEM algorithm. It is shown that, under standard regularity conditions, the observed-data log-likelihood is nondecreasing along the iterates and that accumulation points of the parameter sequence are stationary points of the likelihood. Building on these general results, linear mixed models are presented as a concrete illustration of EM methodology. In particular, the proposed implementation is shown to take the form of an Expectation–Conditional Maximization (ECM) algorithm, where the M–step is decomposed into a sequence of exact conditional maximizations. It is verified that this ECM formulation satisfies the assumptions of Theorem 1.3, thereby ensuring convergence of the algorithm under the stated conditions.

While Section 1 provides a rigorous convergence analysis for the EM/ECM algorithm in the linear mixed model setting, the theoretical results do not account for the potential ill-posedness of the fixed-effects estimation problem in highly correlated designs. In such regimes, unconstrained EM updates may lead to unstable or unreliable estimates of the fixed-effects parameter β . This limitation motivates the developments in Section 2, where regularization is introduced through the partial least squares (PLS) framework in order to stabilize the EM updates while preserving their favorable convergence properties.

2 Regularisation of the EM algorithm with Partial least squares

2.1 Ill-posed Linear Mixed Models

In this section, we work with a population/sample formulation of the linear mixed model; its relation to the stacked notation of Section 1 is made explicit when needed.

Let $\mathbf{X} \in \mathbb{R}^p$ denote a random fixed-effects covariate vector, let $Z \in \mathbb{R}^q$ denote the corresponding random-effects design vector, and let $\mathbf{U} \in \mathbb{R}^q$ and $\varepsilon \in \mathbb{R}$ denote the random effect and the observational noise, respectively. The population mixed model is

$$\mathbf{Y} = \mathbf{X}^\top \beta^* + Z^\top \mathbf{U} + \varepsilon,$$

where

$$\mathbf{U} \sim \mathcal{N}(0, D^*), \quad \varepsilon \sim \mathcal{N}(0, \sigma^{2*}), \quad \mathbf{U} \perp \varepsilon.$$

Here $\beta^* \in \mathbb{R}^p$, $D^* \in \mathbb{S}_{++}^q$, and $\sigma^{2*} > 0$.

The *conditional distribution* of $\mathbf{Y}$ given $\mathbf{X} = x$, $Z = z$, and $\mathbf{U} = u$ is

$$\mathbf{Y} \mid (\mathbf{X} = x, Z = z, \mathbf{U} = u) \sim \mathcal{N}(x^\top \beta^* + z^\top u, \sigma^{2*}).$$

Integrating out the random effect $\mathbf{U}$, the *marginal distribution* of $\mathbf{Y}$ given $\mathbf{X} = x$ and $Z = z$ is

$$\mathbf{Y} \mid (\mathbf{X} = x, Z = z) \sim \mathcal{N}(x^\top \beta^*, z^\top D^* z + \sigma^{2*}).$$

The sample mixed model is

$$y_i = X_i \beta^* + Z_i u_i + \varepsilon_i, \quad i = 1, \dots, m,$$

where

$$u_i \sim \mathcal{N}(0, D^*), \quad \varepsilon_i \sim \mathcal{N}(0, \sigma^{2*} I_{n_i}), \quad u_i \perp \varepsilon_i,$$

and the collections $\{(u_i, \varepsilon_i)\}_{i=1}^m$ are independent across groups.

For each group $i \in \{1, \dots, m\}$, let n_i denote the within-group sample size and define

$$y_i = (y_{i1}, \dots, y_{in_i})^\top \in \mathbb{R}^{n_i}, \quad X_i \in \mathbb{R}^{n_i \times p}, \quad Z_i \in \mathbb{R}^{n_i \times q}, \quad u_i \in \mathbb{R}^q.$$

We assume that the rows $x_{ij}^\top$ of X_i are i.i.d. copies of $\mathbf{X}$. The vectors $z_{ij}^\top$ forming the rows of Z_i are treated as fixed and are determined by the within-group design.

Equivalently, on the row level,

$$y_{ij} = x_{ij}^\top \beta^* + z_{ij}^\top u_i + \varepsilon_{ij}, \quad j = 1, \dots, n_i.$$

The *conditional distribution* of y_i given X_i , Z_i , and u_i is

$$y_i \mid (X_i, Z_i, u_i) \sim \mathcal{N}(X_i \beta^* + Z_i u_i, \sigma^{2*} I_{n_i}).$$

After integrating out u_i , the *marginal distribution* of y_i given X_i and Z_i is

$$y_i \mid (X_i, Z_i) \sim \mathcal{N}(X_i \beta^*, V_i^*), \quad V_i^* := Z_i D^* Z_i^\top + \sigma^{2*} I_{n_i}.$$

Stacking all groups gives

$$y = \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix} \in \mathbb{R}^n, \quad X = \begin{bmatrix} X_1 \\ \vdots \\ X_m \end{bmatrix} \in \mathbb{R}^{n \times p}, \quad Z = \text{blockdiag}(Z_1, \dots, Z_m) \in \mathbb{R}^{n \times mq},$$

and

$$u = \begin{bmatrix} u_1 \\ \vdots \\ u_m \end{bmatrix} \in \mathbb{R}^{mq}, \quad D^* = \text{blockdiag}(D_1^*, \dots, D_m^*), \quad \varepsilon = \begin{bmatrix} \varepsilon_1 \\ \vdots \\ \varepsilon_m \end{bmatrix} \in \mathbb{R}^n, \quad \Sigma^* = \text{blockdiag}(\Sigma_1^*, \dots, \Sigma_m^*).$$

The stacked model reads

$$y = X\beta^* + Zu + \varepsilon, \quad u \sim \mathcal{N}(0, D^*), \quad \varepsilon \sim \mathcal{N}(0, \Sigma^*), \quad \text{Cov}(y) = V^* = ZD^*Z^\top + \Sigma^*.$$

Let

$$\hat{\Sigma}_X := \frac{1}{n} X^\top X.$$

In the regimes of high-dimensional analysis even with fixed p , $\hat{\Sigma}_X$ may be ill-conditioned or nearly singular due to strong collinearities. In such situations, unregularized least-squares or generalized least-squares updates for β are unstable.

To formalize the regularization framework, we extend our model with the latent factor framework. Firstly, we assume that $(\mathbf{X}, \mathbf{Y})$ is centered,

$$\mathbb{E}(\mathbf{X}) = 0, \quad \mathbb{E}(\mathbf{Y} - Z^\top \mathbf{U}) = 0,$$

and has finite second moments

$$\Sigma_{\mathbf{X}} := \mathbb{E}(\mathbf{X}\mathbf{X}^\top) \in \mathbb{R}_{\geq 0}^{p \times p}, \quad \sigma_{\mathbf{X}, \mathbf{Y}} := \mathbb{E}(\mathbf{X}(\mathbf{Y} - Z^\top \mathbf{U})) \in \mathbb{R}^p$$

with $\sigma_{\mathbf{X}, \mathbf{Y}} \neq 0$.

Following the population viewpoint of Finocchio and Krivobokova (2025), we distinguish between directions of the design that are relevant and irrelevant for predicting the response. We postulate that there exists the smallest linear subspace $\mathcal{B} \subset \mathbb{R}^p$ (the *smallest relevant subspace*) such that the component of $\mathbf{X}$ orthogonal to $\mathcal{B}$ is uncorrelated, in population, with both the response and the relevant component of the design.

Let $\Pi_{\mathcal{B}}$ denote the orthogonal projector onto $\mathcal{B}$ and $\Pi_{\mathcal{B}^\perp} = I_p - \Pi_{\mathcal{B}}$. Define the decomposition

$$\mathbf{X}_{\mathcal{B}} := \Pi_{\mathcal{B}} \mathbf{X}, \quad \mathbf{X}_{\mathcal{B}^\perp} := \Pi_{\mathcal{B}^\perp} \mathbf{X}, \quad \mathbf{X} = \mathbf{X}_{\mathcal{B}} + \mathbf{X}_{\mathcal{B}^\perp}.$$

The relevance assumption is encoded by the population orthogonality conditions

$$\mathbb{E}(\mathbf{X}_{\mathcal{B}^\perp} \mathbf{Y}) = 0, \quad \mathbb{E}(\mathbf{X}_{\mathcal{B}^\perp} \mathbf{X}_{\mathcal{B}}^\top) = 0, \quad (3)$$

while $\mathbb{E}(\mathbf{X}_{\mathcal{B}} \mathbf{Y}) \neq 0$. Condition (3) formalizes that directions in $\mathcal{B}^\perp$ are irrelevant for the population signal.

Also note that the assumption (3) can be equivalently written by replacing $\mathbf{Y}$ with $\mathbf{Y} - Z^\top \mathbf{U}$. Indeed,

$$\mathbb{E}(\mathbf{X}_{\mathcal{B}^\perp} (\mathbf{Y} - Z^\top \mathbf{U})) = -\mathbb{E}(\mathbf{X}_{\mathcal{B}^\perp} Z^\top \mathbf{U}) = 0$$

under the independence of $\mathbf{X}$ and $\mathbf{U}$ and $\mathbb{E}(\mathbf{U}) = 0$.

In particular, it implies that the population signal is supported on $\mathcal{B}$, and hence $\beta^* \in \mathcal{B}$.

Under the extended latent factor model of Finocchio and Krivobokova (2025), we assume

$$\Pi_{\mathcal{B}}\mathbf{X} = \mathbf{Q} + \sigma_{\text{lat}}e, \quad \mathbf{Y} = \mathbf{Q}^\top \alpha^* + \varepsilon,$$

where $\mathbf{Q} \in \mathbb{R}^d$ is a centered latent vector satisfying $\mathcal{R}(\Sigma_{\mathbf{Q}}) \subseteq \mathcal{B}$, $\alpha^* \in \mathbb{R}^d$ is the latent coefficient vector governing the linear relationship between the response and the latent factors, and ε is centered and independent of $\mathbf{Q}$.

In the regime of *zero projection noise* ($\sigma_{\text{lat}} = 0$),

$$\Pi_{\mathcal{B}}\mathbf{X} = \mathbf{Q}, \quad \text{rank}(\Sigma_{\mathbf{Q}}) =: d.$$

In this case, the relevant component of the design is exactly low-rank at the population level, which implies that the covariance matrix $\Sigma_{\mathbf{X}}$ is singular.

This may violate the full column-rank assumption on the design matrix imposed in Section 1.4. Nevertheless, the EM updates derived there remain well-defined if the corresponding matrix inverses are replaced by Moore–Penrose pseudoinverses.

However, the zero projection noise assumption is restrictive and rarely holds in practice. It is adopted here as an idealized regime that isolates the effect of irrelevant directions and allows for a transparent characterization of the population behavior of PLS. In particular, it enables exact recovery results that serve as a theoretical benchmark for the regularized EM–PLS procedure described later.

2.2 Existing results on EM regularization

Since our goal is to make use of the latent factor LMM framework by regularizing the M-step update for parameter β , we start this subsection with a literature overview on the EM regularization.

Bondell, Krishna and Ghosh (2010): Joint Variable Selection for Fixed and Random Effects in Linear Mixed-Effects Models. Bondell et al. (2010) propose a simultaneous variable-selection procedure for linear mixed-effects models that targets both the fixed effects and the random-effects structure through a penalized likelihood formulation. Their model the random-effects covariance matrix, $\Psi = D\Gamma^\top\Gamma D$, with the modified Cholesky decomposition, where $D = \text{diag}(d_1, \dots, d_q)$ and Γ is lower triangular with ones on the diagonal.

Selection is induced by the adaptive ℓ_1 -type penalty applied jointly to the fixed-effect coefficients β and to the random-effect scale parameters d in the penalized observed-data log-likelihood. Estimation is carried out by a *constrained EM algorithm* that treats the random effects as latent variables and updates $(\beta, d, \Gamma, \sigma^2)$ by maximizing the penalized complete-data criterion; the penalty allows some components of β and some d_ℓ to be estimated as exactly zero, thereby performing joint model selection.

Theoretical results are derived in a fixed-dimensional regime (fixed p and q) and establish an *oracle* behavior: asymptotically, the procedure identifies the correct zero pattern and the nonzero components attain the same first-order distribution as the corresponding estimator under the true submodel.

However, in contrast to the EM-PLS modification considered here, Bondell et al. modify EM through an explicit penalized likelihood objective with sparsity penalties, so their method fits the “penalized M-step” paradigm.

Yi and Caramanis (2015): Regularized EM Algorithms: A Unified Framework and Statistical Guarantees. Yi and Caramanis (2015) develop a unified high-dimensional analysis of EM where the classical M-step $\beta^{(t)} \in \arg \max_{\beta} Q(\beta \mid \beta^{(t-1)})$ is replaced by a *regularized* update

$$\beta^{(t)} \in \arg \max_{\beta} Q_m^{(t)}(\beta \mid \beta^{(t-1)}) - \lambda_m^{(t)} R(\beta),$$

with a convex regularizer R that is assumed to be a *decomposable norm* with respect to a pair of subspaces $(S, \bar{S})$ and $\beta^* \in S$. Here $Q_m^{(t)}$ denotes the empirical Q -function computed from a block of m i.i.d. observations (see below), and $\lambda_m^{(t)} \geq 0$ is the regularization level at iteration t , typically chosen as a function of the sample size m and the target confidence level. The analysis relies on (i) population-level EM regularity (self-consistency, strong concavity/smoothness, gradient stability) and (ii) sample-level analogues (restricted strong concavity on a cone induced by decomposability, and a statistical error bound in the dual norm).

A key technical device is a *resampling* variant of EM: the sample is split into T disjoint blocks of size $m = n/T$, and iteration t forms $Q_m^{(t)}$ using a fresh block. Under their conditions, Theorem 3.3 yields a finite-iteration “contraction-to-a-ball” guarantee: with probability at least $1 - \delta$, for all $t = 1, \dots, T$,

$$\|\beta^{(t)} - \beta^*\| \leq \kappa^t \|\beta^{(0)} - \beta^*\| + \frac{5}{\gamma_m} \frac{1 - \kappa^t}{1 - \kappa} \Psi(S) \Delta.$$

The quantities in the bound have the following interpretation. The parameter $\gamma_m > 0$ is the *restricted strong concavity* (RSC) constant of the empirical objective $Q_m^{(t)}(\cdot \mid \beta)$ on an appropriate cone determined by decomposability; it plays the role of a curvature modulus and scales like an effective (sample-dependent) strong concavity parameter. The factor $\kappa \in (0, 1)$ is the resulting *contraction coefficient* of the population recursion, obtained by combining the population gradient-stability parameter with the curvature γ_m ; it controls the geometric decay rate κ^t of the optimization error over iterations. The quantity $\Psi(S)$ is the *subspace compatibility constant* of the regularizer with respect to the norm $\|\cdot\|$ used in the error bound, defined by

$$\Psi(S) := \sup_{u \in S \setminus \{0\}} \frac{R(u)}{\|u\|}.$$

It measures how strongly the regularizer scales relative to the norm on vectors restricted to the model subspace S (e.g. it equals $\sqrt{s}$ for ℓ_1 regularization on an s -sparse subspace under the Euclidean norm). The term Δ is a *statistical error level* controlling the discrepancy between empirical and population gradients in a dual norm; it determines the radius of the neighborhood to which the iterates contract. Finally, $r > 0$ is the radius of the local region of attraction $B(r; \beta^*)$ in which the population regularity conditions are assumed to hold; in particular, the initialization $\beta^{(0)}$ is required to lie in this neighborhood.

Hence the iterates converge geometrically up to a neighborhood of radius

$$\limsup_{t \rightarrow \infty} \|\beta^{(t)} - \beta^*\| \leq \frac{5}{\gamma_m(1 - \kappa)} \Psi(S) \Delta,$$

and using their specific choice $\Delta = r\gamma_m/(60\Psi(S))$ this upper bound simplifies to $r/(12(1 - \kappa))$.

Nevertheless, their guarantees require a specific sequence $\lambda_m^{(t)}$ that depends on quantities such as γ_m , $\Psi(S)$, the contraction factor κ , and the (unknown) initial error $\|\beta^{(0)} - \beta^*\|$, as well as a bound Δ on the statistical error. These constants are generally not available in practice, so the tuning rule is primarily a theoretical device enabling the contraction-to-a-ball statement above.

Bazzoli, Lambert-Lacroix and Martinez (2023): Partial least square based approaches for high-dimensional linear mixed models. Bazzoli et al. (2023) propose three PLS-based modifications of the EM algorithm for high-dimensional linear mixed models: *PLS/EM*, where PLS is applied prior to EM; *EM-PLS*, where the M-step least-squares regression is replaced by a PLS regression; and *EMRidge/PLS*, which combines ridge-regularized EM with a final PLS step. All three approaches are motivated algorithmically as dimension-reduction techniques for fixed effects in the $p \gg n$ regime and are evaluated through extensive simulation studies and a real-data application (yeast cell-cycle gene expression).

The paper exclusively focuses on empirical performance. No population-level formulation of the algorithms is provided, and no theoretical guarantees are established for convergence, statistical consistency, or contraction behavior of the iterates. In particular, the proposed procedures are not derived as optimizers of a regularized Q -function with a well-defined penalty structure, and the analysis does not address convergence properties.

2.3 Sample EM-PLS algorithm

The following algorithm is the finite-sample implementation of the EM algorithm, where the M-step update for β is replaced by a PLS-based regularized update.

The regularization addresses the potential ill-posedness of the normal equations. In the unregularized case, the M-step for β at iteration $t \in \mathbb{N}$ corresponds to solving the least-squares problem

$$\min_{\beta \in \mathbb{R}^p} \|y - X\beta - Zm^{(t)}\|_2^2,$$

whose normal equations are

$$\hat{\Sigma}_X \beta = \hat{\sigma}_{X,Y}^{(t)},$$

where

$$\hat{\Sigma}_X := \frac{1}{n} X^\top X \in \mathbb{R}^{p \times p}, \quad \hat{\sigma}_{X,Y}^{(t)} := \frac{1}{n} X^\top (y - Zm^{(t)}) \in \mathbb{R}^p,$$

$$m^{(t)} = \mathbb{E}_{\theta^{(t)}}(u \mid y),$$

and $\theta^{(t)}$ denotes the current parameter estimate at iteration t . When $\hat{\Sigma}_X$ is ill-conditioned or singular, this problem becomes unstable or non-identifiable.

To stabilize the estimation, we restrict the optimization to a low-dimensional subspace. Specifically, instead of minimizing over $\mathbb{R}^p$, we solve the least-squares problem over a Krylov subspace of dimension K .

For a given pair (A, b) , with

$$A \in \mathbb{R}^{p \times p}, \quad b \in \mathbb{R}^p,$$

the Krylov subspace of order K is defined as

$$\mathcal{K}_K(A, b) := \text{span}\{b, Ab, \dots, A^{K-1}b\}.$$

In the EM-PLS setting, the matrices A and b correspond to the (normalized) normal-equation objects at iteration t , namely

$$A = \hat{\Sigma}_X, \quad b = \hat{\sigma}_{X,Y}^{(t)}.$$

The PLS estimator of β at iteration t is therefore defined as

$$\beta_{\text{PLS}}^{(t+1)} \in \arg \min_{\beta \in \mathcal{K}_K(\hat{\Sigma}_X, \hat{\sigma}_{X,Y}^{(t)})} \|y - X\beta - Zm^{(t)}\|_2^2.$$

Algorithm 1 Sample EM-PLS algorithm for the linear mixed model

Require: Observations $\{(y_i, X_i, Z_i)\}_{i=1}^m$, initial parameters $\theta^{(0)} = (\beta^{(0)}, D^{(0)}, \sigma^{2(0)})$, number of PLS components K

Ensure: Sequence of parameter estimates $\{\theta^{(t)}\}_{t \geq 0}$

1: **for** $t = 0, 1, 2, \dots$ until convergence **do**

2: **E-step:**

3: **for** $i = 1, \dots, m$ **do**

4: Compute

$$m_i^{(t)} = \mathbb{E}_{\theta^{(t)}}(u_i \mid y_i), \quad C_i^{(t)} = \text{Var}_{\theta^{(t)}}(u_i \mid y_i)$$

5: **end for**

6: Stack the conditional means as

$$m^{(t)} := ((m_1^{(t)})^\top, \dots, (m_m^{(t)})^\top)^\top$$

7: **CM-step 1 (fixed effects):**

8: Form the EM normal-equation objects

$$A^{(t)} = \frac{1}{n} X^\top X, \quad b^{(t)}(\theta^{(t)}) = \frac{1}{n} X^\top (y - Zm^{(t)})$$

9: Construct the Krylov subspace

$$\mathcal{S}_t = \mathcal{K}_K(A^{(t)}, b^{(t)}(\theta^{(t)}))$$

10: Update

$$\beta^{(t+1)} \in \arg \max_{\beta \in \mathcal{S}_t} Q((\beta, D^{(t)}, \sigma^{2(t)}) \mid \theta^{(t)})$$

▷ equivalently, solve $A^{(t)}\beta = b^{(t)}$ restricted to $\mathcal{S}_t$

11: **CM-step 2 (random-effect covariance):**

12: Update

$$D^{(t+1)} = \frac{1}{m} \sum_{i=1}^m (C_i^{(t)} + m_i^{(t)} m_i^{(t)\top})$$

13: **CM-step 3 (error variance):**

14: Update $\sigma^{2(t+1)}$ in closed form using $\beta^{(t+1)}$ and $\{m_i^{(t)}, C_i^{(t)}\}_{i=1}^m$

15: Set

$$\theta^{(t+1)} = (\beta^{(t+1)}, D^{(t+1)}, \sigma^{2(t+1)})$$

16: **end for**

This formulation emphasizes that the algorithm replaces the potentially ill-posed full M-step for β by a lower-dimensional conditional maximization, while keeping the remaining variance-component updates in closed form.

The immediate consequence of regularizing the M-step is a loss of monotonicity, that is, we cannot guarantee that

$$Q(\beta_{\text{PLS}}^{(t+1)}, D^{(t)}, \sigma^{2(t)} \mid \theta^{(t)}) \geq Q(\beta^{(t)}, D^{(t)}, \sigma^{2(t)} \mid \theta^{(t)}).$$

Therefore, the resulting procedure is no longer a standard ECM algorithm, and the convergence results from Section 1.3 cannot be applied directly. A new convergence framework must therefore be developed.

2.4 Population EM-PLS algorithm

We begin the theoretical analysis of Algorithm 1 by introducing its *population analogue*. This separation between population-level behavior and sample-level fluctuations follows a standard strategy in the analysis of EM-type procedures, motivated by the regularized EM framework of Yi and Caramanis (2015). There, the population recursion is used to isolate the intrinsic stability properties of the algorithm, while the sample-based iterates are subsequently studied as perturbations of this idealized dynamics.

Under this formulation, the population EM-PLS recursion is defined by replacing the sample quantities in Algorithm 1 by their population counterparts. The E-step is given by the population conditional moments

$$m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z) = \mathbb{E}_{\theta_{\text{pop}}^{(t)}}(\mathbf{U} \mid \mathbf{X}, \mathbf{Y}, Z), \quad C_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z) = \text{Var}_{\theta_{\text{pop}}^{(t)}}(\mathbf{U} \mid \mathbf{X}, \mathbf{Y}, Z).$$

The fixed-effects update is obtained by maximizing the population Q -function over a Krylov subspace constructed from the corresponding population normal-equation objects. Specifically, define

$$A_{\text{pop}}^{(t)} := \mathbb{E}(\mathbf{X}\mathbf{X}^\top), \quad b_{\text{pop}}^{(t)}(\theta^{(t)}) := \mathbb{E}(\mathbf{X}(\mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z))).$$

For a fixed number of components K , the population Krylov subspace is

$$\mathcal{K}_K^{(t)} := \mathcal{K}_K(A_{\text{pop}}^{(t)}, b_{\text{pop}}^{(t)}(\theta^{(t)})) = \text{span}\left\{b_{\text{pop}}^{(t)}(\theta^{(t)}), A_{\text{pop}}^{(t)}b_{\text{pop}}^{(t)}(\theta^{(t)}), \dots, (A_{\text{pop}}^{(t)})^{K-1}b_{\text{pop}}^{(t)}(\theta^{(t)})\right\}.$$

The population least-squares problem

$$\arg \min_{\beta \in \mathbb{R}^p} \mathbb{E}\left[(\mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z) - \mathbf{X}^\top \beta)^2\right] \quad (4)$$

admits the solution set

$$\left\{\beta : \Sigma_{\mathbf{X}}\beta = \sigma_{\mathbf{X}, \mathbf{Y}}^{(m, t)}\right\}, \quad \sigma_{\mathbf{X}, \mathbf{Y}}^{(m, t)} := \mathbb{E}(\mathbf{X}(\mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z))).$$

We denote by

$$\beta_{\text{LS}} := \Sigma_{\mathbf{X}}^\dagger \sigma_{\mathbf{X}, \mathbf{Y}}^{(m, t)}$$

the minimum- L^2 -norm solution.

For $K = d$, define the population Krylov subspace generated by $(\Sigma_{\mathbf{X}}, \sigma_{\mathbf{X}, \mathbf{Y}}^{(m, t)})$ as

$$\mathcal{K}_K(\Sigma_{\mathbf{X}}, \sigma_{\mathbf{X}, \mathbf{Y}}^{(m, t)}) := \text{span}\left\{\sigma_{\mathbf{X}, \mathbf{Y}}^{(m, t)}, \Sigma_{\mathbf{X}}\sigma_{\mathbf{X}, \mathbf{Y}}^{(m, t)}, \dots, \Sigma_{\mathbf{X}}^{K-1}\sigma_{\mathbf{X}, \mathbf{Y}}^{(m, t)}\right\}. \quad (5)$$

We now show that the structural assumptions formulated with $\mathbf{Y} - Z^\top \mathbf{U}$ remain valid if $\mathbf{U}$ is replaced by its conditional mean $m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z)$.

First, by the tower property,

$$\mathbb{E}(Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z)) = \mathbb{E}(Z^\top \mathbb{E}_{\theta_{\text{pop}}^{(t)}}(\mathbf{U} \mid \mathbf{X}, \mathbf{Y}, Z)) = \mathbb{E}(Z^\top \mathbf{U}).$$

Since $\mathbb{E}(Z^\top \mathbf{U}) = 0$, it follows that

$$\mathbb{E}(\mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z)) = \mathbb{E}(\mathbf{Y} - Z^\top \mathbf{U}) = 0.$$

Second, the relevance/irrelevance orthogonality condition transfers as well. Indeed,

$$\begin{aligned} \mathbb{E}(\mathbf{X}_{\mathcal{B}^\perp}(\mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z))) &= \mathbb{E}(\mathbf{X}_{\mathcal{B}^\perp} \mathbf{Y}) - \mathbb{E}(\mathbf{X}_{\mathcal{B}^\perp} Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z)) \\ &= \mathbb{E}(\mathbf{X}_{\mathcal{B}^\perp} \mathbf{Y}) - \mathbb{E}(\mathbf{X}_{\mathcal{B}^\perp} Z^\top \mathbf{U}), \end{aligned}$$

where the second equality again follows from the tower property. Therefore, whenever

$$\mathbb{E}(\mathbf{X}_{\mathcal{B}^\perp}(\mathbf{Y} - Z^\top \mathbf{U})) = 0,$$

it follows that

$$\mathbb{E}(\mathbf{X}_{\mathcal{B}^\perp}(\mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z))) = 0.$$

Finally, the required cross-moment is finite. Since the map $u \mapsto \|u\|^2$ is convex and

$$m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z) = \mathbb{E}_{\theta_{\text{pop}}^{(t)}}(\mathbf{U} \mid \mathbf{X}, \mathbf{Y}, Z),$$

Jensen's inequality gives

$$\|m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z)\|^2 \leq \mathbb{E}(\|\mathbf{U}\|^2 \mid \mathbf{X}, \mathbf{Y}, Z).$$

Taking expectations and applying the tower property yields

$$\mathbb{E}\|m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z)\|^2 \leq \mathbb{E}\|\mathbf{U}\|^2 < \infty.$$

Hence

$$\mathbb{E}((Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z))^2) < \infty$$

as Z is fixed. Assume $\mathbb{E}(\mathbf{Y}^2) < \infty$ and apply Cauchy–Schwarz:

$$\mathbb{E}\|\mathbf{X}(\mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z))\| \leq (\mathbb{E}\|\mathbf{X}\|^2)^{1/2} (\mathbb{E}(\mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z))^2)^{1/2} < \infty.$$

Consequently, both the orthogonality condition and the required moment assumptions remain valid when $\mathbf{Y} - Z^\top \mathbf{U}$ is replaced by $\mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z)$.

Therefore, a key adaptivity result (Finocchio and Krivobokova, 2025, Lemma 3.1) states that

$$\mathcal{K}_K(\mathbb{E}(\mathbf{X}\mathbf{X}^\top), \mathbb{E}(\mathbf{X}(\mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z)))) = \mathcal{K}_K(\mathbb{E}(\mathbf{X}_{\mathcal{B}}\mathbf{X}_{\mathcal{B}}^\top), \mathbb{E}(\mathbf{X}_{\mathcal{B}}(\mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\cdot))))). \quad (6)$$

Thus, although the PLS construction does not explicitly know the decomposition $\mathbf{X} = \mathbf{X}_{\mathcal{B}} + \mathbf{X}_{\mathcal{B}^\perp}$, it automatically ignores irrelevant directions at the population level.

Now let β_{PLS} denote the population PLS coefficient obtained by restricting (4) to $\mathcal{K}_K(\mathbf{X}, \mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z))$. For $K = d$, Lemma 3.3 of Finocchio and Krivobokova (2025) yields

$$\beta_{\text{PLS}, d} = \Pi_{\mathcal{B}} \alpha^* = \beta_{\text{LS}} = \Sigma_{\mathbf{X}}^\dagger \sigma_{\mathbf{X}, \mathbf{Y}}^{(m, t)}.$$

Since the transformed pair $(\mathbf{X}, \mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z))$ satisfies the same structural assumptions as the original pair, the above PLS/least-squares identity applies to the population EM fixed-effects criterion as well. Therefore, for $K = d$, the population EM–PLS fixed-effects update coincides with the unregularized population EM update from Section 1.4.

The remaining two updates, namely those for D and σ^2 , are identical in the population EM and population EM–PLS recursions by construction. Consequently, under zero projection noise and with $K = d$, the entire population EM–PLS recursion coincides with the population counterpart of the EM algorithm introduced in Section 1.4.

Therefore, the convergence results established there apply directly. In particular, every accumulation point of the population EM–PLS recursion

$$\{\theta_{\text{pop}}^{(t)}\}_{t \in \mathbb{N}}$$

is a stationary point of the population observed-data log-likelihood $\ell(\theta)$.

Moreover, if the set of stationary points

$$\mathcal{S}(\ell) = \{\theta \in \Theta : \nabla \ell(\theta) = 0\}$$

is a singleton $\{\theta^*\}$, then the population EM–PLS recursion converges to θ^* . Under correct model specification, this limit coincides with the true data-generating parameter.

2.5 Sample EM–PLS algorithm

Algorithm 1 describes the EM–PLS procedure applied to a finite sample of m independent groups. While the population formulation is written at the level of a single generic observation $(\mathbf{X}, \mathbf{Y}, Z, \mathbf{U})$, the sample is generated through a grouped linear mixed model.

In contrast to the population recursion studied in the previous subsection, the sample EM–PLS algorithm replaces population expectations by empirical group averages and uses sample conditional moments $m_i^{(t)}$ and $C_i^{(t)}$ computed from the observed data. As a consequence, the sample update map

$$T_m : \Theta \rightarrow \Theta, \quad T_m(\theta^{(t)}) := \theta^{(t+1)},$$

constitutes a random perturbation of the deterministic population update map

$$T_{\text{pop}} : \Theta \rightarrow \Theta, \quad T_{\text{pop}}(\theta_{\text{pop}}^{(t)}) := \theta_{\text{pop}}^{(t+1)}.$$

The goal of this subsection is to formalize the relationship between the sample and population EM–PLS recursions and to establish that, for any fixed iteration index t , the sample iterates track their population counterparts as $m \rightarrow \infty$. The analysis is carried out under the group-level asymptotic regime introduced earlier, where the number of groups m grows while the within-group sample sizes n_i remain uniformly bounded. Under this regime, consistency of the sample EM–PLS algorithm follows from a combination of (i) law-of-large-numbers arguments for the groupwise quantities entering the E- and CM-steps, (ii) continuity of the population and sample update maps, and (iii) stability of the Krylov/PLS subspace under small perturbations of the normal-equation objects.

The assumptions stated below are grouped according to their role in the analysis. Assumptions B1–B2 concern the observed sample and numerical feasibility of the algorithm, while Assumptions C2–C1 provide the links between the sample-level quantities and their population analogues that are required to control the deviation of the sample recursion from the population recursion.

Assumptions D1–D3 are the regularity assumptions that allow to relax the assumptions C2 and C3.

The technical lemmas used in the main result are proven in Appendix A. The used assumptions are:

B1 (Observed sample). We observe m independent groups $\{(y_i, X_i, Z_i)\}_{i=1}^m$, with within-group lengths n_i uniformly bounded. We assume that the matrices Z_i have uniformly bounded entries and fixed low dimension (for instance, a random intercept/slope structure). Writing the rows of X_i as

$$X_i = \begin{pmatrix} x_{i1}^\top \\ \vdots \\ x_{in_i}^\top \end{pmatrix},$$

we assume that the row vectors x_{ij} are i.i.d. draws from a generic design vector x , and that x has finite moments of second order.

B2 (Feasibility and numerical stability along the algorithm path). Let $t \in \mathbb{N}$ be arbitrary. If the parameter update produces either $D^{(t)}$ or $\sigma^{2(t)}$ outside a prescribed interval due to finite-sample or numerical effects, it is post-processed by a deterministic repair operator $\mathcal{R}$.

For $D^{(t)}$, define the symmetrization

$$\Pi(\tilde{M}) := \frac{1}{2}(\tilde{M} + \tilde{M}^\top).$$

Let $\Pi(\tilde{M}) = U\Lambda U^\top$ be its eigendecomposition, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_q)$. For fixed constants $0 < \lambda_{\min} < \lambda_{\max} < \infty$, define

$$\mathcal{R}_D(\tilde{M}) := U \text{diag}(\min\{\lambda_{\max}, \max\{\lambda_i, \lambda_{\min}\}\}) U^\top.$$

For the noise variance, define the scalar repair operator

$$\mathcal{R}_\sigma(\tilde{\sigma}^2) := \min\{\lambda_{\max}, \max\{\tilde{\sigma}^2, \lambda_{\min}\}\}.$$

This operation enforces positive definiteness of $D^{(t)}$ and ensures that both $D^{(t)}$ and $\sigma^{2(t)}$ remain in the prescribed compact parameter range.

C1 (Stability of the Krylov/PLS subspace under perturbations). Assume a non-degeneracy condition ensuring continuity of the Krylov space map: there exists $C > 0$ such that

$$d(\mathcal{K}_K(A, b), \mathcal{K}_K(A_{\text{pop}}, b_{\text{pop}})) \leq C(\|A - A_{\text{pop}}\| + \|b - b_{\text{pop}}\|),$$

where $d(\cdot, \cdot)$ denotes an operator-norm distance between the corresponding projection matrices.

We further assume that the K -dimensional Krylov space is non-degenerate and that this rank is preserved under sufficiently small perturbations.

D1 (Independent groups with bounded sizes). The groups $\{(y_i, X_i, Z_i)\}_{i \geq 1}$ are independent and satisfy $\sup_i n_i < \infty$.

D2 (Bounded random-effects design). The random-effects design has fixed dimension q and uniformly bounded entries: $\sup_i \|Z_i\|_{\max} < \infty$.

D3 (Moment bounds). The row vectors x_{ij} of the fixed-effects design are i.i.d. draws from a generic design vector x with finite second moments, and the data-generating model satisfies

$$\mathbb{E}\|u_i\|^2 < \infty, \quad \mathbb{E}\|\varepsilon_i\|^2 < \infty.$$

D4 (Subgaussian design). The population fixed-effects covariate vector $\mathbf{X} \in \mathbb{R}^p$ is subgaussian, i.e., there exists a constant $K \geq 1$ such that for all $v \in \mathbb{R}^p$,

$$\|\langle \mathbf{X}, v \rangle\|_{\psi_2} \leq K \|\langle \mathbf{X}, v \rangle\|_{L^2},$$

where for a scalar random variable W ,

$$\|W\|_{\psi_2} := \inf \left\{ c > 0 : \mathbb{E} \exp(W^2/c^2) \leq 2 \right\}.$$

Under these assumptions, the following result can be proven.

Theorem 2.1: Population consistency and pointwise tracking of EM-PLS

Consider the sample EM-PLS recursion $\{\theta^{(t)}\}_{t \geq 0}$ generated by Algorithm 1 from m independent groups, and its population counterpart $\{\theta_{\text{pop}}^{(t)}\}_{t \geq 0}$. Write $\theta = (\beta, D, \sigma^2) \in \Theta$ and let

$$V_i^{(t)} = Z_i D^{(t)} Z_i^\top + \sigma^{2(t)} I_{n_i}$$

for any fixed $t \in \mathbb{N}$.

Assume the assumptions of Section 2.1 as well as B1–B2, C1, and D1–D3 hold. Assume moreover that the population recursion is well-defined on Θ and that the set of stationary points of the population observed-data log-likelihood is a singleton:

$$\mathcal{S}(\ell) = \{\theta \in \Theta : \nabla \ell(\theta) = 0\} = \{\theta^*\},$$

where the population criterion is the expected per-group log-likelihood under the true parameter θ^* ,

$$\ell(\theta) := \mathbb{E}_{\theta^*}[\ell_{\text{pop}}(\theta)], \quad \ell_{\text{pop}}(\theta) := \log f_\theta(\mathbf{Y} \mid \mathbf{X}, Z),$$

and $f_\theta(\cdot \mid X_i, Z_i)$ is the Gaussian marginal density implied by the mixed model,

$$\mathbf{Y} \mid \mathbf{X}, Z \sim \mathcal{N}(\mathbf{X}\beta, V), \quad V = Z D Z^\top + \sigma^2 I.$$

Finally, assume the initialization is consistent:

$$\|\theta^{(0)} - \theta_{\text{pop}}^{(0)}\| \xrightarrow[m \rightarrow \infty]{p} 0,$$

and allow the sample recursion to apply the deterministic repair operator $\mathcal{R}$ in the variance-component update (as in Assumption B2).

Then the following statements hold.

- (i) (*Consistency of the population EM-PLS recursion*) The population EM-PLS recursion $\{\theta_{\text{pop}}^{(t)}\}_{t \geq 0}$ converges to the true parameter θ^* as $t \rightarrow \infty$, that is,

$$\theta_{\text{pop}}^{(t)} \longrightarrow \theta^*, \quad t \rightarrow \infty.$$

- (ii) (*Pointwise tracking of sample iterates*) For every fixed iteration index $t \in \mathbb{N}$,

$$\|\theta^{(t)} - \theta_{\text{pop}}^{(t)}\| \xrightarrow[m \rightarrow \infty]{p} 0.$$

That is, the sample EM-PLS iterates track their population counterparts pointwise in the iteration index.

Proof. **Proof of (i): Consistency of the population EM–PLS algorithm.**

This follows from the equivalence between the population EM–PLS and population EM recursions established in Section 2.4, together with the convergence result for the population EM algorithm.

Proof of (ii): Pointwise tracking of sample iterates. We can apply Lemma 3.1 to conclude for every fixed iteration index $t \in \mathbb{N}$ that

$$\|\theta^{(t)} - \theta_{\text{pop}}^{(t)}\| \xrightarrow{p, m \rightarrow \infty} 0.$$

That is, the sample EM–PLS iterates track their population counterparts pointwise in the iteration index. □

As a consequence of parts (i) and (ii) of Theorem 2.1, there exists a sequence of iteration indices $t_m \rightarrow \infty$ such that

$$\theta^{(t_m)} \xrightarrow{p, m \rightarrow \infty} \theta^*.$$

Fix $\varepsilon > 0$. By part (i), the population EM–PLS recursion satisfies

$$\theta_{\text{pop}}^{(t)} \longrightarrow \theta^* \quad \text{as } t \rightarrow \infty.$$

Hence there exists a finite iteration index $T = T(\varepsilon)$ such that

$$\|\theta_{\text{pop}}^{(T)} - \theta^*\| < \varepsilon.$$

By part (ii), for this fixed T ,

$$\theta^{(T)} \xrightarrow{p, m \rightarrow \infty} \theta_{\text{pop}}^{(T)}.$$

Combining the two displays and using the triangle inequality yields

$$\|\theta^{(T)} - \theta^*\| \leq \|\theta^{(T)} - \theta_{\text{pop}}^{(T)}\| + \|\theta_{\text{pop}}^{(T)} - \theta^*\| \xrightarrow{p, m \rightarrow \infty} \varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, choosing $t_m := T(\varepsilon_m)$ for any sequence $\varepsilon_m \downarrow 0$ produces a sequence $t_m \rightarrow \infty$ such that

$$\theta^{(t_m)} \xrightarrow{p, m \rightarrow \infty} \theta^*.$$

2.6 Simulation study

The simulation study is designed to empirically validate the finite-horizon behavior predicted by the theoretical analysis of the EM–PLS algorithm. In particular, for a fixed number of EM iterations t , the theory predicts that increasing the number of independent groups m improves the accuracy of the parameter estimates through averaging effects. The simulations therefore investigate how the estimation errors for the fixed effects β , the random-effects covariance D , and the noise variance σ^2 evolve across EM iterations as m increases. In addition, the study examines the behavior of the modified Q -function and the impact of dimension reduction via partial least squares in an ill-posed regime.

Data-generating process. Data are generated from a linear mixed-effects model with $m = 83$ independent groups (trajectories), each containing $n_i = 60$ observations, yielding a total sample size of $n = 60m$. The fixed-effects design matrix $X \in \mathbb{R}^{n \times p}$ is high-dimensional with $p = 3200$, rendering the estimation problem severely ill-posed.

The fixed-effects signal is assumed to lie in a low-dimensional latent subspace of dimension $q_{\text{lat}} = 25$. Specifically, the p covariates are split into a relevant and an irrelevant block of equal size $p/2$. An orthonormal basis $P \in \mathbb{R}^{(p/2) \times q_{\text{lat}}}$ is drawn uniformly at random, and latent scores $Q_{\text{lat}} \in \mathbb{R}^{n \times q_{\text{lat}}}$ are generated with independent standard normal entries and exponentially decaying variances. The signal component of the design is constructed as $X_{\text{signal}} = Q_{\text{lat}} P^\top$.

To control the degree of ill-posedness, additional *projection noise* is added in directions orthogonal to the signal subspace. Let $P_\perp$ denote an orthonormal basis of the orthogonal complement of $\text{span}(P)$. Independent Gaussian noise is projected onto this complement and scaled by a factor $\sigma_{\text{proj}} = 0.3$, yielding $X_{\text{proj}} = \sigma_{\text{proj}} Q_{\text{proj}} P_\perp^\top$. The observed relevant design is then $X_{\text{rel}} = X_{\text{signal}} + X_{\text{proj}}$. An additional block of $p/2$ irrelevant covariates is generated independently with matching marginal variance. The final design matrix is given by $X = (X_{\text{rel}}, X_{\text{irr}})$.

The true fixed-effects vector β^* lies entirely in the signal subspace, with $\beta^* = (P\alpha^*, 0)$ for $\alpha^* \in \mathbb{R}^{q_{\text{lat}}}$.

Random effects follow a random-intercept and random-slope structure. For group i , the random-effects design matrix is

$$Z_i = \begin{pmatrix} 1 & t_{i1} \\ \vdots & \vdots \\ 1 & t_{in_i} \end{pmatrix}, \quad t_{ij} = j/n_i,$$

and the random effects are generated as $u_i \sim \mathcal{N}(0, D_{\text{true}})$ independently across groups, where

$$D_{\text{true}} = \begin{pmatrix} 1 & 0.3 \\ 0.3 & 0.6 \end{pmatrix}.$$

Observation noise is Gaussian with variance $(\sigma^2)^* = 0.5$.

The response is generated according to

$$y_i = X_i \beta^* + Z_i u_i + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, (\sigma^2)^* I_{n_i}),$$

independently across groups. Both y and X are centered prior to model fitting.

Algorithms and simulation results. For each of 50 independent simulation runs (seeds), both the unconstrained EM algorithm and the EM-PLS algorithm are applied to the same generated dataset. In the unconstrained EM implementation, the fixed effects are updated via a generalized least-squares step; to ensure numerical stability in the ill-posed setting, a small computational ridge penalty of 10^{-8} is added to the normal equations. In EM-PLS, the fixed-effects M-step is restricted to a Krylov subspace generated by partial least squares, with the number of components fixed to $K = 25$. All other model components are updated identically across the two algorithms.

Each algorithm is run for a fixed number of EM iterations, and performance metrics are evaluated at every iteration and averaged across simulation runs.

As can be seen in Figure 3, both EM and EM-PLS achieve comparable prediction correlation at convergence, which gradually decreases with iteration index, reflecting the increasing generalization of the prediction. One can also see a faster convergence of EM-PLS compared to EM. At the same time, the ℓ_2 error of $\hat{\beta}^{(t)}$ decreases over iterations, indicating continued refinement of the fixed-effects estimate. Due to identifiability and collinearity in the design, the unconstrained EM algorithm

EM vs EM-PLS: Estimation and Prediction Performance




Figure 3: Average training-set and test-set prediction correlations of the marginal fixed-effects predictor $\hat{y} = X\hat{\beta}^{(t)}$ and ℓ_2 error of the parameters across EM iterations. The plot of ℓ_2 error of $\hat{\beta}^{(t)}$ is log-scaled in y-axis.

attains a greater ℓ_2 error in β , therefore EM-PLS yields more stable estimates by restricting updates to a low-dimensional Krylov subspace.

For the noise variance σ^2 , the ℓ_2 error increases over iterations for both methods before stabilizing. This indicates that, although the fixed-effects estimate improves, the allocation of variability between fixed effects, random effects, and residual noise changes across iterations. In mixed models with large p , these variance components are strongly coupled and only partially identifiable. As the fixed-effects estimate stabilizes, the model redistributes variance across components, which may lead to an increase in the estimation error of σ^2 even though prediction performance remains stable.

In addition, the effect of the number of PLS components on predictive performance is studied. For this simulation, $q_{lat} = 40$ was applied. As shown on Figure 4, both train and test marginal correlations improve rapidly for small values of K , followed by a clear saturation. Beyond approximately $K \approx 20$, additional components yield only marginal gains, even though the latent signal dimension is $q_{lat} = 40$. This behavior is consistent with the simulation design, where the latent score variances decay exponentially, so that most of the predictive signal is concentrated in the leading latent directions. Consequently, the relevant structure is captured within a substantially lower-dimensional Krylov subspace. This illustrates the dimension-adaptive nature of the PLS construction.

In addition, we investigate the effect of the number of groups m on the final performance of the two algorithms. For this experiment, $q_{lat} = 40$, $K = 40$, $n_i = 60 \ \forall \ i = 1, \dots, m$, $p = 3200$ were fixed and the number of groups varied over $m \in \{20, 50, 100, 200\}$. For each value of m , the results are averaged over 10 independent simulation runs. We then evaluate the errors in parameters and the predictive performance at the last iteration ($T = 8$).

As shown in Figure 5, the estimation accuracy of the variance components D and σ^2 improves




Figure 4: Train and test marginal correlation of EM-PLS in the simulation study as a function of the number of PLS components K , with latent dimension fixed at $q_{\text{lat}} = 40$.

substantially as m increases, which is consistent with the theory developed in Section 1.3. In contrast, the behavior of the fixed-effects estimator differs markedly between the two methods. While the ℓ_2 error of $\hat{\beta}$ for the unconstrained EM algorithm remains large and even increases for moderate values of m , the EM-PLS estimator remains stable across all regimes. This reflects the ill-posedness of the normal equations, since increasing m does not improve the conditioning of $\hat{\Sigma}_X$, and highlights the regularizing effect of restricting the update to a Krylov subspace.

Despite these differences in parameter estimation, both methods achieve very similar predictive performance, as measured by the marginal correlation on both training and test sets. The prediction quality improves with m , reflecting the increasing amount of information available for estimating the shared structure across groups. Overall, these results indicate that EM-PLS provides substantially more stable fixed-effects estimates without sacrificing predictive accuracy, particularly in the settings with large p .




Figure 5: Final estimation error and predictive performance of EM and EM-PLS as a function of the number of groups m . Shaded areas indicate variability across simulation runs.




Figure 6: Ion (purple) inside the gramicidin A channel.

2.7 Data Description and Exploratory Analysis

The empirical analysis in this thesis is based on molecular dynamics simulations of ion permeation through a biological ion channel, provided by apl. Prof. Dr. Bert de Groot (Max Planck Institute for Multidisciplinary Sciences). Each simulation records the trajectory of a single ion passing through the channel along the membrane axis, resulting in a total of $m = 83$ independent trajectories. A schematic illustration of the channel geometry and a representative ion configuration is shown in Figure 6.

Each trajectory consists of a sequence of observations recorded at consecutive time points. For every time point, a vector of molecular descriptors is extracted, yielding a design matrix with a number of covariates far exceeding the number of observations per trajectory. The response variable corresponds to the position of the ion along the membrane axis.

The fixed-effects design matrix $X \in \mathbb{R}^{7704 \times 1656}$ is high-dimensional and exhibits severe multicollinearity with the 2-norm condition number $\kappa(X) = 19326.2$. Figure 7 displays the eigenvalue spectrum of the empirical Gram matrix $X^\top X$; the spectrum decays rapidly over several orders of magnitude, with a small number of dominant eigenvalues followed by a long tail of near-zero eigenvalues. This indicates that the effective dimension of the design is substantially smaller than the nominal number of covariates, and that many directions in parameter space are only weakly identified. As a consequence, the associated generalized least-squares problem is highly ill-conditioned, rendering naive estimation of the fixed effects unstable and sensitive to noise.

The number of observations in each group differs due to varying ion permeation speeds. Figure 8 illustrates the distribution of group sizes.

Application of EM and EM-PLS algorithms to the dataset. We compare a classical EM implementation for linear mixed models, in which the fixed effects are updated via an unconstrained generalized least-squares (GLS) step (see Section 1.4), with a Krylov-restricted variant (EM-PLS, see Section 2.3), where the M-step for the fixed effects is constrained to a partial least squares (PLS) subspace. Throughout the analysis, the number of PLS components is fixed to $K = 25$. Empirically, the results are stable with respect to this choice, and no substantial differences are observed for $K \in \{20, \dots, 40\}$.

Both algorithms are implemented under the simplifying assumption that the random-effects covari-




Figure 7: First 100 eigenvalues of the empirical Gram matrix $X^\top X$ on a logarithmic scale. The rapidly decaying spectrum, spanning several orders of magnitude, indicates strong collinearity and a pronounced separation between a low-dimensional signal subspace and a large number of weakly identified directions.

ance matrix and the noise covariance are scalar multiples of the identity, that is,

$$D = \tau^2 I_q, \quad \Sigma = \sigma^2 I_{n_i},$$

with the same parameters shared across all trajectories. This specification is consistent with the population-level model introduced in Section 2.4 and isolates the effect of ill-posedness in the fixed-effects estimation.

To account for temporal dependence, the dataset is augmented with the first lag of the response, which leads to the exclusion of the first observation in each trajectory. As a result, the final design matrix has dimension 7621×3312 . To match the assumptions of the population analysis in Section 2.4, both the response vector y and the columns of the design matrix X are centered.

Initial parameters are chosen independently of the data, in accordance with Lemma 3.1. Specifically, the algorithms are initialized with

$$\beta^{(0)} = 0_p, \quad \tau^{2(0)} = 1, \quad \sigma^{2(0)} = 1.$$

The 83 ion trajectories are split into a training set (ions 1–47) and a test set (ions 48–83). Since the ions were randomly shuffled prior to dataset construction, this split does not introduce any systematic bias and may be viewed as an independent hold-out sample.

As a convergence criterion, we monitor the relative change in the parameter $\beta^{(t)}$. The algorithm is terminated once

$$\left| \frac{\beta^{(t+1)} - \beta^{(t)}}{\beta^{(t)}} \right| < 10^{-3},$$

or after a maximum of 50 iterations.

- **Unconstrained EM algorithm.** Due to the severe ill-conditioning of the empirical Gram matrix $X^\top X$, the GLS update for β is numerically unstable. To ensure invertibility, a small




Figure 8: Bar plot of group sample size n_i of group $i \in 1, \dots, 83$. The minimum number of observations per group is $\min_i n_i = 19$, while the maximum is $\max_i n_i = 320$.

Table 1: Marginal prediction performance of EM and EM-PLS. Correlation and RMSE are reported for the training data (fit) and the test data (prediction), based on the fixed-effects component $\hat{y} = X\hat{\beta}$.

Method	Dataset	Correlation	RMSE
EM-PLS	Train	0.737	0.516
EM-PLS	Test	0.539	0.568
EM	Train	0.752	0.476
EM	Test	0.382	0.685

ridge penalty of magnitude 10^{-8} is added to the normal equations. The algorithm doesn't converge even after $T = 50$ iterations, but the parameters' estimates and evaluation metrics eventually stabilize. Figure 9 and Figure 10 display the evolution of correlation and RMSE for the marginal fit on the training data and the marginal prediction on the test data.

- **EM-PLS algorithm.** The EM-PLS algorithm converges after $T = 25$ iterations. Figures 9 and 10 show that EM-PLS exhibits a markedly more stable evolution of both correlation and RMSE, with substantially improved predictive performance on the test data.

The final performance metrics of both algorithms at convergence are summarized in Table 1.

The improvement of the fit along the iterations of both algorithms can be seen of Figures 9-10.

On the present dataset, the unconstrained EM algorithm exhibits clear signs of overfitting, achieving a better fit on the training data while generalizing poorly to held-out observations. In contrast, EM-




Figure 9: Evolution of marginal correlation over EM iterations for the training data (left) and the test data (right). The EM-PLS algorithm converges in few iterations and attains a worse fit on the training data, but exhibits superior generalization.




Figure 10: Evolution of marginal RMSE over EM iterations for the training data (left) and the test data (right). While the unconstrained EM algorithm gradually minimizes the training error, this is accompanied by the decrease in the predictive accuracy. The EM-PLS algorithm achieves higher stability and consistently lower prediction error, reflecting the regularizing effect of the Krylov subspace restriction.

PLS yields substantially more stable estimates and improved predictive performance, highlighting the regularizing effect of the Krylov subspace restriction in ill-posed settings.

Figure 11 displays observed versus fitted values obtained from the EM-PLS estimator. For the training data (red curve), predictions are based on the conditional mixed-model fit,

$$\hat{y}_i^{\text{train}} = X_i \beta^{(t)} + Z_i u_i^{(t)},$$

where $u_i^{(t)}$ denotes the empirical best linear unbiased predictor (BLUP) of the random effects. Under the Gaussian mixed model, the BLUP has the closed-form expression

$$u_i^{(t)} = D_i^{(t)} Z_i^\top \left(Z_i D_i^{(t)} Z_i^\top + \Sigma_i^{(t)} \right)^{-1} (y_i - X_i \beta^{(t)}).$$

This corresponds to the conditional expectation $u_i^{(t)} = \mathbb{E}(u_i \mid y_i; \theta^{(t)})$.

For the test data (blue curve), predictions are obtained marginally,

$$\hat{y}_i^{\text{test}} = X_i \beta^{(t)},$$

since the corresponding random effects are not observed and no BLUP can be computed. Consequently, the training fit reflects both fixed and random effects, while the test fit reflects only the fixed-effects component.

The difference between the red and blue curves illustrates the substantial contribution of the random effects to the conditional fit. On the training data, the BLUP adjustment $Z_i u_i^{(t)}$ captures group-specific deviations in intercept and slope, allowing the fitted values to track the observed responses more closely. In contrast, the marginal prediction on the test set is necessarily smoother and exhibits reduced variability, as it relies solely on the population-level fixed-effects structure. This distinction highlights the role of the random effects in modeling within-group heterogeneity and explains the improved alignment of the conditional training fit with the observed data.

Figure 12 displays the contribution of the estimated random effects to the conditional (training) predictions. Specifically, the plotted quantity corresponds to

$$Z_i u_i^{(t)},$$

which represents the random-effect component in the BLUP fit $\hat{y}_i = X_i \beta^{(t)} + Z_i u_i^{(t)}$.

In the present model, the random effects consist of a group-specific intercept and slope. Writing

$$u_i^{(t)} = \begin{pmatrix} u_{i,0}^{(t)} \\ u_{i,1}^{(t)} \end{pmatrix}, \quad Z_i = \begin{pmatrix} 1 & t_{i1} \\ \vdots & \vdots \\ 1 & t_{in_i} \end{pmatrix},$$

the random-effect contribution takes the form

$$Z_i u_i^{(t)} = u_{i,0}^{(t)} + t_{ij} u_{i,1}^{(t)}, \quad j = 1, \dots, n_i,$$

that is, a group-specific affine adjustment consisting of a vertical shift (intercept) and a linear trend (slope). The variability visible in the figure reflects the heterogeneity across groups captured by these two random-effect components.

Finally, Figure 13 shows the influence of the number of PLS components $K \in \{1, \dots, 40\}$ on predictive performance for the real dataset. Both train and test marginal correlations increase substantially for small values of K , followed by a gradual saturation as K becomes larger. Compared




Figure 11: Observed versus fitted values for EM-PLS. The red curve (training data) shows the conditional fit $X_i\beta^{(t)} + Z_i u_i^{(t)}$ using BLUP random effects, while the blue curve (test data) shows the marginal prediction $X_i\beta^{(t)}$.




Figure 12: Contribution of the estimated random effects $Z_i u_i^{(t)}$ to the training BLUP predictions. The model includes a random intercept and random slope. Gray vertical lines depict the first observation in each group.




Figure 13: Train and test marginal correlation of EM-PLS on the real dataset as a function of the number of PLS components K .

to the simulation scenario, the curves exhibit noticeably higher variability, which is expected due to sampling noise and the absence of a controlled latent structure. Nevertheless, the overall pattern remains qualitatively similar: most predictive gains are achieved within the first moderate number of components, while additional components yield only marginal improvements. This suggests that, also in the real-data setting, the effective predictive dimension is considerably smaller than the maximal number of components considered.

2.8 Summary

Section 2 develops an EM regularization strategy based on the PLS/Krylov framework for ill-posed linear mixed models. The section begins by introducing a latent-factor linear mixed model in which the fixed-effects signal is supported on a low-dimensional subspace, while random effects and within-group error covariance capture group-specific dependence. This structural assumption motivates an EM-type procedure in which the fixed-effects update is stabilized by restricting it to a K -dimensional Krylov subspace, yielding the EM-PLS algorithm. A clear distinction is made between the *sample* EM-PLS algorithm, constructed from empirical group averages, and its *population* analogue, obtained by replacing these averages with expectations with respect to a generic group. This population-sample separation provides a natural theoretical benchmark and forms the basis for subsequent consistency arguments.

A central methodological issue addressed in this section is that the PLS restriction breaks the classical EM property that each M-step globally maximizes the Q -function. As a consequence, monotonicity of the observed-data log-likelihood is no longer guaranteed at the sample level. To place EM-PLS within a principled convergence framework, the section reviews existing results on regularized EM algorithms and identifies conditions under which population-level monotonicity and stability can be recovered. In particular, under zero projection noise and when the Krylov dimension K is at least as large as the latent signal dimension, the population PLS update coincides with the oracle latent least-squares solution. This restores the effective maximization property required for monotonicity at the population level.

Building on these insights, the population EM-PLS recursion is shown to converge to the true parameter. The second theoretical component links the sample EM-PLS iterates to their population counterparts. Section 2 formulates structural and regularity assumptions ensuring well-posed covariance updates, continuity of the EM mappings, and uniqueness of the stationary point of the population observed-data log-likelihood. Using laws of large numbers for groupwise normal-equation objects and covariance updates, together with continuity arguments, a finite-horizon tracking result is established: for any fixed iteration index t , the sample iterates converge in probability to the population iterates as the number of groups $m \rightarrow \infty$. Finally, the theoretical findings are complemented by simulation studies and an application to real data, illustrating the predictive performance and numerical stability of EM-PLS in ill-posed settings with large p .

3 Discussion

The theoretical analysis of EM-PLS is carried out under a number of simplifying assumptions that warrant discussion.

First, the population results are derived under a zero projection noise assumption for the latent factor representation of the relevant component of the design. While this assumption is unlikely to hold exactly in practice, it serves as a tractable idealization that isolates the interaction between the Krylov restriction and the EM recursion. A natural relaxation is to allow for sufficiently small projection noise, in the sense that the population iterate $\beta_{\text{pop}}^{(t)}$ remains within a neighborhood of the minimum-norm least-squares solution. If this neighborhood is small enough to preserve local monotonicity of the population Q -function, the qualitative behavior of the population recursion is expected to remain unchanged. Empirically, the simulation results in Section 2.6 indicate that the sample EM-PLS algorithm continues to converge even in the presence of nonzero projection noise, suggesting that the zero-noise assumption is not critical for practical performance.

Second, the convergence analysis emphasizes that the accuracy of the sample iterate $\theta^{(t+1)}$ relative to its population counterpart $\theta_{\text{pop}}^{(t+1)}$ depends primarily on the closeness of $\theta^{(t)}$ to $\theta_{\text{pop}}^{(t)}$. This reflects the standard perturbation-based viewpoint for EM-type algorithms. At the same time, it is reasonable to expect that improved fixed-effects estimates obtained in an early PLS-regularized M-step lead to more accurate conditional moments in the subsequent E-step, which in turn yield improved variance-component updates. This iterative refinement mechanism provides an intuitive justification for performing multiple EM-PLS iterations even when the number of independent groups is limited.

Finally, while the asymptotic guarantees are derived under the regime $m \rightarrow \infty$, the main contribution of the present framework lies in its structural perspective rather than in a direct replacement of classical EM for linear mixed models. The EM-PLS formulation provides a principled way to combine dimension reduction via Krylov subspaces with latent-variable inference. In more complex models, such as those arising in high-dimensional biological or protein-structure applications where the relationship between predictors and response may not be strictly linear, developing an unconstrained EM algorithm from scratch can be challenging. In such settings, extending an EM framework with a PLS-type regularization may offer a more accessible path for both algorithm design and theoretical analysis.

References

- Bazzoli, C., Lambert-Lacroix, S., and Martinez, M.-J. (2023). Partial least square based approaches for high-dimensional linear mixed models. *Statistical Methods & Applications*.
- Bondell, H. D., Krishna, A., and Ghosh, S. K. (2010). Joint variable selection for fixed and random effects in linear mixed-effects models. *Biometrics*, 66(4):1069–1077.
- Boyles, R. A. (1983). On the convergence of the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 45(1):47–50.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38.
- Finocchio, G. and Krivobokova, T. (2025). An extended latent factor framework for ill-posed generalised linear regression. *arXiv preprint arXiv:2307.08377*.
- Horn, R. A. and Johnson, C. R. (2013). *Matrix Analysis*. Cambridge University Press, Cambridge, 2 edition.
- Laird, N. M. (1982). Computation of variance components using the EM algorithm. *Journal of Statistical Computation and Simulation*, 14(4):295–303.
- McLachlan, G. J. and Krishnan, T. (2008). *The EM Algorithm and Extensions*. Wiley, 2 edition.
- Meng, X.-L. and van Dyk, D. A. (1997). The EM algorithm: An old folk-song sung to a fast new tune. *Journal of the Royal Statistical Society. Series B (Methodological)*, 59(3):511–567.
- Vershynin, R. (2018). *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Wei, M. (1989). The perturbation of consistent least squares problems. *Linear Algebra and its Applications*, 112:231–245.
- Wu, C. F. J. (1983). On the convergence properties of the EM algorithm. *The Annals of Statistics*, 11(1):95–103.
- Yi, X. and Caramanis, C. (2015). Regularized EM algorithms: A unified framework and statistical guarantees. *arXiv preprint arXiv:1511.08551*.

Acknowledgements

I would like to express my sincere gratitude to my colleagues at the Medical University of Vienna for their insightful discussions and valuable feedback throughout the course of this thesis. Their support and scientific exchange have greatly contributed to the development of this work.

I am especially grateful to my wife, Polina, and her family for their unwavering encouragement, patience, and continuous support during this time.

Appendix A

In Theorem 2.1 we use the result

$$\|\theta^{(t)} - \theta_{\text{pop}}^{(t)}\| \xrightarrow{p} 0.$$

To prove this result, we additionally introduce the following two assumptions.

C2 (Pointwise LLNs). For each fixed iteration index t , the following laws of large numbers hold:

$$\begin{aligned} \left\| \frac{1}{m} \sum_{i=1}^m \left(C_i^{(t)} + m_i^{(t)} m_i^{(t)\top} \right) - \mathbb{E}_{\theta^*} \left[C_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z) + m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z) m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z)^\top \right] \right\| &\xrightarrow{p} 0, \\ \left| \frac{1}{m} \sum_{i=1}^m R_i^{(t)} - \mathbb{E}_{\theta^*} \left[R_{i,\text{pop}}^{(t)} \right] \right| &\xrightarrow{p} 0, \\ \left\| A_m^{(t)} - A_{\text{pop}}^{(t)} \right\| + \left\| b_m^{(t)}(\theta) - b_{\text{pop}}^{(t)}(\theta) \right\| &\xrightarrow{p} 0, \end{aligned}$$

where

$$A_m^{(t)} := \frac{1}{n} X^\top X, \quad b_m^{(t)}(\theta) := \frac{1}{n} X^\top (y - Z m^{(t)}),$$

the sample residual criterion is

$$R_i^{(t)} := \frac{1}{\sigma_{2(t)}} \|y_i - X_i \beta^{(t+1)} - Z_i m_i^{(t)}\|_2^2 + \frac{1}{\sigma_{2(t)}} \text{tr}(Z_i C_i^{(t)} Z_i^\top),$$

and the corresponding population residual criterion is

$$R_{i,\text{pop}}^{(t)} := \frac{1}{\sigma_{\text{pop}}^{2(t)}} (\mathbf{Y} - \mathbf{X}^\top \beta_{\text{pop}}^{(t+1)} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z))^2 + \frac{1}{\sigma_{\text{pop}}^{2(t)}} Z^\top C_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z) Z.$$

Population expectations are taken with respect to the joint law of $(\mathbf{X}, \mathbf{Y}, Z)$, with conditional moments defined under the model $\mathbf{Y} \mid \mathbf{X}, Z$.

C3 (Continuity of EM-PLS update maps). For every fixed iteration index $t \in \mathbb{N}$, the following maps are continuous on Θ :

- (a) The population normal-equation vector $\theta \mapsto b_{\text{pop}}^{(t)}(\theta)$ and its sample analogue $\theta \mapsto b_m^{(t)}(\theta)$ are continuous for every m .
- (b) (*Variance-component update maps*) The population variance-component update map

$$(\theta, \beta) \mapsto G_{\text{pop}}(\theta, \beta),$$

and its sample analogue

$$(\theta, \beta) \mapsto G_m(\theta, \beta),$$

are continuous on $\Theta \times \mathbb{R}^p$ for every fixed m .

Later, we show that Assumptions C2–C3 can be replaced by Assumptions D1–D4.

Lemma 3.1: Finite-horizon tracking of sample EM-PLS iterates

Fix $t \in \mathbb{N}$. Let

$$\theta_{\text{pop}}^{(t+1)} = T_{\text{pop}}(\theta_{\text{pop}}^{(t)}) \quad \text{and} \quad \theta^{(t+1)} = T_m(\theta^{(t)})$$

denote the population and sample EM-PLS recursions, respectively.

Let the parameter be partitioned as $\theta = (\beta, D, \sigma^2)$. Write the corresponding population and sample updates as

$$\beta_{\text{pop}}^{(t+1)} = \Psi_K(A_{\text{pop}}^{(t)}, b_{\text{pop}}^{(t)}(\theta_{\text{pop}}^{(t)})), \quad \beta^{(t+1)} = \Psi_K(A_m^{(t)}, b_m^{(t)}(\theta^{(t)})),$$

and

$$(D_{\text{pop}}^{(t+1)}, \sigma_{\text{pop}}^{2(t+1)}) = G_{\text{pop}}(\theta_{\text{pop}}^{(t)}, \beta_{\text{pop}}^{(t+1)}), \quad (D^{(t+1)}, \sigma^{2(t+1)}) = G_m(\theta^{(t)}, \beta^{(t+1)}).$$

Assume C1–C3 hold and that

$$\|\theta^{(0)} - \theta_{\text{pop}}^{(0)}\| \xrightarrow{p} 0.$$

Then

$$\|\theta^{(t)} - \theta_{\text{pop}}^{(t)}\| \xrightarrow{p} 0.$$

Proof. For $t \geq 0$, define

$$e_t := \|\theta^{(t)} - \theta_{\text{pop}}^{(t)}\|, \quad e_{t,\beta} := \|\beta^{(t)} - \beta_{\text{pop}}^{(t)}\|, \quad e_{t,D,\sigma^2} := \|(D^{(t)}, \sigma^{2(t)}) - (D_{\text{pop}}^{(t)}, \sigma_{\text{pop}}^{2(t)})\|.$$

We prove by induction that for each fixed $t \in \mathbb{N}$,

$$e_t \xrightarrow{p} 0.$$

Base step. By assumption,

$$\|\theta^{(0)} - \theta_{\text{pop}}^{(0)}\| \xrightarrow{p} 0.$$

Induction step. Fix $t \geq 0$ and assume $e_t \xrightarrow{p} 0$.

(a) *β -component.* Recall that

$$\beta_{\text{pop}}^{(t+1)} = \Psi_K(A_{\text{pop}}^{(t)}, b_{\text{pop}}^{(t)}(\theta_{\text{pop}}^{(t)})), \quad \beta^{(t+1)} = \Psi_K(A_m^{(t)}, b_m^{(t)}(\theta^{(t)})),$$

where Ψ_K denotes the K -PLS update map.

To make the dependence of Ψ_K explicit, fix A and define the Krylov matrix

$$H_K(A, b) := [b, Ab, \dots, A^{K-1}b] \in \mathbb{R}^{p \times K}, \quad \mathcal{K}_K(A, b) = \text{range}(H_K(A, b)).$$

The K -PLS update is the minimizer of

$$\beta \mapsto \frac{1}{2} \beta^\top A \beta - b^\top \beta$$

over $\mathcal{K}_K(A, b)$. Writing $\beta = H_K(A, b)c$, this reduces to the finite-dimensional least-squares problem

$$c_K(A, b) = \arg \min_{c \in \mathbb{R}^K} \left\{ \frac{1}{2} c^\top H_K(A, b)^\top A H_K(A, b) c - c^\top H_K(A, b)^\top b \right\}, \quad \Psi_K(A, b) = H_K(A, b) c_K(A, b).$$

Under the standing nondegeneracy assumption that $\dim \mathcal{K}_K(A, b) = K$ and that this rank is preserved under small perturbations, the Gram matrix $H_K(A, b)^\top A H_K(A, b)$ is invertible.

Add and subtract $\Psi_K(A_m^{(t)}, b_m^{(t)}(\theta_{\text{pop}}^{(t)}))$ to obtain

$$\begin{aligned} \|\beta^{(t+1)} - \beta_{\text{pop}}^{(t+1)}\| &\leq \underbrace{\|\Psi_K(A_m^{(t)}, b_m^{(t)}(\theta^{(t)})) - \Psi_K(A_m^{(t)}, b_m^{(t)}(\theta_{\text{pop}}^{(t)}))\|}_{(I)} \\ &\quad + \underbrace{\|\Psi_K(A_m^{(t)}, b_m^{(t)}(\theta_{\text{pop}}^{(t)})) - \Psi_K(A_{\text{pop}}^{(t)}, b_{\text{pop}}^{(t)}(\theta_{\text{pop}}^{(t)}))\|}_{(II)}. \end{aligned}$$

For term (II), Assumption C2 yields

$$\|A_m^{(t)} - A_{\text{pop}}^{(t)}\| + \|b_m^{(t)}(\theta_{\text{pop}}^{(t)}) - b_{\text{pop}}^{(t)}(\theta_{\text{pop}}^{(t)})\| \xrightarrow{p} 0.$$

Therefore, by Lemma 3.2,

$$(II) \xrightarrow{p} 0.$$

For term (I), by the induction hypothesis $e_t \xrightarrow{p} 0$ and Assumption C3(a),

$$\|b_m^{(t)}(\theta^{(t)}) - b_m^{(t)}(\theta_{\text{pop}}^{(t)})\| \xrightarrow{p} 0.$$

Since $A_m^{(t)}$ is fixed in both arguments and the reduced least-squares representation of $\Psi_K(A, b)$ is continuous in b on the nondegenerate set, it follows that

$$(I) \xrightarrow{p} 0.$$

Combining the two bounds yields

$$e_{t+1, \beta} = \|\beta^{(t+1)} - \beta_{\text{pop}}^{(t+1)}\| \xrightarrow{p} 0.$$

(b) (D, σ^2) -component. Add and subtract $G_m(\theta_{\text{pop}}^{(t)}, \beta_{\text{pop}}^{(t+1)})$:

$$\begin{aligned} e_{t+1, D, \sigma^2} &= \|(D^{(t+1)}, \sigma^{2(t+1)}) - (D_{\text{pop}}^{(t+1)}, \sigma_{\text{pop}}^{2(t+1)})\| \\ &\leq \underbrace{\|G_m(\theta^{(t)}, \beta^{(t+1)}) - G_m(\theta_{\text{pop}}^{(t)}, \beta_{\text{pop}}^{(t+1)})\|}_{(III)} + \underbrace{\|G_m(\theta_{\text{pop}}^{(t)}, \beta_{\text{pop}}^{(t+1)}) - G_{\text{pop}}(\theta_{\text{pop}}^{(t)}, \beta_{\text{pop}}^{(t+1)})\|}_{(IV)}. \end{aligned}$$

Term (IV) depends on empirical group averages appearing in the variance-component updates and therefore converges to zero in probability by Assumption C2. Moreover, since

$$(\theta^{(t)}, \beta^{(t+1)}) \xrightarrow{p} (\theta_{\text{pop}}^{(t)}, \beta_{\text{pop}}^{(t+1)}),$$

and G_m is continuous by Assumption C3, the continuous mapping theorem yields

$$(III) \xrightarrow{p} 0.$$

Hence

$$e_{t+1, D, \sigma^2} \xrightarrow{p} 0.$$

Finally,

$$e_{t+1} \leq e_{t+1, \beta} + e_{t+1, D, \sigma^2} \xrightarrow{p} 0,$$

which completes the induction. $\square$

In finite samples, the empirical variance-component update $G_m(\theta^{(t)}, \beta^{(t+1)})$ may fall outside the prescribed parameter space. In this case, we apply the repair operator introduced in Assumption B2 and set

$$(D^{(t+1)}, \sigma^{2(t+1)}) = \mathcal{R}\left(G_m(\theta^{(t)}, \beta^{(t+1)})\right).$$

Since the population update $(D_{\text{pop}}^{(t+1)}, \sigma_{\text{pop}}^{2(t+1)})$ belongs to the interior of the parameter space and $\mathcal{R}$ is non-expansive with respect to the product norm, it follows that

$$\|\mathcal{R}(D^{(t+1)}, \sigma^{2(t+1)}) - (D_{\text{pop}}^{(t+1)}, \sigma_{\text{pop}}^{2(t+1)})\| \leq \|(D^{(t+1)}, \sigma^{2(t+1)}) - (D_{\text{pop}}^{(t+1)}, \sigma_{\text{pop}}^{2(t+1)})\| \xrightarrow{p} 0.$$

Therefore, applying $\mathcal{R}$ does not affect convergence in probability.

Verification of Assumptions C2–C3 and proof of the Lipschitz stability lemma

In this subsection, we show that the technical Assumptions C2 and C3 can be replaced by a collection of standard conditions on the group-level sample-generating mechanism.

We focus on the first LLN in Assumption C2. Fix an iteration index $t \in \mathbb{N}$ and consider the deterministic population iterate

$$\theta_{\text{pop}}^{(t)} = (\beta_{\text{pop}}^{(t)}, D_{\text{pop}}^{(t)}, \sigma_{\text{pop}}^{2(t)}).$$

For each group i , define the Gaussian LMM conditional moments computed at $\theta_{\text{pop}}^{(t)}$,

$$m_{i,\text{pop}}^{(t)} := \mathbb{E}_{\theta_{\text{pop}}^{(t)}}(u_i \mid y_i, X_i, Z_i), \quad C_{i,\text{pop}}^{(t)} := \text{Var}_{\theta_{\text{pop}}^{(t)}}(u_i \mid y_i, X_i, Z_i),$$

and set

$$S_{i,\text{pop}}^{(t)} := C_{i,\text{pop}}^{(t)} + m_{i,\text{pop}}^{(t)} m_{i,\text{pop}}^{(t)\top} \in \mathbb{R}^{q \times q}.$$

The claim is that

$$\left\| \frac{1}{m} \sum_{i=1}^m S_{i,\text{pop}}^{(t)} - \mathbb{E}_{\theta^*} [S_{i,\text{pop}}^{(t)}] \right\| \xrightarrow{p} 0, \quad (7)$$

where the expectation is taken with respect to the conditional distribution of $y_i \mid X_i, Z_i$, and $\|\cdot\|$ is any matrix norm on $\mathbb{R}^{q \times q}$ (equivalently, $\|\cdot\|_F$).

Verification of the first LLN in C2. Recall that the parameter space Θ is restricted by the uniform bounds

$$\lambda_{\min} I \preceq D \preceq \lambda_{\max} I, \quad \lambda_{\min} \leq \sigma^2 \leq \lambda_{\max},$$

for fixed constants $0 < \lambda_{\min} < \lambda_{\max} < \infty$. In particular, for every fixed iteration index t we have $\theta_{\text{pop}}^{(t)} \in \Theta$, so the corresponding covariance objects are uniformly well-conditioned.

Proof of (7). By (D1), the groups $\{(y_i, X_i, Z_i)\}_{i \geq 1}$ are independent. Since t is fixed and $\theta_{\text{pop}}^{(t)}$ is deterministic, the matrices

$$S_{i,\text{pop}}^{(t)} = C_{i,\text{pop}}^{(t)} + m_{i,\text{pop}}^{(t)} m_{i,\text{pop}}^{(t)\top}$$

form an independent sequence in i . Therefore, it suffices to show that $\mathbb{E}_{\theta^*} \|S_{i,\text{pop}}^{(t)}\|_F < \infty$ uniformly in i , and then apply a weak law of large numbers entrywise.

Finite first moment of $S_{i,\text{pop}}^{(t)}$. Write

$$\theta_{\text{pop}}^{(t)} = (\beta_{\text{pop}}^{(t)}, D_{\text{pop}}^{(t)}, \sigma_{\text{pop}}^{2(t)})$$

and set

$$V_{i,\text{pop}}^{(t)} := Z_i D_{\text{pop}}^{(t)} Z_i^\top + \sigma_{\text{pop}}^{2(t)} I_{n_i}.$$

Since $\theta_{\text{pop}}^{(t)} \in \Theta$, we have

$$\|D_{\text{pop}}^{(t)}\|_{\text{op}} \leq \lambda_{\max}, \quad \lambda_{\min} \leq \sigma_{\text{pop}}^{2(t)} \leq \lambda_{\max}.$$

Moreover, for any vector $v \in \mathbb{R}^{n_i}$,

$$v^\top V_{i,\text{pop}}^{(t)} v = v^\top Z_i D_{\text{pop}}^{(t)} Z_i^\top v + \sigma_{\text{pop}}^{2(t)} \|v\|^2 \geq \sigma_{\text{pop}}^{2(t)} \|v\|^2 \geq \lambda_{\min} \|v\|^2,$$

so $V_{i,\text{pop}}^{(t)} \succeq \lambda_{\min} I_{n_i}$ and hence

$$\|(V_{i,\text{pop}}^{(t)})^{-1}\|_{\text{op}} \leq \lambda_{\min}^{-1} \quad \text{uniformly in } i.$$

For the Gaussian linear mixed model, the conditional moments admit the standard forms

$$m_{i,\text{pop}}^{(t)} = D_{\text{pop}}^{(t)} Z_i^\top (V_{i,\text{pop}}^{(t)})^{-1} r_{i,\text{pop}}^{(t)}, \quad C_{i,\text{pop}}^{(t)} = D_{\text{pop}}^{(t)} - D_{\text{pop}}^{(t)} Z_i^\top (V_{i,\text{pop}}^{(t)})^{-1} Z_i D_{\text{pop}}^{(t)},$$

where

$$r_{i,\text{pop}}^{(t)} := y_i - X_i \beta_{\text{pop}}^{(t)}.$$

Bounding $C_{i,\text{pop}}^{(t)}$. Using submultiplicativity, $\|D_{\text{pop}}^{(t)}\|_{\text{op}} \leq \lambda_{\max}$, $\|(V_{i,\text{pop}}^{(t)})^{-1}\|_{\text{op}} \leq \lambda_{\min}^{-1}$, and (D2), we obtain

$$\|C_{i,\text{pop}}^{(t)}\|_{\text{op}} \leq \|D_{\text{pop}}^{(t)}\|_{\text{op}} + \|D_{\text{pop}}^{(t)}\|_{\text{op}}^2 \|Z_i\|_{\text{op}}^2 \|(V_{i,\text{pop}}^{(t)})^{-1}\|_{\text{op}} \leq K_C$$

for some finite constant K_C independent of i . Since q is fixed,

$$\|C_{i,\text{pop}}^{(t)}\|_F \leq \sqrt{q} \|C_{i,\text{pop}}^{(t)}\|_{\text{op}},$$

hence

$$\mathbb{E}_{\theta^*} \|C_{i,\text{pop}}^{(t)}\|_F < \infty.$$

Bounding $m_{i,\text{pop}}^{(t)} m_{i,\text{pop}}^{(t)\top}$. Similarly,

$$\|m_{i,\text{pop}}^{(t)}\| \leq \|D_{\text{pop}}^{(t)}\|_{\text{op}} \|Z_i\|_{\text{op}} \|(V_{i,\text{pop}}^{(t)})^{-1}\|_{\text{op}} \|r_{i,\text{pop}}^{(t)}\| \leq K_m \|r_{i,\text{pop}}^{(t)}\|$$

for some finite constant K_m . Therefore,

$$\|m_{i,\text{pop}}^{(t)} m_{i,\text{pop}}^{(t)\top}\|_F = \|m_{i,\text{pop}}^{(t)}\|^2 \leq K_m^2 \|r_{i,\text{pop}}^{(t)}\|^2.$$

It remains to show that $\mathbb{E}_{\theta^*} \|r_{i,\text{pop}}^{(t)}\|^2 < \infty$. Under the data-generating model,

$$r_{i,\text{pop}}^{(t)} = y_i - X_i \beta_{\text{pop}}^{(t)} = X_i (\beta^* - \beta_{\text{pop}}^{(t)}) + Z_i u_i + \varepsilon_i.$$

Since n_i is uniformly bounded by (D1), the rows x_{ij} are i.i.d. with finite second moments by (D3), and $\beta_{\text{pop}}^{(t)}$ is fixed, it follows that

$$\mathbb{E}_{\theta^*} \|X_i (\beta^* - \beta_{\text{pop}}^{(t)})\|^2 < \infty.$$

Moreover, by (D2) and (D3),

$$\mathbb{E}_{\theta^*} \|Z_i u_i\|^2 < \infty, \quad \mathbb{E}_{\theta^*} \|\varepsilon_i\|^2 < \infty.$$

Hence

$$\mathbb{E}_{\theta^*} \|r_{i,\text{pop}}^{(t)}\|^2 < \infty,$$

and consequently

$$\mathbb{E}_{\theta^*} \|m_{i,\text{pop}}^{(t)} m_{i,\text{pop}}^{(t)\top}\|_F < \infty.$$

Combining the bounds yields

$$\mathbb{E}_{\theta^*} \|S_{i,\text{pop}}^{(t)}\|_F \leq \mathbb{E}_{\theta^*} \|C_{i,\text{pop}}^{(t)}\|_F + \mathbb{E}_{\theta^*} \|m_{i,\text{pop}}^{(t)} m_{i,\text{pop}}^{(t)\top}\|_F < \infty.$$

Applying LLN entrywise and concluding in norm. Let $\text{vec}(\cdot)$ denote vectorization. Since $\text{vec}(S_{i,\text{pop}}^{(t)}) \in \mathbb{R}^{q^2}$ is independent across i and uniformly integrable, the weak law of large numbers implies

$$\frac{1}{m} \sum_{i=1}^m \text{vec}(S_{i,\text{pop}}^{(t)}) \xrightarrow{p} \mathbb{E}_{\theta^*} [\text{vec}(S_{i,\text{pop}}^{(t)})].$$

Equivalently,

$$\frac{1}{m} \sum_{i=1}^m S_{i,\text{pop}}^{(t)} \xrightarrow{p} \mathbb{E}_{\theta^*} [S_{i,\text{pop}}^{(t)}].$$

Since q is fixed and the Frobenius norm is continuous, this implies (7). $\square$

Verification of the second LLN in C2. We now verify the *second* law of large numbers stated in Assumption C2. Fix an iteration index $t \in \mathbb{N}$. In the notation of Assumption C2, the sample residual criterion entering the σ^2 -update is

$$R_i^{(t)} := \frac{1}{\sigma_{2(t)}} \|y_i - X_i \beta^{(t+1)} - Z_i m_i^{(t)}\|_2^2 + \frac{1}{\sigma_{2(t)}} \text{tr}(Z_i C_i^{(t)} Z_i^\top),$$

and its population counterpart is

$$R_{i,\text{pop}}^{(t)} := \frac{1}{\sigma_{\text{pop}}^{2(t)}} \|y_i - X_i \beta_{\text{pop}}^{(t+1)} - Z_i m_{i,\text{pop}}^{(t)}\|_2^2 + \frac{1}{\sigma_{\text{pop}}^{2(t)}} \text{tr}(Z_i C_{i,\text{pop}}^{(t)} Z_i^\top),$$

where

$$m_{i,\text{pop}}^{(t)} := \mathbb{E}_{\theta_{\text{pop}}^{(t)}}(u_i \mid y_i, X_i, Z_i), \quad C_{i,\text{pop}}^{(t)} := \text{Var}_{\theta_{\text{pop}}^{(t)}}(u_i \mid y_i, X_i, Z_i),$$

and $\theta_{\text{pop}}^{(t)}$ is deterministic. The claim is that

$$\left| \frac{1}{m} \sum_{i=1}^m R_{i,\text{pop}}^{(t)} - \mathbb{E}_{\theta^*} [R_{i,\text{pop}}^{(t)}] \right| \xrightarrow{p} 0, \tag{8}$$

where the expectation is taken under the true group-level data-generating law.

Proof of (8). We reuse Assumptions (D1)–(D3). Define the population residual vector

$$r_{i,\text{pop}}^{(t)} := y_i - X_i \beta_{\text{pop}}^{(t+1)} - Z_i m_{i,\text{pop}}^{(t)}.$$

Then

$$R_{i,\text{pop}}^{(t)} = \frac{1}{\sigma_{\text{pop}}^{2(t)}} \|r_{i,\text{pop}}^{(t)}\|_2^2 + \frac{1}{\sigma_{\text{pop}}^{2(t)}} \text{tr}(Z_i C_{i,\text{pop}}^{(t)} Z_i^\top).$$

Step 1: independence. By (D1), the groups $\{(y_i, X_i, Z_i)\}_{i \geq 1}$ are independent. Since t is fixed and $\theta_{\text{pop}}^{(t)}$ and $\beta_{\text{pop}}^{(t+1)}$ are deterministic, each $R_{i,\text{pop}}^{(t)}$ is a measurable function of (y_i, X_i, Z_i) . Hence $\{R_{i,\text{pop}}^{(t)}\}_{i \geq 1}$ is an independent sequence. It therefore suffices to show that $\mathbb{E}_{\theta^*} |R_{i,\text{pop}}^{(t)}| < \infty$, uniformly in i , and then apply a weak law of large numbers.

Step 2: integrability. Since $\theta_{\text{pop}}^{(t)} \in \Theta$,

$$\lambda_{\min} \leq \sigma_{\text{pop}}^{2(t)} \leq \lambda_{\max}, \quad \|D_{\text{pop}}^{(t)}\|_{\text{op}} \leq \lambda_{\max}.$$

Set

$$V_{i,\text{pop}}^{(t)} := Z_i D_{\text{pop}}^{(t)} Z_i^\top + \sigma_{\text{pop}}^{2(t)} I_{n_i}.$$

Then $V_{i,\text{pop}}^{(t)} \succeq \lambda_{\min} I_{n_i}$, so

$$\|(V_{i,\text{pop}}^{(t)})^{-1}\|_{\text{op}} \leq \lambda_{\min}^{-1} \quad \text{uniformly in } i.$$

For the Gaussian linear mixed model,

$$m_{i,\text{pop}}^{(t)} = D_{\text{pop}}^{(t)} Z_i^\top (V_{i,\text{pop}}^{(t)})^{-1} (y_i - X_i \beta_{\text{pop}}^{(t)}),$$

and hence, by submultiplicativity together with (D1)–(D2),

$$\|m_{i,\text{pop}}^{(t)}\|_2 \leq K_m (\|y_i\|_2 + \|X_i\|_{\text{op}})$$

for some constant $K_m < \infty$ independent of i .

Therefore,

$$\|r_{i,\text{pop}}^{(t)}\|_2 \leq \|y_i\|_2 + \|X_i\|_{\text{op}} \|\beta_{\text{pop}}^{(t+1)}\|_2 + \|Z_i\|_{\text{op}} \|m_{i,\text{pop}}^{(t)}\|_2 \leq K_r (1 + \|y_i\|_2 + \|X_i\|_{\text{op}})$$

for some constant $K_r < \infty$. Since n_i is uniformly bounded by (D1), the rows x_{ij} are i.i.d. with finite second moments by (D3), and $\mathbb{E}\|u_i\|_2^2 < \infty$, $\mathbb{E}\|\varepsilon_i\|_2^2 < \infty$, it follows that

$$\mathbb{E}_{\theta^*} \|y_i\|_2^2 < \infty, \quad \mathbb{E}_{\theta^*} \|X_i\|_{\text{op}}^2 < \infty,$$

and hence

$$\mathbb{E}_{\theta^*} \|r_{i,\text{pop}}^{(t)}\|_2^2 < \infty.$$

Next, using $\|Z_i\|_{\text{op}} \leq K_Z$ from (D2),

$$|\text{tr}(Z_i C_{i,\text{pop}}^{(t)} Z_i^\top)| \leq n_i \|Z_i C_{i,\text{pop}}^{(t)} Z_i^\top\|_{\text{op}} \leq n_i \|Z_i\|_{\text{op}}^2 \|C_{i,\text{pop}}^{(t)}\|_{\text{op}}.$$

As in the proof of (7), $\|C_{i,\text{pop}}^{(t)}\|_{\text{op}}$ is uniformly bounded, while n_i is uniformly bounded by (D1). Therefore

$$|\text{tr}(Z_i C_{i,\text{pop}}^{(t)} Z_i^\top)| \leq K'_C$$

for some constant $K'_C < \infty$, uniformly in i .

Combining the bounds,

$$|R_{i,\text{pop}}^{(t)}| \leq \frac{1}{\sigma_{\text{pop}}^{2(t)}} \|r_{i,\text{pop}}^{(t)}\|_2^2 + \frac{1}{\sigma_{\text{pop}}^{2(t)}} |\text{tr}(Z_i C_{i,\text{pop}}^{(t)} Z_i^\top)|,$$

so $\mathbb{E}_{\theta^*} |R_{i,\text{pop}}^{(t)}| < \infty$.

Step 3: LLN. Since $\{R_{i,\text{pop}}^{(t)}\}_{i \geq 1}$ is independent and integrable, the weak law of large numbers yields

$$\frac{1}{m} \sum_{i=1}^m R_{i,\text{pop}}^{(t)} \xrightarrow{p} \mathbb{E}_{\theta^*} [R_{i,\text{pop}}^{(t)}],$$

which proves (8). $\square$

Verification of the third LLN in C2. Fix an iteration index $t \in \mathbb{N}$ and consider the population parameter

$$\theta_{\text{pop}}^{(t)} = (\beta_{\text{pop}}^{(t)}, D_{\text{pop}}^{(t)}, \sigma_{\text{pop}}^{2(t)}) \in \Theta.$$

For each group i , define the groupwise normal-equation contributions for the EM- β step by

$$A_i^{(t)} := \frac{1}{n_i} X_i^\top X_i \in \mathbb{R}^{p \times p}, \quad b_i^{(t)}(\theta) := \frac{1}{n_i} X_i^\top (y_i - Z_i m_i^{(t)}) \in \mathbb{R}^p,$$

where

$$m_i^{(t)} := \mathbb{E}_{\theta_{\text{pop}}^{(t)}}(u_i \mid y_i, X_i, Z_i)$$

denotes the conditional mean from the E-step.

Let

$$A_m^{(t)} := \frac{1}{n} X^\top X = \frac{1}{n} \sum_{i=1}^m X_i^\top X_i, \quad b_m^{(t)}(\theta) := \frac{1}{n} X^\top (y - Z m^{(t)}) = \frac{1}{n} \sum_{i=1}^m X_i^\top (y_i - Z_i m_i^{(t)}),$$

and define the population counterparts by

$$A_{\text{pop}}^{(t)} := \Sigma_{\mathbf{X}} = \mathbb{E}_{\theta^*}(\mathbf{X} \mathbf{X}^\top), \quad b_{\text{pop}}^{(t)}(\theta) := \mathbb{E}_{\theta^*}(\mathbf{X}(\mathbf{Y} - Z^\top m_{\text{pop}}^{(t)}(\mathbf{X}, \mathbf{Y}, Z))).$$

Under Assumptions (D1)–(D4), it holds that

$$\|A_m^{(t)} - A_{\text{pop}}^{(t)}\| + \|b_m^{(t)}(\theta) - b_{\text{pop}}^{(t)}(\theta)\| \xrightarrow{m \rightarrow \infty} 0.$$

Proof. Fix $t \in \mathbb{N}$.

Step 1: convergence of $A_m^{(t)}$. By Assumption (D4), the rows of X are i.i.d. subgaussian with covariance $\Sigma_{\mathbf{X}}$. Therefore, (Vershynin, 2018, Theorem 9.2.4) yields a non-asymptotic operator-norm bound for the empirical covariance matrix, and in particular

$$\left\| \frac{1}{n} X^\top X - \Sigma_{\mathbf{X}} \right\|_{\text{op}} \xrightarrow{p} 0.$$

Since $A_m^{(t)} = \frac{1}{n} X^\top X$ and $A_{\text{pop}}^{(t)} = \Sigma_{\mathbf{X}}$, it follows that

$$\|A_m^{(t)} - A_{\text{pop}}^{(t)}\|_{\text{op}} \xrightarrow{p} 0.$$

Step 2: convergence of $b_m^{(t)}(\theta)$. Define the groupwise block contributions

$$B_i^{(t)}(\theta) := X_i^\top (y_i - Z_i m_i^{(t)}(\theta)) \in \mathbb{R}^p.$$

Then

$$b_m^{(t)}(\theta) = \frac{1}{n} \sum_{i=1}^m B_i^{(t)}(\theta), \quad n = \sum_{i=1}^m n_i.$$

By Assumption (D1), the groups are independent, hence the vectors $\{B_i^{(t)}(\theta)\}_{i \geq 1}$ are independent across i .

We now verify a uniform second-moment bound. Since

$$\|B_i^{(t)}(\theta)\|_2 \leq \|X_i\|_{\text{op}} \|y_i - Z_i m_i^{(t)}(\theta)\|_2,$$

it suffices by Cauchy–Schwarz to control the second moments of $\|X_i\|_{\text{op}}$ and $\|y_i - Z_i m_i^{(t)}(\theta)\|_2$.

For the Gaussian linear mixed model,

$$m_i^{(t)} = D_{\text{pop}}^{(t)} Z_i^\top (V_i^{(t)})^{-1} (y_i - X_i \beta_{\text{pop}}^{(t)}), \quad V_i^{(t)} := Z_i D_{\text{pop}}^{(t)} Z_i^\top + \sigma_{\text{pop}}^{2(t)} I_{n_i}.$$

Because $\theta_{\text{pop}}^{(t)} \in \Theta$ and Assumption (D2) holds, there exists a constant $C_t < \infty$ such that

$$\|(V_i^{(t)})^{-1}\|_{\text{op}} \leq \lambda_{\min}^{-1}, \quad \|Z_i\|_{\text{op}} \leq K_Z, \quad \|D_{\text{pop}}^{(t)}\|_{\text{op}} \leq \lambda_{\max},$$

uniformly in i . Consequently,

$$\|Z_i m_i^{(t)}\|_2 \leq C_t (\|y_i\|_2 + \|X_i\|_{\text{op}}),$$

and therefore

$$\|y_i - Z_i m_i^{(t)}\|_2 \leq C'_t (\|y_i\|_2 + \|X_i\|_{\text{op}})$$

for some constant $C'_t < \infty$, uniformly in i .

Since $\sup_i n_i < \infty$, finite second moments of the rows x_{ij} imply

$$\sup_i \mathbb{E}_{\theta^*} \|X_i\|_{\text{op}}^2 < \infty,$$

and Assumption (D3) yields

$$\sup_i \mathbb{E}_{\theta^*} \|y_i\|_2^2 < \infty.$$

Hence

$$\sup_i \mathbb{E}_{\theta^*} \|y_i - Z_i m_i^{(t)}\|_2^2 < \infty,$$

and consequently

$$\sup_i \mathbb{E}_{\theta^*} \|B_i^{(t)}(\theta)\|_2^2 < \infty.$$

Now write

$$b_m^{(t)}(\theta) - \frac{1}{n} \sum_{i=1}^m \mathbb{E}_{\theta^*} B_i^{(t)}(\theta) = \frac{1}{n} \sum_{i=1}^m \left(B_i^{(t)}(\theta) - \mathbb{E}_{\theta^*} B_i^{(t)}(\theta) \right).$$

Since the centered block contributions are independent and have uniformly bounded second moments,

$$\mathbb{E}_{\theta^*} \left\| \frac{1}{m} \sum_{i=1}^m \left(B_i^{(t)}(\theta) - \mathbb{E}_{\theta^*} B_i^{(t)}(\theta) \right) \right\|_2^2 \leq \frac{C_t}{m}$$

for some constant $C_t < \infty$. Therefore

$$\frac{1}{m} \sum_{i=1}^m \left(B_i^{(t)}(\theta) - \mathbb{E}_{\theta^*} B_i^{(t)}(\theta) \right) \xrightarrow{p} 0.$$

Since $1 \leq n_i \leq \sup_i n_i < \infty$, we have $m/n \in [c, C]$ for constants $0 < c \leq C < \infty$, and thus

$$\frac{1}{n} \sum_{i=1}^m \left(B_i^{(t)}(\theta) - \mathbb{E}_{\theta^*} B_i^{(t)}(\theta) \right) \xrightarrow{p} 0.$$

Finally, by the definition of the row-level population target $b_{\text{pop}}^{(t)}(\theta)$, the deterministic term satisfies

$$\frac{1}{n} \sum_{i=1}^m \mathbb{E}_{\theta^*} B_i^{(t)}(\theta) \rightarrow b_{\text{pop}}^{(t)}(\theta).$$

Hence

$$\|b_m^{(t)}(\theta) - b_{\text{pop}}^{(t)}(\theta)\|_2 \xrightarrow{p} 0.$$

Combining the two parts gives

$$\|A_m^{(t)} - A_{\text{pop}}^{(t)}\| + \|b_m^{(t)}(\theta) - b_{\text{pop}}^{(t)}(\theta)\| \xrightarrow{p} 0.$$

□

Verification of C3 (continuity of update maps). (a) *Continuity of the normal-equation objects.*

For each fixed group i , $m_i^{(t)}(\theta)$ is an explicit continuous function of θ in the Gaussian LMM, so $\theta \mapsto b_i^{(t)}(\theta)$ is continuous as well. Hence, for every fixed m , the map

$$\theta \mapsto b_m^{(t)}(\theta)$$

is continuous as finite average of continuous functions.

Continuity of the population maps follows from dominated convergence. Indeed, if $\theta_n \rightarrow \theta$ in Θ , then

$$b_i^{(t)}(\theta_n) \rightarrow b_i^{(t)}(\theta)$$

almost surely by pointwise continuity, and the integrable envelopes established in the proof of Assumption C2 yield

$$\lim_{n \rightarrow \infty} \mathbb{E}_{\theta^*} [b_i^{(t)}(\theta_n)] = \mathbb{E}_{\theta^*} [b_i^{(t)}(\theta)].$$

Hence

$$b_{\text{pop}}^{(t)}(\theta_n) \rightarrow b_{\text{pop}}^{(t)}(\theta)$$

so the map is continuous on Θ .

(b) *Continuity of the variance-component update maps.* Write $G_{\text{pop}}(\theta, \beta)$ for the population variance-component update and $G_m(\theta, \beta)$ for its sample analogue. In the Gaussian LMM, both updates are constructed from the groupwise quantities $m_i^{(t)}(\theta)$, $C_i^{(t)}(\theta)$, and the residual criterion

$$R_i^{(t)}(\theta, \beta) = \frac{1}{\sigma^2} \|y_i - X_i \beta - Z_i m_i^{(t)}(\theta)\|_2^2 + \frac{1}{\sigma^2} \text{tr}(Z_i C_i^{(t)}(\theta) Z_i^\top),$$

all of which depend continuously on (θ, β) . Hence, for each fixed m , the sample map $G_m(\theta, \beta)$ is continuous.

For the population map, continuity follows again by dominated convergence, using the same integrability bounds as in the LLN proofs for $R_i^{(t)}$ and $C_i^{(t)} + m_i^{(t)} m_i^{(t)\top}$. Thus

$$(\theta, \beta) \mapsto G_{\text{pop}}(\theta, \beta)$$

is continuous on $\Theta \times \mathbb{R}^p$.

This verifies Assumption C3.

Stability of the PLS update via least-squares perturbation. We use the following auxiliary result based on Theorem 1.1 of Wei (1989).

Lemma 3.2: Stability of the PLS update

Let $A_{\text{pop}} \in \mathbb{R}^{p \times p}$ be symmetric and let $b_{\text{pop}}(\theta) \in \mathbb{R}^p$. For each m , let A_m and $b_m(\theta)$ satisfy

$$\|A_m - A_{\text{pop}}\|_{\text{op}} \xrightarrow{P} 0, \quad \|b_m(\theta) - b_{\text{pop}}(\theta)\|_2 \xrightarrow{P} 0.$$

Fix $K \in \mathbb{N}$ and define the Krylov matrices

$$H_K(A, b) := [b, Ab, \dots, A^{K-1}b] \in \mathbb{R}^{p \times K}.$$

Assume that the Krylov space is non-degenerate, i.e.

$$\text{rank } H_K(A_{\text{pop}}, b_{\text{pop}}(\theta)) = K,$$

and that this rank is preserved under sufficiently small perturbations.

Then

$$\|\Psi_K(A_m, b_m(\theta)) - \Psi_K(A_{\text{pop}}, b_{\text{pop}}(\theta))\|_2 \xrightarrow{P} 0.$$

Proof. Write the PLS update in Krylov coordinates. For A and b , define

$$H_K(A, b) = [b, Ab, \dots, A^{K-1}b] \in \mathbb{R}^{p \times K}, \quad \mathcal{K}_K(A, b) = \text{range}(H_K(A, b)).$$

Every $\beta \in \mathcal{K}_K(A, b)$ can therefore be written as

$$\beta = H_K(A, b)c, \quad c \in \mathbb{R}^K.$$

Hence the restricted quadratic minimization problem

$$\Psi_K(A, b) = \arg \min_{\beta \in \mathcal{K}_K(A, b)} \left\{ \frac{1}{2} \beta^\top A \beta - b^\top \beta \right\}$$

reduces to

$$c(A, b) = \arg \min_{c \in \mathbb{R}^K} \left\{ \frac{1}{2} c^\top H_K(A, b)^\top A H_K(A, b) c - c^\top H_K(A, b)^\top b \right\},$$

and

$$\Psi_K(A, b) = H_K(A, b)c(A, b).$$

Define the reduced matrices and right-hand sides

$$M_K(A, b) := H_K(A, b)^\top A H_K(A, b) \in \mathbb{R}^{K \times K}, \quad h_K(A, b) := H_K(A, b)^\top b \in \mathbb{R}^K.$$

Then $c(A, b)$ is the minimum-norm least-squares solution of the consistent linear system

$$M_K(A, b)c = h_K(A, b).$$

Now set

$$M_{K,m} := M_K(A_m, b_m(\theta)), \quad h_{K,m} := h_K(A_m, b_m(\theta)),$$

and

$$M_{K,\text{pop}} := M_K(A_{\text{pop}}, b_{\text{pop}}(\theta)), \quad h_{K,\text{pop}} := h_K(A_{\text{pop}}, b_{\text{pop}}(\theta)).$$

Since $H_K(A, b)$ depends polynomially on (A, b) , we have

$$\|H_K(A_m, b_m(\theta)) - H_K(A_{\text{pop}}, b_{\text{pop}}(\theta))\|_{\text{op}} \xrightarrow{P} 0.$$

Consequently,

$$\|M_{K,m} - M_{K,\text{pop}}\|_{\text{op}} \xrightarrow{P} 0, \quad \|h_{K,m} - h_{K,\text{pop}}\|_2 \xrightarrow{P} 0.$$

By assumption,

$$\text{rank } H_K(A_{\text{pop}}, b_{\text{pop}}(\theta)) = K,$$

and this rank is preserved under sufficiently small perturbations. Hence, for all large m , the reduced systems have the same rank K . Since the minimization problem above is equivalent to the normal equations

$$M_{K,m}c = h_{K,m}, \quad M_{K,\text{pop}}c = h_{K,\text{pop}},$$

and both systems are consistent, Theorem 1.1 of Wei (1989) applies and yields

$$\|c(A_m, b_m(\theta)) - c(A_{\text{pop}}, b_{\text{pop}}(\theta))\|_2 \xrightarrow{P} 0.$$

Finally,

$$\begin{aligned} & \|\Psi_K(A_m, b_m(\theta)) - \Psi_K(A_{\text{pop}}, b_{\text{pop}}(\theta))\|_2 \\ &= \|H_K(A_m, b_m(\theta))c(A_m, b_m(\theta)) - H_K(A_{\text{pop}}, b_{\text{pop}}(\theta))c(A_{\text{pop}}, b_{\text{pop}}(\theta))\|_2 \\ &\leq \|H_K(A_m, b_m(\theta)) - H_K(A_{\text{pop}}, b_{\text{pop}}(\theta))\|_{\text{op}} \|c(A_m, b_m(\theta))\|_2 \\ &\quad + \|H_K(A_{\text{pop}}, b_{\text{pop}}(\theta))\|_{\text{op}} \|c(A_m, b_m(\theta)) - c(A_{\text{pop}}, b_{\text{pop}}(\theta))\|_2. \end{aligned}$$

Since both terms converge to zero in probability and $\|c(A_m, b_m(\theta))\|_2$ is bounded in probability,

$$\|\Psi_K(A_m, b_m(\theta)) - \Psi_K(A_{\text{pop}}, b_{\text{pop}}(\theta))\|_2 \xrightarrow{P} 0.$$

This completes the proof. □

Declaration of Authorship

I hereby declare that I have written this Master's thesis independently and that I have not used any sources or aids other than those indicated. All passages that are taken verbatim or in substance from published or unpublished sources have been marked as such.

All figures, tables, and data used in this thesis are either my own work or are properly cited.

AI-based tools (ChatGPT) were used exclusively for linguistic support, including grammar correction and formulation of text. These tools were not used to generate scientific ideas, analyses, results, or interpretations. All scientific content of this thesis originates from my own work.

Abstract

Diese Arbeit untersucht einen regularisierten Expectation–Maximization-(EM)-Algorithmus für lineare gemischte Modelle, bei denen das klassische Maximum-Likelihood-Schätzverfahren aufgrund der Schlechtgestelltheit der Normalgleichungen für die Fixed Effects instabil wird. In solchen Situationen ist die Designmatrix stark kollinear, was zu einer schlechten Konditionierung der empirischen Kovarianzmatrix und zu unzuverlässigen Parameterschätzungen führt. Daher wird in dieser Arbeit ein EM–PLS-Algorithmus vorgeschlagen und analysiert, bei dem das M-Step-Update für die Fixed-Effects-Parameter durch ein auf Partial Kleinst Quadrate (PLS) basierendes regularisiertes Update ersetzt wird.

Die theoretische Analyse erfolgt durch eine Trennung in eine populationsbasierte Rekursion und deren Stichprobengegenstück. Auf Populationsebene wird gezeigt, dass der EM–PLS-Algorithmus unter einem latenten Faktormodell ein stabiles Verhalten aufweist, wobei das Signal in einem niedrigdimensionalen Unterraum liegt. In Abwesenheit von Projektionsrauschen rekonstruiert der durch das PLS-Verfahren erzeugte Krylov-Unterraum exakt den relevanten Signalraum. Daraus folgt, dass die kleinst-Quadrate-Lösung mit derjenigen Lösung übereinstimmt, die durch Minimierung der Zielfunktion über dem Krylov-Unterraum der Dimension K , erzeugt durch die Normalgleichungsgrößen, erhalten wird. Somit entspricht der populationsbasierte EM–PLS-Algorithmus genau dem populationsbasierten Gegenstück einer GEM-Implementierung für lineare gemischte Modelle. Daraus folgt insbesondere die Monotonie der Populations-Parameterupdates, welche die Konvergenz zum wahren Parameter unter geeigneten Regularitätsannahmen sicherstellt.

Um die populationsbasierte Analyse mit dem Stichprobenalgorithmus zu verknüpfen, wird ein Ergebnis zur Nachverfolgung über einen endlichen Iterationshorizont gezeigt. Es wird bewiesen, dass für jede feste EM-Iteration t die stichprobenbasierten EM–PLS-Iterate in Wahrscheinlichkeit gegen ihre populationsbasierten Gegenstücke konvergieren, wenn die Anzahl unabhängiger Gruppen wächst. Dieses Resultat basiert auf einer Kombination aus punktwisen Gesetzen der großen Zahlen, Stetigkeitseigenschaften der Update-Abbildungen sowie Stabilität der Krylov-Unterräume gegenüber Störungen.

Neben den theoretischen Ergebnissen wird eine umfangreiche Simulationsstudie durchgeführt, um die Leistungsfähigkeit des EM–PLS-Algorithmus im Vergleich zum klassischen EM-Algorithmus zu evaluieren. Die Experimente betrachten synthetische Designs mit latenter niedrig-rangiger Struktur und abklingenden Eigenwerten sowie unterschiedlichen Signal-Rausch-Verhältnissen. Die Ergebnisse zeigen, dass EM–PLS deutlich stabilere Schätzungen der Fixed Effects liefert, gemessen am L2-Fehler, während gleichzeitig eine mit dem unregularisierten EM-Algorithmus vergleichbare Vorhersagegüte erreicht wird. Insbesondere reduziert die durch den Krylov-Unterraum induzierte Regularisierung effektiv die Auswirkungen schlechter Konditionierung, ohne die Out-of-Sample-

Vorhersageleistung zu verschlechtern. Gleichzeitig zeigt sich, dass die Konvergenzgeschwindigkeit von EM-PLS diejenige des klassischen EM-Algorithmus übertrifft.

Der Einfluss der Anzahl der PLS-Komponenten sowie die Modellangemessenheit wird zudem anhand von Protein-Permeationsdaten untersucht, die vom Max-Planck-Institut für Multidisziplinäre Naturwissenschaften bereitgestellt wurden. Die Ergebnisse zeigen, dass die Vorhersageleistung bereits bei einer relativ kleinen Anzahl von PLS-Komponenten sättigt, was die Konzentration des Signals in einem niedrigdimensionalen latenten Raum widerspiegelt. Während der klassische EM-Algorithmus eine bessere Anpassung an die Trainingsdaten erreicht, zeigt er eine schlechtere Leistung auf Testdaten, was auf eine höhere Anfälligkeit für Overfitting im Vergleich zum EM-PLS-Ansatz hinweist.

Insgesamt zeigt die Arbeit, dass die Kombination des EM-Algorithmus mit einer PLS-basierten Dimensionsreduktion einen rechnerisch effizienten und theoretisch fundierten Ansatz für lineare gemischte Modelle darstellt. Die vorgeschlagene Methode erreicht ein günstiges Gleichgewicht zwischen statistischer Stabilität und Vorhersagegenauigkeit und stellt damit ein vielversprechendes Werkzeug für Erweiterungen auf komplexere Daten mit wahrscheinlich nicht-linearem Zusammenhang dar.